

Protecting Personal Medical and Health Information in the Era of Generative Artificial Intelligence

Shan Chen

Xinjiang University of Finance and Economics, Urumqi, Xinjiang, China

Abstract: Generative artificial intelligence is increasingly used in assisted consultation, record generation, image interpretation, and health advice, shifting personal medical and health information from closed processing inside medical institutions to large-scale circulation across multiple actors and scenarios. Such information is sensitive personal information under the Personal Information Protection Law, carrying personality interests that can rarely be restored once infringed. This Article asks how the existing normative system can accommodate the logic of large models in medical settings. The study adopts a doctrinal and typological method, tracing the full processing cycle of corpus aggregation, training, interactive use, and data feedback. It then reviews the applicable norms layer by layer, drawing on comparative experience under the European Union's General Data Protection Regulation and Artificial Intelligence Act. The findings show that rules resting on informed consent as the legal basis, anonymisation as the risk-blocking mechanism, and tort damages as the ex post remedy face structural difficulty: separate consent cannot be obtained from every subject, purpose limitation and minimum necessity fail to constrain cross-scenario reuse, and de-identified data remain exposed to re-identification. Protection must therefore shift from prior consent to risk-based governance of the entire process.

Keywords: Generative Artificial Intelligence; Medical and Health Information; Sensitive Personal Information; Informed Consent; Risk-Based Governance

1. Introduction

In practice, patients upload test reports and past medical histories into general-purpose large models to seek a treatment plan. Subjectively this is a private consultation; objectively, however,

the content may be retained by the service provider, used for further training, or even transferred across borders. The patient's rights to informed consent and to informational self-determination are left unprotected, and the provider's duties of disclosure, confidentiality, and security have not been clarified in response to the change of scenario. Large models can perform health education, assisted triage, medical record writing, and management decisions, and are expected to alleviate the unequal distribution of quality medical resources [1]; the medical community has formed a preliminary consensus on prudent application, safety review, algorithmic transparency, and ethical compliance [2]. At the same time, growing numbers of ordinary patients turn directly to general-purpose large models for medical advice, and risks have surfaced: models give erroneous advice and the content of conversations is collected and exploited [3]. Medicine is the field in which the data risks of generative artificial intelligence are most concentrated. Model training depends on massive multi-source data, and user data continue to accumulate during operation, so that the unlawful acquisition of corpora, the leakage of sensitive information, and algorithmic bias run through the entire process [4]. Medical and health information records disease history, the course of diagnosis and treatment, physiological characteristics, and even genetic and reproductive information; it is highly sensitive, lasts a lifetime, and cannot be altered. Once it is infringed, even deletion cannot restore the status quo ante. China has formed a normative system whose skeleton consists of the Book of Personality Rights of the Civil Code, the Personal Information Protection Law, the Data Security Law, the Cybersecurity Law, and the Basic Medical and Health Care and Health Promotion Law, which is supplemented by the Interim Measures for the Management of Generative Artificial Intelligence Services, the Regulations on the Management of Network Data Security, and the Measures for the Administration of Compliance Audits of Personal Information

Protection, and which is further elaborated by instruments such as the Measures for Labelling AI-Generated Synthetic Content. Whether legal norms can adapt to the technological logic of large models, and how they should properly apply in medical scenarios, is the question this Article seeks to answer.

2. How Generative AI Processes Personal Medical and Health Information

2.1 The Full-Cycle Structure of Processing

Generative artificial intelligence processes health information throughout the life cycle of the model. The first stage is pre-training: developers assemble corpora through web scraping, dataset purchase, and access to the information systems of medical institutions, and it is difficult to avoid mixing in medical record summaries, consultation records, and accounts of medical treatment that patients have shared online. The second is fine-tuning and feedback: the model is fine-tuned with labelled data and instructions and reinforced through human feedback so that it adapts to medical question answering. The third is interaction in operation: the user inputs symptoms, test results, and medication, the model answers on that basis, and the provider retains the conversation under default settings. The fourth is output feedback: the model may reproduce personal information contained in the corpus, and user inputs flow back into training data. This structure has three features. The purpose of processing can hardly be specified at the time of collection; it becomes known only at the time of use. Actors are dispersed along the chain of developer, deployer, and user. The model memorises the corpus in a statistical manner, the process is non-transparent, and those whose information is used can hardly know what is happening. Health information circulates among multiple actors and across scenarios: the same piece of information moves between consultation, research, public health, and commercial development, the risk varies with the scenario, and massive corpora are aggregated across institutions and borders and used repeatedly. Rules that look only to a single act of processing cannot cope with such a complex reality.

