



Governing Platform Data in the Generative AI Era: Personal Information Protection Lessons from China's Didi Cybersecurity Review Case

Yuanyuan Wang¹, Shiyi Xu², Yang Gao³, Yanjun Xu*

¹Big Data Application and Information Monitoring Center of Hohhot Public Security Bureau, Hohhot, China

² Big Data Application and Information Monitoring Center of Hohhot Public Security Bureau, Hohhot, China

³Big Data Application and Information Monitoring Center of Hohhot Public Security Bureau, Hohhot, China

*(Corresponding Author) Inner Mongolia Honder College of Arts and Sciences, Hohhot, China

Email: 16446284@qq.com; gaoyang992008@163.com; 18947926204@163.com
13500691220@139.com

Abstract: Generative artificial intelligence has transformed personal information protection from a narrow compliance issue into a broader problem of platform data governance. Although much current debate focuses on model outputs, hallucination, content moderation, or deepfake misuse, this study argues that a deeper source of risk lies in the data infrastructures that precede AI deployment. Using China's Didi cybersecurity review and administrative penalty as a qualitative single-case study, the paper examines how large-scale platform data accumulation, sensitive personal information processing, algorithmic inference, and weak internal governance may create systemic privacy and security risks in the generative AI era. The Didi case is not a direct generative AI case. Rather, it provides a pre-generative-AI lesson: trustworthy AI governance depends on lawful, secure, transparent, and auditable platform data governance. Drawing on doctrinal legal analysis and regulatory document analysis, the study develops a platform-data-infrastructure framework that links data aggregation, algorithmic inference, AI reuse, and governance response. It shows that Didi's violations involved excessive and unlawful processing of large volumes of personal information, including facial recognition information, precise location information, identity card numbers, and inferred travel-related data. The case reveals three governance challenges: platform data aggregation risk, heightened vulnerability of sensitive personal information, and the expansion of privacy risk from direct collection to algorithmic inference. The paper argues that generative AI governance should not begin only at the model layer. It should integrate training data compliance, sensitive information protection, platform accountability, data security audits, and lifecycle-based risk assessment. The study contributes to AI governance scholarship by linking platform data regulation with generative AI privacy protection and by highlighting upstream data governance as a precondition for responsible AI development.

Keywords: Generative artificial intelligence; personal information protection; platform data governance; data security;

1. Introduction

This study examines China's Didi cybersecurity review and administrative penalty case as a single-case study of platform data governance. In July 2022, the Cyberspace Administration of China imposed an administrative fine of RMB 8.026 billion on Didi Global Inc. for violations of the Cybersecurity Law, the Data Security Law, and the Personal Information Protection Law [13]. The official explanation stated that Didi had illegally processed 64.709 billion items of personal information, including sensitive personal information such as facial recognition information, precise location information, and identity card numbers [14]. The case also involved excessive collection of clipboard information, application list information, passenger facial recognition information, family relationship information, home and company ride-hailing addresses, and inferred information such as travel intention, resident city, and business or tourism status [14]. These facts make the Didi case one of the most significant platform data governance enforcement cases in China.

The Didi case is not, strictly speaking, a generative AI case. It occurred before the commercial expansion of large-scale generative AI services and before China adopted the Interim Measures for the Management of Generative Artificial Intelligence Services in 2023 [18]. Nevertheless, the case is highly relevant to generative AI governance because it reveals the upstream conditions under which AI-related privacy risks are produced. Generative AI risks are not located only at the model layer. They also originate from platform data accumulation, sensitive personal information processing, algorithmic inference, weak internal controls, and insufficiently documented data reuse. If platform data are collected excessively, stored insecurely, or processed without adequate transparency, subsequent AI training and deployment may inherit and amplify these defects.



The central argument of this paper is that platform data governance is a precondition for responsible generative AI governance. The Didi case demonstrates that large digital platforms may create systemic privacy risks through continuous data aggregation, multi-app data integration, opaque inference, and insufficient compliance mechanisms. In the generative AI era, these risks become more serious because platform data can be reused for model training, personalized generation, intelligent recommendation, automated decision-making, or retrieval-augmented generation. Personal information protection must therefore move beyond notice-and-consent formalism toward lifecycle-based governance covering data collection, storage, analysis, transfer, model use, audit, deletion, and accountability. The paper makes three contributions: first, it reinterprets the Didi case not merely as a conventional personal information violation, but as a foundational case for understanding platform data risks in the generative AI era; second, it connects Chinese platform regulation with global concerns about training data privacy, data minimization, dataset provenance, model memorization, and algorithmic inference; third, it proposes a conceptual framework linking platform data infrastructure, aggregation and inference, AI reuse and amplification, and governance response.

