KHAZANAH RESEARCH INSTITUTE

VIEWS 23/24 | 24 OCTOBER 2024

Data relationality: privacy in the AI age

Khoo Wei Yang

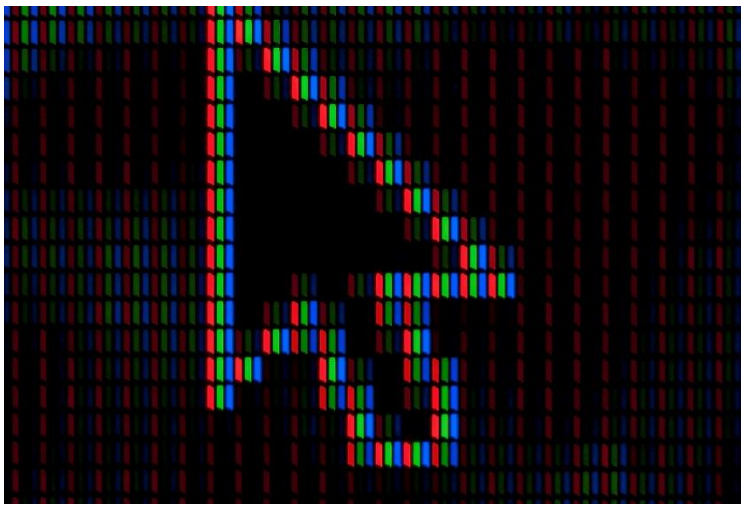

Introduction

Existing privacy protections are not sufficient to curtail big tech automated decision-making. Automated decisions are increasingly widespread and can have harmful impacts.

Artificial Intelligence (AI) relies on vast amounts of data. Data's social or relational properties can reveal information about individuals that wasn't directly provided. This reduces the meaningful control individuals have over their data.

This article explores the tension between data production practices and privacy protection in the AI age.

This view was prepared by Khoo Wei Yang, a researcher from the Khazanah Research Institute (KRI). The author is grateful for the valuable comments from Dr Rachel Gong, Dr Jun-E Tan, Gregory Ho Wai Son, and Nik Syafiah Anis Nik Sharifulden.

Author's email address:
khoo.weiyang@krinstitute.org

Automated decision-making directly impacts our lives

Companies increasingly deploy AI systems to make automated decisions about millions of customers and workers.

Consider how digital ride-hailing platforms dispatch great numbers of ride matches at a time[1] and tailor dynamic pricing[2] using AI. Often, these decisions are presented as a choice for the users, but in a significantly constrained sense[3]. Drivers, in particular, may be harmed by automated decisions about their tasks and remuneration, leaving them with little control over their work[4]. At a sufficiently large scale, collective harm can occur when people with similar conditions are affected by the same decisions[5].

In advanced economies, the harms of automated decisions by AI systems were reported in healthcare insurance[6], unemployment benefits[7], and more. AI systems' underperformance in their operational contexts may be a risk, but there are other problems associated with automated decisions that go beyond accidental harm.

How is automated decision-making done?

In commercial applications, automated decision-making is based on actionable insights derived from predictive analytics. AI systems enable analytics to be automated and scaled up. AI systems can combine vast amounts of data from various sources, process data and make decisions autonomously for millions of users (or cases) at a time.

For example, in recommender systems, machine learning (ML) models predict users' preferred products, movies, or music by learning from a dataset of other users with similar browsing or purchasing histories. The model outputs a recommended list as an actionable insight.

ML models rely on an immense volume of data[8]. The more data an ML model is trained on, the better the accuracy of the output[9]. As ML models' performance depends on the size and quality of data, companies are clamouring to expand the scale of data production[10].

This has led to the mushrooming of the data production industry dedicated to the collection, processing, storage, and circulation of data. Not only are individuals now subjected to collection

---

[1] Ling (2023)
[2] Lei and Ukkusuri (2023)
[3] For instance, is the choice legitimate when user is only presented with a premediated options instead of the full gamut of options?
[4] Tan and Gong (2024)
[5] Efforts have already been made to regulate the harms of automated decisions. The European Union General Data Protection Regulation (GDPR), for example, provided for the right of data subjects to opt out of automated decisions and profiling that significantly affects them (Art 22).
[6] Lopez (2023)
[7] Lam (2013)
[8] Budach et al. (2022)
[9] Hestness et al. (2017)
[10] *Yahoo Finance* (2023)

of identifiable personal information, but also to an expanding surveillance of their behaviour, turning all aspects of life into data. This is known as dataification[11].