2.2 A Typology of Processing Risks

One commentator observes that in this setting the risks of processing sensitive information are characterised by concealed collection, dynamic

processing, automated inference, and diffusion, so that the traditional path built on separate consent and individual control is steadily losing its effect [5]. In the medical field these risks can be classified into five types.

First, unlawful collection and aggregation at the training stage. Health information scattered across cyberspace is frequently incorporated into corpora without permission; multiple fragments, once aggregated and assembled, are sufficient to reconstruct a complete health profile. Large-scale aggregation is itself a high-risk condition, and in serious cases may trigger the crime of infringing citizens' personal information [6].

Second, over-collection and cross-border transfer at the operation stage. Users tend to disclose their real conditions in order to obtain accurate advice, yet voluntary disclosure does not mean that they have been adequately informed [7]. For some offshore large models the servers and training activities are located abroad, and health information may leave the country without a security assessment.

Third, reproduction and re-identification at the output stage. The model may reproduce verbatim the text of medical records contained in its corpus, which constitutes training data. Even after de-identification, a person may still be re-identified through associative inference combined with external data, so that the expectation of anonymisation is defeated.

Fourth, the generation of sensitive information and discrimination at the inference stage. From neutral information it has obtained—consumption records, exercise data, linguistic expression—the model can infer disease propensity and generate 'inferred sensitive data' that the user never disclosed. One study divides health information into four layers: clinical diagnosis and treatment; biometric and genetic; physiological and metabolic; and derived and inferred. The last is produced by algorithmic aggregation and predictive modelling, and its black-box character and discriminatory effect make it the most concentrated site of privacy challenges [8]. Once such data are used in fields such as employment screening, discrimination results, and the disposal of the data and the distribution of their benefits also give rise to disputes over entitlement [9].

Fifth, the risk of harm caused by content distortion. The model may fabricate clinical guidelines or drug indications that do not exist; a user who seeks treatment or takes medication on

that basis endangers health and even life.

3. The Difficulties of Applying the Current Protection Framework

3.1 The Structure of the Protective Norms

The current legal norms can be divided into four layers. Article 111 of the Civil Code establishes the general rule of personal information protection; Article 1034 brings health information within personal information; and Article 1226 imposes a duty of confidentiality on medical personnel. Article 28 of the Personal Information Protection Law classifies medical and health information as sensitive personal information; Articles 29 and 30 require separate consent and special notification; and Article 55 requires a prior impact assessment. Article 7 of the Interim Measures requires lawful sources of training data, and Article 9 confirms the status of the service provider as a processor; Article 19 of the Regulations on the Management of Network Data Security strengthens the security management of training data.

3.2 The Validity Crisis of Informed Consent as a Legal Basis

Notice and consent is the primary legal basis established by Article 13(1) of the Personal Information Protection Law; for sensitive personal information such as medical and health information, Article 29 further requires separate consent. Consent takes effect under Article 14(1), and its application presupposes that the subject can be identified, the purpose determined, and the will formed freely. Large-model scenarios satisfy none of these three requirements.

First, the subject is not identifiable and consent cannot be obtained. The volume of the corpus is enormous, health information is scattered across medical records, forums, and reviews of medical treatment, and obtaining consent from each person is factually impossible. Second, the purpose is indeterminate and consent loses its point. Model training is exploratory, and a general privacy policy can hardly satisfy the requirement of Article 14 that consent be fully informed, voluntary, and explicit; party autonomy and health data privacy have themselves been listed among the foremost challenges [10]. Third, the will is not free and consent becomes a formality. Patients have almost no room to negotiate, and 'no consent means no service' reduces consent to a mechanical click. The doctor–patient relationship

is itself marked by information asymmetry, and a formal tick is not enough to carry genuine authorisation.