2. Literature Review

2.1 Generative AI and Privacy Risk

Existing research on generative AI privacy has focused heavily on model memorization, training data extraction, inference attacks, and leakage of personally identifiable information. Carlini et al. demonstrated that large language models may memorize and reproduce training data, including names, phone numbers, email addresses, and other identifiable sequences, when queried in specific ways [1]. Later research further showed that memorization tends to increase with model capacity, repeated exposure to training examples, and longer prompts [2]. These studies challenge the assumption that once personal data are absorbed into statistical models, they become practically inaccessible. Instead, generative models may preserve traces of training data in ways that create new privacy risks.

Recent surveys similarly divide large language model privacy risks into risks during training and risks during inference [3], [4]. Training-stage risks include unlawful data collection, insufficient data cleaning, inclusion of sensitive or confidential information, and weak provenance controls. Inference-stage risks include prompt leakage, model output of private data, re-identification, user interaction logging, and downstream misuse of generated content. This literature is important because it shows that generative AI privacy cannot be limited to the output layer. The lawfulness, quality, provenance, and sensitivity of training data are equally central.

Privacy engineering and dataset governance research reinforces this point. Privacy-by-design literature emphasizes that privacy protection should be embedded in system architecture, rather than added after deployment [5]. Dataset governance scholarship similarly argues for documentation of dataset composition, collection purpose, recommended use, limitations, and maintenance responsibilities [6]. These approaches are directly relevant to generative AI, where training data, fine-tuning data, prompt logs, vector databases, and evaluation datasets may all contain personal information. However, much of the technical literature pays limited attention to the platform environments in which data are accumulated before model training. The Didi case helps fill this gap by showing that platform data governance failures may precede and later shape AI privacy risks.

2.2 Platform Data Governance and Accountability

A second body of literature concerns platform power and data governance. Digital platforms operate as data-intensive infrastructures that mediate transport, commerce, communication, finance, and public services. Their value depends not only on direct user input but also on observed behavior, inferred preferences, predictive analytics, and cross-service integration. Zuboff describes this as an economic logic based on the extraction and prediction of behavioral data [7]. Van Dijck, Poell, and de Waal similarly argue that platforms reorganize social and economic activities by controlling data flows, interfaces, and algorithmic visibility [8].

For privacy governance, this means that personal information risk is not limited to isolated data points. The more serious risk lies in aggregation. A single location record may reveal little. Continuous location histories, ride-hailing destinations, device identifiers, family relationships, and payment records can reveal work routines, medical visits, social relations, consumption patterns, and other intimate aspects of daily life. In platform environments, personal information protection must therefore address aggregation risk and inferred identity, not only direct collection.

International legal developments also reflect this shift. The General Data Protection Regulation emphasizes lawfulness, fairness, transparency, purpose limitation, data minimization, accuracy, storage limitation, integrity, confidentiality, and accountability [9]. The EU Artificial Intelligence Act adds a risk-based approach to AI governance, including obligations related to data governance, technical documentation, transparency, human oversight, and risk management for high-risk AI systems [10]. Although these instruments differ from China's regulatory approach, they show a shared concern: AI governance requires not only output control but also accountable data practices.

2.3 Personal Information Protection in China

China's personal information protection regime combines individual rights protection with broader concerns over cybersecurity, data security, national security, and platform governance. The Personal Information Protection Law, effective from November 1, 2021, defines personal information as information related to an identified or identifiable natural person and requires processing to follow principles of lawfulness, necessity, good faith, purpose limitation, openness, transparency, and minimum scope [12]. It also provides specific rules for sensitive personal information, automated decision-making, cross-border provision, individual rights, and processor obligations.

Academic work has noted that China's privacy regime differs from approaches that frame privacy primarily as an individual autonomy issue. Large-scale personal information processing can become a matter of both individual rights and

public regulatory concern [11], [12]. The Didi case illustrates this combined logic. It was not treated merely as an app privacy violation but as a cybersecurity-review-related administrative penalty involving personal information protection, data security, and national security risk.