Legal scholars have pointed out the incompatibility of privacy with data production in the AI age, owing to the fact that data is social or relational in nature. The advancement in statistical tools and AI has changed the ways in which data are processed and used. To understand this incompatibility, we must learn how value is derived from data relationality.

The value of social data

Data is social

In 2018, Cambridge Analytica harvested user data through their app "thisisyourdigitallife" to develop predictive psychological profiles used to target users with similar profiles for political advertisements on Facebook[12]. In this case, most Facebook users did not disclose their data to Cambridge Analytica but accurate prediction had exposed them to ad-targeting.

What the incident has demonstrated is a problem of privacy. The ability of Cambridge Analytica's algorithm to make predictions about one group based on information collected elsewhere suggested that information reveals relationships between people.

Consider a financial services platform that uses an ML model trained on user data such as browsing histories, socio-economic class, and financial product preferences. Suppose Alice shares only her browsing history with this platform. The model infers sensitive information about her, such as socio-economic class and financial interests, from her browsing data. Suppose the platform uses this inferred information to target her for advertisements of financial products. In that case, Alice is affected by the data of others, independent of her choice in disclosing the target information.

Salome Viljoen (2021) called this the "relationality" of data[13]. Relationality refers to the phenomena where information about others has the potential to reveal information about us when processed or aggregated[14].

Data production is motivated by the social nature of data

Individual datum is not useful in itself; it is only by relating one datum to another that meaningful links are derived to inform valuable insights[15]. According to Viljoen (2021),

> *In the digital economy, data isn't collected solely because of what it reveals about us as individuals. Rather, data is valuable primarily because of how it can be aggregated and*

[11] Mehta (2023)
[12] Rehman (2019)
[13] Viljoen (2021)
[14] Parsons and Viljoen (2023). Relationality in this article shall not be confused with the concept of relational database, which is a type of database that organises data in predefined relationships.
[15] Ashraf (2020)

> *processed to reveal things (and inform actions) about groups of people. Dataification, in other words, is a social process, not a personal one.*[16]

Companies and organisations now voraciously collect data to produce predictive analytics about users. More data give better approximations about groups and relationships between the features linked to users.

Machine learning (ML) aims to "automatically detect meaningful patterns in training data" to make predictions about new data[17]. This ability to gain insights and automate decisions is crucial for deriving value from data. The goal is to develop a prediction rule that approximates the relationship between pieces of information, such as correlations between input features and target variables[18].

In a way, models construct identities at an aggregated level, sometimes called "profiles."[19] For example, "women earning below median wage" is an input variable or a profile that groups individuals based on similar characteristics[20]. The prediction rule approximates the relationship between profiles and a target variable, such as the likelihood of women earning below the median wage in taking loans. This is a target function or "pattern". ML seeks to predict the target variable in the new data based on the patterns modelled in the training data.

A subset of users' data is selected as training data to train a predictive model. These are often data of users who disclose some target information like gender, or earnings. A prediction rule is modelled between the target information and some readily available auxiliary information like browsing history, clicks, and latency. The prediction rule modelled from this pool of data is then used to infer new data from the rest of the users, even if they haven't explicitly disclosed the target information (Figure 1). This prediction is produced as an actionable insight to either make automated decisions for users, such as ad targeting, or aid in decision-making.

[16] Viljoen (2021)
[17] Shalev-Shwartz and Ben-David (2014)
[18] van Otterlo (2013)
[19] Ibid.
[20] Taylor (2017); Wachter (2022)

