GENERATIVE AI AND DATA SUBJECT RIGHTS UNDER TURKISH DATA PROTECTION LAW

Assist. Prof. Dr Osman Gazi GÜÇLÜTÜRK*

1. Introduction

Generative Artificial Intelligence (**AI**) has established itself as a groundbreaking innovation in contemporary technology, demonstrating the ability to autonomously generate diverse forms of content, including but not limited to textual narratives, visual imagery, and computer code. This transformative capability has not only expanded the horizons of creativity and productivity but has also introduced a myriad of ethical, legal, and regulatory challenges. As generative AI systems, particularly Large Language Models (**LLMs**), are increasingly integrated into various industries—such as healthcare, education, finance, and entertainment—they bring into focus critical issues regarding their alignment with data protection frameworks and the safeguarding of individual rights.

This paper aims to explore the complex interplay between generative AI systems and data protection law, specifically within the context of Turkish data protection legislation, primarily Law No. 6698 on the Protection of Personal Data (**TDPL**)[1]. By examining the potential conflicts and synergies between generative AI and the rights of data subjects, it seeks to illuminate the challenges and considerations involved in ensuring compliance with applicable legal rules, principles, and the decisions of the Turkish Data Protection Board, along with drawing some analogies between the TDPL and the European Union's General Data Protection Regulation (**GDPR**)[2]. Particular emphasis is placed on understanding how fundamental data

* Galatasaray University, Faculty of Law, Department of IT Law. ogucluturk@gsu.edu.tr. This article is prepared as a part of the roundtable and article series, co-organised by Meta and Galatasaray University, and finalised after the roundtable discussion took place at Galatasaray University on 5 Dcemver 2024. All URLs in this paper are accessed on 20 December 2024, unless explicitly stated otherwise.

[1] Law No. 6698 on the Protection of Personal Data. The official Turkish version of the law can be found at https://mevzuat.gov.tr/mevzuat?MevzuatNo=6698&MevzuatTur=1&MevzuatTertip=5. An unofficial translation of the initial version of the law by the Turkish Data Protection Authority can be found at: https://www.kvkk.gov.tr/Icerik/6649/Personal-Data-Protection-Law .

[2] Regulation (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), https://eur-lex.europa.eu/eli/reg/2016/679/oj/eng.

subject rights—such as the right to access, rectify, or erase personal data—can be meaningfully exercised in the context of generative AI.

Additionally, the paper also examines practical examples of these challenges, such as ensuring transparency in data processing by LLMs, mitigating the risks of unintentional data retention or misuse, and addressing potential conflicts between innovation and individual rights. By grounding the discussion in the Turkish legal landscape, the paper offers a comprehensive and context-specific analysis that also has broader implications for global data protection practices. Naturally, before proceeding with the details on these points, a brief introduction is required on what generative AI is and how it functions.

2. The Concept of Generative AI

While there is no universally accepted legal definition of generative AI, it can broadly be described as algorithms designed to produce new content by identifying and replicating patterns from existing data[3]. Unlike traditional AI systems, which are primarily designed for classification or prediction tasks, generative models[4] or systems excel in creating seemingly original and contextual outputs that closely resemble the data they were trained on.

The concept of an algorithm, or a piece of software, generating content is not inherently novel. While the public perception of AI has evolved significantly over the past few decades, the foundational concept of AI has existed for over 50 years. Earlier iterations of today's AI applications, such as primitive chatbots or basic translation tools, were also capable of generating some form of content. What distinguishes modern generative AI—particularly advanced models such as LLMs—is not the fundamental function of content generation but rather the remarkable quality, coherence, and sophistication of the outputs. For instance, generative image models today can produce visuals that are virtually indistinguishable from

[3] OECD defines generative AI with a reference to AI as "*a category of AI that can create new content such as text, images, videos, and music*". Organisation for Economic Co-operation and Development (OECD), 'Generative AI' https://www.oecd.org/en/topics/sub-issues/generative-ai.html.

[4] There is a difference between AI systems and models, and these must not be used interchangably. The term *model* refers to the software framework embodying the rules and patterns learned during the training phase. While not all AI systems are model-based, the most advanced ones almost invariably rely on such structures. For further details, see Osman Gazi Güçlütürk, *Yapay Zeka ve Verinin Kullanımı* (On İki Levha Yayıncılık, 2022) 66 fn 168.

those created by human artists[5], while advanced LLMs can engage in conversations with a level of fluency that mimics human interaction[6].

The current discourse on generative AI centres primarily on these advanced applications and their enhanced generation capabilities. This focus is driven by the paradigm shift in deep learning, reinforcement learning, and other modern AI techniques. Although the technical intricacies of these advancements are beyond the scope of this paper[7], it is critical to note that the quality of content produced by a generative model is highly dependent on the quality and volume of the data used during training. To develop a model capable of engaging in human-like conversation, vast datasets comprising human-generated texts and dialogues are required[8]. Inevitably, such datasets often include fragments of personal data, raising significant privacy and data protection concerns.

Generative AI systems can be designed to produce a wide range of outputs, including articles, summaries, conversational responses, realistic images, design prototypes, and musical compositions. The specific type of content a model aims to generate significantly influences its development phase, including data collection processes, algorithm design, and the associated legal and regulatory considerations. For example, text-based LLMs inherently present more significant user-level privacy risks[9] compared to general-purpose image or audio generation models. This heightened risk arises from the nature of human conversations, which tend to be customised and often require users to input personal information. Even when there is no direct conversation, the tasks performed by LLMs frequently necessitate extensive textual input, increasing the likelihood of processing personally identifiable information.

Given these risks and the widespread adoption of LLMs, this paper focuses on these models to analyse their unique privacy implications. Unlike image or code generation tools, LLMs

[5] For a detailed analysis on whether humans can distinguish human-generated art from AI-generated art,see Anna Yoo Jeong Ha and others, 'Organic or Diffused: Can We Distinguish Human Art from AI-generated Images?' (Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security, Salt Lake City, UT, USA, 2024) https://doi.org/10.1145/3658644.3670306.

[6] On LLMs' capacity to engage in human-like conversation, see J Ou and others, 'DialogBench: Evaluating LLMs as Human-like Dialogue Systems' (2023) ArXiv https://arxiv.org/abs/2311.01677.

