


RESEARCH ARTICLE

Generative AI and data protection

Hannah Ruschemeier 

Faculty of Law, University of Hagen, Hagen, Germany
Email: hannah.ruscheimer@fernuni-hagen.de

(Received 29 July 2024; revised 26 September 2024; accepted 09 October 2024)

Abstract

Generative artificial intelligence (AI) has catapulted into the legal debate through the popular applications ChatGPT, Bard, Dall-E and others. While the predominant focus has hitherto centred on issues of copyright infringement and regulatory strategies, particularly in the context of the AI Act, a critical but often overlooked issue lies in the friction between generative AI and data protection laws. The rise of these technologies highlights unresolved tension between safeguarding fundamental protection rights and the vast, almost universal, of scale of data processing required for machine learning. Large language models, which scrape nearly the whole Internet rely on and may even generate personal data falling under the GDPR. This tension manifests across multiple dimensions, encompassing data subjects' rights, the foundational principles of data protection and the fundamental categories of data protection. Drawing on ongoing investigations by data protection authorities in Europe, this paper undertakes a comprehensive analysis of the intricate interplay between generative AI and data protection within the European legal framework.

Keywords: generative AI; data protection; GDPR; LLMs; regulation; privacy

1. Prelude: risks and challenges of generative AI

Now that the initial hype around generative artificial intelligence (AI) in the form of large language models and image generators has subsided, legal issues are coming to the fore. In addition to discussions about generative AI and copyright, there is an increasing focus on the friction between generative models and the requirements of data protection law. In the US, several lawsuits are underway against Google and OpenAI regarding potential privacy violations by generative models.¹ Regulators are currently active in the European Union² and the European Data Protection Board has set up a task force to deal with ChatGPT (EDPB, 2023). The Italian data protection authority Garante had already opened a case against OpenAI in 2023, which led to a temporary national ban on ChatGPT (GPDP, 2023). After the proceedings were concluded, the authority found violations of the GDPR (GPDP, 2024). Investigations into data protection violations are also underway in Poland (UODO, 2023). Other countries such as Germany have issued requests for information (LDI NRW,

¹ See the class action against OpenAI: <https://storage.courtlistener.com/recap/gov.uscourts.cand.425482/gov.uscourts.cand.425482.1.0.pdf>.

² For an overview of ongoing investigations of data protection authorities outside the European Union, see Zanfir-Fortuna (2023).

2023),³ and the French data protection authority has developed an action plan (CNIL, 2023). In the case of Maximilian Schrems, the data protection NGO NYOB filed a complaint with the Austrian data protection authority in April 2024 (NOYB, 2024), centred on incorrect information about an individual provided by ChatGPT, which OpenAI neither corrected nor responded to the request for information about what information was processed. These cases make it clear that data protection authorities are already AI regulators and generative AI is a core issue for data protection.⁴

From a legal perspective, generative models introduce a range of distinct issues, which are well-documented across various scholarly sources (Charlotin, 2023; Chen, 2024; Chiara, 2023; El-Mhamdi et al., 2023; Gillotte, 2020; Hazell, 2023; Marcus & Pullin, 2023; Sag, 2023; Weidinger et al., 2022). In particular, the foundation models on which the popular large language models (LLMs) are built pose new security risks and vulnerabilities that need to be addressed. As a result, a socio-technical assessment is needed to understand these risks and the necessary safety mechanisms, including legal and ethical aspects. Understanding the risks posed by LLMs requires a contextual approach: normative rules, like law, always operate in context.

A major concern is the protection of personal data and privacy. Different experiments have shown that it is possible to extract personal and sensitive information about individuals from LLMs (Carlini et al., 2023; El-Mhamdi et al., 2023; Gupta, 2023; Nasr et al., 2023). Researchers have proven that LLMs are able to memorise training data, either through over-application of abundant parameters to small datasets, thus reducing the capacity to generalise to new data, or through optimisation for generalisation in long-tailed data distributions (Novelli et al., 2024). Although this phenomenon most often occurs where duplicates exist in the training data, it will still appear where training data have been partially deduplicated. Larger models with more parameters “remember” more data than smaller models (Carlini et al., 2022). Violations of people’s privacy and right to data protection result from both incorrect information and correct information they do not want published (Mühlhoff & Ruschemeier, 2024a). These risks are exacerbated by unregulated and therefore uncontrolled secondary downstream use of the models. In the case of popular LLMs operated by global technology companies, commercial resale seems remote, as the companies have no interest in giving up their exclusive option for commercial exploitation. The situation is different for smaller, but in some cases no less risky models: Mixtral 8x7B competes with and surpasses GPT 3.5 in some respects, due to a smart architecture that combines eight different expert models, and has been recently made open source (Hacker, 2023). This only highlights the needs for an overview of the purposes for which these models are used and a categorisation that enables a context-based risk assessment (Mühlhoff & Ruschemeier, 2024b).

Data protection law gives rise to its own particular frictions, from the general function and technical specificities of big data applications and generative AI on the one hand, and on the other, the particularities of generative models. Generative models are used in different contexts for different purposes, to generate text, code, video, images, audio, etc. In this article, I will focus on LLMs, which generate text by calculating the probability of word order. Data that are linguistically translatable, i.e., can be understood by human recipients, can clearly also contain personal data as covered by data protection law. For this reason, LLMs are a good example of the problems of how data protection law works in relation to AI generated content.

This article is structured as follows: first, I outline the overarching lines of conflict between data protection law and generative AI. I then go into the specific legal issues of the GDPR: The scope

³The State Commissioner for Data Protection and Freedom of Information Rhineland-Palatinate, LfDI asks about ChatGPT, 24 April 2023. The State Commissioner for Data Protection Schleswig-Holstein sent out a list of questions: https://www.datenschutzzentrum.de/uploads/chatgpt/20230419_Request-OpenAI_ULD-Schleswig-Holstein_IZG.pdf.

⁴Until the beginning of 2024, OpenAI did not have an establishment in Europe, so no member state data protection authority was responsible under Art. 56, 60 GDPR. Therefore, the supervisory authorities of the member states could all act in accordance with Art. 55 GDPR within their areas of competence. OpenAI now operates an office in Dublin, which is designated as the data controller.

and legal basis for authorisation of different steps of data processing by generative AI (Section 2), the principles of data processing (Section 3), the rights of data subjects (Section 4) and questions of responsibility (Section 5). In Section 6, I discuss the transferability of the argument to models that create images, audio and video. The article concludes with an outlook (Section 7).

2. Structural challenges of generative AI for data protection law

Data protection law in the EU is primarily addressed by the GDPR.⁵ The current system of the GDPR is rooted in the Data Protection Directive adopted in 1995,⁶ the right to data protection (article 8 CFR), the right to privacy (article 7 CFR)⁷ and the primary legal foundations in article 16 TFEU.⁸ Article 1 GDPR sets the matter and scope as the processing of personal data with the objective of protecting fundamental rights and freedoms of natural persons. Correspondingly, the understanding of “processing” in terms of personal data is very broad and takes an all-encompassing approach to cover practically any interaction with personal data.⁹ For this reason, when personal data are involved, all stages in the life cycle of an AI model may fall within the scope of the GDPR.