This paper uses four terms in their legal and analytical senses. Personal information refers to information related to an identified or identifiable natural person. Sensitive personal information refers to information that, once leaked or illegally used, may easily harm personal dignity or personal and property safety, such as biometric identification, religious belief, specific identity, medical health, financial accounts, or location tracking information. Important data is a data-security category related to national security, economic operation, social stability, public health, and safety. National data security risk refers to risks that may affect national security or public interest through data processing, cross-border transfer, platform concentration, or systemic exposure. Distinguishing these terms is necessary because the Didi case connects personal information protection with broader data security governance.

3. Regulatory Context

China's regulatory framework for platform data and generative AI is built on several overlapping legal instruments: the Cybersecurity Law, the Data Security Law, the Personal Information Protection Law, the Regulations on Network Data Security Management, and the Interim Measures for the Management of Generative Artificial Intelligence Services. These instruments are not identical in purpose. The Cybersecurity Law focuses on network operation security and critical information infrastructure protection. The Data Security Law extends governance to data processing activities more generally. The Personal Information Protection Law focuses on individual-related data processing. The Network Data Security Management Regulations integrate network data processing obligations, while the Interim Measures target generative AI services provided to the public.

The Cybersecurity Law provides the foundational framework for network operation security, critical information infrastructure protection, and network operator obligations. It requires network operators to adopt technical and organizational measures to protect network security and prevent data leakage, theft, or tampering [15]. In the Didi case, the cybersecurity review was initiated in 2021 to prevent national data security risks and safeguard national security and public interests [14]. This shows that platform data protection in China is embedded in a broader cybersecurity governance system.

The Data Security Law, effective September 1, 2021, extends governance beyond personal information to data processing activities more generally. It requires data processors to establish whole-process data security management systems, conduct education and training, take technical measures, and conduct risk assessment for important data [16]. It also requires data to be collected by lawful and proper means and used within legally prescribed purposes and scopes. This law is important for generative AI because AI systems depend on the lawful collection, classification, storage, and reuse of data.

The Personal Information Protection Law is the central statute governing personal information processing. It requires explicit and reasonable purposes, minimum necessary collection, openness and transparency, and accountability of personal information processors [12]. It also governs automated decision-making and sensitive personal information. In the Didi case, the regulator emphasized not only excessive collection but also the processing of sensitive personal information, including facial recognition information, precise location data, and identity card numbers [14]. These categories are especially relevant to generative AI because they may enable biometric synthesis, mobility profiling, identity simulation, and personalized generation.

The Regulations on Network Data Security Management, effective from January 1, 2025, further integrate network data governance with personal information protection and AI-related data risks. The regulation defines network data processing broadly to include collection, storage, use, processing, transmission, provision, disclosure, and deletion [17]. Article 19 specifically provides that network data processors offering generative AI services shall strengthen the security management of training data and training data processing activities and take effective measures to prevent and address network data security risks [17]. This provision directly connects platform data governance with generative AI governance.

The Interim Measures for the Management of Generative Artificial Intelligence Services, effective August 15, 2023, establish a targeted regulatory framework for generative AI services provided to the public in China [18]. They apply to services that generate text, images, audio, video, or other content. The measures require providers to use data and foundation models from lawful sources, respect intellectual property rights, obtain consent or satisfy other legal bases when personal information is involved, improve training data quality, and comply with the Cybersecurity Law, Data Security Law, and Personal Information Protection Law. Taken together, these instruments show a regulatory trajectory from network security to data security, from data security to personal information protection, and from personal information protection to generative AI governance.

4. Conceptual Framework

This paper proposes a platform-data-infrastructure framework. The framework does not assume that Didi used generative AI. Instead, it explains why a pre-generative-AI platform case can illuminate AI governance. Generative AI systems are not only models; they are socio-technical systems connected to data pipelines, platform databases, user interfaces, organizational decisions, and regulatory obligations. The risks generated by such systems depend on the data environment that precedes and surrounds them.

The framework contains four linked components. First, platform data infrastructure refers to the collection, storage, integration, and management of data across apps, services, and business lines. Second, aggregation and inference refer to

the transformation of provided and observed data into behavioral profiles, predictions, and inferred identities. Third, AI reuse and amplification refers to the possibility that platform data may be used in training, fine-tuning, retrieval-augmented generation, AI agents, recommendation systems, or automated decision-making. Fourth, governance response refers to legal, organizational, and technical mechanisms that control these risks, including data minimization, purpose limitation, sensitive data safeguards, audit trails, risk assessment, and accountability. This framework allows the Didi case to be analyzed without overstating its empirical connection to generative AI. The case is not evidence of generative AI misuse. Rather, it is evidence of how platform data practices can create the upstream conditions for future AI risks. Its value lies in showing that AI governance cannot be separated from the data infrastructures of digital platforms.