Figure 1: Schematic diagram of a typical ML pipeline

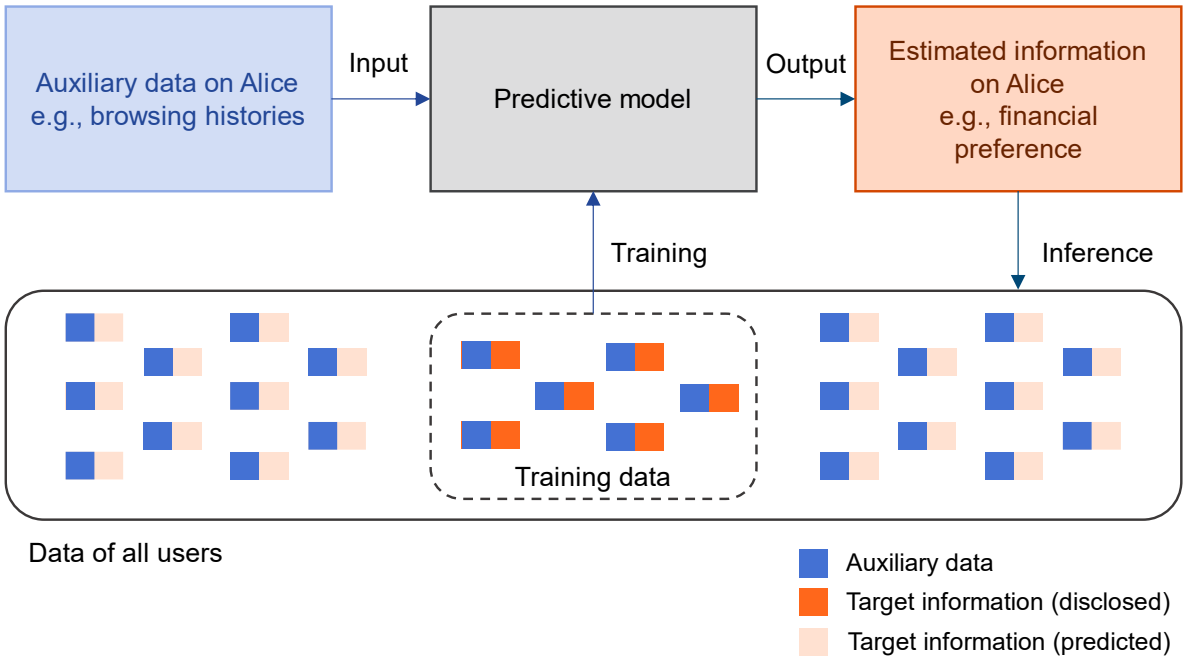

Source: Adapted from Mühlhoff (2023)

In a way, predictive analytics amounts to building profiles, estimating information, and making decisions about them.

The capacity of prediction is independent of whether an individual's target data was part of the training dataset, provided that sufficient auxiliary information is available[21]. While data subjects can control the information they disclose, they cannot control the estimated information about them. This becomes problematic in two ways: (1) when private attributes are inferred from publicly shared information and (2) when the effect on the data subject occurs without their consent.

Current privacy protections are insufficient

The dominant regulatory approach to information flow is a combination of transparency and choice, also known as notice-and-consent or informed consent[22]. The approach "requires that individuals be notified and grant their permission before information about them is collected and used"[23]. The approach also stresses the role of the individual as data subjects and their autonomy in information disclosures. Hence, regulatory efforts often emphasise the protection of personal information or personally identifiable information[24].

[21] Dwork and Roth (2014)
[22] Barocas and Nissenbaum (2014)
[23] Susser (2019)
[24] This view has dominated data protection laws. For example, the Malaysian Personal Data Protection Act (PDPA) requires data users to inform data subjects on processing of personal data under the notice and choice principle

Data relationality undermines privacy protection based on informed consent. Data protection laws protect information at an individual level, whereas AI sidestepped the need for an individual's informed consent to learn information about that individual.

The ability of AI to produce highly accurate predictions about us based on aggregated information of others erodes privacy. Thus, there are constraints to the extent of meaningful control one has over their data.

Privacy disclosures (or privacy notices) that inform how users' data are collected and used are now widely implemented across the web. According to this view, the data subject's privacy is protected so long as people have legitimate control over the permissions they give to disclose their personal information.

In reality, most digital platforms implement opt-in contracts on a "take-it-or-leave-it" basis for their services[25]. These opt-in contracts leave users with little deciding power, as big digital platforms accrue users by undercutting competition from alternative platforms[26].

Helen Nissenbaum posited the impracticality of informed consent in the Internet age[27]. Modern Big Data analytics draw and combine data from various sources. Companies also trade data among each other[28], making it hard for users to assess the trade-offs for giving away their information. The ability of AI to infer private information about us from public auxiliary information such as cookies, clickstreams, latencies, IP addresses, and so on makes drawing boundaries between private and public information a futile exercise and individual privacy calculus[29] infinitely tricky.