From a regulatory perspective, the various steps of data processing are therefore important in the life cycle of an AI model, and for generative models can be differentiated as follows:¹⁰ The first step is the collection of training data, made up of many data points. These data points may comprise personal or non-personal information. In certain instances, this process utilises extremely large datasets, making it challenging, if not impossible, to differentiate between various categories of data. For instance, ChatGPT was developed using copious amounts of data freely available on the Internet. The second step is the actual training of the model using the collected data, resulting in a configured model. The third step is model application, meaning that the trained model is applied to specific cases or individuals, making the model a tool that computes a specific output in response to input data (Mühlhoff & Ruschemeier, 2024b). This quantity of data and the training process mean model output contains information about cases or individuals as well as that of “third parties” that were not part of the training data.

2.1 Quantity

The first problem area relates to how the training of powerful AI models, or the processing of large amounts of data, works in relation to the amount of data processed (Zarsky, 2017). The sheer quantity of data processed by potent AI models is the core, as yet unresolved, problem of AI and data protection.¹¹ Generative AI models are typically trained on billions, if not hundreds of billions, of parameters and require large amounts of training data and computing power (Brown et al., 2020). Data protection law on the other hand is based on the idea that the individual steps of data processing and the data processed can be identified. This concept applies the idea of individual control to empower individuals by allowing them to manage their own personal information (Solove, 2024). But models trained on unprecedentedly large data sets make it impossible to manually identify or

⁵Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance) OJ 2016 No. L119, 4 May 2016.

⁶Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data OJ L 281, 23.11.1995, p. 31–50.

⁷I understand privacy as one of the protected interests under data protection law. More about the relationship: Eichenhofer (2021), p. 47 et seq.).

⁸Additionally, article 39 TEU lays down specific provisions for the area of the Common Foreign and Security Policy.

⁹See V. P. & P. D. Hert, ‘GDPR Art. 4 Abs. 2 Processing’ (2023), par. 3.

¹⁰These steps are developed in Mühlhoff & Ruschemeier (2024a, 2024b).

¹¹Generative models are based on neural networks, usually Generative Adversarial Networks and Transformer Networks (Huang et al., 2024).

even review whether data processing complies with legal requirements, and thus harbours potential for privacy and data protection violations.¹² Furthermore, this approach conflicts with the principle of data minimisation laid down in article 5(1)(c). This mode of operation reveals the problems of governance arising from the systematic design of the GDPR, which, for example, envisages individual consent as the basis for authorisation and presupposes the identification of individual data subjects and the data to be attributed to them.

2.2 Purposes

Privacy and data protection also seem to be at odds with the general concept of generative AI when it comes to the relevance of purposes. Data protection is highly contextual, and its level of protection depends on the type of data processed, by whom, in which settings and for which purposes, article 5(1)(b). LLMs on the other hand, cover a wide range of purposes, applications and operating environments. According to article 3(63) of the new European regulation on AI¹³ (in the following: AI Act), a general-purpose AI model includes AI models trained with a large amount of data using self-supervision at scale, which display significant generality, are capable of competently performing a wide range of distinct tasks regardless of how the model is placed on the market, and that can be integrated into a variety of downstream systems or applications. It does not include AI models used for research, development or prototyping activities before they are placed on the market. This definition is a good description of the current market situation; OpenAI, for example, now offers a wider variety of different GPTs for specific tasks: the laundry buddy for laundry specific questions about stains and laundry settings, the sous chef that provides users with recipes or the negotiator that helps a user to argue in their favour (available under ChatPT 4o at chatgpt.com with a payed subscription). These downstream applications will gain more relevance, as it can be expected that the foundation models will not continue to be used primarily as isolated applications, as has been the case to date, but will be integrated into other models as modular building blocks. As a result, this will increase both desirable and undesirable effects due to the possible scaling of model output. Here, even the design aspect of LLMs is difficult to reconcile with legislation, seemingly conflicting with the GDPR's purpose limitation principle. In particular, when models are made available to numerous third parties via an interface, ensuring compatibility for that model and its data with the purposes for which the data was originally collected (article 6(4)) becomes difficult, if not impossible (Mühlhoff & Ruschemeier, 2024b).

3. Scope of the GDPR and legal basis

The GDPR is open in terms of material and geography, i.e., it extends to the processing of personal data for activities within the EU, even when that processing takes place elsewhere, (article 3.1), and where goods or services are offered to data subjects within the Union (article 3.2). It therefore applies for all generative models in use in the Union.

3.1 Scope of application

3.1.1 Personal data

The GDPR applies to the processing of personal data (article 2(1)) if none of the exceptions in paragraphs 2–3 apply. This processing includes both the collection of training data and the training of the models, as well as the use or sale of the model to generate output based on user requests.

The processing of personal data begins with step one, the collection of vast amounts of data with which to train an LLM. As the effectiveness of LLMs is directly linked to the breadth and variety of

¹²On the relationship between privacy and data protection: Gellert (2023); Kokott and Sobotta (2013).

¹³P9_TA(2024)0138 (COM(2021)0206 – C9-0146/2021 – 2021/0106(COD)).

their datasets, this data is obtained by scraping content from numerous websites. Inevitably this often includes personal data such as names, dates of birth, or other identifying information (Ruscheimer, 2023a). As personal data also includes incomplete, or indirect details which may result in an individual being identified through additional information,¹⁴ this processing is therefore covered by the GDPR, even before the model is trained or released.¹⁵

In the second step of data processing, the training of the model, identifying personal data becomes more challenging, as the final trained model may differ from its training data. An artificial neural network, for example, is represented by a large matrix of numbers, which in turn are determined by weights and other parameters such as activation thresholds (Mühlhoff & Ruschemeier, 2024b, p. 27; Abadi et al., 2016). While the training data may include personal information, the data in the model may not necessarily retain that characteristic: personal data may be anonymised where advanced techniques such as differential privacy and federated machine learning are used during the training process to remove references to the training data (Abadi et al., 2016; Dwork, 2006; El Emam et al., 2015).

A trained model resulting from such anonymisation that makes the reconstruction of training data impossible or highly unlikely is not considered to constitute personal data. However, the current popular large language models tend to persist in producing identifying information, whether by design or by accident. It cannot therefore always be assumed that model data has been fully anonymised: research into this “remembering” phenomenon is ongoing. This is critical from the GDPR standpoint as the storage of the model also constitutes data processing under the GDPR if the model data is not anonymised. In addition, many authors argue that the anonymising of personal data itself is also a processing operation that requires justification under the GDPR (El Emam & Alvarez, 2014; Roßnagel, 2021; Schreurs et al., 2008).¹⁶

In processing step number three, the production of output, the models or application using them can produce personal data. Whether the information provided is correct or not is immaterial: when LLMs produce outputs that contain the names and bibliographical information of real people, they are processing personal data. Additionally, individuals can often be easily identified from the context of the text prompt or text output, or by using search engines. LLMs linked to search engines may also facilitate identification. Particularly in the case of public LLMs, it is likely that many data subjects can be identified for the reasons mentioned above (Pesch & Böhme, 2023). It is important to note that the people in the training data are not theoretically the same as those produced in the output data even where they have the same name, as LLMs can also generate names of existing people, for example, by producing information which can then be assigned by users.

