



AI AND LAW

HOW AUTOMATION IS CHANGING THE LAW

**Aurelia Tamò-Larrieux Clement Guitton,
and Simon Mayer**

A **Chapman & Hall** Book



CRC Press
Taylor & Francis Group

AI and Law

This book provides insights into how AI is changing legal practice, government processes, and individuals' access to those processes, encouraging each of us to consider how technological advances are changing the legal system. Particularly, and distinct from current debates on how to regulate AI, this book focuses on how the progressive merger between computational methods and legal rules changes the very structure and application of the law itself.

We investigate how automation is changing the legal analysis, legal rulemaking, legal rule extraction, and application of legal rules and how this impacts individuals, policymakers, civil servants, and society at large. We show through many examples that a debate on how automation is changing the law is needed, which must revolve around the democratic legitimacy of the automation of legal processes, and be informed by the technical feasibility and tradeoffs of specific endeavors.

Chapman & Hall/CRC Artificial Intelligence and Robotics Series

Series Editor: Roman Yampolskiy

AI in and for Africa

A Humanist Perspective

Susan Brokensha, Eduan Kotzé, Burgert A. Senekal

Artificial Intelligence on Dark Matter and Dark Energy

Reverse Engineering of the Big Bang

Ariel Fernández

Explainable Agency in Artificial Intelligence

Research and Practice

Silvia Tulli

David W. Aha

An Introduction to Universal Artificial Intelligence

Marcus Hutter, Elliot Catt, and David Quarel

AI: Unpredictable, Unexplainable, Uncontrollable

Roman V. Yampolskiy

Transcending Imagination:

Artificial Intelligence and the Future of Creativity

Alexander Manu

Responsible Use of AI in Military Systems

Jan Maarten Schraagen

AI iQ for a Human-Focused Future

Strategy, Talent, and Culture

Seth Dobrin

Federated Learning

Unlocking the Power of Collaborative Intelligence

Edited by M. Irfan Uddin and Wali Khan Mashwan

Designing Interactions with Robots

Methods and Perspectives

Edited by Maria Luce Lupetti, Cristina Zaga, Nazli Cila, Selma Šabanović, and

Malte F. Jung

The Naked Android

Synthetic Socialness and the Human Gaze

Julie Carpenter

For more information about this series please visit: <https://www.routledge.com/Chapman--HallCRC-Artificial-Intelligence-and-Robotics-Series/book-series/ARTILRO>

AI and Law

How Automation Is Changing the Law

Aurelia Tamò-Larrieux, Clement Guitton, and
Simon Mayer



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business
A CHAPMAN & HALL BOOK

Designed cover image: Shutterstock

First edition published 2025

by CRC Press

2385 NW Executive Center Drive, Suite 320, Boca Raton FL 33431

and by CRC Press

4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

CRC Press is an imprint of Taylor & Francis Group, LLC

© 2025 Aurelia Tamò-Larrieux, Clement Guitton, and Simon Mayer

Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyrighted material has not been acknowledged, please write and let us know so we may rectify it in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, access www.copyright.com or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. For works that are not available on CCC please contact mpkbookspermissions@tandf.co.uk

Trademark notice: Product or corporate names may be trademarks or registered trademarks and are used only for identification and explanation without intent to infringe.

ISBN: 9781032480060 (hbk)

ISBN: 9781032464527 (pbk)

ISBN: 9781003386919 (ebk)

DOI: 10.1201/9781003386919

Typeset in Minion

by codeMantra

*To children:
who make writing so much harder
and life so much better!*



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Contents

Preface, xi

Authors, xiii

CHAPTER 1 ■ AUTOMATION OF LAW	1
1.1 WHY THIS BOOK?	2
1.1.1 First Wave of Legal Automation	4
1.1.2 Second Wave of Legal Automation	7
1.1.3 Third Wave of Legal Automation	11
1.2 A GUIDE TO THIS BOOK	12
CHAPTER 2 ■ LAW AND COMPUTER SCIENCE INTERACTIONS	16
2.1 IF-THEN: LAWS AS CODE	19
2.2 ENCODING LEGAL KNOWLEDGE FOR MACHINES	31
2.2.1 Types of Machine Knowledge	34
2.2.2 Exemplifying the Translation of a Legal Norm into Its Machine-Processable Version	43
2.3 ENTERS MACHINE LEARNING	47
NOTES	53
CHAPTER 3 ■ AUTOMATICALLY PROCESSABLE REGULATION	54
3.1 TERMINOLOGIES AND TYPOLOGIES	55
3.1.1 Applying the Typology to Our Initial Examples	60
3.1.2 Challenges When Classifying Applications of Automatically Processable Regulation	61

3.2	EFFICIENCY YOU SAID?	63
3.3	WHEN AND WHEN NOT TO	69
3.3.1	“Involves Calculations”	70
3.3.2	“Can Be Delivered Digitally”	72
3.3.3	“Repeated Process”	73
3.3.4	“Compliance Process”	74
3.3.5	“Factual Information”	77
	NOTES	81
CHAPTER 4 ■ CHALLENGES AND CONTROVERSIES		82
4.1	REPRESENTATIVE AND BALANCED AUTOMATICALLY PROCESSABLE REGULATION	84
4.2	COMPUTER SAYS NO	86
4.3	TRANSPARENT AND CONTESTABLE AUTOMATICALLY PROCESSABLE REGULATION	87
4.4	REPLACEMENT OF THE HUMAN TOUCH AND WORKFORCE	92
4.5	IS THIS FOR REAL?	93
4.6	TOWARDS RESPONSIBLE AUTOMATICALLY PROCESSABLE REGULATION	95
	NOTE	102
CHAPTER 5 ■ NEEDED (PUBLIC) DEBATES		103
5.1	LAW FOR ALL: IS THERE A MANDATE TO MAKE LAW ACCESSIBLE?	104
5.2	TO WHICH EXTENT SHOULD THE STATE STRIVE TO MAKE THE LAW DIGITALLY READY?	108
5.3	HOW TO PROMOTE LEGAL DESIGN THINKING?	113
CHAPTER 6 ■ HOW EDUCATION SHOULD SHIFT		117
6.1	PRIMARY AND SECONDARY EDUCATION	120
6.1.1	Ubiquitous Computers and Ubiquitous Computing	120
6.1.2	Connected Computers	124