[7] For an analysis on how reinforcement learning is being used in AI, see Y Cao and others, 'Reinforcement Learning for Generative AI: A Survey' (2023) ArXiv https://arxiv.org/abs/2308.14328.

[8] See Interface Media, 'Big Data Isn't Big Enough to Train Generative AI' (6 March 2024) https://interface.media/blog/2024/03/06/big-data-isnt-big-enough-to-train-generative-ai/#:~:text=%E2%80%9CThe%20development%20and%20effectiveness%20of,billion%20words%E2%80%94to%20train%20ChatGPT. Also see Wayne Xin Zhao and others, 'A Survey of Large Language Models' (2024) ArXiv https://arxiv.org/abs/2303.18223.

[9] See Yifan Yao and others, 'A Survey on Large Language Model (LLM) Security and Privacy: The Good, The Bad, and The Ugly' (2024) 4(2) High-Confidence Computing 100211 https://doi.org/10.1016/j.hcc.2024.100211.

operate in conversational formats, encouraging users to disclose personal data more readily. Understanding the risks and implications of such advanced models requires examining how content generation operates[10].

At the core of generative AI models like LLMs are sophisticated algorithms and extensive datasets. These models undergo a multi-phase training process that includes, among others, data ingestion in the sense that the assimilation of vast amounts of text and, in some cases, images from diverse sources, learning as in neural networks identifying patterns, statistical relationships, and linguistic structures within the data, and refinement, which means that outputs are improved through techniques like reinforcement learning and user feedback.

LLMs, such as GPT-4, generate text by predicting the next word in a sequence based on contextual cues, enabling them to produce coherent and contextually appropriate responses[11]. Their reliance on large and varied datasets allows them to capture linguistic nuances, including colloquialisms, idiomatic expressions, and cultural references. However, it is critical to emphasise that despite their human-like conversational abilities and features such as humour or simulated "thinking," generative AI models do not "think" as humans do[12]. Their operation is purely statistical, devoid of logic, reasoning, or judgment.

This fundamental distinction between human cognition and generative AI processes carries significant implications. When interacting with generative AI, users may anthropomorphise these systems, assuming they possess human-like capabilities, empathy, or intentions[13]. This misunderstanding can lead to a failure to recognise the potential privacy risks associated with AI's high capacity to process, analyse, and infer data. For instance, users may inadvertently disclose sensitive information, underestimating the model's data retention and processing capabilities.

Taking these elements into account, this paper examines the interaction between LLMs and human users, highlighting the enhanced privacy risks associated with generative AI. By

[10] See Zhiping Zhang and others, '"It's a Fair Game", or Is It? Examining How Users Navigate Disclosure Risks and Benefits When Using LLM-Based Conversational Agents' (Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, Honolulu, HI, USA, 2024) https://doi.org/10.1145/3613904.3642385.
[11] For more information on how LLMs work, see Giovanni Briganti, 'How ChatGPT Works: A Mini Review' (2024) 281(3) European Archives of Oto-Rhino-Laryngology 1565 https://doi.org/10.1007/s00405-023-08337-7.
[12] For a comparative study on human and LLM problem-solving skills, see Yufei Tian and others, 'Thinking Out-of-the-Box: A Comparative Investigation of Human and LLMs in Creative Problem-Solving' (ICML 2024 Workshop on LLMs and Cognition, 2024) https://openreview.net/forum?id=rxkqeYHXy0.
[13] For detailed information on AI and anthropomorphism, see driana Placani, 'Anthropomorphism in AI: Hype and Fallacy' (2024) 4(3) AI and Ethics 691 https://doi.org/10.1007/s43681-024-00419-4.

addressing these considerations, it seeks to provide a nuanced understanding of the implications of advanced AI systems in both technical and regulatory contexts.

3. Generative AI and Its Interplay with Data Protection Rules: Identification of Additional Risks

This section starts exploring the substance of the discussion by focusing on the examination of the interaction between generative AI and data protection rules, identifying unique risks that arise from this relationship. It is essential first to clarify that not all data processed by or for generative AI systems constitutes personal data. Accordingly, data protection rules are only one aspect of the broader regulatory framework governing generative AI. Other significant regulatory considerations include intellectual property concerns, such as copyright and trade secrets, which often intersect with the development and application of these systems.

Moreover, processing personal data by or for generative AI is fundamentally a subset of data processing by or for AI systems in general. Thus, the broader concerns, principles, and evaluations applicable to AI-related data processing are equally relevant to generative AI. Under the TDPL, the applicability of data protection rules to generative AI hinges on whether the datasets involved contain personal data.

To comply with the TDPL, generative AI-related data processing must adhere to several foundational principles and legal requirements, including:

1. **General Principles (Article 4)**: Data processing must be lawful, fair, and conducted transparently. It should also respect the principles of purpose limitation, data minimisation, accuracy, and accountability.
2. **Legal Grounds for Processing (Article 5 and Article 6)**: The processing of personal data must be based on one of the lawful grounds outlined in the law, such as the explicit consent of the data subject, performance of a contract, compliance with legal obligations, or protection of vital interests. For sensitive data, stricter requirements under Article 6 apply.
3. **Data Retention and Destruction (Article 7)**: Personal data must be destroyed when the purpose of processing is achieved, or the lawful basis ceases to exist.

4. **Domestic and Cross-Border Data Transfers (Articles 8 and 9)**: Data transfers must comply with specific rules, including obtaining consent for international transfers unless an exception applies.
5. **Disclosure Obligations (Article 10)**: Data controllers must provide clear and transparent information to data subjects regarding the purpose, legal basis, and methods of data processing.
6. **Data Subject Rights (Article 11)**: Data subjects have the right to access, rectify, erase, and object to the processing of their personal data, among others.
7. **Data Security Measures (Article 12)**: Controllers must implement appropriate technical and administrative measures to ensure data security and protect against unauthorised access, alteration, or loss.

These conditions are cumulative, and the absence of any of them renders a personal data processing activity unlawful. Additionally, there is an obligation for data controllers to register with the National Data Controller Registry (**VERBIS**)[14] if they meet certain criteria or engage in specific types of personal data processing activities.