3.1.2 Territorial scope of application

Article 3(1) of the GDPR states that the regulation applies “to the processing of personal data in the context of the activities of an establishment of a controller or a processor in the Union, regardless of whether the processing takes place in the Union or not.” Thus, the processing of personal data does not have to take place in the Union itself, but can be performed on servers that are for example based in the USA or other third-countries. As mentioned, *lex loci solutionis* is principle of the GDPR, article 3(2) GDPR, means the requirements apply even if the data processor is not located in the EU, but

¹⁴The Hamburg data protection authority, on the other hand, believes that LLMs do not store personal data, without explaining how this relates to the scraping of information that undoubtedly constitutes personal data: (https://Datenschutz-Hamburg.de/Fileadmin/User_upload/HmbBfDI/Datenschutz/Informationen/240715_Diskussionspapier_HmbBfDI_KI_Modelle.Pdf, n.d.).

¹⁵In its judgement on the Breyer case (CJEU C-582/14, ECLI:EU:C:2016:779), the ECJ also considers whether the identification process is lawful in the context of the question of the identifiability of persons, i.e., whether personal data are involved. This, however, contradicts the protective direction of the GDPR, which aims to protect data subjects from unlawful processing of their data. For this: Hacker (2021).

¹⁶Article 29 Working Party, Opinion 05/2014 on Anonymisation Techniques (WP 216) 0829/14/EN. On this dispute in the context of predictive analytics: Mühlhoff and Ruschemeier (2024a), p. 19.

offers their services to EU citizens. This therefore brings global technologies such as LLMs and other AI models such as Chat-GPT, Bard, Gemini, etc., squarely under the GDPR where these are accessible from the European Union.

3.2 Legal basis for data processing

All processing of personal data within the scope of the GDPR requires a legal basis, article 6(1) GDPR. The question of the legal basis for data processing across the life cycle of a generative AI system poses different problems, as it depends on the stage of the data processing. As argued before, it is essential to distinguish between the different steps of data processing when analysing AI and data protection (Mühlhoff & Ruschemeier, 2024a).

3.2.1 Collection of training data

The first step in the life cycle of a generative model is the collection of training data. In the case of LLMs like GPT-4 or Bard, this step consists of scraping data from the internet. The indiscriminate scouring of almost the entire internet logically excludes the legal basis of consent, article 6(1)(a). In the absence of legal obligations or contractual relationships between the operators of LLMs and all internet users worldwide, the scraping of training data can only rely on the legal basis of legitimate interest provided in article 6(1)(f) GDPR (Borgesius et al., 2017; Zarsky, 2017).

Article 6(1)(f) GDPR states that data processing is lawful if it is necessary for the purposes of pursuing the legitimate interests of the data controller or by a third party, provided these interests do not override the fundamental rights and freedoms of the data subjects requiring protection (Ruscheimer, 2023a). The ECJ has clarified that this provision lays down three cumulative conditions for the processing of personal data covered by that provision to be lawful: (1) the pursuit of a legitimate interest by the controller or by a third party; (2) the processing of personal data must be necessary to pursue that legitimate interest; and (3) that the legitimate interest of the controller or of a third party are not outweighed by the interests or fundamental freedoms and rights of the data subject.¹⁷

The fact that this constitutes the only plausible legal basis exposes the structural problem of data protection law in relation to data-intensive technologies, not least because whether article 6(1)(f) provides a sufficient legal basis must be determined on a case-by-case basis (Donnelly & McDonagh, 2019; Novelli et al., 2024). There are indications that general interest may outweigh the purpose of processing, or that this can be assumed, if data subjects could reasonably expect their data be processed for training purposes. Ultimately, it depends on the individual case. However, the nature of mass data scraping works makes it almost impossible to identify individual interests at all, and therefore cannot provide satisfactory answers in terms of current legal doctrine and legal systems.

3.2.2 Legitimate interest

The term “legitimate interests” is deliberately broad, to encompass legal, economic or idealistic interests, excluding only hypothetical and public interests. The collection of data to train a generative model for commercial use is initially a legitimate economic interest and protected by the freedom to conduct a business under article 16 ECFR. The argument that the ECJ also cited freedom of information (article 11 (2) ECFR) as a legitimate interest in processing transferable to generative model training (Krönke, 2023) in the Google Spain case¹⁸ does not apply to models that are only accessible for a fee. Furthermore, search engines and generative models operate differently, and are therefore not comparable. Source references in search engines can be deleted or corrected, whereas LLMs generate a unique new text for each question for which a new probability is calculated. If the output text is incorrect, it cannot be corrected for future outputs.

¹⁷Case C-252/21 Meta ECLI:EU:C:2023:537, par. 106; Case C-597/19, EU:C:2021:492, para. 106.

¹⁸Case C-131/12 Google Spain, ECLI:EU:C:2014:317.

3.2.3 Necessity

The necessity test under article 3 (1)(f) this provision means that the processing of personal data must be a proportionate means of achieving the legitimate interests. Processing is considered necessary if the processing of personal data is essential in order to achieve the objective of the processor's legitimate interest, in this case, a trained AI model and that these interests do not outweigh the rights of the data subject. In rare cases, where only anonymised data is sufficient to train the model, that training may not require personal data. However, anonymised data alone is generally not adequate for training generative models even were such anonymisation possible in the training phase.

3.2.4 Balancing of interests

Article 6(1)(f) GDPR requires a balancing of the conflicting rights and interests between processor and data subject, which must also consider the rights of data subjects under articles 7 and 8 ECFR. Their interests are particularly affected when AI collects, links and contextualises personal data available on the Internet in response to user queries.

The fact that powerful generative models require a large amount of training data to achieve a certain level of performance, e.g., to generate word sequences that correspond to human language, speaks in favour of the interests of the operators. However, it is not absolutely necessary to scrape data on a scale that covers almost all publicly accessible resources on the internet to develop generative models: data sets can be generated in other ways, such as through data donations, effective consent solutions or data collection by the data controller itself. Although, none of these alternative options would be able to create the required breadth of data. The question is therefore which specific interest of the data processor is worth protecting. Meta, for example, has publicly admitted that the acquisition of licences for copyrighted material would have made the development of generative models considerably more difficult, simply because it would have made them more expensive. The same argument was made against collecting training data in a privacy-compliant manner: the approach would have required considerable resources. However, it is unlikely that cost savings can constitute a legitimate interest, and at any rate, an interest based on structural infringements has considerably less protective value.

In the Meta case,¹⁹ the ECJ also ruled that the personalisation of content – Meta's core business model – was not necessary for the operation of a social network.

The ECJ went on, stating that legitimate interests do not adequately justify Meta's practices of tracking and profiling individuals for the purpose of conducting its behavioural advertising business across its social platforms:

it is important to note that, despite the fact that the services of an online social network such as Facebook are free of charge, the user of that network cannot reasonably expect that the operator of the social network will process that user's personal data, without his or her consent, for the purposes of personalised advertising. In those circumstances, it must be held that the interests and fundamental rights of such a user override the interest of that operator in such personalised advertising by which it finances its activity, with the result that the processing by that operator for such purposes cannot fall within the scope of point (f) of the first subparagraph of Article 6(1) of the GDPR²⁰.