Table 1. Conceptual framework linking platform data governance and generative AI governance

Component	Analytical focus	Generative AI relevance
Platform data infrastructure	Collection, storage, integration, and management of data across apps and services	Determines whether AI systems are built on lawful and auditable data foundations
Aggregation and inference	Transformation of provided and observed data into profiles, predictions, and inferred identities	Creates privacy risks beyond direct data collection and formal consent
AI reuse and amplification	Reuse of platform data in training, fine-tuning, RAG, AI agents, or automated decision-making	May amplify unlawful, excessive, or sensitive data practices
Governance response	Data minimization, purpose limitation, audits, risk assessment, and accountability	Provides the institutional basis for responsible AI governance

5. Methodology

This study adopts a qualitative single-case study design, supported by doctrinal legal analysis and regulatory document analysis. The single-case approach is appropriate because the Didi cybersecurity review and administrative penalty represents an unusually rich and influential case of platform data governance in China. Rather than treating the case as an isolated compliance failure, the paper uses it as an analytical site for examining how large-scale platform data processing may generate personal information, data security, and national security risks in the generative AI era.

The case is selected for four reasons. First, the penalty amount was exceptional: the Cyberspace Administration of China imposed an RMB 8.026 billion fine on Didi Global Inc., making the case one of China's most significant data governance enforcement actions [13]. Second, the data types involved were complex, including facial recognition information, precise location information, identity card numbers, clipboard information, app list information, and inferred travel-related data [14]. Third, the case connected personal information protection with broader data security and national security concerns. Fourth, Didi operated across multiple apps and business lines, making the case typical of platform-based data governance problems.

The primary materials used in this study are official regulatory documents, including the cybersecurity review announcement, the multi-agency review notice, the administrative penalty decision, and the official Q&A issued by the Cyberspace Administration of China [13], [14], [19], [20]. These materials were analyzed in three steps. First, the paper extracted factual elements, including dates, responsible entities, penalty amount, violation categories, data types, and regulatory justifications. Second, the paper mapped these elements onto analytical categories derived from the literature: aggregation risk, sensitive personal information, inferred personal information, platform accountability, and lifecycle governance. Third, the paper connected these categories to generative AI governance through the conceptual framework developed above.

The methodology has limitations. Official regulatory materials are authoritative for case facts and legal basis, but they also reflect the regulator's framing and do not provide full access to internal company documents, technical systems, or user experiences. The paper therefore does not claim to provide a complete empirical reconstruction of Didi's internal data practices. Nor does it claim that Didi's conduct involved generative AI. Its claim is narrower: the case illuminates platform data conditions that become central to generative AI governance.

6. Case Background

The Didi cybersecurity review and administrative penalty is a landmark case in China's platform data governance regime. The case began on July 2, 2021, when the Cybersecurity Review Office announced that it would conduct a cybersecurity review of Didi Chuxing. The stated purpose was to prevent national data security risks, safeguard national security, and protect public interests. During the review period, Didi Chuxing was required to suspend new user registration in order to prevent the further expansion of risks [19]. On July 16, 2021, the Cyberspace Administration of China, together with the Ministry of Public Security, the Ministry of State Security, the Ministry of Natural Resources, the Ministry of Transport, the State Taxation Administration, and the State Administration for Market Regulation, entered Didi Chuxing Technology Co., Ltd. to conduct the cybersecurity review [20]. This multi-agency process shows that the case was treated not merely as a routine app privacy issue, but as a major platform data security matter.

On July 21, 2022, the Cyberspace Administration of China issued the administrative penalty decision. According to the official decision, Didi Global Inc. had violated the Cybersecurity Law, the Data Security Law, and the Personal Information Protection Law. The regulator imposed a fine of RMB 8.026 billion on Didi Global Inc. and imposed separate

finer of RMB 1 million each on Cheng Wei, the company's chairman and CEO, and Liu Qing, the company's president [13]. The penalty decision described the violations as clear in fact, conclusive in evidence, serious in circumstance, and severe in nature. This language indicates a regulatory judgment that the case involved systematic failures rather than minor or accidental defects.