In response to these difficulties, some scholars propose supplementing consent with an information fiduciary duty [11]; one view expands fair use so as to leave room for training [12]; another study divides consent into four levels—strict, strong, weak, and not required—according to the degree of sensitivity, supplemented by dynamic consent [13]; and a further doctrine would impose fiduciary duties on processors and construe consent as authorisation grounded in reliance [14]. Each of these proposals, when applied to health information, must first pass review under the proportionality principle; the relaxed logic applicable to ordinary information may not be transplanted into health information by analogy.

3.3 Purpose Limitation versus Data Reuse

Article 6 of the Personal Information Protection Law requires that the purpose be explicit and reasonable, that the method of least impact be adopted, and that collection be confined to the minimum scope; where the purpose changes, consent must be obtained afresh. The tension is that the value of a model derives precisely from cross-purpose reuse: the same corpus may be used in turn for training, alignment, evaluation, and commercial development. Fixing a single purpose at the time of collection is neither economical nor consistent with the logic of the technology, while obtaining consent at each stage is factually impossible.

To ease the conflict, one study argues for 'weak protection' of already public personal information, narrowing elements such as the individual's express refusal and significant impact so as to leave room for training [15]. That view may have some force for ordinary information, but it cannot be applied to health information: patients disclose their conditions for the specific purpose of seeking medical advice, not as a blanket authorisation of secondary use. Once the logic of weak protection extends to health information, Article 29 is hollowed out.

3.4 The Failure of the Anonymisation Assumption and Re-Identification Risk

The Personal Information Protection Law excludes anonymised information from personal information, and processors often claim exemption on the ground that the data have been

de-identified. Large models unsettle that assumption. On the one hand, model-based correlation analysis lowers the threshold of re-identification: genomic data can identify a subject by the uniqueness of the genotype, and rare diseases, because cases are few, are easy to pin down. Even where each source has been anonymised, the combination of several sources may still permit re-identification through algorithmic correlation. On the other hand, the model may reproduce the original corpus through its memory mechanism, and anonymisation is insufficient to prevent the downstream leakage of personal health information. Requiring complete anonymisation across the board, however, would damage research value. A foreign study once found that the combined use of paroxetine and pravastatin could raise blood glucose; once the data were thoroughly anonymised the patients could not be contacted, it could not be determined whether they had been hyperglycaemic to begin with, and the conclusion could not be reached. The dilemma of anonymisation being at once too strict and ineffective persists.

3.5 Evidentiary Obstacles in Tort Liability and Rights Remedies

Large models form a structure of provider, deployer, and user, and health data circulate through successive layers of entrusted processing. Under Article 21, the entrusting party must agree the matters of processing and supervise the entrusted party; one study construes the supervisory duty as a hybrid of a third-party duty in administrative law and a duty of care in tort law, holding that upon discovering unlawful processing the entrusting party bears duties to prevent and to report it, with the formal boundaries of recording the entrustment and retaining the record for three years[16]. In practice, entrusting parties often replace supervision with a single sheet of paper, and once rights are harmed, any claim the rights-holder wishes to bring encounters obstacles.

The first is that the duty-holder and the form of liability are unclear. Article 9 of the Interim Measures establishes the service provider's status as processor, but where a medical institution deploys an open-source model on its own and trains it locally, how responsibility is to be shared with the provider is not clear, and the attribution of criminal liability is also disputed. In civil terms, those who jointly determine the purposes and means of processing are jointly and severally

liable under Article 20, while those who process independently must be distinguished under Article 1168 et seq. of the Civil Code according to the form of the tort.

The second is the difficulty of proving fault and causation. The algorithmic black box makes it hard for the rights-holder to prove which duty was breached and how causation is established. Article 69(1) adopts a presumption of fault, but it reverses only fault; the unlawful act, the damage, and causation still have to be proved by the rights-holder. Taking the state of the art as the objective standard of fault [17] is quite workable, but it has not been elevated into a rule of adjudication.