This raises substantial doubts about the ability of companies like OpenAI to defend the processing of vast amounts of personal data to establish a commercial generative AI enterprise, particularly given that such tools pose numerous emerging risks to identified individuals, including issues like disinformation, defamation, identity theft and fraud, among others.

¹⁹Case C-252/21 Meta ECLI:EU:C:2023:537, para. 102 et seq.

²⁰Case C-252/21 Meta ECLI:EU:C:2023:537 para.117.

Context is therefore critical in safeguarding privacy and data protection (Ruscheimer, 2023a). The public accessibility of data on the internet, even where disclosed by the data subjects themselves, this does not completely negate their legitimate interest in its protection. As recital 47 states, the interests and fundamental rights of the data subject may in particular override the interest of the data controller where personal data are processed in circumstances, or in ways that data subjects do not reasonably expect. Although it is now public knowledge that data posted on the internet may be processed in ways other than initially thought, it is also a question of the specific purposes of the processing. For example, the legitimate expectation of privacy means that decades-old or deleted posts, personal websites and entries cannot be used in perpetuity to train commercial models. It is reasonable to suggest that the typical Internet user does not expect, nor intend, for their data to be utilised as training material for LLMs for the financial gain of others. Therefore, the use of the data for training these models represents a secondary purpose. In most instances, it is unlikely that a data subject made their data publicly available to serve as a dataset for the financial gain of LLM providers, making the use of such publicly available data an infringement on contextual privacy (Nissenbaum, 2011).

Moreover, a legitimate interest must be determined within the broader European and national regulatory context. The broad scope of scraping also means that an unmanageable number of people are affected, which opens the claim of legitimacy to questions of proportionality. According to the German constitutional doctrine of the Federal Constitutional Court, this can intensify interventions if a particularly large number of people being affected without cause can affect the claim of legitimacy. This impact is referred to as “scattering width,” and is a line of argumentation also used by the ECJ.²¹ Such effects also arise in the case of universal data processing, as almost all Internet users are affected.

Additionally, the legitimate interest must also be lawful, meaning it should conform to all applicable laws and regulations, including the principles and other provisions of data protection law. This includes ensuring that processing aligns with the expectations of the data subject based on their relationship with the controller, adheres to the principles of data minimisation and implements appropriate safeguards. In the case of broad-scale web scraping, individual interests are difficult to identify. However, there were concerns about the legality of scraping from the outset, including with regard to potential copyright infringements. An interest pursued through structural infringements cannot be legitimate.

Compatibility with the principles of data protection law under article 5 GDPR also plays a role in the balancing of interests. This requires legitimate interests be evaluated in terms of the fairness of processing (article 5(1)(a)), purpose limitation (article 5(1)(b)), data minimisation (article 5(1)(c)), and data accuracy (article 5(d)).

Therefore, legitimate interests cannot be assumed across all training data. The matter is made more complex given it is extremely difficult, if not impossible to comprehensively exclude personal data pertaining to minors or special categories of personal data according to article 9(1) from training data. To complicate this matter further, the point at which the processing of personal data “reveals” special categories of personal data under article 9(1) GDPR has not yet been conclusively clarified.

3.2.5 Training of the model

It is worth taking a chronological review of the various data processing operations used for training the model. A key consideration is whether data anonymisation occurs during model training. The prevailing view is that an authorisation basis is required. In this context, anonymisation is understood as normative rather than technical, in line with the ECJ ruling that holds that data has been

²¹ Case C-511/18 u.a., ECLI:EU:C:2020:791, Rn. 143 f. – La Quadrature du Net; Case C-293/12, ECLI:EU:C:2014:238, para. 57 ff. – Digital Rights Ireland Ltd; Case C-203/15, ECLI:EU:C:2016:970, para 105 f. – Tele2; Case C-140/20, ECLI:EU:C:2022:258, Rn. 66 – Commissioner of An Garda Síochána.

anonymised even if it is technically possible, but unlikely, that the controller could carry out identification with the means available including additional information.²² In addition, according to the court, data are considered anonymous under the GDPR if re-identification is illegal.²³

In principle, the anonymisation of personal data is generally easy to justify under article 6 GDPR. The practice aligns with the principle of data minimisation and storage limitation, effective and permanent anonymisation can serve the interests of both data subjects and data controllers: the former are protected from unauthorised interference with their fundamental data protection rights, while the latter are freed from some of the perceived burdens of complying with the stringent requirements of data protection law (Hornung & Wagner, 2020). However, this argument struggles to hold in light of the the volume of data, as effective consent from the data subjects pursuant to articles 6(1)(a) and 7 GDPR, cannot be obtained in practice.

While it is conceivable to institute legal obligations to anonymise training data under Article 6(1)(c), this is not yet relevant in practice. This means that the legal basis of legitimate interest in article 6(1)(f) may also apply to anonymisation. Generally speaking, this provision may provide adequate results as anonymisation will typically be in the interest of the data subjects themselves, at the very least a conflicting interest is improbable. In the case of large LLMs, it is equally unlikely that a data subject will have an individual interest in non-anonymisation, and moreover, that even where such interest of an individual data subject exists, that it would outweigh other relevant interests, e.g., of the other data subjects.

Assessing the processing of special categories of personal data under article 9 is more difficult. If anonymisation as processing requiring justification would require a case be made under article 9(2). As described above, the processing of special categories of data cannot be ruled out for LLMs. Although the hurdles of article 9(2) are high, but the “made public” provision of article 9(2)(e) can also be considered here. Others argue for a teleological reduction of article 9(1) for anonymisation (Hornung & Wagner, 2020). Neither variant constitutes an infringement of data subjects’ rights where training data has been anonymized. These complex considerations alone show that there are gaps between the individual-based approach of the GDPR and the tools required for adequately regulating generative AI.

3.2.6 Generating output

The output of generative language models may constitute the processing of personal data. Here, a distinction must be made here between the processing of scraped training data and the processing of user data in the form of prompts entered while using the model. There is no legitimate interest in processing user data, e.g., in the context of input prompts when using LLMs. Instead, effective consent, pursuant to article 6(1)(a) must be obtained, a legal measure that needs to be critically evaluated in the digital space (Ben-Shahar & Schneider, 2011; Borgesius et al., 2017; De & Imine, 2020; Nguyen et al., 2022; Mühlhoff & Ruschemeier, 2024b, p. 31). Open AI had to update its privacy policy for EU users after being investigated by the Italian data protection authority. It now states: “We use the content you provide to improve our services, such as to train the models that run our services. Read our instructions on how to opt out of the use of your content to train our models” (OpenAI, 2023). However, consent is only a valid basis for prompts containing personal information about the user themselves. If users generate prompts that include personal data from other persons, they cannot validly consent on their behalf (Novelli et al., 2024).

When a generative model is capable of producing output which includes personal data, the issue of how training data was collected remains relevant throughout its life cycle. If there was no legal basis for collecting the training data, there is no legal basis for using it to generate output. Theoretically,

²²Case C-582/14, ECLI:EU:C:2016:779 – Breyer, para. 45–49.