The official Q&A released on the same day provided more detailed information about the investigation and the violations. According to that document, the regulator conducted inquiries, technical evidence collection, review of submitted materials, and analysis of evidence, while also hearing Didi's views and protecting its procedural rights [14]. The Q&A stated that Didi had sixteen illegal facts, which were summarized into eight main categories. These included illegally collecting 11.9639 million items of screenshot information from users' mobile phone albums; excessively collecting 8.323 billion items of clipboard information and app list information; excessively collecting passengers' facial recognition information, age group information, occupational information, family relationship information, and home and company ride-hailing address information; and excessively collecting precise location information in several service scenarios [14]. The case also involved sensitive identity and inference-related data. The regulator stated that Didi excessively collected drivers' education information and stored 57.8026 million driver identity card numbers in plaintext. It also analyzed 53.976 billion items of passenger travel intention information, 1.538 billion items of resident city information, and 304 million items of intercity business or tourism information without clearly informing passengers [14]. These findings are important because platform personal information risk is not limited to data directly submitted by users. It also includes observed data generated during service use and inferred data produced through platform analytics.

The official materials further explained the basis for identifying Didi Global Inc. as the responsible entity. Didi's domestic business lines included ride-hailing, hitch, two-wheel mobility, automobile-related services, and other operations. Its products included 41 apps, such as Didi Chuxing, Didi Driver, Didi Hitch, and Didi Enterprise Edition. The regulator found that Didi Global Inc. had ultimate decision-making authority over major matters in domestic business lines and that its internal governance rules applied across those lines [14]. This point is central to the regulatory logic of the case: responsibility was attributed to the platform-level entity because the violations were connected with unified decision-making, supervision, and management.

The penalty was based on several considerations. The regulator emphasized the nature of the violations, their long duration, their harm to users' personal information rights, the enormous quantity of illegally processed personal information, and the involvement of multiple apps and multiple types of unlawful processing. It also stated that Didi's illegal processing involved 64.709 billion items of personal information, including facial recognition information, precise location information, identity card numbers, and other sensitive personal information [14]. The regulatory logic therefore combined individual rights protection, platform accountability, data security, and national security.

7. Analysis

7.1 Platform Data Accumulation and Aggregation Risk

The Didi case demonstrates that platform data risk does not arise only from the unlawful collection of isolated pieces of personal information. Its deeper significance lies in the accumulation and aggregation of data across services, applications, user groups, and business scenarios. In platform ecosystems, personal information is continuously generated, captured, linked, inferred, and reused. Ride-hailing platforms are especially data-intensive because their core service depends on identity verification, location tracking, route planning, payment processing, driver-passenger matching, risk control, customer service, and operational optimization. Each of these functions may appear necessary when viewed separately, but their combination can create a highly detailed picture of individual behavior.

Aggregation changes the nature of privacy risk. A single ride record may only reveal where a user traveled at a specific time. Repeated ride records, however, can reveal residential location, workplace, commuting patterns, medical visits, family relationships, social contacts, nightlife habits, and other intimate aspects of daily life. When such records are combined with device information, clipboard data, app list information, identity information, facial recognition data, and payment-related data, the platform gains the ability to construct a multi-dimensional profile of the user. The risk is therefore not limited to whether one item of data was lawfully collected; it also concerns whether the platform has created a data environment capable of continuous profiling and behavioral prediction.

This risk also exposes the limits of consent-based governance. In theory, users may be informed of personal information processing rules through privacy policies and consent interfaces. In practice, however, users rarely understand the full scope of data collection, the logic of data combination, or the future uses of their information. Consent becomes especially fragile when users depend on a platform for daily mobility, employment, or business operations. Drivers and passengers may have little real bargaining power. They may accept broad data practices because refusal means exclusion from an essential service.

For this reason, the Didi case suggests that personal information protection in platform environments must move beyond formal consent. Governance should focus on data necessity, purpose limitation, access control, internal accountability, and data lifecycle management. The key question is not only whether users clicked agree, but whether the platform can justify why each category of data is collected, how long it is retained, who can access it, whether it is linked with other data, and whether it is later used for new purposes.

