

The universal implications of artificial intelligence (AI) demand a radical reconsideration of how we approach all fields of study and research. AI law and policy is no different. Grappling with AI law and policy means considering how AI—and topics surrounding its regulation and litigation arising from its use—affects all aspects of society and government. This can involve theories of justice and governance structures, social and political movements, the digital economy and political economy, knowledge production and ethics, labor and employment, rights and privileges, freedom of expression and association, power and defense, security and sovereignty, and other frameworks relevant to the rule of law.

AI law and policy is an evolving field of legal inquiry that intersects with a wide range of other legal fields. Depending on one's legal perspective, AI law and policy can be seen as either encompassing or being encompassed by other areas of technology law and policy: data privacy and data protection, information security and cybersecurity, cybercrime and cyber law, algorithms and the law, intellectual property, Internet law and the Internet of Things, social media and platform regulation, First Amendment law, criminal procedure, and other areas of law. AI is embedded in a wide range of systems, such as surveillance, financial technology or fintech, and biometrics, to name just a few. Because AI systems are embedded within other systems unfolding in an information society and digital economy, AI law and policy is part of the DNA of other legal organisms and other legal ecosystems.

Additionally, AI architecture and the uses of AI are evolving. As the technology adapts, so, too, the laws, definitions, and standards surrounding AI are also adapting. The data market ecosystems and enabling applications, the digital infrastructures and other foundational data and technological structures, and other AI systems and structures are also accommodating the speed and nature of innovation. Scientists, corporations, nongovernmental organizations, and governmental bodies are setting different standards and establishing different definitions for AI and machine learning to address their unique needs and goals.

Reconciling law and technology—assessing AI's legal impact while also simultaneously assessing AI's technological impact—is an enormously complex undertaking for those in the field of AI law and policy. Attorneys, academics, policymakers, regulators, legislative and judicial bodies, industry leaders, civil society experts, and others are struggling to grasp

the technological capacities and impact of AI. Often this struggle occurs between preexisting legal frameworks and newly proposed regulatory regimes and litigation theories. Because of the emerging nature of the law and policy surrounding AI, which is developing as the technology is developing, the law surrounding AI is far from clear. Where AI regulations are lacking, and when AI technology has already been deployed, the significant legal implications of AI necessitate an evaluation of what laws currently exist that might be available in the future to redress AI's harms and risks, leaving AI law and policy often in a reactive state.

Consequently, there is no consensus on what constitutes the field of AI law and policy. This project does not purport to state definitively the parameters of the field of AI law and policy. Rather, it serves a more modest goal: It offers the scaffolding of a pedagogical tool. It presents one method for how to approach a discussion on the state of an evolving field that is referred to as AI law and policy.

Casebooks customarily reflect the law as it is understood in definitive sources, such as case decisions, statutes and regulations, and constitutions. AI law and policy, by contrast and necessity, incorporates a survey of complaints and soft law, such as policy guidance documents, due to the inchoate nature of the field. Representative complaints have been included in this book to reflect newly emerging legal theories in AI cases. Complaints do not reflect conclusions of law. Similarly, guidance documents have been included to reflect theories of how to enforce voluntary forms of AI oversight. Guidance documents also do not reflect instruments that possess the force of law. The field of AI law is awaiting legislative and judicial developments to set more conclusively the parameters of the field. Future editions will be adjusted accordingly.

This casebook, as a snapshot in a moment of time, provides a range of primary sources—complaints, cases, statutes, regulations, guidance documents, congressional reports, government press releases, and other materials—to contextualize how laws and policies are developing as we encounter new consequences of the integration of AI into our lives and digital economy. Because AI law and policy is still a work in progress, only a limited number of illustrative primary sources have been included. The excerpts of these primary sources reflect a case study approach to introduce the subject matter. This does not represent a comprehensive survey, but, rather, an introduction to AI governance concepts through the vehicle of select primary sources and source summaries. Any casebook or textbook is page-limited. The page constraints required that in some instances, only summaries of the sources are included—for example, a press release or a public announcement. In other instances, I intentionally selected summaries and excerpts to increase accessibility, such as news articles over academic articles. The full text of many of the sources can be found in the statutory supplement and elsewhere.

Finally, the focus of this book project, intended for a survey course that can be offered in a wide range of disciplines beyond schools of law and policy, is on primary sources. Citations to secondary sources are few. References to the rich academic debates surrounding AI law and policy, and an engagement with ongoing and important scholarly research in this exciting field, have been limited. Similarly, the sometimes conflicting goals of both simplicity and breadth, and the reality of page limitations, led to the development of a structure that is more akin to a sourcebook and that offers only limited introductions to the sources. This allows for the pedagogical focus to be trained on the excerpted sources. Each chapter is intended to offer introductory concepts only and, to exemplify select introductory concepts, provide an overview of limited case studies to serve illustrative purposes. I recognize that this book—and the case study approach of this fast-moving field—is not exhaustive and may incentivize a search for greater context. This is available in supplementary materials provided by this publisher and others.

Editorial decisions to exclude certain sources in any casebook project are arduous. The aims of the project undertaken here made these decisions particularly difficult, given the scale and tempo of the subject matter. The editorial team and I acknowledge how contested these decisions might be in a field of law as nascent as AI law and policy. I invite a conversation on revisions and expansions in future editions.