²³Case C-582/14, ECLI:EU:C:2016:779 – Breyer, para. 46. Rightfully critical: Hacker (2021).

legitimate interest could also be considered here under article 6(1)(f), but must be assessed on a case-by-case basis according to the criteria described above. However, LLMs make individual assessments difficult because of the quantity of data they process. In addition, generative models are scalable in terms of their output, which means that false information can be disseminated to a large number of users and third parties.

Output processing is also problematic in cases where models infer or disclose special categories of personal data under article 9(1) GDPR. It has been shown that models can memorise and reproduce private and personal information such as phone numbers, addresses and medical documents (Carlini et al., 2023). In the age of big data, it is now potentially possible to infer sensitive information from almost any data, especially if one includes the boundless category of political opinions, included in article 9(1) GDPR. This means that “normal” personal data can reveal special categories of personal data covered by article 9(1), although the criteria for distinguishing between general and sensitive data remain contested (Mühlhoff & Ruschemeier, 2024a, p. 22). One proposed criterion goes to the intention behind the data processing. Scenarios involving context-specific information, the purpose of the evaluation could lead to the generation of sensitive data depending on the purpose of evaluation (Matejek & Mäusezahl, 2019; Schulz, 2018; Georgieva & Kuner, 2020, p. 373; Wachter & Mittelstadt, 2019).²⁴ Court rulings tend to support this assumption: the ECJ seemed to interpret “revealing” broadly in the *Meta* case,²⁵ and in another ruling, the Court decided that the disclosure of a spouse, partner or cohabitee’s name could potentially indicate the sexual orientation of the applicant.²⁶ The court has established minimal criteria for what constitutes the “revealing” sensitive data: the act of an “intellectual operation involving comparison or deduction” is deemed sufficient to extend the special protection regime meant for sensitive data to personal data that are not inherently sensitive. However, this judgement was not directly related to big data, leaving the distinction somewhat ambiguous (Mühlhoff & Ruschemeier, 2024a, p. 24).

Consequently, in many instances involving big data, merely being able to potentially infer sensitive information may subject processes such as AI training to the provisions of article 9, and there is little likelihood that LLMs satisfy the exceptions in article 9(2). For instance, the research exemption under article 9(2)(j), for instance, is restricted to the development of models for research purposes and does not permit their commercial exploitation, as indicated in Recitals 159 and 162 (Novelli et al., 2024).

Another important distinction is whether LLM output can be used to infer sensitive information about individuals that they have not made public themselves. Even if certain indicators, e.g., of political orientation, are available on the Internet, LLM output may aggregate this information. As such, article 9(2)(e) does not constitute a legal basis for this type of derivation.²⁷

Data accuracy requirements (article 5(1)(d) GDPR) also apply to LLM output. LMs have been shown to “hallucinate” and produce incorrect information, including incorrect personal data (Wachter et al., 2024). Under the GDPR, operators are responsible for ensuring data accuracy, articles 5(2), 24, 25(1) GDPR.²⁸ Although all popular applications provide disclaimers to inform users that the models may not always be correct, the effect of such notifications are questionable given the *automation bias* (Hondrich & Ruschemeier, 2023; Ruschemeier, 2023b; Ruschemeier & Hondrich, 2024). Even if the current error rate of LLMs (Metz, 2023) does not justify generally prohibiting such applications on the basis of ensuring data accuracy, it does affect data subjects’ rights. The right to data

²⁴ Question No. 2 in ECJ Case C- 252/21 *Meta Platforms and Others* [2021] (ECLI:EU:C:2022:704), dismissed by the Advocate general in his Opinion, Par. 41.

²⁵ Case C-252/21, *Meta*, ECLI:EU:C:2023:537, para. 73.

²⁶ Case C-184/20 ECLI:EU:C:2022:601.

²⁷ See also Case C-252/21, *Meta*, ECLI:EU:C:2023:537, para. 73.

²⁸ As the ECJ already stated in 2008 with regard to the maintenance of data in the Central Register of Foreigners – similar to an actor acting under public authority, in order to delete or correct incorrect information without delay. Case C-524/06, 2008 I-09705. Krönke (2023).

accuracy becomes even more significant if the right to rectification or erasure cannot be effectively enforced.

4. Data subject's rights

As with other areas of data-intensive technology application, there are problems with the enforcement of data subjects' rights in the case of generative models (Solove, 2023). In general, the actors involved mean many data driven AI technologies are developed, promoted, sold and used by a handful of big-tech-companies, which establishes an informational power asymmetry between the powerful processors and the users. As a result, privacy rights alone are insufficient to address the issue of data disempowerment.²⁹ Individuals typically possess limited capacity to manage their personal data, as there is a fundamental limit to the control they can exert. While rights can afford a modest degree of influence in certain isolated cases, this influence is too sporadic and disjointed to significantly safeguard privacy. Ultimately, rights function primarily as a minor element within a broader framework (Solove, 2023).

The sheer quantity of the data processed from various sources seems to make it impossible to identify and inform individuals of the processing, or of the processor of their data, effectively ruling out compliance with the right to information, and making it difficult for data subjects to assert their rights. In practice, reporting indicates that companies such as OpenAI and Midjourney have not responded to requests for information from people who had found themselves in the training data (Harlan & Brunner, 2023).

The prerequisite for the exercise of the data subjects' rights under the GDPR provided in articles 13–22 is, first and foremost, that the data subject is aware of the data processing. Users providing input into an AI model in the form of prompts are covered by article 13 GDPR. However, article 14 GDPR also comes into play where data has not been collected from the data subjects themselves. According to both standards, data subjects must be informed about who has processed which data (categories of data), for what purposes, on what legal basis and whether this data has been disclosed to third parties. These transparency provisions have the specific purpose of enabling data subjects to exercise their other rights, such as the right to erasure or rectification. Article 14(5) states the exception that the transparency obligation does not apply where and insofar as the provision of such information proves impossible or would involve a “disproportionate effort, in particular for processing for archiving purposes in the public interest, scientific or historical research purposes or statistical purposes [...] In such cases the controller shall take appropriate measures to protect the data subject's rights and freedoms and legitimate interests, including making the information publicly available.”

Again, it depends on the individual case, although it seems doubtful whether LLM operators can invoke unreasonableness if they already knew before the model was developed that individual requests for information could not be enforced. In any case, the principle of responsibility in article 5(2) GDPR means a failure to respond to such requests or a reference to general impossibility is not sufficient.

Practice has shown LLMs and possibly other generative AI models that produce content operate almost universally, not just at an individual level. This near-universal infringement reflects the profound mismatch between data-intensive models and the individual rights approach to data protection taken by data protection laws. As a result of this universality, other rights of data subjects such as the right to rectification (article 16 GDPR) and the right to erasure (article 17 GDPR) exist on paper,

²⁹We have explained in detail elsewhere that systemic, supra-individual solutions and a collective understanding of data protection are needed: Mühlhoff and Ruschemeier (2024c); Mühlhoff and Ruschemeier (2024a); Mühlhoff and Ruschemeier (2024b); Mühlhoff (2023).

but become unenforceable in practice (Novelli et al., 2024).³⁰ Furthermore, removal requests from an individual data subject cannot produce the intended outcome, particularly in cases where the same information has been disseminated by multiple users interacting with the LLM (Brown et al., 2020; Novelli et al., 2024). In essence, simply deleting data from a training dataset offers only a superficial remedy, as it does not guarantee the elimination of the ability to retrieve that data or extract related information embedded within the model's parameters (Novelli et al., 2024). As the output of certain machine learning models is shaped by the data used during the training phase, the original training data or information related to removed data may be deduced or "leaked" (Novelli et al., 2024).