6.1.3	Data Trails and Their Privacy Implications	130
6.1.4	Data Analytics: Understanding Statistics and Biases	136
6.1.5	Pioneer?	138
6.2	SPECIALIZED EDUCATION	140
	NOTES	145
CHAPTER 7 ■ EXERCISES		146
7.1	DESIGNING A LEGAL DECISION TREE	146
7.2	TURNING NATURAL LANGUAGE INTO CONTROLLED LANGUAGE	148
7.3	MODELING A RULE	149
7.4	CLASSIFYING AUTOMATICALLY PROCESSABLE REGULATION PROJECTS	150
7.4.1	Rates Rebate	151
7.4.2	Mes Aides	152
7.4.3	Victor	155
7.4.4	DoNotPay	155
7.4.5	Overtime Regulation	157
7.5	IDENTIFYING OPEN-TEXTURED TERMS	157
7.6	DEBATING ABOUT ISSUES OF AUTOMATING LEGAL PROCESSES	160
	NOTES	161
EPILOGUE, 162		
	PERSONALIZING LAW?	162
	REGULATING AUTOMATICALLY PROCESSABLE REGULATION	164
ACKNOWLEDGMENTS, 166		
GUIDING APPROACHES FOR SOLUTIONS, 168		
	EXERCISE 7.1: DESIGNING A LEGAL DECISION TREE	168

EXERCISE 7.2: TURNING NATURAL LANGUAGE INTO CONTROLLED LANGUAGE	172
EXERCISE 7.3: MODELING A RULE	174
EXERCISE 7.4: CLASSIFYING AUTOMATICALLY PROCESSABLE REGULATION PROJECTS	175
EXERCISE 7.5: IDENTIFYING OPEN-TEXTURED TERMS	176
EXERCISE 7.6: DEBATING ABOUT ISSUES OF AUTOMATING LEGAL PROCESSES	176
REFERENCES, 177	
INDEX, 191	

Preface

COMING FROM EACH OF the different disciplinary perspectives, law, political science, and computer science, we approached the topic of automation and law distinctly. From a legal perspective, Aurelia built upon her expertise on privacy by design approaches to analyze the bi-directional relationship of law and automation, meaning how law can mandate certain automated processes and how automation can be leveraged to fulfill regulatory goals. From a political science perspective, Clement analyzed the structural sociopolitical changes in policymaking that the automation of law generates and could further generate, as well as how feasible (technically and from a social acceptability point of view) those changes can take place, including depending on how responsible they are. Finally, from a computer science perspective, Simon built upon his expertise in increasingly autonomous cyber-physical systems and the deployment of such systems in ubiquitous computing environments where they interact with people and society. These different “starting points” or viewpoints of the authors were key for the writing of this book and the many joint research papers that form its basis. The differences in approach kickstarted numerous engaging discussions, were the basis for reaching out to the community through the organization of workshops, and formed a central part of the interdisciplinary work presented in this book. Working across disciplines truly attempting to understand and integrate disciplinary perspectives is a lot of work that does not immediately bear fruit. Keep at it: It is important to understand each other, and it keeps giving back.

As we will show throughout this book, automatically processable regulation is an ideal topic to discuss from our three disciplinary perspectives: Encoding regulation gives not only a chance to revisit how people gain knowledge about the legal system in which they operate, and which concerns them in so many ways. From a political science perspective, it offers many insights on how we can think of reshaping the social contract with

the state by, for instance, making sure that access to legal processes is not left to the few but to the many. Automation of legal processes could work towards enabling laypeople to either know how to apply the law or at least get access to legal processes in a way that suits their needs, and they could lead to products that respect the prevailing regulation. From a societal perspective, we hence face a situation where, to understand future regulatory environments, individuals will require a fundamental understanding of both law and automation, which poses a formidable challenge to our education system.

To ensure that the law is automated *responsibly*, we would love to see more dialogs emerge that are as interdisciplinary as this book. We view this as a necessity not only in research but also across the broader society.

Prof. Aurelia Tamò-Larrieux*

Dr. Clement Guitton*

Prof. Simon Mayer

Switzerland, 2024

* Aurelia and Clement share the first authorship.

Authors

Aurelia Tamò-Larrieux is an associate professor at the University of Lausanne (UNIL), Faculty of Law, heading the team of Digital and Computational Law and leading the Legal Design & Code Lab. Previously, she was an assistant professor at Maastricht University and a member of the Law & Tech Lab. She has a background in law and economics and specializes in research at the intersection of law and digital technologies with numerous publications and projects in the field of privacy, data protection, design approaches, transparency of automated decision-making, automatically processable regulation, and trust in automation.

Clement Guitton is currently a postdoctoral researcher at the University of St. Gallen, focusing on topics around law and technology, or more generally on understanding and developing policy solutions around technology. He previously worked on cyber security, bringing together the fields of politics and technology. He has published dozens of peer-reviewed journal articles, and *AI and Law* is his fourth book. He has degrees in telecommunication engineering, international affairs, and finance. In the past, he worked in counter-espionage, consulting at the International Telecommunication Union and as a political analyst for a large reinsurance company.