Although the above requirements apply universally to personal data processing, the unique characteristics of generative AI introduce additional considerations. This paper narrows its focus to two specific aspects of this intersection:

1. **Additional Privacy and Data Protection Risks Associated with Generative AI**: Generative AI systems, particularly LLMs, often involve the large-scale ingestion and processing of text-heavy datasets that may include fragments of personal data. These risks are amplified by the interactive nature of generative AI, which encourages users to input potentially sensitive information. Risks such as unintended data retention, inferences about personal characteristics, and the potential misuse of outputs necessitate close scrutiny. These risks must be considered while training an LLM or developing a generative AI system.

[14] VERBIS is a publicly accessible registry available at https://verbis.kvkk.gov.tr/. As a general rule, registration with VERBIS is mandatory for all data controllers, although the Board has the authority to introduce exemptions. In addition to sector- and profession-specific exemptions, the Board has established two cumulative criteria—based on the number of employees and annual turnover—that may exempt a data controller from the registration requirement. For more information, see Ömer Ekmekçi and others, *Kişisel Verilerin Korunması Hukuku* (On İki Levha Yayıncılık 2024), p. 219. For the decision establishing the most recent threshold see https://www.resmigazete.gov.tr/eskiler/2023/07/20230725-5.pdf.

2. **Interplay Between Data Subject Rights and Generative AI**: Generative AI systems pose challenges to the effective exercise of data subject rights. For instance, ensuring transparency and accountability in data processing by LLMs can be complex due to the opacity of their underlying algorithms and training processes. Additionally, the right to erasure or rectification may be challenging to implement when personal data is incorporated into large, unstructured datasets. These concerns, on the other hand, must be taken into account while engaging with the users and facilitating them to exercise their data subject's rights to the extent that they are applicable.

By focusing on these two dimensions, this paper seeks to illuminate the unique challenges generative AI poses to privacy and data protection frameworks, with particular attention to the issues that must be taken into account during the training and deployment phase by the controllers on the one hand and implications for data subject rights under the TDPL on the other.

4. Additional Privacy Risk Posed by Generative AI and Potential Mitigations

The functioning of generative AI, particularly LLMs and LLM-powered systems, introduces unique privacy risks that distinguish these technologies from other AI applications. These risks arise from both the design of these systems and the behavioural patterns they encourage in their users, making them a complex area of concern for data protection.

4.1. Privacy Risks Stemming from Training Data

Generative AI models are typically trained on vast datasets, which often include personal data—either directly or in fragmented forms. These datasets may originate from publicly available sources, proprietary repositories, or even web scrapes, often without the explicit consent of the data subjects involved. While the aim is to create a model capable of understanding patterns and producing coherent outputs, the sheer volume of processed data heightens the likelihood of unintentional inclusion of sensitive or personally identifiable information[15].

[15] See Xiaodong Wu, Ran Duan, and Jianbing Ni, 'Unveiling Security, Privacy, and Ethical Concerns of ChatGPT' (2024) 2(2) Journal of Information and Intelligence 102.

For example, an LLM trained on datasets containing online forums, emails, or publicly posted documents may inadvertently retain recognisable fragments of actual conversations or personal identifiers. Even anonymised data poses risks if the model, through its advanced processing, reconstructs or infers identities by combining disparate data points.

4.2. Privacy Risks in Interaction and Input

Once trained, the practical use of generative AI often requires continuous input from users. This interaction is more sophisticated and contextual than other AI systems, as LLMs are designed to generate responses that are tailored to the nuances of a user's input.

In a chatbot interface, for instance, users often provide detailed and personalised information to achieve more customised and relevant outputs. For example:

- A user requesting an LLM to draft a professional email might input real names, job titles, or sensitive corporate details.
- Another user seeking advice on health or legal matters may inadvertently share personal health data or confidential information.

These interactions create a feedback loop where the system, by design, encourages further personalisation. This phenomenon not only increases the volume of personal data processed but also makes the data more identifiable and sensitive over time.

4.3. Behavioral Nudging and Design Features

A unique aspect of generative AI systems, particularly in chatbot form, is their ability to emulate human-like interaction. Design features such as conversational tone, the use of colloquialisms, artificial "thinking" pauses, and even voice-enabled interactions enhance the illusion of a human counterpart. Users may begin to view these systems as trustworthy confidants, sharing intimate details they might not reveal to a real person they have just met[16].

For example:

[16] See N Mireshghallah and others, 'Trust No Bot: Discovering Personal Disclosures in Human-LLM Conversations in the Wild' (2024) ArXiv https://arxiv.org/abs/2407.11438.

- A voice-enabled LLM may mimic natural conversational flow, complete with hesitations or affirmations, making users feel they are engaging with a thoughtful and empathetic individual.
- A text-based chatbot equipped with humour or expressions of concern ("I understand how you feel") may encourage users to disclose emotionally charged or sensitive information.

Such behavioural nudges amplify privacy risks by lowering users' guard and promoting the sharing of personal details.

4.4. Profiling and Data Linkage

Another significant risk arises from the aggregation of user inputs over time. Information fragments provided across multiple interactions can be linked to create a comprehensive user profile. Even seemingly innocuous details, when combined, can reveal sensitive insights about a person's identity, preferences, or behaviours[17]. For instance:

- A user seeking travel recommendations across different sessions may disclose their location, favourite destinations, and travel habits, which can collectively form a detailed personal profile.
- A user discussing their daily routine in various contexts may inadvertently reveal details about their work schedule, family dynamics, and leisure activities.

This ability to link inputs into coherent profiles makes LLMs inherently more prone to privacy violations compared to other AI systems.

4.5. Potential Mitigation Measures

Considering the practical use of LLMs in light of the additional risks mentioned above, it can be stated that generative AI users often treat interactions with LLMs more casually than they do with formal systems or even with other humans. The perceived anonymity and the human-like nature of these tools lower the psychological barriers to sharing sensitive information. For

[17] Such linkages may result in or faciliate reidentification. For more information, see A Nyffenegger, M Stürmer, and J Niklaus, 'Anonymity at Risk? Assessing Re-Identification Capabilities of Large Language Models' (2023) ArXiv https://arxiv.org/abs/2308.11103.

example, a user might experiment by typing hypothetical personal dilemmas or use the system to draft sensitive communications, unaware that their input could be stored or analysed.

This behaviour underscores the need for clear, proactive measures to mitigate privacy risks. Addressing the privacy risks posed by generative AI, particularly LLMs, requires a multi-faceted approach. Given the flexibility and diversity of AI systems, effective mitigation measures must be tailored to their specific technical characteristics, operational context, and associated risks. The foundational step in this process is conducting a comprehensive risk assessment before any AI system begins processing personal data[18]. Such assessments help identify the unique challenges and vulnerabilities of a given system, allowing for the development of targeted and effective mitigation mechanisms.

