AI and healthcare
A case of an omelette without eggs?

Francisco Arga e Lima

Introduction

The adoption of artificial intelligence ("AI") in the healthcare sector has the potential to improve medical diagnosis and treatment, thereby improving public health efforts, equitable access to care and ultimately contributing to the advancement of the public good.

However, the training of these AI models depends on personal data gathered, for example, during patient interactions with healthcare services. The processing of this information is subject to heightened protection under the General Data Protection Regulation ("GDPR"), creating significant limitations to AI developers in processing it in a way that respects the principle of lawfulness.

It is with these limitations in mind that this essay aims to bring clarity on a possible way forward for AI developers to lawfully process health data in the training process of their models.[1]

Health data and the GDPR

In order to have a complete overview of the steps AI developers need to adopt, it is important to clarify what makes data personal, as well as the concrete limitations derived from the principle of lawfulness to the processing of health data.

[1] In this regard, we will limit ourselves to Article 9(2) and 6(4) of the GDPR, assuming AI developers act as controllers. Considerations over applicable legal bases under Article 6(1) of the GDPR will, therefore, not be the main scope of this essay.

1. Personal data and health data

The GDPR defines personal data in Article 4(1) as all forms of data, whether objective or subjective,[2] whose content, purpose or effect has a direct or indirect connection to an individual.[3] Importantly, this individual needs to be identified or, at least, identifiable. Here, the GDPR adopts a contextual approach, where the main criteria for identifiability are the means – i.e. additional information held by third parties – reasonably likely to be used to identify the data subject.[4] Within this broad concept, Article 4(15) of the GDPR sets out a specific classification for data concerning health, which is to be interpreted as including more than purely medical data, i.e. hospital admission records, appointment schedules, invoices for healthcare services, and social security numbers.[5] Recital 35 GDPR also broadens its temporal dimension, encompassing personal data about past, present, and future health conditions.

An important requirement for the definition of personal and health data is, therefore, the need for the information to relate, at the very least, to an identifiable individual. This brings us to the distinction between personal data and non-personal data – in particular anonymised data – important since the GDPR does not apply to situations where the re-identification of the data subject becomes practically impossible or excessively difficult. Two caveats must, however, be made.

First, assessing anonymisation involves evaluating whether reasonable means (i.e. financial, technical, and human resources) exist to either (1) re-identify individuals from data held by a controller or (2)

---

[2] CJEU, Case C-434/16, Nowak, para. 34; CJEU, Case C-413/23 (Opinion Advocate-General), EDPB v. SRB, para. 29.

[3] CJEU, Case C-434/16, Nowak, para. 35; CJEU, Case C-413/23 (Opinion Advocate-General), EDPB v. SRB, para. 30; CJEU, Joined cases C-92/09 and C-93/09, Schnecke, para. 53.

[4] CJEU, Case C-582/14, Breyer, para. 43 and 46.

[5] EUROPEAN DATA PROTECTION SUPERVISOR, Opinion of the European Data Protection Supervisor on the Proposal for a Directive of the European Parliament and of the Council on the Application of Patients' Rights in Cross-Border Healthcare (2009), para. 15.

access the information necessary to re-identify the data subject.[6] In this context, a line has to be drawn between anonymised data and pseudonymised data. According to guidance from the European Data Protection Board ("EDPB"), pseudonymisation involves replacing identifying information with pseudonyms.[7] Still, pseudonymised data are viewed as personal data since the pseudonymisation process is reversible,[8] through (1) the identification of the data subject, (2) reconnection of pseudonymised data to the original data, or (3) recreation of original data via additional information held by the controller.[9] By contrast, anonymisation is not reversible through reasonably available methods.[10] Thus, for personal data to be considered anonymised, controllers must assess what trumps their anonymisation efforts, considering not only their capabilities but also third-parties' and technological progress.[11]

Second, while anonymised data itself escapes the GDPR's scope, the anonymisation process does not.[12] In this regard, anonymisation operations, generally further processing activities, must respect Article 6(4) of the GDPR. This is in most cases not too high of a standard to achieve, as confirmed by the Article 29 Working Party ("WP29").[13]

---

[6] CJEU, Case C-582/14, Breyer, para. 46; Recital 26 GDPR.

[7] TOSONI, L.; "Article 4(5). Pseudonymisation". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 135; ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 20; CJEU, Case C-413/23 (Opinion Advocate-General), EDPB v. SRB, para. 47.

[8] EUROPEAN DATA PROTECTION BOARD, Guidelines 01/2025 on Pseudonymisation (2025), p. 11.

[9] Idem. See also CJEU, Case C-413/23 (Opinion Advocate-General), EDPB v. SRB, para. 47-57.