Simon Mayer is a professor at the Institute of Computer Science at the University of St. Gallen, Switzerland, where he heads the Chair for Interaction- and Communication-based Systems. Previously, he worked at Siemens Corporate Technology in Berkeley, USA, most recently as senior key expert for Smart and Interacting Systems, and afterward headed the Cognitive Products research group at the Austrian COMET research center Pro2Future at Graz University of Technology. Simon holds degrees in Computer Science, Economics and Business Administration,

and Secondary and Higher Education; he has co-authored more than 200 journal, magazine, conference, and workshop publications on the topics of ubiquitous computing, the Web of Things, autonomous systems, and interoperability, and is serving on the steering committee of the International Conference on the Internet of Things as well as the editorial board of the *IEEE Pervasive Computing* magazine.

Automation of Law

MORE AND MORE, LEGAL experts are focusing on creating legal rules that computers can understand and process automatically. To do so, they leverage new technologies to extract and apply legal rules. With this change towards the automation of law, an abundance of terms—from *computer-encoded regulation*, *computable* or *computational* law, *legal informatics*, *legal tech*, and *legal artificial intelligence* (AI)—have emerged, leading to an array of academic discussions that have provided us with a rich set of different examples of and viewpoints on what we refer to as *automatically processable regulation* (we will explore the meaning of the term more in details in Chapter 3: Automatically Processable Regulation). These examples are useful to take a bottom-up approach to understand how we can create legal rules that computers can process and how we can leverage automation to extract and apply legal rules. In addition, the focus on concrete and real-world use cases of automatically processable regulation enables us to disentangle the various concerns that come with encoding regulation. In this book, we address the question of why everyone, from legal experts to policymakers, to engineers, to educators to civil society, should pay attention to the developments around automation and law. Throughout the book, we showcase through multiple use cases what is at stake for individuals and society when automation meets the law.

Science fiction movies sometimes capture these developments, albeit often in dark terms, and can be used to illustrate the changes: In *Elysium*, the protagonist, played by the actor Matt Damon, must interact with an automated civil servant, namely a robot parole officer. The robot officer is portrayed as an AI embodied in a metallic, human-like interface (with a

suit and a tie) which is programmed to know the legal rules applicable and can perceive the tone and behavior of the protagonist. After a frustrating interaction with the robot officer, one scene ends with the question, “Would you like to talk to a human?” and Matt Damon answers in a tone imitating a robot: “No. I am OK. Thank you”.—indicating his discontent with the automated process. Automating legal processes might be more efficient in handling individual complaints, yet by no means leads to systems that are more accepted by citizens, and debates with the general public about how much automation of the law we want are crucial even if currently insufficient (see Chapter 5: Needed (Public) Debates). While the movie *Elysium* plays in the future, such interactions with “legally knowledgeable” machines are not that futuristic. In fact, current research has shown the breadth of domains where automation of legal rules occurs already today: In Brazil, we witness an array of new applications of automated systems within the judiciary, for instance to assess the merit of cases presented to the Supreme Federal Court; around the world systems to assess re-offense risks of criminals have been deployed; and in China, so-called *Internet Courts* that include “non-human judges” are handling cases that are today mostly related to e-commerce (Vasdani, 2020). Aside from these examples, where legal processes handled by the state are being automated, we see multiple private initiatives to make legal rules more accessible to laypeople. In the UK and in the US, for instance, the famous *DoNotPay* application was heavily used to contest parking tickets, and endless legal software to automate various often repetitive tasks such as contract drafting, or the filing of mandates has emerged under the *LegalTech* heading. More recently, with the increased public access to large language models, natural language processing has come to the forefront of the discussion, especially with a recent version of a Generative Pre-trained Transformer (GPT) model, OpenAI’s GPT-4, being able to solve bar exam questions better than the average human bar exam taker in 2023 (Katz, 2023; Martínez, 2024). This raises the question of how far off the sci-fi future described in *Elysium* and similar movies is, and how we want to shape it to benefit society.

1.1 WHY THIS BOOK?

In 2024, a great many people are talking about AI and law. European regulators have proposed regulation in 2021 (and adopted it in May 2024) on how to tackle risky applications of AI systems—focusing on the discriminative AI systems that were abundant at that time. These models are able to classify inputs along a decision boundary and to compute regressions—in statistical terms, they are able to map a given input to a class label

(e.g., hot/cold; car/bicycle/bus) conditional on observed probabilities in training samples—i.e., they capture conditional probabilities, $P(y|x)$. With the rise of generative AI, which generates new data based on learned patterns of training data and became highly visible worldwide in the form of ChatGPT at the end of 2022, policymakers around the world have started deliberating about the best ways to enable further research within the field while mitigating its risks to fundamental rights, such as balancing the freedom to information and democracy when seeking to curtail the rapid spread of disinformation and deepfakes, or in the context of challenges to copyrights of authors. While the title of this book—*AI and Law*—seems to immediately link to such new debates on how to regulate AI, this book takes a different perspective and discusses how the progressive merger between computational methods and legal rules *changes the very structure and application of the law itself*. We uncover how automation, including current developments in the field of AI, changes the legal field, breaking away from traditional and normative analysis of the role of law in society. We do not ask “How should law regulate automation?” but instead investigate the question: “How is automation changing the legal analysis, legal rulemaking, legal rule extraction, and application of legal rules, and how does this impact individuals, policymakers, civil servants, and society at large?”