A significant challenge lies in the technical limitations of current advanced AI models, particularly their inability to provide clear interpretability or explainability regarding how individual data points influence outputs. This opacity, sometimes referred to as the "black box" problem, makes it nearly impossible to trace the impact of specific data points once they are incorporated into a model. Consequently, regulatory efforts should focus on the entry points of data processing—primarily the data collection and pre-training phases. This approach is reflected in frameworks such as the EU AI Act[19], which emphasises data governance and documentation requirements to ensure transparency and accountability from the outset.

4.5.1. Measures Targeting Data Collection and Pre-Training

Given that data collection and pre-training are critical phases in the lifecycle of generative AI, strict controls must be implemented to mitigate risks at this stage. First, data sourcing must prioritise selectivity and ethical considerations. For instance, only data from publicly accessible and reputable sources should be used, avoiding the scraping of websites requiring user authentication or payment, as these often contain sensitive or private information. Additionally, websites known to host personal data should be excluded from the data collection process unless explicit consent is obtained. In cases where personal data is processed, particularly for systems or models that are not explainable by design, data subjects may be offered customised

[18] It must be noted that there is no mandatory risk or impact assessment under the TDPL that may be considered comparable to data protection impact assessments under the GDPR.

[19] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng.

disclosure texts about the nature of the data processing or the choice of selecting types of personal data that they want to keep out from the training data

Data deduplication techniques[20] should also be employed rigorously during pre-training. By identifying and eliminating duplicate entries in training datasets, the risk of the model memorising specific personal information is reduced. For example, if a training dataset includes multiple instances of the same sensitive email address or contact detail, deduplication ensures that such information does not become disproportionately embedded in the model's parameters.

4.5.2. Addressing User Behaviour and Design Risks

As previously explained, the conversational nature of LLMs, coupled with human-like design features, nudges users to disclose personal information. To counteract this, AI systems should be designed with privacy-preserving mechanisms that reduce the likelihood of users unknowingly sharing sensitive details. One approach is to incorporate real-time input monitoring that detects and flags potentially sensitive or identifiable information during user interactions. For instance, if a user enters a credit card number or personal address, the system could issue a warning, discouraging the submission of such data.

Transparency and user education are equally critical. Users must be clearly informed about how their data is processed and the potential privacy implications of their interactions. This can be achieved through concise and accessible disclaimers integrated into the user interface, as well as detailed privacy policies that outline data handling practices. Furthermore, anonymisation techniques can be applied to inputs and outputs to ensure that identifiable information is not retained or displayed.

4.5.3. Enhancing Accountability and Oversight

Mitigation measures must also extend to ongoing oversight and accountability mechanisms. Regular audits of AI models and their training data can help identify and address potential privacy risks that emerge after deployment. For instance, audits can detect whether the model

[20] See Kushal Tirumala and others, 'D4: Improving LLM Pretraining via Document De-Duplication and Diversification' in A Oh and others (eds), *Advances in Neural Information Processing Systems*, vol 36 (Curran Associates, Inc. 2023) 53983 https://proceedings.neurips.cc/paper_files/paper/2023/file/a8f8cbd7f7a5fb2c837e578c75e5b615-Paper-Datasets_and_Benchmarks.pdf.

unintentionally memorises sensitive information and implement corrective measures, such as retraining or fine-tuning with privacy-focused modifications.

Organisations deploying generative AI should establish robust governance frameworks that include clearly defined roles and responsibilities for data protection. This may involve appointing data protection officers to oversee compliance with applicable laws and standards, as well as implementing internal policies to ensure adherence to data minimisation principles.

4.5.4. Balancing Learning Capabilities and Privacy

The challenge lies in balancing the learning capabilities of generative AI models with the need to protect privacy. For example, while reducing dataset diversity might limit privacy risks, it could also diminish the model's ability to generate nuanced and contextually rich outputs. A potential compromise is the use of synthetic data—data generated artificially to simulate real-world patterns—during training. While synthetic data cannot fully replace real-world data, it can reduce dependence on sensitive information and enhance privacy safeguards[21].

In sum, addressing the privacy risks associated with generative AI requires a combination of proactive data governance, user-focused design, robust oversight, and innovative techniques like synthetic data. These measures not only mitigate risks but also ensure that generative AI systems operate within ethical and legal boundaries, fostering trust among users and stakeholders alike.

5. Interplay Between Data Subject Rights and LLMs

5.1. Overview of Data Subject Rights under the Turkish Data Protection Law

Data subject rights, as outlined in Article 11 of the TDPL, provide data subjects whose personal data are processed with a comprehensive framework for controlling the processing of their personal data. These rights include[22]:

1. The right to learn whether personal data is being processed.

[21] For more information o using syntetic data see T Marwala and S Stinckwich, 'The Use of Synthetic Data to Train AI Models: Opportunities and Risks for Sustainable Development' (2023) ArXiv https://arxiv.org/abs/2309.00652.

[22] For more information on data subject rights in general under the TDPL, see Ömer Ekmekçi and others, *Kişisel Verilerin Korunması Hukuku* (On İki Levha Yayıncılık, 2024).

2. The right to request information about the processing of personal data.

3. The right to understand the purpose of personal data processing and whether it aligns with the stated objectives.

4. The right to know the third parties, both domestic and international, to whom personal data is transferred.

5. The right to request the correction of incomplete or inaccurate personal data.

6. The right to request the erasure or destruction of personal data under the conditions specified in Article 7 of the TDPL.

7. The right to request that operations regarding correction and erasure of personal data be reported to third parties to whom the data has been transferred.

8. The right to object to decisions based solely on automated data processing that produce results detrimental to the individual.

9. The right to seek compensation for damages caused by unlawful personal data processing.

5.2. An Overview of Challenges Posed by LLMs in the Exercise of Data Subject Rights

The integration of generative AI, particularly LLMs, complicates the effective exercise of these rights in several ways. While the fundamental principles governing these rights remain unchanged, the technical characteristics and operational complexities of LLMs create new obstacles that demand careful examination.