(required by means of written notice Article 7 Act 709, 2010). The EU GDPR, one of the strongest data protection law, goes a further step in requiring "controllers", or data processor to uphold "data subject rights". The regulation buttresses principles of notice-and-consent by requiring meaningful consent and transparent information on personal data processing. The regulation further enforces data minimisation, data portability, integrity and confidentiality, as well as accountability clauses (see Burgress, 2020) .

[25] Nissenbaum (2011); Guirguis and Howarth (2019)

[26] Feiner (2024)

[27] Nissenbaum (2010); (2011)

[28] Cyphers and Gebhart (2019)

[29] Privacy calculus refers to the decision making process of trading off privacy for the benefits of disclosing information. For example, an internet user may trade-off her contact information for a service.

Conclusion

In the age of information flow, where data collection, processing, and use are everywhere, data governance is crucial. The crux of data governance is about managing the tension in "balancing data openness and control"[30]. Because data brings about essential benefits in the public interest, improved access to and broader sharing of data are crucial to expanding the reach of benefits that raise living standards[31]. Conversely, data misuse and unjust outcomes can arise from loose data flow[32].

Protection of privacy has been one of the critical principles for data handling to strengthen trust in information systems[33]. However, the current regimes of privacy protection rely on individualist notions of information control. This may not be sufficient to safeguard society from harms derived from an economy driven by social predictions based on shared data.

References

Ashraf, Shaharudin. 2020. "Open Government Data: Principles, Benefits and Evaluations." Discussion Paper. Kuala Lumpur: Khazanah Research Institute. http://www.krinstitute.org/Discussion_Papers-@-Open_Government_Data-;_Principles,_Benefits_and_Evaluations.aspx.

Barocas, Solon, and Helen Nissenbaum. 2014. "Big Data's End Run around Anonymity and Consent." In *Privacy, Big Data, and the Public Good*, edited by Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum, 1st ed., 44–75. Cambridge University Press. https://doi.org/10.1017/CBO9781107590205.004.

Budach, Lukas, Moritz Feuerpfeil, Nina Ihde, Andrea Nathansen, Nele Noack, Hendrik Patzlaff, Felix Naumann, and Hazar Harmouch. 2022. "The Effects of Data Quality on Machine Learning Performance." https://arxiv.org/abs/2207.14529.

Burgess, Matt. 2020. "What Is GDPR? The Summary Guide to GDPR Compliance in the UK." *Wired*, March 24, 2020. https://www.wired.com/story/what-is-gdpr-uk-eu-legislation-compliance-summary-fines-2018/.

Cyphers, Bennett, and Gennie Gebhart. 2019. "Behind the One-Way Mirror: A Deep Dive Into the Technology of Corporate Surveillance." Electronic Frontier Foundation. https://www.eff.org/wp/behind-the-one-way-mirror.

Dwork, Cynthia, and Aaron Roth. 2014. "The Algorithmic Foundations of Differential Privacy." *Foundations and Trends® in Theoretical Computer Science* 9 (3–4). Now Publishers, Inc.:211–407. https://doi.org/10.1561/0400000042.

[30] OECD (2022)
[31] Ashraf (2020)
[32] Marcucci et al. (2023)
[33] Ibid.

Feiner, Lauren. 2024. "Judge Rules That Google 'Is a Monopolist' in US Antitrust Case." *The Verge*, August 5, 2024. https://www.theverge.com/2024/8/5/24155520/judge-rules-on-us-doj-v-google-antitrust-search-suit.

Guirguis, Ayman, and David Howarth. 2019. "ACCC's Digital Platforms Report: Market Power in Advertising, Search Services & Media & Privacy Implications." *K&L Gates* (blog). August 12, 2019. https://www.klgates.com/ACCCs-Digital-Platforms-Report-Market-Power-in-Advertising-Search-Services--Media--Privacy-Implications-08-12-2019.

Hestness, Joel, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. 2017. "Deep Learning Scaling Is Predictable, Empirically." arXiv. https://doi.org/10.48550/arXiv.1712.00409.