To understand the impact of these developments, it is important to have an understanding of how the field of automatically processable regulation has evolved over time. An important parenthesis first: We use the term *automatically processable regulation* throughout the book to refer to any regulation that is expressed in a form that makes it accessible to being processed automatically and where there is a will and intention of encoding said regulation for a specific purpose (e.g., efficiency gains) as automatically processable regulation (Guitton, Tamò-Larrieux, & Mayer, 2022b). We will expand on this definition when looking at how to classify different projects that fall within the term in Chapter 3. But before we get there, let us first go back to the development of automatically processable regulation: Contissa and Sartor (2022) have proposed to describe these developments in *three seminal waves* that shaped the field. While the earliest forms of automation’s impact on law can be thought of as going back to the writing of laws on material artifacts like wood or paper (Hildebrandt, 2015), these three seminal waves describe contemporary automation processes, such as the electronic storing and retrieval of legal texts for rapid search and retrieval, the semi-automatic application of legal rules, and the modeling of the law through machine-learning techniques to predict legal outcomes. By flying

over some of the technical details of each of these waves, it will already be possible to distinguish different kinds of limitations to automating the law (which are also further explored in Chapter 2: Law and Computer Science Interactions).

1.1.1 First Wave of Legal Automation

The first wave started in the 1960s with the increased digital access to legal resources (Contissa & Sartor, 2022). These legal resources typically contain metadata and are structured in available formats that enable querying them. Text (or information) retrieval systems may then be applied to digitally available documents that exhibit structured data, such as statutes, regulations, and case law. Lawyers and legal scholars work daily with such information; these systems enable them to rapidly find relevant information, such as precedents. Generally, an information retrieval system starts by *indexing* available texts, which means that an organized database is created that can map terms to the locations where they occur within the indexed documents; indexing is relevant for enabling the efficient retrieval of information from documents since for a large number of documents, searching all of them for each new prompt that is given by a user is impractical. After indexing (which sometimes continues in the background), users are given the option to formulate queries—for instance, as keywords, phrases, or in the form of complex queries. Complex queries can be created, for instance, by using logical operations (i.e., AND, OR, NOT) to join simple clauses—these are commonly referred to as *Boolean* searches (hinting at George Boole, who conceived of Boolean algebra as a cornerstone in the development of digital electronics); alternatively, systems might make use of a *structured query language* that is more powerful than Boolean search. After users provide a query, the information retrieval system attempts to match this query with the indexed documents, aiming to achieve high recall (a recall score of 1.0 indicates that all indexed documents that are relevant to the query have been returned; see Figure 1.1 and Box 1.1) as well as high precision (a precision score of 1.0 indicates that all documents that have been returned are relevant to the query).

Because systems that only optimize one of these metrics are highly problematic (e.g., a perfect recall score of 1.0 can be attained by simply returning *all* documents in the database), the quality of information retrieval systems is often expressed as an F-score that is computed from both the precision and recall scores: The F_1 score, a type of F-score, achieves this by computing the harmonic mean of the *precision* and *recall* scores:

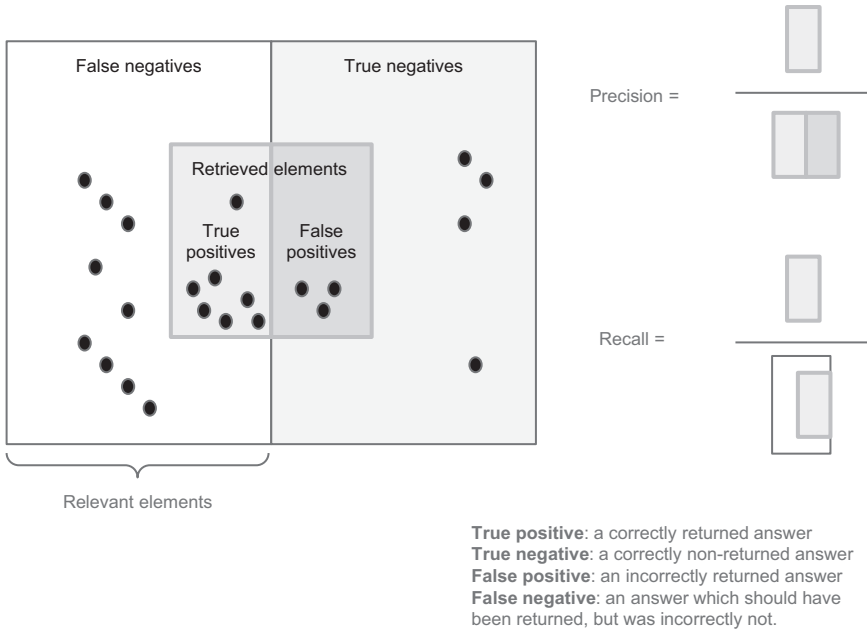


FIGURE 1.1 Depiction of precision and recall used in the computation of F1-scores.

$$F_1 = 2 * \frac{\text{precision} * \text{recall}}{\text{precision} + \text{recall}} = \frac{2 * \text{true positives}}{2 * \text{true positives} + \text{false positives} + \text{false negatives}}$$

In Figure 1.1 above, precision corresponds to 7/10 (with 7 retrieved true positives divided by 10 retrieved elements) while recall corresponds to 7/16 (with 7 retrieved true positives divided by 16 relevant elements), and hence the F_1 score is 0.538. A perfect information retrieval system would return all 16 relevant elements and none of the 7 irrelevant elements, yielding an F_1 score of 1.0.

BOX 1.1 EXAMPLE OF COMPUTING SCORES WHEN PERFORMING A RETRIEVAL TASK

Imagine you are working on a legal document search engine, and you are evaluating its performance on a set of search results for a specific query. You have a set of relevant documents for that query, and your search engine has returned a set of documents as well.