For instance, the purpose of processing personal data is a cornerstone of transparency and accountability[23]. However, for a general-purpose generative AI model, defining or communicating this purpose can be inherently challenging. Unlike narrowly tailored AI systems, LLMs are trained for broad applications and are designed to generate outputs based on user prompts. This generality often makes it difficult for data controllers to articulate

[23] Under the TDPL, the purpose must be specific, clear, and legitimate. For more information on these, see Aksoy Retornaz EE and Güçlütürk OG, 'Yapay Zekanın Kişisel Veri Kavramı ve Kişisel Verilerin İşlenmesinde Temel İlkelerle İlişkisi' in Aksoy Retornaz EE and Güçlütürk OG (eds), *Gelişen Teknolojiler ve Hukuk II: Yapay Zeka* (2021).

specific purposes for data processing, complicating their ability to comply with the requirements of Article 11(3).

Similarly, as will be detailed below, the provision of information to data subjects (Article 11(2)) becomes increasingly complex in the context of generative AI. Explaining how a user's personal data is processed, and its influence on the model's outputs may require disclosing technical details that are difficult to interpret for non-experts. For example, while an LLM might process data to improve conversational coherence, explaining the intricate neural network operations underpinning this improvement could overwhelm or confuse data subjects, undermining transparency.

To address these challenges, several measures can be adopted to facilitate the exercise of data subject rights in the context of LLMs. Data controllers must invest in developing tools and interfaces that enhance transparency, such as user-friendly dashboards that allow individuals to view how their data has been processed and its influence on AI models. Additionally, anonymisation and pseudonymisation techniques should be applied to training data wherever feasible to minimise the risks associated with data retention.

Regulators, on the other hand, may need to establish clearer guidelines for defining the purpose of data processing in generative AI contexts and set practical standards for erasure requests that account for the technical realities of LLMs. Collaboration between AI developers, legal experts, and data protection authorities is essential to align technical capabilities with regulatory requirements and ensure the meaningful exercise of data subject rights.

Having explored the general overview of challenges as well as certain possible mitigation measures, the remainder of this chapter will proceed with a detailed exploration of challenges concerning a set of specific rights that become more problematic to exercise with the involvement of generative AI.

5.3. The Right to Request Information on Whether, How, and for What Purposes Personal Data Are Processed

The right to request information about whether, how, and for what purposes personal data is processed is a cornerstone of data subject rights under Article 11 of the TDPL. This right encompasses the ability to learn whether personal data is processed, to obtain details about the processing, and to understand whether the processing aligns with its intended purpose. While

these rights are straightforward in conventional data processing contexts, the involvement of generative AI systems, particularly LLMs, introduces unique complexities that must be carefully analysed.

The first complexity arises in determining whether personal data is being processed in the first place. When an LLM is trained, personal data present in the training dataset is undoubtedly processed, making the activity fall squarely within the scope of the TDPL. Similarly, when data subjects provide prompts to an LLM, any personal data included in these prompts is actively processed to generate outputs. However, the situation becomes less clear when considering whether the potential for a model to generate outputs containing personal data qualifies as ongoing processing.

Generative AI models encode statistical relationships rather than storing personal data in retrievable forms[24]. Thus, the mere possibility that a model could generate outputs containing personal data does not necessarily constitute personal data processing. This paper argues that personal data processing should be recognised only at the moment a model actually generates outputs containing identifiable information. Adopting a contrary view would lead to impractical results, such as categorising all generative AI models as perpetually processing personal data, which would impose unreasonable and unmanageable regulatory obligations.

The second layer of complexity concerns the level of detail that data controllers are required to provide to data subjects exercising their right to request information. Modern AI systems, particularly LLMs, often function as black boxes, where the influence of individual data points on model outputs is opaque even to the developers themselves[25]. This lack of transparency creates significant challenges in determining what information should be provided to data subjects. Should controllers limit their disclosures to personal data collected and stored in conventional formats, such as forms, emails, or other structured records? Or are they also obligated to disclose inferred data or details about how the model encodes personal data into its parameters?

To address these questions, it is crucial to consider the purpose of the right to request information. This right is designed to ensure transparency and accountability in data processing

[24] Also see, 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' (Adopted 17 December 2024) https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf.
[25] Güçlütürk (n 4) 332.

and to give data subjects a certain degree of control over their personal data[26]. However, it does not grant data subjects proprietary rights over personal data or the AI systems that process it. Granting excessive access—such as requiring disclosure of training algorithms or the internal structure of a model—would risk exposing trade secrets and intellectual property, undermining the economic value of the AI system. This aligns with the approach under the GDPR, where Recital 63 emphasises that the right of access must not adversely affect the rights or freedoms of others, including intellectual property and trade secrets.

The application of these principles can be observed in case law and regulatory decisions. Under GDPR, the Court of Justice of the European Union (**CJEU**) has addressed the scope of the right of access in cases such as C-141/12[27] and C-434/16[28], emphasising the need for a balanced approach. In C-141/12, the Court held that extending the right of access to include a legal analysis in a residence permit application would exceed the scope of the right of access to personal data but would fall within the scope of the right of access to administrative documents[29]. However, in C-434/16, the Court ruled that access to examination answers and comments was justified because it served the purpose of providing transparency in the processing of personal data[30]. These cases highlight the importance of context when determining the scope of the right to request or access information.

---

[26] Güçlütürk (n 4) 332.

[27] YS v Minister voor Immigratie, Integratie en Asiel and Minister voor Immigratie, Integratie en Asiel v M and S, C-141/12, The Court of Justice of the European Union, https://curia.europa.eu/juris/liste.jsf?language=en&num=C-141/12.

[28] Peter Nowak v Data Protection Commissioner, C-434/16, The Court of Justice of the European Union, https://curia.europa.eu/juris/liste.jsf?num=C-434/16.