5. Responsibility

In addition to the various steps of data processing in generative models, multiple parties could potentially be considered as data controllers under the GDPR due to their various levels of involvement. The GDPR establishes three categories of responsibility for data processing in relation to the data subject: controller, processor and third parties.

The data controller is primarily responsible for compliance with the provisions of the GDPR (article 5(2)). As defined by article 4(7) GDPR according to which the data controller is a "natural or legal person [...], which, alone or jointly with others, determines the purposes and means of the processing of personal data." Article 4(8) goes on to define the processor as a "natural or legal person [...] to which the personal data are disclosed, whether a third party or not." Third parties on the other hand are actors other than the data subject, controller or processor, article 4(10).

Prima facie, the legal companies that develop and deploy generative models are the data controllers. However, a differentiated picture emerges in the various steps of data processing. Indisputably, companies like OpenAI and Google act as data controllers in relation to the processing steps involved in establishing the parameters for foundational training and storing of the model, given they exclusively determine the modalities of data processing, such as the decision to release a freely accessible LLM. However, in terms of output production, generative models process data based on the prompts from their users. Whether this can be used to establish providers and users as joint controllers within the meaning of article 26 GDPR remains an open question.

Joint controllership under article 26 GDPR refers to the situation where two or more controllers jointly determine the purposes and means of data processing. In contrast, the relationship between a controller and a processor (articles 4 (7,8), 28 GDPR) is different, as in this constellation, the processor processes data *on behalf* of, and subject to the instructions of, the controller. Joint responsibility is thus a relationship of equality, whereas the data processor is subject to the instructions of the data controller. Whether these evaluations can be transferred to the relationship between providers of generative models and users is questionable.

Users are not considered processors within the meaning of article 28 GDPR, as while there is a contract between them and the providers they do not have the obligations of a processor, especially those imposed by article 28(3) GDPR, as they may generate prompts at will, and are not processing data according to instruction.. The purpose of generative models is to enable users to freely use the model for their own defined purposes, free from instructions from the providers.

Users and providers could therefore be joint controllers, but this would require them to jointly define the purposes of data processing, and set out transparent and mutual obligations. This classification is supported by the fact that users and providers both influence the purposes of data processing: the providers of generative models set the basic framework within which their models are used, while users specify the purposes according their individual needs. Consequently, both users and providers are interdependent and have a reciprocal effect on data processing and are also dependent on each

³⁰Novelli et al. (2024) rightfully point out, that invoking the right to erasure (and correction under article 21 GDPR) depends on whether the LLM itself is personal data. This depends whether the training data is anonymised or not, see 3.2.3.

other. However, this is contradicted by the fact that users also tend to be data subjects, and according to article 26(3) GDPR, data subjects have the right to bring claims against any of the joint controllers. Although the law does not require joint controllers hold the same level of responsibility, mere contributory causation without cooperative action is not sufficient for joint responsibility.³¹ Additionally, users' limited influence over data processing means users of generative models cannot effectively be held responsible to third parties as users have no ability to grant rights of access to providers or to delete personal data from the training data.

The relationship between users and providers of generative AI presents therefore a special case that cannot be seamlessly subsumed under the categories of the GDPR. On the one hand, users are more than just data subjects, as their active inputs are required to generate and shape the model's output. On the other hand, they are neither data processors nor joint controllers, as they have no influence over the fundamental modalities of data processing. For instance, providers are able to simply deactivate models or make them subject to fees (as in the case of ChatGPT). The ECJ considers the extent to which data controllers participate in the joint data processing and the specific processing phases to be crucial.³² In this case, users only participate in output generation, which is significantly dependent on the previous steps, such as training. The purpose and aim of the regulations concerning joint controllership is to counteract a diffusion of responsibility among multiple participants. Affected individuals should be able to clearly identify who is collecting their personal data, and for what purpose (recital 58). Therefore, while providers may be responsible for user-generated content in the case of generative models, but the inverse does not apply. This follows from the reasoning and fundamental rights protection of the GDPR provisions regarding responsibility and also corresponds to the technical and economic reality.³³

6. Images, audio and videos as personal data

The considerations for LLMs are not always transferable to generative models that produce audio, images and video. This is because the aim of these models is not to generate information, which may be incorrect, but to generate new audio or visual material. The primary aim, of generating new content, has given rise to many copyright issues arising in these cases. Images and videos can also be personal data if they can be used to identify the person, something easily achieved today, through image searches.

A major problem is the significant increase in deepfakes in the digital context, which now affects not only public figures but also the general population. Women in particular are often victims of deepfake pornography, where explicit images and videos are generated using their images, without their consent. This is an unlawful processing of personal data that violates the GDPR and, in many cases, national regulations. The photographed image of a person in these cases constitutes personal data if the person is still alive, regardless of whether the data are fake or not. The purpose of deepfakes is to disparage or discredit a specific individual, thus fulfilling the decisive characteristic of article 4(1) GDPR, namely that the person is identified or identifiable. Voices may also constitute personal data, if the person is identifiable, visual or acoustic identification methods recorded using pattern recognition, such as facial or voice recognition (speaker recognition), can even be considered biometric data under article 4(14) GDPR (Weichert, 2024).

The latest addition to the data protection cavalry, the AI Act only imposes a labelling obligation on deepfakes (article 50(4)), leaving considerable doubt as to whether there is an adequate level of legal protection at European level (Fallis, 2021; Kumkar & Rapp, 2022; Mania, 2024).

³¹ Case C-2010/16 ECLI:EU:C:2018:388, para. 45.

³² Ibid.

³³ Similiar: Case C- 131/ 12, ECLI:EU:C:2014:317 – Google Spain, para. 33.

7. Conclusion and outlook

The popular use cases of generative AI models show that data protection law is reaching its limits when it comes to regulating data-intensive technologies. In addition to the problems highlighted here, further questions arise regarding the principle of purpose limitation for data processing for data-intensive models and their downstream applications. The use of LLMs in decision-making situations raises questions about the scope of the prohibition in article 22 GDPR.

Structural problems of almost universal concern exist between the GDPR's focus on individual protection and the volume of data processed for training purposes, and in terms of a structural enforcement deficit, particularly regarding data protection principles and data subjects' rights.

