AI Compliance – Challenges of Bridging Data Science and Law

PHILIPP HACKER, European New School of Digital Studies, European University Viadrina, Germany
FELIX NAUMANN, Hasso Plattner Institute, University of Potsdam, Germany
TOBIAS FRIEDRICH, Hasso Plattner Institute, University of Potsdam, Germany
STEFAN GRUNDMANN, Humboldt University of Berlin, Germany
ANJA LEHMANN, Hasso Plattner Institute, University of Potsdam, Germany
HERBERT ZECH, Weizenbaum Institute, Humboldt University of Berlin, Germany

This vision article outlines the main building blocks of what we term *AI Compliance*, an effort to bridge two complementary research areas: computer science and the law. Such research has the goal to model, measure, and affect the quality of AI artifacts, such as data, models, and applications, to then facilitate adherence to legal standards.

CCS Concepts: • **Applied computing → Law**; • **Information systems → Data cleaning**; • **Security and privacy → Human and societal aspects of security and privacy**; • **Social and professional topics → Intellectual property;**

Additional Key Words and Phrases: AI Act, compliance, liability, privacy, transparency, information quality

ACM Reference format:
Philipp Hacker, Felix Naumann, Tobias Friedrich, Stefan Grundmann, Anja Lehmann, and Herbert Zech. 2022. AI Compliance – Challenges of Bridging Data Science and Law. *J. Data Inform. Quality* 14, 3, Article 21 (June 2022), 4 pages.
https://doi.org/10.1145/3531532

1 COMPLIANCE FROM A LEGAL AND A TECHNICAL PERSPECTIVE

We propose the notion of *AI Compliance* to bridge two complementary research fields, data science and law, and outline the associated challenges. Interdisciplinary analyses in data science have focused mostly not on law but on other "hard" sciences, such as biology or medicine. Whereas there is now a growing trend toward computational social science [8], and fruitful collaborations have been developed between data science and ethics [2, 9], this cannot be said, to the same extent, for the law. The rapid advance and deployment of artificial intelligence (AI) tools in all sectors of society, however, clearly calls for a new interdisciplinary paradigm: lawyers need to team up with computer scientists to answer some of the most pressing technical and policy questions our society

Authors' addresses: P. Hacker, European New School of Digital Studies, European University Viadrina, Große Scharrnstr. 59, 15230 Frankfurt (Oder), Germany; email: hacker@europa-uni.de; F. Naumann, T. Friedrich, and A. Lehmann, Hasso Plattner Institute, University of Potsdam, Prof.-Dr.-Helmert-Str. 2-3, 14482 Potsdam, Germany; emails: {felix.naumann, tobias.friedrich, anja.lehmann}@hpi.de; S. Grundmann, Humboldt University of Berlin, Unter den Linden 6, 10099 Berlin, Germany; email: stefan.grundmann@rewi.hu-berlin.de; H. Zech, Weizenbaum Institute, Humboldt University of Berlin, Hardenbergstraße 32, 10623 Berlin, Germany; email: herbert.zech@rewi.hu-berlin.de.

1936-1955/2022/06-ART21 $15.00
https://doi.org/10.1145/3531532

faces. Bridging computer science and law seems all the more urgent as governments around the globe actively pursue a strategy on regulating AI, as outlined, for example, in the EU White Paper on AI (WPAI) [3] and the proposed EU AI Act [4].

The term compliance is sometimes understood as a subdiscipline of corporate law, dealing with the various tasks and regulations that compliance officers face in companies. From an information systems perspective, compliance usually describes various aspects of data governance. We take a more encompassing perspective, entailing two conceptual dimensions. First, AI compliance means that AI systems must follow the law; hence, the relevant law must be analyzed and data science tools must be developed to facilitate such compliance. Second, we submit that laws should also be critically scrutinized and re-designed, where appropriate, so that they can be effectively and efficiently complied with. Law's governance also of data use and retrieval is paramount.