[10] ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 24.

[11] CJEU, Case C-582/14, Breyer, para. 43.

[12] ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 7.

[13] Idem.

2. The principle of lawfulness and health data

When the development of AI models requires the processing of personal data, the applicability of the GDPR and the need for a legal basis is, therefore, unequivocal. Hence, the path forward will depend on whether AI training is the initial or a subsequent processing purpose.[14]

2.1. Personal data is collected for AI training purposes

If personal data is initially collected for AI training, given the sensitivity inherent to health data, developers will require an exception under Article 9(2) of the GDPR to apply.[15] Generally, two options are put forward.

Firstly, Article 9(2)(a) of the GDPR allows the processing of health data based on the explicit consent of the data subject.[16] Despite being theoretically possible to obtain it, consent has clear practical problems, due to, i.e. information and explicitly requirements, as well as the need to ensure a proper right to withdraw consent.[17]

The second option comes through Article 9(2)(e) of the GDPR, that allows for the processing of special categories of data if the data subject had manifestly made them public. However, here we also have practical difficulties, since controllers will need to demonstrate the data subject's intentionality in doing so, which is hard to achieve at scale. In fact, it must be evident that the data subject intended to make this information public on a case-by-case basis, considering i.e. the medium used

---

[14] EUROPEAN DATA PROTECTION BOARD, Guidelines 03/2020 on the processing of data concerning health for the purpose of scientific research in the context of the COVID-19 outbreak, p. 6.

[15] GEORGIEVA, L., KUNER, C.; "Article 9. Processing of special categories of personal data". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 376-377.

[16] Idem.

[17] Idem.

to disclose the information, its standards settings, accessibility, etc.[18]

Given these challenges, the GDPR provides another route. This comes through Article 9(2)(j), where the GDPR allows for the processing of health data for scientific research purposes.[19] Three cumulative conditions determine the applicability of this exception: (1) the processing purposes must relate to scientific research, excluding, i.e. the subsequent commercial exploitation of the AI model;[20] (2) it must be based in EU or national law; and (3) the processing must be limited to what is strictly necessary, with suitable safeguards to mitigate risks towards data subjects.[21]

2.2. AI training is a subsequent data processing purpose

If AI training is a subsequent purpose from the one for which health data was initially collected, the GDPR allows it if it is compatible with the initial purpose.[22]

Here, a joint reading of Articles 5(1)(b), 6(4), and 9(2)(j) of the

---

[18] EUROPEAN DATA PROTECTION BOARD, Guidelines 8/2020 on the Targeting of Social Media Users (2020), p. 34–36.

[19] VERHENNEMAN, G.; "AI and Healthcare Data". In: The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence. Smuha, N. A. (ed). Cambridge: Cambridge University Press (2025) Cambridge, p. 312; KOTSCHY, W.; "Article 6. Lawfulness of processing". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 342.

[20] SVANBERG, C. W.; "Article 89. Safeguards and derogations relating to processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 1246.

[21] GEORGIEVA, L., KUNER, C.; "Article 9. Processing of special categories of personal data". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 381.

[22] An important discussion is whether compatible further processing still requires a new legal basis under Article 6 and Article 9 of the GDPR, if applicable. Although this is a central topic also for the training of AI in healthcare, since the common denominator for Article 6(4) and Article 9(2) is scientific research, we find that this discussion does not bear significant consequences for the discussion in this essay. See also VERHENNEMAN, G.; "AI and Healthcare Data". In: The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence. Smuha, N. A. (ed). Cambridge: Cambridge University Press (2025) Cambridge, p. 312;

GDPR leads us to conclude that this can be the case for scientific research. However, the assessment under Article 6(4) GDPR is still required,[23] with special attention to the sensitive nature of health data, as required by Article 6(4)(c) of the GDPR.[24] Additionally, and under Article 89(1) of the GDPR, this requires additional safeguards to be put in place (i.e. pseudonymisation and anonymisation) to minimise risks towards data subjects.[25]

2.3 Scientific research under EU law

We thus arrive at scientific research as a common denominator for the processing of health data for AI training purposes, which is broadly defined by the GDPR as including technological development and privately funded research.[26-27] This way, to qualify as scientific research, the activity must align with two core elements:

a. Firstly, there must be a systematic activity based on the collection and analysis of structured data aimed at increasing a *corpus* of knowledge.[28] Additionally, the EDPS, looking at the

---

[23] KOTSCHY, W.; "Article 6. Lawfulness of processing". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 342.

[24] ARTICLE 29 WORKING PARTY, Opinion 03/2013 on Purpose Limitation (2014), p. 25.