As important as the structure of data protection law is for the protection of fundamental rights, new solutions are needed for the structural challenges posed by generative AI and other data-intensive technologies. These may also lie outside data protection law. To address these challenges, it is important to recognise the structural dimension of AI as a socio-technical development (Mühlhoff, 2020). As a result, there is a need for structural solutions that go beyond the enforcement of individual rights (Mühlhoff & Ruschemeier, 2024a, 2024b). Unfortunately, these issues remain unaddressed, as the AI Act does not contribute solutions to remedy the structural and specific challenges for data protection law posed by generative AI. Despite its stated goal of protecting fundamental rights including data protection, the structure of the AI Act follows product safety law parameters and as such has a fundamentally different approach than legal frameworks aimed at protecting fundamental rights, as found in the GDPR. The AI Act does establish certain obligations for high-risk AI-systems such as data governance, transparency requirements and standards for human oversight. However, these provisions do not address the privacy and data protection of users. This leaves a great need for legal regulation.

Competing interests. The author declares none.

References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., & Zhang, L. (2016). Deep learning with differential privacy. In Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security - CCS (pp. 308–318).
- Ben-Shahar, O., & Schneider, C. E. (2011). The failure of mandated disclosure. *University of Pennsylvania Law Review*, 159, 647–749.
- Borgesius, F. J. Z., Kruikemeier, S., Helberger, N., & Helberger, N. (2017). Tracking walls, take-it-or-leave-it choices, the GDPR, and the ePrivacy regulation. *European Data Protection Law Review*, 3(3), 353–368.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... Amodei, D. (2020). Language models are few-shot learners. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (pp. 1877–1901).
- Carlini, N., Hayes, J., Nasr, M., Schwag, V., Tramèr, F., Balle, B., ... Wallace, E. (2023). *Extracting Training Data from Diffusion Models*.
- Carlini, N., Ippolito, D., Jagielski, M., Lee, K., Tramer, F., & Zhang, C. (2022). Quantifying memorization across neural language models. *arXiv preprint*. doi:10.48550/arXiv.2202.07646
- Charlotin, D. (2023). Large language models and the future of law. *SSRN Electronic Journal*, Retrieved from <https://www.ssrn.com/abstract=4548258>.
- Chen, S. (2024). Potential applications and safety of large language models in healthcare. *Interdisciplinary Humanities and Communication Studies*, 1(6), 1–5.
- Chiara, P. G. (2023). Italy: Italian DPA v. OpenAI's ChatGPT: The reasons behind the investigation and the temporary limitation to processing. *European Data Protection Law Review*, 9(1), 68–72.
- CommissioN Nationale de l'Informatique et des Libertés (CNIL) (2023, May 16). Artificial intelligence: The action plan of the CNIL. Retrieved from <https://www.cnil.fr/en/artificial-intelligence-action-plan-cnil>.
- De, S. J., & Imine, A. (2020). Consent for targeted advertising: The case of Facebook. *AI & SOCIETY*, 35(4), 1055–1064.
- Donnelly, M., & McDonagh, M. (2019). Health research, consent and the GDPR exemption. *European Journal of Health Law*, 26(2), 97–119.
- Dwork, C. (2006). Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone & I. Wegener (Eds.), *Automata, Languages and Programming: 33rd International Colloquium, ICALP, Venice, Italy, July 10–14, 2006, Proceedings, Part II*, (pp. 1–12). Springer.

- Eichenhofer, J.** (2021). *e-Privacy. Theorie Und Dogmatik Eines Europäischen Privatheitsschutzes Im Internet-Zeitalter: Jus Publicum*, vol. 301.
- El Emam, K., & Alvarez, C.** (2014). A critical appraisal of the article 29 working party opinion 05/2014 on data anonymization techniques. *International Data Privacy Law*, 5(1), 73–87.
- El Emam, K., Rodgers, S., & Malin, B.** (2015). Anonymising and sharing individual patient data: *BMJ (Clinical research ed.)*. *BMJ (Clinical Research Ed.)*, 350, h1139.
- El-Mhamdi, E.-M., Farhadkhani, S., Guerraoui, R., Gupta, N., Hoang, L.-N., Pinot, R., ... Stephan, J.** (2023). *On the impossible safety of large AI models*.
- European Data Protection Board (EDPB)** (2023, April 13). *EDPB resolves dispute on transfers by meta and creates task force on chat GPT*. Retrieved from https://www.edpb.europa.eu/news/news/2023/edpb-resolves-dispute-transfers-meta-and-creates-task-force-chat-gpt_en.
- Fallis, D.** (2021). The epistemic threat of deepfakes. *Philosophy & Technology*, 34, 623–643.
- Garante per la protezione dei dati personali (GPDP)** (2023, March 31). *Intelligenza artificiale: Il Garante blocca ChatGPT. Raccolta illecita di dati personali. Assenza di sistemi per la verifica dell'età dei minori*. Retrieved from <https://www.garanteprivacy.it:443/home/docweb/-/docweb-display/docweb/9870847>.
- Garante per la protezione dei dati personali (GPDP)** (2024, January 29). *ChatGPT: Garante privacy, notificato a OpenAI fatto di contestazione per le violazioni alla normativa privacy*. Retrieved from <https://gdpd.it/web/guest/home/docweb/-/docweb-display/docweb/9978020>.
- Gellert, R., & Gutwirth, S.** (2013). The legal construction of privacy and data protection. *Computer Law & Security Review*, 29(5), 522–530.
- Georgieva, L., & Kuner, C.** (2020). Art. 9 Processing of special categories of personal data. In C. Kuner, L. A. Bygrave, C. Docksey & L. Drechsler (Eds.), *The EU General Data Protection Regulation (GDPR)* (pp. 365–384). Oxford University Press.
- Gillotte, J. L.** (2020). Copyright infringement in AI-generated artworks. *UC Davis Law Review*, 53(5), 2655–2691..
- Gupta, M., Akiri, C., Aryal, K., Parker, E., & Praharaj, L.** (2023). From chatGPT to threatGPT: Impact of generative AI in cybersecurity and privacy. *IEEE Access*, 11, 80218–80245.
- Hacker, P.** (2021). A legal framework for AI training data—from first principles to the Artificial Intelligence Act. *Law, Innovation and Technology*, 13(2), 257–301.
- Hacker, P.** (2023, December 13). *What's missing from the EU AI Act: Addressing the four key challenges of large language models*. *Verfassungsblog*. Retrieved from <https://verfassungsblog.de/whats-missing-from-the-eu-ai-act>.
- Harlan, E., & Brunner, K.** (2023, July 7). *We are all raw material for AI*. Retrieved from <https://interaktiv.br.de/ki-trainingsdaten/en>.
- Hazell, J.** (2023, December 14). *Spear phishing with large language models*. Retrieved from <https://www.governance.ai/research-paper/llms-used-spear-phishing>.
- Hondrich, L., & Ruschmeier, H.** (2023). Addressing automation bias through verifiability. In *EWAF'23: European Workshop on Algorithmic Fairness*.
- Hornung, G., & Wagner, B.** (2020). *Anonymisierung als datenschutzrelevante verarbeitung? Rechtliche anforderungen und grenzen für die anonymisierung personenbezogener daten*. *ZD*, 223–228. Retrieved from https://www.bfdi.bund.de/SharedDocs/Downloads/DE/Konsultationsverfahren/1_Anonymisierung/Stellungnahmen/Universit%C3%A4t-Kassel.pdf?__blob=publicationFile&v=5.
- Huang, K., Wang, Y., & Zhang, X.** (2024). Foundations of Generative AI. In K. Huang, Y. Wang, B. Goertzel, Y. Li, S. Wright & J. Ponnappalli (Eds.), *Generative AI Security: Theories and Practices* (pp. 3–30). Springer.
- Kokott, J., & Sobotta, C.** (2013). The distinction between privacy and data protection in the jurisprudence of the CJEU and the ECtHR. *International Data Privacy Law*, 3(4), 222–228.
- Krönke, C.** (2023, April 14). *Attention is all you need: ChatGPT und die DSGVO*. *Verfassungsblog*. Retrieved from <https://verfassungsblog.de/attention-is-all-you-need>.
- Kumkar, L. K., & Rapp, J. P.** (2022). Deepfakes: Eine herausforderung für die rechtsordnung. *Zeitschrift Für Digitalisierung Und Recht (Zfdr)*, 3, 199–228.
- Landesbeauftragte für Datenschutz und Informationssicherheit NRW (LDI NRW)** (2023, October 30). *Prüfung von ChatGPT geht in die nächste Runde*. Retrieved from <https://www.badische-zeitung.de/der-neue-sc-freiburg-coach-julian-schuster-kickt-beim-trainingsaufakt-direkt-selbst-mit>.
- Mania, K.** (2024). Legal protection of revenge and deepfake porn victims in the European Union: Findings from a comparative legal study. *Trauma, Violence, & Abuse*, 25(1), 117–129.
- Marcos, H., & Pullin, M.** (2023, October 11). Large language models and EU data protection: Mapping (some) of the problems — The digital constitutionalist (October 2023). Retrieved from <https://digi-con.org/large-language-models-and-eu-data-protection-mapping-some-of-the-problems>.
- Matejek, M., & Mäusezahl, S.** (2019). Gewöhnliche vs. sensible personenbezogene Daten. Abgrenzung und Verarbeitungsrahmen von Daten gem. Art. 9 DS-GVO: *ZD*, 551–556.