Table 2. Three layers of platform data and their privacy implications

Data layer	Examples in platform services	Main risk
Provided data	Name, phone number, identity information, driver credentials	Users knowingly provide data, but may not understand later reuse
Observed data	Location, route, time, device information, app behavior	Continuous monitoring may reveal behavioral patterns
Inferred data	Travel intention, resident city, business or tourism status	Users may not know that such profiles are created

7.2 Sensitive Personal Information and Algorithmic Inference

The second major lesson from the Didi case concerns sensitive personal information and algorithmic inference. Sensitive personal information requires stricter protection because misuse may cause serious harm to personal dignity, property security, physical safety, reputation, or social equality. In the Didi case, the relevant sensitive data included facial recognition information, precise location information, identity card numbers, and other information that could directly or indirectly identify individuals and expose intimate aspects of their lives.

Precise location data is especially sensitive in a mobility platform. Location data is not merely technical service data. It can reveal where a person lives, works, studies, receives medical treatment, participates in social activities, or meets others. When location records are accumulated over time, they become a behavioral map of the individual. This map may be used for service optimization, but it may also enable excessive profiling, discriminatory pricing, unauthorized monitoring, or exposure of private routines.

Facial recognition information presents another serious risk. Biometric information is difficult to replace once compromised. A password can be changed, but a face cannot. If facial data are collected excessively, stored insecurely, or used beyond the original purpose, individuals face long-term risks of identity misuse, deepfake generation, unauthorized authentication, and biometric surveillance. In the generative AI era, this risk becomes more severe because facial information can be used not only for recognition but also for synthetic content creation. A face can become training material, a target for simulation, or an element in manipulated media.

The case is also significant because it involved inferred personal information. The regulator noted that the platform analyzed passengers' travel intention, resident city, and business or tourism status without sufficiently clear notification [14]. Inferred information is often less visible than directly collected information. Users may know they are providing a phone number or location permission, but they may not know that the platform is using their behavior to infer identity attributes, habits, economic status, social roles, or likely intentions.

Algorithmic inference creates a new category of privacy risk. It allows platforms to produce information about individuals that they never directly disclosed. Inferred data can influence service ranking, pricing, advertising, fraud control, credit assessment, employment opportunities, or access to platform benefits. In this sense, personal information protection should not only cover raw input data. It should also cover algorithmically generated profiles and predictions, especially when such profiles affect user rights or expose sensitive aspects of personal life.

7.3 From Platform Data Governance to Generative AI Governance

The Didi case provides a bridge between traditional platform regulation and generative AI governance. Although the case did not involve a generative AI model, it shows why AI governance cannot begin only at the stage of model output. Before a model generates text, images, predictions, recommendations, or decisions, it depends on data. These data may come from public sources, purchased datasets, user interactions, platform logs, business databases, or internal documents. If the data sources are unlawful, excessive, biased, insecure, or poorly documented, the AI system built on them will be structurally unsafe.

This point is especially relevant to AI training data compliance. Generative AI providers often emphasize model alignment, output moderation, and content filtering. These measures are necessary, but they are not sufficient. A model trained or fine-tuned on improperly obtained personal information may still create legal and ethical risks even if its outputs are later moderated. Platform companies therefore need upstream controls: data source review, dataset documentation, consent verification, sensitive information filtering, retention management, and deletion mechanisms.

The Didi case also suggests that AI governance should treat platform data systems as part of the AI system itself. In modern digital platforms, AI models are rarely standalone tools. They are integrated with user databases, recommendation engines, customer service systems, risk control systems, advertising systems, and operational dashboards. A generative AI assistant deployed inside a platform may access historical records, user profiles, transaction logs, location data, and internal knowledge bases through retrieval or tool-calling mechanisms. Therefore, the boundary between data governance and AI governance is increasingly artificial.

The core lesson is that responsible generative AI requires responsible platform data governance. Model safety, content moderation, and output labeling cannot compensate for an unlawful or opaque data foundation. A trustworthy AI system must be built on data that are lawfully obtained, purpose-limited, minimized, secured, documented, and auditable.