This is the information you have:

- Total number of documents in the collection: 5,000
- Number of query-relevant documents in the collection: 50
- Total number of documents retrieved by the search engine: 80
- Number of query-relevant documents retrieved by the search engine: 40

First, calculate the recall, precision, and F_1 score for the search engine's performance:

- Recall (R) = (Number of Relevant Documents Retrieved) / (Total Number of Relevant Documents) = $40/50 = 0.8$
- Precision (P) = (Number of Relevant Documents Retrieved) / (Total Number of Documents Retrieved) = $40/80 = 0.5$
- F_1 score (F_1) = $2 * (\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall}) = 2 * (0.5 * 0.8) / (0.5 + 0.8) = 0.6154$ (rounded)

Through an upgrade, the search engine can retrieve almost all relevant documents (i.e., 49 instead of 40). However, the number of retrieved documents also increased from 80 to 200. Based on the F_1 score, would you upgrade your search engine?

- Recall (R) = (Number of Relevant Documents Retrieved) / (Total Number of Relevant Documents) = $49/50 = 0.98$
- Precision (P) = (Number of Relevant Documents Retrieved) / (Total Number of Documents Retrieved) = $49/200 = 0.245$
- F_1 Score (F_1) = $2 * (\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall}) = 2 * (0.98 * 0.245) / (0.98 + 0.245) = 0.392$

Therefore, the search engine should not be upgraded based on the F_1 score. If the target application benefits from a higher recall more than the lower precision hurts it, the F_1 score should be adapted or abandoned as a quality metric for this case.

There are many ways how an information retrieval system may match and rank documents given a user query, where typical systems will consider *term frequency* (i.e., the number of occurrences of a term in each document) or *inverse document frequency* (IDF; a measure that emphasizes occurrences of terms that are rare in the overall document body). Advanced information retrieval systems will, in addition, employ natural language

processing techniques to analyze the content of documents semantically (i.e., considering synonyms, related terms, or the context of a term). For *ranking* of returned results, a range of relevant metrics can be taken into account: For instance, the very well-known *BM25* family of scoring functions ranks documents according to an occurrence-based weighing of the inverse document frequency. While *BM25* and similar functions ignore the proximity of different terms in a document, ranking based on inverse document frequency is still today prevalent in the majority of information retrieval systems. Alternatively, and similar to the *PageRank* algorithm that was introduced with the Google search engine, search results might be ranked according to how often they are cited from other parts of a text.

Similar to how the Google search engine has changed how users of the World Wide Web—a *global* information system—find results online, such text retrieval systems changed how legal practitioners or government employees discover relevant legal rules. It has hence already altered the way certain legal materials are made visible (Ashley, 2017). This happened, for instance, through relevance ranking algorithms that highlight certain legal materials over others, where citation-based metrics propel highly cited decisions to the top of the results that an information retrieval query returns; it also happens through semi-summarization which automatically selects fragments that it estimates to be more relevant within a document.

1.1.2 Second Wave of Legal Automation

The *second wave* moves from improving access to sources to the establishment of *automated applications* of formal accounts of the law (Contissa & Sartor, 2022). A lot of focus, already early on, has been set on manual, rule-based representations of legal rules. One famous example has been the formalization of the *British Nationality Act* in 1986 using logic programming (Sergot et al., 1986). The idea behind this formalization was to transform legal rules into *if-then* statements that can be represented as a computer program and executed by a computer. To bring an example, we look at the equivalent law in Switzerland, whose first article stipulates that: “*The following persons are Swiss citizens from birth: a child whose parents are married to each other and whose father or mother is a Swiss citizen [...]*”. This statement can be expressed as if-then rules as shown in Figure 1.2; the condition “X is Swiss” will then evaluate to “True” if it holds that (1) X is a child and that (2) this child’s parents are married and that (3) the child’s father or the child’s mother or both are Swiss (see Box 1.2).

BOX 1.2 EXAMPLE OF IF-THEN MODELING

X is Swiss

IF X is a child

AND X's parents are married to each other

AND (X's father is Swiss OR X's mother is Swiss)

This transformation into rules often relies on propositional logic statements, which are either true or false and will be elaborated in Chapter 2: Law and Computer Science Interactions. However, and as will also be shown in the next chapter, propositional statements can be augmented by predicate logic which introduces variables, or predicates, that can further qualify non-logical objects. The reason why propositional logic statements are so appealing is because of their simplicity, and of the ease with which both humans and machines can understand if-then statements. In fact, research even shows that humans tend to understand legal documents better than the law itself when they are formulated in propositional logic statements (Guitton, Tamò-Larrieux, Mayer, & Zumbrunn, 2025). However, it requires a reformulation of laws to fit the scheme of simple if-then statements (Allen & Engholm, 1978). If-then statements can further be visualized as (logical) circuit diagrams that graphically illustrate the linking between statements within a given legislation. Applied to the extract from Swiss nationality law shown above, this representation is shown in Figure 1.2. There, the two logic gates represent the “AND” and “OR”, respectively, and the output of the circuit (on the very right) will be “True” if the clause is satisfied and “False” otherwise.

As the quest to visualize legal norms in this way is not new (see e.g., Allen & Engholm, 1978), we have already witnessed some of the difficulties

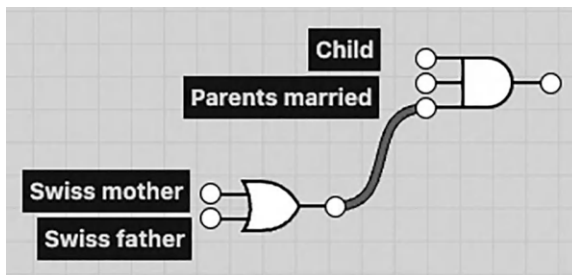


FIGURE 1.2 Example of a legal article modeled with logic gates (created with <https://logic.ly/demo/>).

that come with it. These difficulties are linked with the need to ensure that the representation of legal texts into logical expression is correct and efficient in terms of the resources needed for the translation or mapping. Both correctness and efficiency raise different challenges and are inherently interconnected: For a rule-based system to be “correct” means interpreting a legal rule at a certain time. This might, at times, simply not be feasible due to the vagueness of terms used in a law and due to the spectrum of possible interpretations of the law itself (see Chapter 3: Automatically Processable Regulation). In addition, the law is dynamic, and taking this temporal aspect of law into account is not a trivial task, as legal knowledge bases must be updated with changes in legislation and application of legal norms (Contissa & Sartor, 2022). Moreover, legal rules are at times conflicting and to apply them “correctly”, an automated system must understand how to resolve conflicting norms (e.g., in cases where a specific legal rule has priority over more general legal rules, the so-called *lex specialis* maxim). Here, *defeasible* reasoning comes in (see Box 1.3), which is a type of logic that acknowledges that a certain type of legal norm or legal decision is not absolute but can be challenged and overridden under certain conditions (e.g., in case a more specialized rule applies as for *lex specialis*) (Governatori, Rotolo, & Sartor, 2005; Hage, 2003).