Here, the term *compliance* comprises the many facets by which we can assess an AI system in a specific way that connects requirements from data science and law. While the precise determination of what AI Compliance in fact constitutes is a core challenge, we can already establish its cornerstones. Traditionally, computer science evaluates AI systems based on the correctness of their decisions, for instance using precision, recall, and accuracy measures based on specific test data. In contrast, society, and by extension the law, is interested in and regulates many further aspects: the transparency of the system; the adequacy of training and test data, including non-discrimination and privacy aspects; the appropriateness of its deployment for a certain use case; access rights to, and conversely (intellectual property) protection for, training data and trained models; and finally, its overall added value for the deployer. All these aspects, and possibly more, are suitable extensions to the traditional set of information quality criteria as defined in [13]. Just as Firmani et al. proposed the extension of data quality along the ethical dimension [5], we present the challenges posed by jointly regarding the legal dimension of data and data science and the challenge to ultimately create an *AI compliance framework* guiding the interplay between development and deployment of AI systems on the one hand and the interpretation and design of the corresponding legal fields on the other hand.

2 DIMENSIONS OF AI COMPLIANCE

We believe that six themes serve as unifying concepts to guide research across AI and law: *liability*, *transparency*, *intellectual property protection*, *privacy*, *information quality*, and *cost*. The choice of these six concepts reflects the heritage of existing scholarship at the intersection of AI and law [6, 7, 11]. A key challenge lies in linking hard quantitative thresholds from data science with usually openly textured legal concepts (e.g., fairness, defect, and accuracy) in these different fields.

Liability denotes legally relevant responsibility, usually implying the obligation to pay fines or damages in case harm occurs. In AI Compliance, liability may arise particularly in the context of the General Data Protection Regulation (GDPR), of general contract and tort law, including company law, and of intellectual property (IP) law. For instance, poor training data quality may result in liability of the developers. For training AI models, typically, liability law will take the *cost* of achieving better performance measures into account and balance it against respective safety gains [12, 15]. Finally, in concrete deployment contexts, liability will often hinge on the *quality* of the model, comprising not only performance on a test set, but also the inference performance in the deployment scenario.

IP protection describes legal exclusivity conferred to intangible assets like inventions or works of art. (Training) datasets and trained AI models are distinct yet novel intangible assets; to which extent they are subject to existing IP protection, and whether they should be, is hotly debated [14]. IP questions are conspicuously absent from [4], but should be high on the agenda of interdisciplinary research between AI and the law. Importantly, when *transparency requirements* extend to models themselves, complementary IP rights in the models and their deployment methods may serve as incentives to develop them in the first place.

Transparency denotes the availability of the relevant information for decision-making. There are many ways to achieve transparency of AI in data science, such as access to training data, model, and specific applications. Equally, different instruments to realize it exist in the legal space, e.g., mandatory disclosure or incentives to disclose. These interactions must be developed further, creating a mapping between the respective concepts. In concrete data acquisition or model deployment contexts, *transparency* may compete with *privacy* if the former comprises access to personal data, for example in auditing procedures or litigation. However, without access to application data, models, and potentially even training data, those harmed by AI models in concrete use cases often will not be able to prove causation of harm to enforce their claims. Resolving this tension constitutes a key challenge for real-world AI Compliance.

Data privacy is an aspect of data security that is concerned with the proper handling of (personal) data. Already by nature a strongly interdisciplinary field, it involves not only computer science but also law and social sciences. The advent of data-driven applications and AI has amplified the challenge of how to use and benefit from data, while preserving the users' right to privacy. Hence, privacy is a pivotal term both for legal analysis and technical constraints. Data subjects may, for example, have a right to be erased from the training data (Art. 17 GDPR), forcing technical adaptations to the datasets, such as "machine unlearning" [1, 10].

Information quality permeates many aspects of AI and the law. As quality of data, it comprises the wide range of quality dimensions of a given (training) dataset, such as correctness and completeness [13], but also criteria that we posit to be compliance-specific, such as understandability, origin, and diversity of training data. Notably, the GDPR also contains data quality aspects in the yet undefined principle of data accuracy and freshness. As quality of trained AI models, the concept represents the ability of an AI system to deliver legally compliant results. One particular challenge of AI Compliance research is to unify these facets in a common framework to guide data scientists and legal professionals in building and assessing AI systems. Another is to expand information quality by incorporating metrics guaranteeing the representativeness of the training data for the application context, and a feasible balance of the training dataset between different protected groups. Vague legal terms significantly complicate this endeavor.