[25] KOTSCHY, W.; "Article 6. Lawfulness of processing". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 342-343.

[26] Recital 159 GDPR; SVANBERG, C. W.; "Article 89. Safeguards and derogations relating to processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes". In: The EU General Data Protection Regulation (GDPR): A Commentary. Kuner, C., Bygrave, L. A., Docksey, C. (eds). Oxford: Oxford University Press (2020), p. 1249.

[27] Notwithstanding the GDPR's broad definition, Member States may define additional conditions for processing special categories of data. This means that the conclusions we arrive at will need to be analysed in line with national derogations and definitions of scientific research. See also VERHENNEMAN, G.; "AI and Healthcare Data". In: The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence. Smuha, N. A. (ed). Cambridge: Cambridge University Press (2025) Cambridge, p. 312.

[28] EUROPEAN DATA PROTECTION SUPERVISOR, A Preliminary Opinion on data pro-

definition provided by the Copyright in the Digital Single Market Directive ("CDSMD"), frames scientific research as a methodological pursuit of knowledge with public benefits.[29]

b. Secondly, scientific research requires the adoption of established scientific methods, i.e. the observation of phenomena, hypothesis creation and testing to ascertain their validity, openness to scholarly critique, and transparency of the conclusions arrived at.[30]

How to process health data for AI training in the healthcare sector

Based on this, we find room to conclude that AI training may be conceptualised as scientific research. In fact, it can be argued that an AI model is in itself a means that allows the increase of insights and the available *corpus* of knowledge since it leads to the creation and analysis of information through the rules and patterns it generalises from its training – a systematic and methodological process of arriving at the desired result: a functional AI model.

The Hamburg Regional Court in Kneschke v. LAION arrived at a similar conclusion, albeit when looking at the CDSMD, where it acknowledged the collection of training data as scientific research, establishing that:[31]

a. The concept of research includes preliminary steps and not solely direct knowledge creation, as long as they intend to lead to knowledge advancement.
b. The definition does not hinge upon the success of the research

tection and scientific research (2020), p. 9-10.

[29] Idem. See also Article 3 CDSMD.

[30] EUROPEAN DATA PROTECTION SUPERVISOR, A Preliminary Opinion on data protection and scientific research (2020), p. 10.

[31] LG Hamburg, Urteil vom 27.09.2024 – 310 O 227/23, para. 113.

but rather rests on methodological and systematic intent to achieve these objectives.

Therefore, AI training can be considered as scientific research when that training contributes systematically, methodologically, and transparently to societal knowledge and technological advancement, which seems to be the case with models aimed at being deployed in the healthcare space to i.e. render better diagnosis and medical treatment to patients.

However, to process health data based on scientific research purposes, the conditions described above need to be complied with, namely limiting the data processing to what is strictly necessary, with appropriate data minimization and anonymisation processes to mitigate the risks towards data subjects.

With this in mind, there are three main phases of minimisation and anonymisation AI developers need to consider.

1. Data collection

The healthcare sector processes personal data for multiple purposes, including to validate therapies, diagnosis and other healthcare practices. This information is typically collected and generated through various interactions with patients and kept in electronic health records ("EHRs"), which are essential sources for AI training datasets.[32] However, while unique patient identification is vital to healthcare services, patient names and other identification elements are typically unnecessary in the context of AI training. This means that all personal data that is not necessary to train the AI model must be removed or anonymised, ideally before being collected by the AI developer.

Consequently, to go in line with GDPR standards, developers will need to ensure they collect anonymised datasets from their suppliers to

[32] VERHENNEMAN, G.; "AI and Healthcare Data". In: The Cambridge Handbook of the Law, Ethics and Policy of Artificial Intelligence. Smuha, N. A. (ed). Cambridge: Cambridge University Press (2025) Cambridge, p. 308-310.

the extent possible for the purposes pursued and given the practical difficulties of anonymising EHRs.[33] Indeed, fully anonymising i.e. free-text clinical reports, may prove difficult, leading to residual personal data remaining within the dataset. In such cases, AI developers must also implement additional safeguards in a second phase to manage residual risks of re-identification.