[29] "*As regards those rights of the data subject, referred to in Directive 95/46, it must be noted that the protection of the fundamental right to respect for private life means, inter alia, that that person may be certain that the personal data concerning him are correct and that they are processed in a lawful manner. As is apparent from recital 41 in the preamble to that directive, it is in order to carry out the necessary checks that the data subject has, under Article 12(a) of the directive, a right of access to the data relating to him which are being processed. That right of access is necessary, inter alia, to enable the data subject to obtain, depending on the circumstances, the rectification, erasure or blocking of his data by the controller and consequently to exercise the right set out in Article 12(b) of that directive. In contrast to the data relating to the applicant for a residence permit which is in the minute and which may constitute the factual basis of the legal analysis contained therein, such an analysis … is not in itself liable to be the subject of a check of its accuracy by that applicant and a rectification under Article 12(b) of Directive 95/46. In those circumstances, extending the right of access of the applicant for a residence permit to that legal analysis would not in fact serve the directive's purpose of guaranteeing the protection of the applicant's right to privacy with regard to the processing of data relating to him, but would serve the purpose of guaranteeing him a right of access to administrative documents, which is not however covered by Directive 95/46*." YS v Minister voor Immigratie, Integratie en Asiel and Minister voor Immigratie, Integratie en Asiel v M and S, C-141/12, The Court of Justice of the European Union, https://curia.europa.eu/juris/liste.jsf?language=en&num=C-141/12, paras 44-46.

[30] "*In so far as the written answers submitted by a candidate at a professional examination and any comments made by an examiner with respect to those answers are therefore liable to be checked for, in particular, their accuracy and the need for their retention, within the meaning of Article 6(1)(d) and (e) of Directive 95/46, and*

Under the TDPL, despite the lack of explicit reference to a right of access to data under Article 11, the Turkish Data Protection Board has taken a similar stance and held that the right to request information encompassed the right to access personal data. In terms of the means and methods of access, in its decision No. 2020/13[31], the Board ruled that a data subject's right to request information about phone call recordings did not entitle them to direct access to the recording environment but required the data controller to provide transcripts instead. This reasoning can be extended to generative AI: data subjects should not be granted direct access to the model or its underlying hardware but should receive relevant and meaningful summaries of how their personal data has been processed[32].

A related question is whether data controllers are obligated to provide access to both personal data provided as inputs during the model's development and personal data generated as outputs by the model. During the training phase, data is often formatted and transformed in ways that make it challenging to identify individual data points. Allowing access to this formatted data could compromise the integrity of the model while failing to provide meaningful insights into the data subject. Instead, access should be limited to the original form of the data if it is still retained and legally accessible[33]. For outputs generated by the model, data subjects should be informed about results containing their personal data. Still, they should not be entitled to access every possible outcome calculated by the model, especially when these outcomes are not directly related to the data subject[34].

The complexity of LLMs also raises questions about the extent to which data controllers must explain the operational mechanisms of these models. While transparency is essential, providing detailed technical explanations—such as the mathematical infrastructure of machine learning algorithms—would overwhelm most data subjects and fail to achieve the intended purpose of

---

may be subject to rectification or erasure, under Article 12(b) of the directive, the Court must hold that to give a candidate a right of access to those answers and to those comments, under Article 12(a) of that directive, serves the purpose of that directive of guaranteeing the protection of that candidate's right to privacy with regard to the processing of data relating to him … irrespective of whether that candidate does or does not also have such a right of access under the national legislation applicable to the examination procedure." Peter Nowak v Data Protection Commissioner, C-434/16, The Court of Justice of the European Union, https://curia.europa.eu/juris/liste.jsf?num=C-434/16, para 56.

[31] "*Within the framework of Article 11(1)(b) of Law No. 6698 on the Protection of Personal Data (the Law), the right of the data subject to request information regarding the processing of their personal data also encompasses the right of access to such data. Moreover, the right of access complements the right to request information, thereby enabling the data subject to fully exercise their rights over their personal data by gaining comprehensive knowledge about how their personal data is being processed*.", The Board Decision No 2020/13, https://www.kvkk.gov.tr/Icerik/6698/2020-13.

[32] See, Güçlütürk (n 4) 333.

[33] Güçlütürk (n 4) 333.

[34] Güçlütürk (n 4) 334.

the right[35]. Instead, controllers should focus on explaining the logical steps and implications of data processing in clear, non-technical terms. For instance, a controller could describe how an LLM uses user inputs to generate personalised outputs, highlighting potential risks such as data retention or profiling without delving into the intricacies of neural networks.

Nevertheless, even simplified explanations must be carefully crafted to avoid undermining trade secrets or intellectual property. Striking this balance requires controllers to disclose meaningful information about the logic and significance of processing while avoiding unnecessary details that could harm their commercial interests[36]. For example, it is reasonable for a controller to inform a data subject that their input data may be used to improve the quality of service but unnecessary to detail the exact algorithms or datasets involved.

Finally, it is important to consider the interplay between the right to request information and other data subject rights, such as the right to object to automated decisions under Article 11(1)(g) of the TDPL. Given that generative AI systems inherently involve automated processing, data controllers must be prepared to address objections related to the use of these systems. However, determining whether the processing is "solely automated" often requires a case-by-case analysis, further emphasising the need for contextual evaluation, which will be covered in detail from another perspective below.

In conclusion, the right to request information under the TDPL must be applied in a way that balances the transparency owed to data subjects with the practical and technical realities of generative AI. By focusing on meaningful and relevant disclosures while safeguarding trade secrets and intellectual property, data controllers can fulfil their obligations without compromising the integrity or security of their AI models.

5.4. Interplay Between the Right to Rectification, Erasure, and Generative AI

The rights to rectification, erasure, and anonymisation of personal data present significant challenges when applied to generative AI systems, especially LLMs. Although Turkish law does not explicitly codify the right to be forgotten, as established under the GDPR, this principle is recognised and frequently referred to in Turkish practice. The Turkish Data Protection Board, for example, addressed the right to be forgotten in its decision dated

[35] Güçlütürk (n 4) 332.
[36] Güçlütürk (n 4) 332-334.

23.06.2020 and numbered 2020/481[37], which deals with requests to remove search engine results linked to individuals' names. This decision, along with the other legal grounds referred to therein, demonstrates that the right to be forgotten has practical application under Turkish law, even in the absence of specific codification[38].

Despite the separate definitions of the right to rectification and the right to erasure, these rights are interconnected in practice. They are further linked to the right to request information since a data subject's ability to rectify or erase their data often depends on their understanding of how the data is processed. This interrelation becomes more complex when generative AI systems are involved, raising two critical questions. First, do these rights allow data subjects to request the deletion or retraining of an LLM that "contains" their personal data or generates outputs based on it? Second, how should a data controller handle requests for rectification or erasure if the input data has been deleted but the model itself is retained?