The third is the weak incentive to seek remedies. Moral damage from the leakage of health information is hard to quantify, the standard of compensation is uniform and the scales of adjudication inconsistent [18]; the rights-holder faces high costs and low returns, and the claim for damages lies unused. The personality-rights claim does not require fault, and under Article 997 one may demand cessation of the infringement and even an injunction, but it cannot fill the loss already incurred. Public interest litigation under Article 70 could supplement private litigation, yet medical cases are few and the forms of liability remain to be elaborated [19].

3.6 The Inadequate Supply of Rules Specific to Medical and Health Information

China's unified legislation adopts a one-size-fits-all model for sensitive information, and detailed rules for health and medical data are still lacking; the rules on medical records and on population health information respond inadequately to model training and to the treatment of data as a factor of production [20]. In comparative law, the European Union sets strict conditions for health data under the General Data Protection Regulation and, through the Artificial Intelligence Act, imposes duties of data governance, risk control, and human oversight on high-risk medical systems. China's requirements remain broadly framed, and ownership, attribution, and consent have not been systematised.

4. Pathways to Institutional Reform

In response to these difficulties, reform should follow a risk-based and whole-process approach: the consent mechanism should be configured differently according to scenario and risk level, and whole-process duties, procedural regulation,

and the allocation of responsibility should advance together, so that personality interests are protected without impeding the reasonable use of medical data.

4.1 Reconstructing Consent Rules by Scenario

Consent should not be abandoned wholesale; it should be allocated by stage and by risk level, which accords with the idea of tiered and dynamic consent. At the pre-training stage, corpora containing health information should in principle not be used merely because they are already public, and the weak protection approach should be excluded. At the operation stage, a rule of not used for training by default should be established: a conversation may be used for optimisation only after express opt-in, and separate consent must be obtained for training. In research settings where obtaining consent one by one is genuinely impossible, a menu-style itemised authorisation can replace package consent, allowing rights-holders to adjust and to withdraw at any time, with ethical review substituting for part of the prior consent. Sensitive data generated by inference should be treated in the same way as data actively provided, and Articles 29 and 30 applied to both.

4.2 Allocating Whole-Cycle Duties through Impact Assessment and Tiered Classification

The impact assessment required by Article 55 should be implemented across the entire life cycle of the model. Before training, the legality of the corpus should be verified, sensitive information filtered, and re-identification testing carried out. During training, priority should be given to de-identification and differential privacy, so that raw data can participate in training without leaving the medical institution, and in operation minimisation of retention and cross-border assessment should be practised; at the output end, medical advice should be prominently labelled and marked as no substitute for a physician's diagnosis, and managed by tiered classification across four levels using the training—fine-tuning—generation—feedback framework [21].

4.3 Procedural Regulation: High-Risk Access and Algorithmic Audit

Drawing on the risk-graded regulation of the European Union's Artificial Intelligence Act, applications involving diagnosis, treatment, prescription, triage, and psychological intervention can be classified as high-risk systems

and subjected to prior assessment, clinical validation, human oversight, and post-market monitoring, with lighter duties in low- and medium-risk scenarios. Before launch, the purpose of the algorithm should be reviewed and unfit data excluded. An independent third-party algorithmic audit system should also be established: medical guidelines can be issued under the Measures for the Administration of Compliance Audits of Personal Information Protection, and model memory testing and cross-border data flows should be made mandatory audit items.

4.4 Allocating Responsibility among Multiple Actors and Improving Remedies

The allocation of responsibility should identify the processor by the power to determine processing: the provider is responsible for the legality of the data and the safety of the output; a medical institution that deploys on its own is responsible for input, retention, and provision; those that determine jointly constitute joint processors and are jointly and severally liable under Article 20. Where the entrusting party neglects its supervisory duty under Article 21 and the entrusted party thereby commits a tort, one study argues for proportionate liability. Whether an omission and an act can be brought within Article 1172 of the Civil Code depends on whether each party's conduct was sufficient to cause the corresponding harm; it should not be treated as proportionate across the board. Records of public-law supervision can serve as *prima facie* evidence of fault in private law.