Lam, Khoa. 2013. "Incident Number 373: Michigan's Unemployment Benefits Algorithm MiDAS Issued False Fraud Claims to Thousands of People." Edited by Khoa Lam. *AI Incident Database*. Responsible AI Collaborative. https://incidentdatabase.ai/cite/373.

Lei, Zengxiang, and Satish V. Ukkusuri. 2023. "Scalable Reinforcement Learning Approaches for Dynamic Pricing in Ride-Hailing Systems." *Transportation Research Part B: Methodological* 178 (December):102848. https://doi.org/10.1016/j.trb.2023.102848.

Ling, Bo. 2023. "Innovative Recommendation Applications Using Two Tower Embeddings at Uber." Uber Blog. July 26, 2023. https://www.uber.com/blog/innovative-recommendation-applications-using-two-tower-embeddings/.

Lopez, Ian. 2023. "UnitedHealthcare Accused of AI Use to Wrongfully Deny Claims (1)." *Bloomberg Law*, November 15, 2023. https://news.bloomberglaw.com/health-law-and-business/unitedhealthcare-accused-of-using-ai-to-wrongfully-deny-claims.

Marcucci, Sara, Natalia Gonzalez Alarcon, Stefaan G. Verhulst, and Elena Wullhorst. 2023. "Mapping and Comparing Data Governance Frameworks: A Benchmarking Exercise to Inform Global Data Governance Deliberations." arXiv. https://doi.org/10.48550/arXiv.2302.13731.

Mehta, Mita. 2023. "Monitoring Algorithm for Datafication and Information Control for Data Optimization." In *ICT with Intelligent Applications*, edited by Jyoti Choudrie, Parikshit N. Mahalle, Thinagaran Perumal, and Amit Joshi, 1–7. Singapore: Springer Nature. https://doi.org/10.1007/978-981-99-3758-5_1.

Mühlhoff, Rainer. 2023. "Predictive Privacy: Collective Data Protection in the Context of Artificial Intelligence and Big Data." *Big Data & Society* 10 (1). SAGE Publications Ltd:20539517231166890. https://doi.org/10.1177/20539517231166886.

Nissenbaum, Helen. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford University Press.

———. 2011. "A Contextual Approach to Privacy Online." *Daedalus* 140 (4):32–48. https://doi.org/10.1162/DAED_a_00113.

OECD. 2022. "Going Digital: Guide to Data Governance Policy Making." Paris: OECD. https://doi.org/10.1787/49a65317-en.

Otterlo, Martijn van. 2013. "A Machine Learning View on Profiling." In *Privacy, Due Process and the Computational Turn*. Routledge.

Parsons, Amanda, and Salome Viljoen. 2023. "Valuing Social Data." *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4513235.

Rehman, Ikhlaq ur. 2019. "Facebook-Cambridge Analytica Data Harvesting: What You Need to Know." *Library Philosophy and Practice (e-Journal)*, January, 2497.

Shalev-Shwartz, Shai, and Shai Ben-David. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781107298019.

Susser. 2019. "Notice After Notice-and-Consent: Why Privacy Disclosures Are Valuable Even If Consent Frameworks Aren't." *Journal of Information Policy* 9:37. https://doi.org/10.5325/jinfopoli.9.2019.0037.

Tan, Jun-E, and Rachel Gong. 2024. "Algorithmic Management and Societal Relations: The Plight of Platform Workers in Southeast Asia." Kuala Lumpur: Khazanah Research Institute.

Taylor, Linnet. 2017. "Safety in Numbers? Group Privacy and Big Data Analytics in the Developing World." In *Group Privacy: New Challenges of Data Technologies*, edited by Linnet Taylor, Luciano Floridi, and Bart van der Sloot, 13–36. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-46608-8_2.

Viljoen, Salome. 2021. "A Relational Theory of Data Governance." *Yale Law Journal* 131 (2):370–781.

Wachter, Sandra. 2022. "The Theory of Artificial Immutability: Protecting Algorithmic Groups Under Anti-Discrimination Law." *Tulane Law Review* 97 (2):149.

Yahoo Finance. 2023. "Global Datafication Market Report 2023-2028 Featuring IBM, Oracle, Microsoft, SAP, Google, AWS, HPE, SAS Institute, Teradata, and Dell," November 22, 2023. https://finance.yahoo.com/news/global-datafication-market-report-2023-125300773.html/.