The right to rectification becomes particularly intricate in the context of generative AI because these systems can generate new, inference-based personal data using existing data. This inferred data may not accurately reflect the reality of the data subject, creating a basis for rectification requests. For example, a generative AI system trained on partial or outdated information about an individual might generate an inaccurate profile or predictions about that individual's behaviour. In such cases, data subjects should have the right to request rectification of the inaccurate outputs. However, extending this right to include retraining or rebuilding the model itself raises practical and legal challenges. Retraining an LLM would require repeating the entire data processing pipeline, which may be prohibitively costly or practically impossible[39]. Moreover, the technical and economic burdens of retraining would likely outweigh the benefits of correcting isolated inaccuracies, making such requests disproportionate.

---

[37] The Board Decision No. 2020/481, https://www.kvkk.gov.tr/Icerik/6776/2020-481.

[38] "*… while the concept of the right to be forgotten is not explicitly recognised in our legislation, it is evident that there are mechanisms in our legal framework aimed at safeguarding this right. These mechanisms include, for instance, the provision under Law No. 5651 regarding the restriction of access to content on the grounds of privacy protection, as well as Article 7 of the [TDPL] regulating the deletion of data. Accordingly, based on the aforementioned explanations, it is concluded that the right to be forgotten can be addressed within the scope of domestic law through the regulations stipulated in Article 20(3) of the Constitution, Articles 4, 7, and 11 of Law No. 6698, and Article 8 of the Regulation on the Deletion, Destruction, or Anonymization of Personal Data, without necessitating its explicit recognition as an independent right.*" The Board Decision No. 2020/481, https://www.kvkk.gov.tr/Icerik/6776/2020-481.

[39] Güçlütürk (n 4) 335.

The right to erasure faces similar complexities. Personal data used during the training of generative AI models is typically transformed into a specific format that cannot be directly traced within the model's parameters. Under the TDPL, data subjects may request the deletion of identifiable personal data used in training datasets if the conditions outlined in Article 7 are met. However, this right does not extend to the deletion or destruction of the model itself unless the model is specifically designed to generate highly personalised outputs that pose a clear risk to the data subject's privacy[40]. The effect of an individual data point on a model like an LLM diminishes as the size of the training dataset increases. For example, a single name scraped from a social media profile has such a negligible impact on the model's performance that it is practically ineffective. Therefore, unless a clear and demonstrable risk exists, the right to erasure should not be interpreted as granting data subjects the ability to demand the deletion or retraining of a generative AI model.

A related issue arises when personal data is provided as input to a pre-trained model. In such cases, data subjects can exercise their right to request the deletion of the input data itself. However, outputs generated by the model require a separate assessment to determine whether they qualify as personal data. Even if personal data is used as input, the model's output may not necessarily contain identifiable information. For example, a generative AI system might use personal data to train its language structure but generate outputs that are entirely generic and non-identifiable. Data subjects cannot demand the deletion of all outputs produced by the model merely because their personal data was part of the input dataset.

When a data subject requests rectification or erasure and the underlying personal data is deleted, but the model itself is retained, the handling of such requests becomes more nuanced. The response of the data controller depends on whether the retained model continues to generate outputs containing personal data. If the model does not generate or leak personal data and there is no reasonable indication to the contrary, the controller can process the request without altering the model. However, if the model generates identifiable personal data, the nature and scope of this generation or leakage may necessitate additional actions, such as informing the data subject, deploying safeguards, or even retraining the model in extreme cases[41].

[40] Güçlütürk (n 4) 334.
[41] For a similar analysis under the GDPR, see Osman Gazi Güçlütürk, 'GDPR Data Access Requests' (OECD AI Wonk Blog) https://oecd.ai/en/wonk/gdpr-data-access-requests.

Given the technical and economic limitations of retraining or deleting generative AI models, data controllers must consider alternative measures to uphold the spirit of the rights to rectification and erasure. One effective approach is the use of content filters or additional screening layers that assess model outputs before they are presented to users. These filters, which are already employed in many LLM systems to block harmful or inappropriate outputs, can be adapted to detect and suppress outputs containing personal data. For instance, when a data subject submits an erasure request, such filters can evaluate whether the model's outputs contain the requested data and prevent these outputs from being displayed to users. While these measures do not prevent the model from generating outputs containing personal data internally, they effectively mitigate privacy risks by ensuring such outputs are not exposed.

Dynamic adjustments to the model's outputs can further enhance privacy protection. For example, if a data controller receives a rectification or erasure request, they can deploy mechanisms to monitor the model's behaviour and suppress outputs that rely on the deleted data. Logging and monitoring tools can also provide valuable oversight, allowing data controllers to track how data is processed and verify compliance with rectification and erasure requests.

Another practical safeguard is the application of advanced anonymisation techniques to training datasets. Techniques such as differential privacy[42] or data masking can reduce the likelihood that individual data points influence the model in a recognisable way. These methods enhance privacy protection while preserving the model's overall functionality.

The balance between respecting data subject rights and maintaining the functionality of generative AI systems is delicate. While data subjects must have mechanisms to ensure the accuracy and lawful processing of their data, granting excessive rights—such as the ability to demand model retraining—could impose disproportionate burdens on data controllers and undermine the commercial viability of AI systems. Transparency is crucial in managing this balance. Data controllers should clearly explain to data subjects how their data is processed, the limits of rectification and erasure rights, and the measures in place to mitigate privacy risks. By combining clear communication with robust technical safeguards, data controllers can

[42] For more information on differential privacy, see Cynthia Dwork, 'Differential Privacy' in Michele Bugliesi and others (eds), *Automata, Languages and Programming* (Springer Berlin Heidelberg 2006) 1.

uphold the rights of data subjects while preserving the integrity and value of generative AI systems.

5.5. The Question of the Right to Object under Article 11

The last right that will be covered in this paper is the right to object to the occurrence of an adverse result against the data subject by analysing the data processed solely through automated means. To begin with, while this right may seem to resemble the GDPR's Article 22, there are significant differences. Unlike the GDPR, the TDPL does not explicitly grant data subjects the right to object to the exclusive use of automated systems for processing their data. This reflects a less interventionist approach in Turkish law, emphasising the freedom of data controllers in determining how to process data. The right provided here is not a right to object to the processing but one to object only to the emergence of a certain type of result based on such processing.