The principle of attribution must be handled with care. Whether a large model can be regarded as a product remains to be argued, so legislation should be the basis. Causation can follow a sequence of low-threshold *prima facie* showing, presumption by the court, and rebuttal by the processor: once the rights-holder proves the harm and the existence of the processing activity, the court presumes causation, and the processor who cannot rebut is not exempted; a processor who refuses to provide log records is additionally subject to the rule on obstruction of evidence.

As to remedies, preventive relief may take the form of a personality-rights injunction under Article 997, combined with a demand under Article 47 for the deletion of training data and parameters. Compensatory relief should activate public interest litigation under Article 70, bringing suit over large-scale unlawful collection

and export; punitive damages under Article 69(2) should apply to intentional violations, and moral damage should be assessed in tiers according to sensitivity. Criminal liability should be exercised with restraint, focusing on the illicit chain of unlawfully obtaining and selling health information.

4.5 A Governance Structure Combining Legal Regulation and Industry Supervision

Risk in this field runs through data, algorithms, and applications; no single instrument can cover it [22], and legal rules and industry supervision must work together. At the level of legal norms, the specialised norms on health and medical data should specify the boundaries of legality for training use, the tiered catalogue, and the cross-border rules, and should connect with the Basic Medical and Health Care and Health Promotion Law and the medical records regime. At the level of industry supervision, national standards, industry guidelines, and ethics committees should translate privacy-enhancing technologies and ethical review into compliance checklists, and medical institutions should plan protection and use in tandem [23]. As for the claim of fiduciary duty, its division into loyalty and care may serve as a doctrinal reference for raising the standard of conduct, but personal information lacks the independence of trust property, so direct analogy to the Trust Law remains obstructed, and the proposition that breach of the duty of loyalty entails no-fault liability also conflicts with Article 69; it should first be transformed through industry codes and medical ethics. Under the treatment of data as a factor of production, ownership should also be clarified, confirming both the patient's personality interests and the rights of the data source [24].

5. Conclusion

Generative artificial intelligence has not changed the position that medical and health information, as sensitive personal information, receives strict protection; what it has changed is the structure of information processing and the distribution of risk. The massive aggregation of corpora, the continuous shifting of purposes of use, and the growing capacity of models to memorise and to re-identify heighten the risk that individuals will be identified, and the traditional rules centred on notice and consent struggle to meet the problem. The legal response can neither stop at heavier penalties nor rest at ethical exhortation. It should

stand on a risk-based footing: restructure the rules of processing through scenario-specific consent and whole-process impact assessment, restructure the regulatory system through high-risk access and algorithmic audit, and restructure the remedial system through plural attribution and public interest litigation, striking a balance between safeguarding the patient's personality interests and releasing the public value of medical data.

References

- [1] Huatang Zeng, Xiatong Ke, Ping Xu, et al. Exploring the Application and Challenges of Generative Artificial Intelligence in Primary Healthcare Services Based on ChatGPT. *Chinese Journal of Hospital Administration*, 2023, 39(10):791-794.
- [2] Guanghui Xia, Yanlin Cao, Bingshu Chen, et al. Expert Consensus on the Application of Large Model Artificial Intelligence Technology in Medical Services. *China Health Law*, 2023, 31(5):124-126.
- [3] Wenjing Ding. "Assistant" or "Killer": A Risk Analysis of Consulting ChatGPT for Medical Advice. *Medicine & Philosophy*, 2023, 44(23):22-25.
- [4] Xin Zhang. Data Risks and Governance Pathways of Generative Artificial Intelligence. *Science of Law: Journal of Northwest University of Political Science and Law*, 2023, 41(5):42-54.
- [5] Xiaodong Tou. Risk and Control: On Personal Information Protection in the Application of Generative Artificial Intelligence. *Journal of Political Science and Law*, 2023(4):59-68.
- [6] Tong Ge and Zexin Zhou. On the Protection of Personal Medical and Health Information in China from the Perspective of the Personal Information Protection Law. *China Health Law*, 2024, 32(3):34-39.
- [7] Chuan Li. The Dilemma of Criminal Regulation and the Path of Attribution for Personal Information in the Context of Generative Artificial Intelligence. *Journal of East China University of Political Science and Law*, 2025, 28(6):22-34.
- [8] Ruibo Yue and Li Xue. The Legal Regulatory Path for Personal Health Information Governance from the Perspective of Smart Healthcare. *Journal of Jinzhou Medical University (Social Science Edition)*, 2026, 24(3):13-19.