Table 3. Traditional platform issues and generative AI implications

Governance stage	Traditional platform issue	Generative AI implication
Data collection	Excessive permissions, unclear purpose, overcollection	Unlawful or unnecessary data may enter AI pipelines
Data storage	Weak access control, plaintext storage, poor segregation	Sensitive data may be exposed through model access or retrieval

Governance stage	Traditional platform issue	Generative AI implication
Data analysis	User profiling and inferred attributes	AI may generate or amplify inferred personal information
Data reuse	Secondary use beyond original purpose	Training, fine-tuning, or RAG may violate purpose limitation
Data audit	Limited internal accountability	AI risk assessment becomes unreliable without data audit trails

8. Discussion

Generative artificial intelligence has shifted privacy protection from a narrow consent-based issue to a broader concern over data infrastructure governance. Large language models and multimodal AI systems rely on extensive data collection, storage, annotation, model training, fine-tuning, and user interactions. Consequently, protecting personal information in the generative AI era requires not only attention to model outputs but also careful oversight of the underlying platform data ecosystems, ensuring that they are lawful, necessary, secure, auditable, and accountable. The Didi cybersecurity review and administrative penalty case exemplifies the challenges of platform-level data governance. In July 2022, the Cyberspace Administration of China fined Didi Global Inc. RMB 8.026 billion for violations of the Cybersecurity Law, Data Security Law, and Personal Information Protection Law, citing illegal processing of over 64 billion items of personal information, including sensitive data such as facial recognition, location, and identity information. This case demonstrates that large digital platforms can generate systemic privacy risks through continuous data accumulation, excessive processing of sensitive personal information, and creation of inferred user profiles. Traditional notice-and-consent mechanisms are insufficient to address these risks, highlighting the need for lifecycle-based governance.

For generative AI governance, the Didi case suggests several critical institutional responses. Training data compliance should be central, with rigorous review of the legality, necessity, and provenance of data sources. Sensitive personal information must be subject to heightened safeguards, including strict access control, encryption, retention limits, and exclusion from AI training unless legally justified. Platform accountability is essential: parent entities with unified data governance responsibilities must be held responsible for multi-app and multi-department practices, while data security audits should become routine and cover all stages of data processing, training, and AI deployment.

Risk assessment must integrate personal information protection with broader concerns of data security, public interest, and national security. The Didi case shows that weak platform data governance can propagate systemic AI risks if left unchecked. However, the Chinese governance model also has limitations. Enforcement-centered regulation demonstrates strong regulatory capacity and deterrence but may lack transparency, detailed proportionality, and due process. Balancing individual rights with national security framing is crucial, ensuring that privacy harms are not obscured and that international lessons are adapted thoughtfully.

The Didi case illustrates that effective generative AI governance cannot be restricted to monitoring model outputs or content moderation. It must address the underlying data infrastructures and operational practices that feed AI systems. A responsible governance framework should combine platform data compliance, sensitive information protection, lifecycle-based risk management, algorithmic accountability, auditability, independent review, and clear allocation of platform responsibility. Implementing these mechanisms is essential for building trustworthy AI systems that respect privacy while mitigating systemic risks.

9. Conclusion

This study has examined the Didi cybersecurity review and administrative penalty as a foundational case for understanding personal information protection in the generative AI era. Although the Didi case did not directly involve generative AI services, it reveals a core governance problem that is highly relevant to AI development: model-level risks are often rooted in platform-level data practices. Generative AI systems depend on large-scale data collection, storage, classification, retrieval, training, fine-tuning, and user interaction. If the underlying data infrastructure is unlawful, excessive, opaque, or weakly protected, AI systems built upon it may reproduce and amplify those risks. Therefore, platform-level governance constitutes the precondition for any trustworthy generative AI deployment.

The main finding of this study is that platform data governance is essential for mitigating systemic privacy and security risks. The Didi case demonstrates that large digital platforms may accumulate massive volumes of personal information, excessively process sensitive personal data, and generate inferred user profiles that users are often unaware of. Traditional notice-and-consent mechanisms are insufficient in such complex ecosystems because users cannot meaningfully track how their data are collected, combined, analyzed, reused, or transformed. Consequently, personal information protection must evolve from formal compliance toward comprehensive lifecycle-based governance, encompassing all stages of data handling.

Finally, this study emphasizes that AI governance cannot be limited to model outputs, content moderation, or user-facing safeguards. Effective generative AI governance requires training data compliance, dataset provenance, sensitive information filtering, internal access control, retention policies, deletion mechanisms, and independent security audits. The Didi case highlights the bridge between platform regulation and AI regulation, illustrating that governance must start before model training and continue throughout the full data lifecycle. While the study is limited by its focus on a single Chinese case and reliance on official regulatory materials, it offers a clear lesson: trustworthy AI cannot be built on weak data governance. Integrating platform data compliance, sensitive information protection, algorithmic accountability,

security auditing, and lifecycle-based risk management is critical for the responsible development and deployment of generative AI.