While the content and the implications of this right are already unclear[43], LLMs complicate the exercise of this right even further. The training process for generative models often involves human intervention during the formatting and selection of datasets, but personal data provided as input to a pre-trained model is typically processed exclusively by automated systems. In such cases, the right to object may become relevant. However, it must be noted that as long as there is meaningful human intervention in the processing, this right does not apply[44].

Determining the presence of human intervention in generative AI systems is more complex than in other AI contexts. In narrowly tailored systems designed for specific tasks, such as an AI tool used to screen resumes, there may be a human operator actively involved in monitoring the model's operations. In these cases, human involvement can prevent the application of this right. However, for LLMs, where millions of simultaneous interactions occur across diverse users, it is neither practical nor realistic to expect continuous human oversight. Instead, an alternative approach could involve mechanisms allowing users to escalate issues to a human reviewer. Whether this constitutes sufficient human intervention depends on the extent of the reviewer's influence over the system's outputs. In most scenarios, such human-led moderation

[43] For more information, see Güçlütürk (n 4) 335.
[44] Güçlütürk (n 4) 337.

occurs post hoc, meaning the AI system continues to operate autonomously during the interaction. In these cases, the right to object remains applicable[45].

Taking a step further in terms of the exercise and practical implications of this right, the lack of clarity in the TDPL and its related materials regarding the outcomes of objections adds another layer of complexity. Neither the law nor the Board's decisions specify what should happen when a data subject objects to automated processing. If the objection is tied to an unlawful outcome, the data subject would already have grounds to seek compensation under Article 11(1)(ğ). However, the scope of this right should not be limited to monetary compensation alone. Ideally, an objection should trigger a re-evaluation of the relevant decision, this time incorporating human intervention to ensure a fairer outcome[46]. For example, decisions on job applications, credit approvals, or salary calculations initially determined solely by machine-learning systems should ideally be reviewed by human participants.

This approach, however, becomes less straightforward in systems relying entirely on automated closed-loop processes. Human intervention in such systems may be prohibitively costly or technically infeasible. In these situations, the appropriate course of action remains unclear. If the decision-making mechanism of the AI model is explainable, the data controller or relevant personnel could review the system's logic and render an independent decision. This alternative, while imperfect, would align with the broader goals of transparency and fairness.

The challenges associated with exercising this right are even more pronounced in the context of LLMs. The outputs of these models often lack the direct traceability found in more narrowly focused systems. In instances where an LLM serves merely as an additional interface to an underlying decision-making system, objections should lead to a re-evaluation of the decision by human reviewers, provided the right to object applies. However, when an LLM itself plays a substantive role in generating a harmful decision or outcome, reevaluating the entire conversation or process manually may not be practical. In such cases, this right must be considered in conjunction with the right to seek compensation under Article 11(1)(ğ), as the re-evaluation process may not provide a sufficient remedy or may not even be practical.

In sum, the right to object under the TDPL, while relevant in the context of generative AI, requires careful interpretation and practical adaptation. Data controllers must evaluate the

[45] Güçlütürk (n 4) 336-337.
[46] Güçlütürk (n 4) 337.

extent of automation and the feasibility of human intervention in each case. While objections should ideally lead to fairer outcomes through human oversight, the technical realities of generative AI systems may necessitate alternative remedies, such as explainable decision reviews or compensation mechanisms, to uphold the rights of data subjects effectively.

6. Conclusion

The interplay between generative AI and data subject rights under TDPL underscores the transformative yet complex nature of modern artificial intelligence systems. As generative AI, particularly LLMs, continues to redefine how personal data is processed, stored, and utilised, it simultaneously challenges the foundational principles of data protection law, creating a pressing need for nuanced interpretations and practical adaptations of existing rights.

Throughout this paper, we have seen how the capabilities and operational complexities of generative AI systems complicate the exercise of fundamental data subject rights under the TDPL. The right to request information, for example, becomes difficult to enforce due to the opaque, "black box" nature of LLMs and their inability to disclose the specific influence of individual data points on outputs. The right to rectification, erasure, and anonymisation is similarly constrained, as these models transform personal data into statistical representations that cannot easily be reversed or erased. Even the right to object, while seemingly aligned with the unique challenges posed by automated decision-making systems, reveals gaps in clarity and feasibility when applied to highly automated and large-scale generative AI systems.

The inherent risks posed by generative AI—such as unintended data retention, profiling, and behavioural nudging—further exacerbate these challenges. As these systems increasingly encourage users to disclose sensitive information, their capacity to process and infer personal data raises significant privacy concerns. These concerns are amplified by the difficulty of identifying when personal data is actually processed, particularly when outputs are generated based on probabilistic patterns rather than identifiable information.

Addressing these challenges requires a dual approach. On the one hand, data controllers must adopt proactive measures to enhance transparency, accountability, and privacy protection. These measures include real-time input monitoring, the application of anonymisation and pseudonymisation techniques, the deployment of content filters, and the development of user-friendly tools that empower data subjects to exercise their rights effectively. On the other hand, regulatory bodies must establish clearer guidelines tailored to the unique characteristics of

generative AI systems. These guidelines should address issues such as the scope of data subject rights in the context of probabilistic data generation, the obligations of data controllers in providing meaningful explanations of AI systems, and the mechanisms for ensuring fair outcomes in cases of automated decision-making.

The paper also highlights the necessity of balancing competing interests—namely, the rights of data subjects to control their personal data and the legitimate interests of AI developers and operators in preserving the functionality, efficiency, and commercial value of their models. Granting excessive access or imposing disproportionate obligations, such as mandatory model retraining, could jeopardise the viability of generative AI technologies, undermining their transformative potential.

While the TDPL provides a robust framework for data protection, it must be interpreted and applied flexibly to accommodate the novel challenges posed by generative AI. This flexibility should be grounded in the law's underlying principles, including transparency, accountability, and proportionality. Drawing insights from international frameworks such as the GDPR, Turkish regulators, lawmakers, and stakeholders must collaborate to refine the legal and ethical standards governing generative AI, ensuring that they remain relevant in an era of rapid technological advancement.

Ultimately, the evolution of generative AI presents not only legal challenges but also opportunities to rethink and enhance data protection frameworks. By fostering innovation while safeguarding individual rights, Turkish data protection law can serve as a model for addressing the complexities of generative AI in a way that promotes trust, fairness, and accountability. Through continued dialogue, technical innovation, and legal refinement, a balance can be struck that benefits both society and the development of transformative technologies.