- Metz, C.** (2023, November 6). Chatbots may “Hallucinate” More Often Than Many Realize. Retrieved from <https://www.nytimes.com/2023/11/06/technology/chatbots-hallucination-rates.html>.
- Mühlhoff, R.** (2020). Human-aided artificial intelligence: Or, how to run large computations in human brains? Toward a media sociology of machine learning. *New Media & Society*, 22(10), 1868–1884.
- Mühlhoff, R.** (2023). Predictive privacy: Collective data protection in times of AI and big data. *Big Data & Society*, 10, 1–14. doi:10.1177/20539517231166886
- Mühlhoff, R., & Ruschemeier, H.** (2024a). Predictive analytics and the collective dimensions of data protection. *Law, Innovation and Technology*, 16(1), 261–292.
- Mühlhoff, R., & Ruschemeier, H.** (2024b). Regulating AI with purpose limitation for models. *Journal of AI Law and Regulation*, 1(1), 24–39.
- Mühlhoff, R., & Ruschemeier, H.** (2024c). Predictive analysis und DSGVO. *Kommunikation & Recht (K&R)*, 12, 15–26.
- Nasr, M., Carlini, N., Hayase, J., Jagielski, M., Cooper, A. F., Ippolito, D., ... Lee, K.** (2023). Scalable extraction of training data from (production) language models.
- Nguyen, T. T., Backes, M., & Stock, B.** (2022). Freely given consent? Studying consent notice of third-party tracking and its violations of GDPR in android apps. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, (New York, NY, USA: Association for Computing Machinery, 2022) (pp. 2369–2383).
- Nissenbaum, H.** (2011). A contextual approach to privacy online. *Daedalus*, 140(4), 32–48.
- None of Your Business (NOYB) – Europäisches Zentrum für digitale Rechte** (2024, April 29). *ChatGPT provides false information about people, and OpenAI can't correct it*. Retrieved from <https://noyb.eu/en/chatgpt-provides-false-information-about-people-and-openai-cant-correct-it>.
- Novelli, C., Casolari, F., Hacker, P., Spedicato, G., & Floridi, L.** (2024). *Generative AI in EU law: Liability, privacy, intellectual property, and cybersecurity*. Retrieved from <https://papers.ssrn.com/abstract=4694565>.
- OpenAI** (2023, December 15). *Datenschutzerklärung*. Retrieved from <https://openai.com/de/policies/eu-privacy-policy>.
- Pesch, P., & Böhme.** (2023). Verarbeitung personenbezogener daten und datenrichtigkeit bei großen sprachmodellen. *Multimedia Und Recht*, 16(12), 917–923.
- Roßnagel, A.** (2021). Datenlöschung und anonymisierung. Verhältnis der beiden datenschutzinstrumente nach der DSGVO. *ZD*, 11, 188–192.
- Ruscheimer, H.** (2023a, April 7). *Squaring the circle*. Retrieved from <https://verfassungsblog.de/squaring-the-circle>.
- Ruscheimer, H.** (2023b). The problems of the automation bias in the public sector – A legal perspective. In W. Institute (Ed.), *Weizenbaum Conference Proceedings: AI, Big Data, Social Media and People on the Move, (Rochester, NY, 2023b)* (pp. 59–69).
- Ruscheimer, H., & Hondrich, L.** (2024). Automation bias in public administration - an interdisciplinary perspective from law and psychology. *Government Information Quarterly*, 41(3), 101953.
- Sag, M.** (2023). Copyright safety for generative AI. *Houston Law Review*, 61(2), 295.
- Schreurs, W. J., Hildebrandt, M., Kindt, E., & Michaël, V.** (2008). Cogitas, Ergo Sum. The role of data protection law and non-discrimination law in group profiling in the private sector. In M. Hildebrandt & S. Gutwirth (Eds.), *Profiling the European Citizen* (pp. 241–270). Springer.
- Schulz, S.** (2018). Art. 9 DS-GVO, para. 13.
- Solove, D. J.** (2023). The limitations of privacy rights. *Notre Dame Law Review*, 98, 975.
- Solove, D. J.** (2024). Artificial Intelligence and Privacy.
- Urząd ochrony danych osobowych (UODO)** (2023, September 9). Technologia musi być zgodna z Rodo. Retrieved from <https://uodo.gov.pl/pl/138/2823>.
- Wachter, S., & Mittelstadt, B.** (2019). A right to reasonable inferences: Re-thinking data protection law in the age of big data and AI. *Columbia Business Law Review*, 2019(2), 494–620.
- Wachter, S., Mittelstadt, B., & Russell, C.** (2024). Do large language models have a legal duty to tell the truth?
- Weichert, T.** (2024). DS-GVO, Art. 4 Nr. 14 biometrische Daten, marg. 3.
- Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., ... Gabriel, I.** (2022). Taxonomy of risks posed by language models. In *2022 ACM Conference on Fairness, Accountability, and Transparency, (Seoul Republic of Korea: ACM)* (pp. 214–229).
- Zanfir-Fortuna, G.** (2023, September 12). *How data protection authorities are de facto regulating generative AI - Future of privacy forum*. Retrieved from <https://fpf.org/blog/how-data-protection-authorities-are-de-facto-regulating-generative-ai>.
- Zarsky, T. Z.** (2017). Incompatible: The GDPR in the age of big data. *Seton Hall Law Review*, 47, 995–1020.