References

- [1] N. Carlini, F. Tramèr, E. Wallace, M. Jagielski, A. Herbert-Voss, K. Lee, A. Roberts, T. Brown, D. Song, Ú. Erlingsson, A. Oprea, and C. Raffel, “Extracting Training Data from Large Language Models,” in Proc. 30th USENIX Security Symposium, 2021, pp. 2633–2650.
- [2] N. Carlini, D. Ippolito, M. Jagielski, K. Lee, F. Tramèr, and C. Zhang, “Quantifying Memorization Across Neural Language Models,” in Proc. 11th Int. Conf. Learning Representations, 2023.
- [3] H. Kibriya, W. Z. Khan, A. Siddiq, and M. K. Khan, “Privacy issues in large language models: A survey,” Computers and Electrical Engineering, vol. 120, Part A, Art. no. 109698, 2024, doi: 10.1016/j.compeleceng.2024.109698.
- [4] B. C. Das, M. H. Amini, and Y. Wu, “Security and Privacy Challenges of Large Language Models: A Survey,” arXiv:2402.00888, 2024.
- [5] J.-H. Hoepman, “Privacy Design Strategies,” in ICT Systems Security and Privacy Protection, IFIP Advances in Information and Communication Technology, vol. 428, 2014, pp. 446–459, doi: 10.1007/978-3-642-55415-5_38.
- [6] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé III, and K. Crawford, “Datasheets for Datasets,” Communications of the ACM, vol. 64, no. 12, pp. 86–92, 2021, doi: 10.1145/3458723.
- [7] S. Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York, NY, USA: PublicAffairs, 2019.
- [8] J. van Dijck, T. Poell, and M. de Waal, *The Platform Society: Public Values in a Connective World*. Oxford, U.K.: Oxford University Press, 2018.
- [9] European Parliament and Council of the European Union, “Regulation (EU) 2016/679 of 27 April 2016 on the protection of natural persons with regard to the processing of personal data,” Official Journal of the European Union, L 119, pp. 1–88, May 4, 2016.
- [10] European Parliament and Council of the European Union, “Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence,” Official Journal of the European Union, L, Jul. 12, 2024.
- [11] I. Calzada, “Citizens' Data Privacy in China: The State of the Art of the Personal Information Protection Law,” Smart Cities, vol. 5, no. 3, pp. 1129–1150, 2022, doi: 10.3390/smartcities5030057.
- [12] Standing Committee of the National People's Congress, “Personal Information Protection Law of the People's Republic of China,” adopted Aug. 20, 2021, effective Nov. 1, 2021.
- [13] Cyberspace Administration of China, “Decision on Administrative Penalties Related to the Cybersecurity Review of Didi Global Inc.,” Jul. 21, 2022. [Online]. Available: https://www.cac.gov.cn/2022-07/21/c_1660021534306352.htm
- [14] Cyberspace Administration of China, “Official Answers to Reporters' Questions on the Administrative Penalty Decision Related to the Cybersecurity Review of Didi Global Inc.,” Jul. 21, 2022. [Online]. Available: https://www.cac.gov.cn/2022-07/21/c_1660021534364976.htm
- [15] Standing Committee of the National People's Congress, “Cybersecurity Law of the People's Republic of China,” adopted Nov. 7, 2016, effective Jun. 1, 2017.
- [16] Standing Committee of the National People's Congress, “Data Security Law of the People's Republic of China,” adopted Jun. 10, 2021, effective Sept. 1, 2021.
- [17] State Council of the People's Republic of China, “Regulations on Network Data Security Management,” State Council Order No. 790, promulgated Sept. 24, 2024, effective Jan. 1, 2025.
- [18] Cyberspace Administration of China et al., “Interim Measures for the Management of Generative Artificial Intelligence Services,” Order No. 15, Jul. 10, 2023, effective Aug. 15, 2023.
- [19] Cybersecurity Review Office, “Announcement on Initiating Cybersecurity Review of Didi Chuxing,” Jul. 2, 2021. [Online]. Available: https://www.cac.gov.cn/2021-07/02/c_1626811521011934.htm
- [20] Cyberspace Administration of China, “CAC and Six Other Departments Enter Didi Chuxing Technology Co., Ltd. to Conduct Cybersecurity Review,” Jul. 16, 2021. [Online]. Available: https://www.cac.gov.cn/2021-07/16/c_1628023601191804.htm