Finally, the **cost** of an AI system comes in several guises, usually trading off with *data quality*. First, human and data cost is incurred by configuring, training, deploying, and maintaining AI systems. Second, cost occurs in terms of runtime performance of a system or method. And finally, legal risks can turn into monetary costs via damages and fines. The lack of *IP protection*, however, may engender an incentive problem for the generation of high-quality, discrimination-sensitive AI training data, given the significant cost such data preparation entails, e.g., in credit scoring.

3 OUTLOOK

To tackle the challenges of bridging data science and the law, future research should pursue a *two-way research strategy*. First, technical solutions may be optimized under legal constraints while holding current legal requirements constant. To achieve this, legal constraints must be translated to measurable dimensions of data and data science pipeline components. For instance, this concerns the operationalization of legal concepts for AI training data quality contained in Article 10 of the proposed AI Act [4], such as "relevance," "representativeness," and "freedom from errors." Conversely, the legislator should seek to optimize legal requirements while assuming certain technological capabilities or limitations. For example, IP law needs an update to decide who ought to own rights in AI creations (copyright) or AI inventions (patent). In this way, new compliance strategies can be developed for existing law; and new regulation can be designed with a view to current and future technical capabilities.

REFERENCES

[1] Yinzhi Cao and Junfeng Yang. 2015. Towards making systems forget with machine unlearning. In *IEEE Symposium on Security and Privacy (SP'15)*. 463–480.

[2] Mark Coeckelbergh. 2020. *AI Ethics*. MIT Press, Cambridge.

[3] European Commission. 2020. On artificial intelligence - a European approach to excellence and trust. White Paper, COM (2020) 65.

[4] European Commission. 2021. Proposal for a regulation of the European parliament and of the council laying down harmonised rules on artificial intelligence, COM (2021) 206 final. https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52021PC0206.

[5] Donatella Firmani, Letizia Tanca, and Riccardo Torlone. 2019. Ethical dimensions for data quality. *Journal on Data and Information Quality* 12, 1 (2019).

[6] Philipp Hacker. 2020. A legal framework for AI training data – From first principles to the artificial intelligence act. *Law, Innovation and Technology* 13 (2020). forthc., https://ssrn.com/abstract=3556598.

[7] Margot E. Kaminski. 2018. Binary governance: Lessons from the GDPR's approach to algorithmic accountability. *Southern California Law Review* 92 (2018).

[8] D. M. J. Lazer, A. Pentland, D. J. Watts, S. Aral, S. Athey, N. Contractor, D. Freelon, S. Gonzalez-Bailon, G. King, H. Margetts, A. Nelson, M. J. Salganik, M. Strohmaier, A. Vespignani, and C. Wagner. 2020. Computational social science: Obstacles and opportunities. *Science* 369 (2020), 1060–1062.

[9] B. D. Mittelstadt, P. Allo, M. Taddeo, S. Wachter, and L. Floridi. 2016. The ethics of algorithms: Mapping the debate. *Big Data & Society 3.2*. DOI: https://doi.org/10.1177/2053951716679679

[10] Sebastian Schelter. 2020. "Amnesia" – Machine learning models that can forget user data very fast. In *Proceedings of the Conference on Innovative Data Systems Research (CIDR'20)*.

[11] Andrew D. Selbst. 2020. Negligence and AI's human users. *Boston University Law Review* 1315 (2020). https://ssrn.com/abstract=3350508.

[12] Gerhard Wagner. 2019. Robot liability. In *Münster Colloquia on EU Law and the Digital Economy IV, Liability for Artificial Intelligence and the Internet of Things*, Lohsse/Schulze/Staudenmayer (Eds.). Nomos, Baden-Baden, 25–62.

[13] Richard Y. Wang and Diane M. Strong. 1996. Beyond accuracy: What data quality means to data consumers. *Management of Information Systems* 12(4) (1996), 5–34.

[14] Herbert Zech. 2019. Artificial intelligence: Impact of current developments in IT on intellectual property. *Gewerblicher Rechtsschutz und Urheberrecht, Internationaler Teil (GRUR Int.)* (2019), 1145–1147.

[15] Herbert Zech. 2019. Liability for autonomous systems: Tackling specific risks of modern IT. In *Münster Colloquia on EU Law and the Digital Economy IV, Liability for Artificial Intelligence and the Internet of Things*, Lohsse/Schulze/Staudenmayer (Eds.). Nomos, Baden-Baden, 185–200.

Received November 2020; revised December 2021; accepted March 2022