BOX 1.3 EXAMPLE OF DEFEASIBLE REASONING

General Rule: *All vehicles are prohibited in the park.*

Specialized Rule: *Emergency vehicles are allowed in the park.*

We can express these rules using predicates:

- $V(x)$ represents the condition “ x is a **vehicle**”
- $P(x)$ represents the condition “ x is **prohibited** in the park”
- $E(x)$ represents the condition “ x is an **emergency vehicle**”
- $A(x)$ represents the condition “ x is **allowed** in the park”

To address the conflict between the general and specialized rule, an automated system uses defeasible reasoning. The conflict can be expressed as:

Rule 1: $\forall x \quad [V(x) \rightarrow P(x)]$ (All vehicles are prohibited in the park. Please note that \forall is read “for all” and will be explained further in the next chapter).

However, since the general rule that $\forall x \quad [E(x) \rightarrow V(x)]$ (All emergency vehicles are vehicles) now would lead to the prohibition of emergency vehicles, we require an additional, defeasing statement:

Rule 2: $\forall x \ [E(x) \rightarrow A(x)]$ (All emergency vehicles are allowed in the park.)

Thus, Rule 2 contradicts Rule 1. In a classical reasoner, this would lead to the reasoner rejecting the logic program. Hence, defeasible reasoning is required to represent this circumstance.

Aside from clear conflicts in law, such as the example given in Box 1.3, even in a given specific context without apparent conflict of norms, it regularly happens that legal scholars, lawyers, and judges themselves disagree with each other about how a specific regulation should be interpreted. As Schafer (2010) on the automatic interpretation of wills and testaments states: “If it is possible that even the top experts rationally hold mutually contradictory opinions, what exactly is the ‘knowledge’ that the computer models, and on what basis is the decision what to include, or which one to choose, taken?” (p. 387).

These difficulties related to the correctness of the legal knowledge representation are linked to the ones on efficiency, as formalizing legal texts into computable models “is a highly specialized skill, requiring both legal and technical expertise. The process of formalizing legal texts requires a considerable expenditure of resources (time, workforce) to overcome the difficulties that systems face in scaling up from small prototypes to large applications. Further resources are needed to keep the system up to date” (Contissa & Sartor, 2022, p. 30)—in addition, only recently a labor market for individuals with this combination of skills has been emerging (e.g., with “legal engineers”). However, despite these limitations, logic and rule-based approaches are still useful when automating legal reasoning. Specifically, rule-based systems are inherently more transparent than statistical approaches (e.g., neural networks, see below): For each decision that our above system takes about assigning Swiss citizenship, it is immediately possible to ask the system *why* it has assigned citizenship in a specific case, and in the opposite case it is possible to ask the system which of the conditions would need to change in order to assign citizenship. Upon such an investigation, the system is then able to give a precise response about which rules have been combined with what facts to yield the decision, which permits not only scrutiny but also straightforward repairing of facts or knowledge (or the program code of the system itself) if the system yields a wrong or biased result. Societally highly important, the system’s decisions can hence be *interpreted* and *explained*, including, crucially, to individuals who are subjected to the system’s judgment and who might disagree.

1.1.3 Third Wave of Legal Automation

Finally, the *third wave* has been propelled by the deployment of machine learning in the legal domain, enabling, through the use of natural language processing techniques, legal text analytics, and extraction of legal semantics (i.e., meaning) in court cases, contracts, legislation, or legal scholarship (Contissa & Sartor, 2022). Machine learning techniques have been deployed to automate the summarization of legal documents (e.g., Ash, Kesari, Naidu, Song, & Stambach, 2024), discover correlations across various documents to support the identification of relevant information needed in legal practice (Ashley, 2017), and converting legislation into machine-interpretable rules (Xu & Ashley, 2023). Likewise, a broad area of interest has been the application of machine learning to court cases, and especially for predicting outcomes based on the features and facts of a case. Within this domain, commercial applications have emerged partially also with the vision to democratize the law by facilitating access to the law through affordable AI tools or so-called “robot-lawyers” (Guitton, Tamò-Larrieux, & Mayer, 2022b). One example of such commercial applications is *ROSS Intelligence*, a company created in 2014, which is focused on the US legal system and leverages natural language processing techniques to understand, retrieve, and rank legally relevant information. Many other examples of tools exist today that leverage machine learning algorithms to analyze legal texts, reason about specific conclusions and legal arguments, and learn from new legal sources. Overall, such applications can help to answer legal questions and explain legal decisions and assessments to legal practitioners, judges, or laypeople (Contissa & Sartor, 2022). A compelling example of the support that such systems may provide is *Claudette* (see Figure 1.3), which takes its acronym from *Clause Detect* and analyzes consumer contracts to flag unfair clauses within these legal documents (Lippi et al., 2019). To do so, the researchers identified categories of *potentially*



FIGURE 1.3 CLAUDETTE from <http://claudette.eui.eu/demo/>.

unfair clauses based on the literature and case law, such as clauses that establish jurisdiction for disputes in a country different than the consumer's residence, limitations of liability, or the provider's right to unilaterally modify the contract. To identify possible violations, they extracted a corpus of 50 relevant online consumer contracts (i.e., *Terms of Service* of online platforms) and manually annotated the contracts based on the different categories of clauses by grouping them into three categories: *Clearly fair clauses*, *potentially unfair clauses*, or *clearly unfair clauses*. Upon the annotated corpus, they studied how machine learning can be applied to first detect sentences with (potentially) unfair clauses and second predicting the category an unfair clause belongs to. After testing different results depending on the models used for the analysis, the researchers released their prototype to the public (Lippi et al., 2019).