2. Data pre-processing

After the data collection phase, developers will need to pre-process the collected information and remove any unnecessary personal data by using appropriate anonymisation processes. Here, two primary methods can be deployed:

a. Randomisation: This approach modifies data attributes, altering their accuracy to obscure the link to data subjects. One of the main randomisation techniques is noise addition, which adjusts data values just enough to retain the same distribution without allowing re-identification.[34]
b. Generalisation: Its purpose is to dilute data's specificity, adjusting magnitudes or broadening its attributes. Aggregation, for example, aims to combine data subjects into sufficiently large groups to effectively prevent identification at an individual level.[35]

After adopting an anonymisation process, developers will need to test and ensure it is effective, considering both the technical feasibility and expertise required to reverse it.[36] This must be done taking into

[33] Idem.
[34] ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 12.
[35] Idem., p. 16-17.
[36] Idem., p. 9.

account risks of singling out,[37] linkability,[38] and inference,[39] as mentioned by the WP29.[40] The results of this assessment will inherently vary according to factors such as dataset attributes (i.e. their uniqueness), limitations on the access to the training datasets, availability of additional data allowing re-identification and the effort, cost, and required technological expertise necessary to reverse the anonymisation, which the developer needs to consider.[41]

This means that developers needing to anonymise personal data must combine different anonymisation strategies. Neither randomisation nor generalisation alone offer complete protection, so incorporating different methods is required for an efficient anonymisation process.[42] Moreover, developers should exclude rare or unique attributes to the extent that it doesn't undermine the training of the model, as well as ensure the underlying personal data is properly deleted.[43]

3. Model training

When personal data is necessary for the model's training, developers must adopt adequate safeguards to minimise its inherent risks, particularly those linked to the memorisation and output of health data.

In fact, the EDPB concluded that AI models trained on personal data cannot automatically be considered anonymous. Instead, given the high likelihood of extraction and inference, a case-by-case risk

---

[37] Risk of individual records being isolated within the dataset.

[38] Risk of linking multiple records related to one data subject within the same database or across separate databases.

[39] Risk of deducing sensitive attributes with significant probability from other information available within the dataset.

[40] ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 11-12.

[41] CJEU, Case C-479/22 P, OC, para. 50; EUROPEAN DATA PROTECTION BOARD, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (2024), p. 16.

[42] ARTICLE 29 WORKING PARTY, Opinion 05/2014 on Anonymisation Techniques (2014), p. 23.

[43] Idem., p. 9 and 25.

assessment must be undertaken by the developer to reach that conclusion.[44] This assessment should include considerations on the training data's characteristics, the model's architecture, training procedures, contextual use once deployed, cost, effort, and available technological resources.[45] Importantly, developers must determine whether third parties can reasonably identify data subjects from interactions with the model or access to it, if it is i.e. open-sourced.[46]

Therefore, AI developers must pay particular attention to the following when training their models on health data:[47]

a. Data Protection Impact Assessment ("DPIA"): Given the substantial risks connected with health data, conducting a DPIA is not just mandatory but also allows developers to systematically identify, evaluate and mitigate risks towards data subjects.
b. Privacy by design and by default: Developers should evaluate the impact different design and training choices have on the identifiability of data subjects from the model's parameters and output. Additionally, the EDPB recommends training methods involving i.e. regularisation, which improves model generalisation and reduces risks linked to overfitting. Privacy-preserving technologies such as differential privacy should be evaluated and applied when suitable.
c. Data sources and preparation: Developers must ensure only the strictly necessary personal data is used to train the model. Thus, and following the steps outlined in this essay, they must adopt relevant criteria when selecting the training data, assess the adequacy of data sources with the training objective, and exclude

[44] EUROPEAN DATA PROTECTION BOARD, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (2024), p. 14-16

[45] Idem.; CJEU, Case C-479/22 P, OC, para. 50-51.

[46] EUROPEAN DATA PROTECTION BOARD, Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models (2024), p. 16 et seqs.

[47] Idem.

unnecessary or inadequate sources. Additionally, they should adhere to the data pre-processing methodologies discussed previously, using anonymised or at least pseudonymised data whenever possible.

d. Testing for vulnerabilities and residual risks: Testing should be conducted throughout the training stage to address common and the most critical privacy attacks and risks, including attribute inference, membership inference, model inversion, memorisation, reconstruction, exfiltration, and regurgitation of training data.

Conclusion

This essay aimed to outline a three-phase approach enabling AI developers in the healthcare sector to lawfully process health data to train AI models as scientific research. Initially, they must obtain anonymised datasets, thus avoiding the application of the GDPR from the start. To the extent that this is not possible, developers must anonymise or remove any remaining personal data before training the model that is not needed at the training stage. If it is, developers must adopt a third layer of protection, by mitigating risks to data subjects through the pseudonymisation of personal data, and the adoption of an architecture and training process that limits the memorisation and output of personal data and is resilient to common and severe attacks.

While this essay aims to provide another route for AI developers to process special categories of data in this sector, other obligations will also need to be accounted for, such as the respect towards data subject rights, namely when it comes to the right to be informed and to be forgotten, whose practical implementations are far from easy to achieve.