- [9] Jinming Yang, Na Wang, Yexun Hu, et al. Preventing Legal Risks in Artificial Intelligence Medicine. *Journal of Sichuan University (Medical Sciences)*, 2025, 56(1):143-148.
- [10] Jiayou Shi and Jingyi Xu. Legal Challenges and Governance of the Application of Medical Artificial Intelligence. *Journal of Northwest University (Philosophy and Social Sciences Edition)*, 2024, 54(2):91-103.
- [11] Xiaoming Shou and Xianxin Cao. Risks, Logic and Norms of Personal Information Protection in Generative Artificial Intelligence Scenarios. *Journal of Kunming University of Science and Technology (Social Sciences Edition)*, 2024, 24(1):21-29.
- [12] Hongtao Hu. On the Fair Use of Personal Information by Generative Artificial Intelligence Service Providers. *Jiangsu Social Sciences*, 2024(2):185-192.
- [13] Lele Liu. Improving the Civil Law Protection Mechanism for Personal Information in the Context of Generative Artificial Intelligence. LL.M. Thesis. Shandong University of Political Science and Law, 2026.
- [14] Hui Deng and Hui Sun. Justification and Institutional Construction of the Fiduciary Duty of Processors of Personal Medical and Health Data. *Journal of East China University of Political Science and Law*, 2025, 28(1):76-90.
- [15] Ruobing Wang. Weak Protection of Already Public Personal Information in Generative Artificial Intelligence Scenarios. *Journal of Renmin University of China*, 2025, 39(6):75-87.
- [16] Mingzhi Fan and Yue Li. The Nature and Liability of Supervision over Entrusted Processing of Personal Information in the Context of Artificial Intelligence. *Journal of Chongqing University of Posts and Telecommunications (Social Science Edition)*, advance online publication, 2026-07-20. <https://link.cnki.net/urlid/50.1180.C.20260720.1140.002>.
- [17] Ruobing Wang. Determining the Fault of Service Providers in Generative Artificial Intelligence Infringement: Taking the State of the Art as the Standard. *Journal of Comparative Law*, 2023(5):20-33.
- [18] Zhao Wang. Legal Protection of Medical and Health Information in the Big Data Era. LL.M. Thesis. Northwest Normal University, 2025.
- [19] Qian Wang, Rui Sun, Man Xia. The Dilemma and Path Optimisation of Personal Information Protection for Patients in Smart Healthcare. *Medicine & Philosophy*, 2025, 46(19):62-67.
- [20] Kaixin Meng, Xiaona Zhao, Gaofang Cao. Current Practice and Reflections on the Protection of Personal Health and Medical Information. *Chinese Health Service Management*, 2023, 40(11):858-860,866.
- [21] Yuheng Chen. Constructing a Whole-Process Compliance System for Personal Information Protection in Generative Artificial Intelligence. *Journal of East China University of Political Science and Law*, 2024, 27(2):37-51.
- [22] Xi Zheng and Surong Zhu. Legal Risks and Regulation of Generative Artificial Intelligence. *Changbai Journal*, 2023(6):80-88.
- [23] Di Zhou, Mao Wang, Jingwei Wang. Protection and Use of Medical and Health Information in the Construction of Smart Hospitals. *Clinical Focus*, 2024, 39(9):833-836.
- [24] Zhe Wang and Yong Deng. Constructing a System for the Attribution of Rights in Health and Medical Data: Current Situation, Necessity and Key Concerns. *Soft Science of Health*, 2024, 38(3):25-28,34.