Tools like Claudette also show the potential for consumer-empowering tools that can be generated through the automation of legal rule extraction and reasoning, especially when coupled with newer machine learning approaches and natural language processing techniques. At the same time, such changes in how law or legally relevant terms are applied and analyzed also require legal scholars, lawyers, judges, and government officials to gain a better understanding of what is at stake when AI and law meet. As developments within the field of AI continue to evolve, educators, policymakers, and practitioners need to understand the underlying factors and systemic shifts that go hand in hand with this evolution. This book provides a gateway to this interdisciplinary field of research and domain of application by building upon key developments and projects in the field of automatically processable regulation and debating the tradeoffs involved when law meets automation.

1.2 A GUIDE TO THIS BOOK

The aim of this book is to gather diverse information on how automation and AI are altering the legal field and make it accessible to a wider audience beyond interdisciplinary researchers in the field. In particular, our focus rests on educating the next generation of individuals interested in interdisciplinary research situated at the intersection of law, computer science, and political science. We believe that to be well-equipped for the future of legal practice, policymaking, and public service, it is key that *students* across disciplines are aware of how automation is changing these fields and how to leverage such changes to the benefit of individuals and society, while mitigating the challenges that arise within this process. This book also

targets *policymakers* who are required to grasp these developments to promote an informed public discourse around AI and law. Such a discourse needs to revolve around the *democratic legitimacy* of the automation of legal processes, informed by the *technical feasibility* of specific endeavors. As for *legal practitioners*, this book provides insights into how the AI and law field is changing legal practice, encouraging them to analyze how technological advances are altering their profession. Finally, *developers* working on automating law should be able to grasp the potential challenges and issues that their technologies will face and cause. Developers require a good understanding not only of the techniques they use to create a legal automation solution, but they also need to develop sensitivity about the decisions they make (partially implicitly) while implementing this solution and need to consider the overall ecological viability of their products.

The remainder of this book is structured around five main questions that will uncover diverse scholarship and perspectives: First, what is technically speaking automatically processable regulation? Second, what use cases of automatically processable regulation are well described, and what can we learn from them? Third, what challenges and issues arise from those use cases, and what can we deduce from them? Fourth, what public debates are needed to address the challenges and issues identified? And fifth, what does the future of AI and the law hold and how can we shape future (educational) developments? More specifically:

Chapter 2 (entitled *Law and Computer Science Interactions*) presents a current picture of how automation is changing the law to a broad target audience. To cater to this audience, which we do not require to be trained lawyers or legal scholars nor to bring a background in computer science, we review in this chapter the most important technical foundations that are required to put the content of this book into perspective, in a way similar to how we have presented the three waves in this first chapter. We start by understanding simple modeling of the law in if-then rules. We show how such rule-based models can be created manually and how approaches to make law digitally ready can be promoted through the use of controlled language, which builds upon propositional and predicate logic. Upon this basis, we discuss how machines may encode legal meaning through semantic technologies and how ontologies and semantic networks provide us with the language and format to turn regulation into an automatically processable form. Because encoding legal knowledge is a time-consuming task, many researchers have turned to machine learning approaches to support rule extraction or reasoning. Within the last part of Chapter 2, we

explain how those machine learning approaches are currently changing the field of AI and law.

Chapter 3 (entitled *Automatically Processable Regulation*) dives into what automatically processable regulation is on the ground by taking concrete use cases to illustrate the different approaches that exist in Europe and worldwide. These projects range from the creation of robot-judges to the deployment of tools that provide citizens with access to social benefits. Such use cases or projects are key to understand the various facets of automatically processable regulation and provide a starting point to describe the terminologies encountered in the field of AI and law and typologies that have emerged to enable comparing different use cases with each other. It also provides the basis to discuss whether automatically processable regulation really leads to increased efficiency of legal processes overall, as well as the underlying specifications of law that lend themselves better to being processed automatically. A typical example of such an underlying specification of the law is norms that involve calculation, a task in which rule-based machines are clearly able to outperform humans. Within Chapter 3 we discuss these different aspects of automatically processable regulation in order to also understand when societal challenges due to the implementation of automated legal processes arise, thus building the foundation for the discussions in Chapter 4.

Chapter 4 (entitled *Challenges and Controversies*) tells real stories of automatically processable regulation that went wrong. By means of those stories, we disentangle different categories of issues, ranging from a lack of transparency over how legal processes, when automated, are rendered and a lack of contestability of the decisions to issues tied to the often-vague nature of law and the need to balance conflicting interests. To stay with the latter set of issues, many scholars have pointed out the risks of freezing certain interpretations of the law in time if rendered automatically. As the law is drafted in ways that enable its evolution (e.g., by including terms that are open-textured, meaning that they are likely to have more than one meaning depending on the context and person interpreting the norm), hardcoding an interpretation to a machine to execute is bound to be too restrictive and likewise not be open to evolve when cultural perceptions and contexts change. Having a deep understanding of these issues and tradeoffs at hand is important to develop *responsible* guidelines on how to develop and deploy automatically processable regulation in the real world. Within Chapter 4, we provide a framework that can be applied by multiple stakeholders—from developers to government bodies

developing automatically processable regulation—to ensure an informed debate around the challenges that arise and mitigation measures that can be applied to minimize the risk of arising controversies .

Chapter 5 (entitled *Needed (Public) Debates*) builds upon the different issues discussed and posits that public debates are needed. We argue that a central discussion must occur around the topic of whether policymakers have a duty to make law more accessible to their citizens and, therefore, in part a duty to automate—in a responsible manner—certain legal processes for their citizens (e.g., access to social benefits, access to justice). While this debate in itself is not new, the discussed second and third waves of legal automation have changed the scope of what we can consider “accessible” beyond just physical access to a legal document (The Engine Room, 2019). These discussions on accessibility are occurring on a policy level with initiatives promoting the development and deployment of digitally ready legislation. While different terminologies can be used, digitally ready laws at their core need to be drafted in a manner that simplifies their transformation into a digital process. Such initiatives connect to the field of legal design thinking, which has developed methodologies and tools to put the user at the center of the legal (automated) solutions that are being developed.

Chapter 6 (entitled *How Education Should Shift*) looks into how these developments have changed the ways we educate our future generations. As different types of literacy are needed for a functioning democratic society—from functional literacy (that pertains to reading and writing) to civic literacy—it is central that we rethink not only specialized education (in particular in law and public governance) but also primary and secondary education. Students will need a diverse set of skills to navigate in an increasingly digitalized environment in which legally relevant decisions can be taken in milliseconds. Current regulations, from the *AI Act* to the *Data Act* in the European Union, acknowledge the need for greater AI and data governance literacy overall. These initiatives are welcomed and set a legal basis to rethink and promote AI and data literacy in the European Union. We describe the educational shifts required and elaborate on initiatives already taken in this domain. We couple the discussion with **Chapter 7** (entitled *Exercises*), where we provide exercises that can be used or adapted to critically examine the different topics discussed throughout the book. We hope that these exercises provide a baseline to further engage the reader to think about ways that AI and law interact!

Law and Computer Science Interactions

LAW AND COMPUTER SCIENCE are in constant interaction and represent two tools in the toolbox used to govern: As Lawrence Lessig famously pointed out, software and technical architectures (in short, “code”) govern how computers and networks operate and, in doing so how we can interact with technology—hence his famous statement, “code is law” (Lessig, 2003). He refers to this as *West Coast Code*, as digital technologies have been driven by Silicon Valley. *East Coast Code*, on the other hand is the more traditional mode of governance by legal rules and government institutions that apply them. Both forms of governance shape human behavior and impact societal norms. We often think of law as regulating *how* technology is developed and implemented, such as within the European Union’s AI Act that prohibits certain implementations of AI systems, whereby law *constrains* certain developments or applications of technology. However, the law has more functionalities than constraining activities, which can be subsumed within the regulatory strategy of command and control (Baldwin, Cave, & Lodge, 2011). In fact, an important functionality of the law is to *level the playing field* within markets (Gasser, 2016). This leveling function can take different forms, such as information disclosures that help generate transparency over how certain systems work (e.g., transparency disclosures with respect to how AI systems have been trained) or rights to demand access to specific data held by companies (e.g., in the European Data Act, end-users may request data that are generated through the use of smart products and

related services). Another function of the law that Gasser (2016) mentions is that of *enabling* specific innovation. Law can shape how technologies are designed and enable new implementations of systems that enhance certain regulatory goals and societal values. Legal norms can combine different functionalities: A case in point here are the regulatory developments towards privacy by design, which can constrain the design of new services and products as well as enable new innovative privacy-friendly solutions by demanding the technical implementation of overarching principles. How these principles are achieved is not constrained.

Privacy by design was popularized as a concept in the year 1990 with Ann Cavoukian, former *Information and Privacy Commissioner* in Ontario, Canada. The original idea was broad, to encompass not only the design of networked Information Technology (IT) infrastructure and IT systems themselves but also organizational practices (Cavoukian, 2009). From it emerged privacy by design principles which stayed rather high-level (e.g., to be proactive about privacy and not reactive, and to embed privacy into the design). While the high-level guidance has been criticized (Klitou, 2014), it also enables creative engagement with how to engage and design for privacy and the concept was picked up by many government agencies within the European Union, the United States of America (Federal Trade Commission, 2012), the Organization for Economic Co-operation and Development (OECD, 2013), and finally found its way into Article 25 of the *General Data Protection Regulation* (GDPR) in Europe, making it binding law in the European Union and beyond through its extraterritorial reach (Henseler & Tamò-Larrieux, 2022).

Article 25, entitled “Data Protection by Design and Default” is a good example of how law and computer science interact. The norm requires companies that determine the means and purposes of the data processing to implement technical and organizational measures to ensure that the requirements of the GDPR are fulfilled. In other words, measures that safeguard that the processing occurs in accordance with the fundamental principles of the law are required to be put in place. The fundamental principles within the European data protection law are formulated broadly and include assuring that the processing is fair, transparent, and limited to the amount needed to achieve a certain purpose. Computer scientists and legal scholars alike have worked on ways to translate these requirements for developers (Hoepman, 2021; Tamò-Larrieux, 2018). For instance, Hoepman (2021) provides us with a set of privacy by design strategies that guide developers through the development of a new service or

product. These design strategies can be summarized with one command: “Minimize, hide, separate, aggregate, inform, control, enforce, and demonstrate”. This stands for *minimizing* personal data that is being processed, *hiding* relationships among data, *separating* data in different databases, *aggregating* personal data to the highest level possible, *informing* individuals about the data processing, providing *controls* to individuals whose data is being processed, *enforcing* the terms in privacy policies, and *demonstrating* compliance with legal requirements. In the end, these strategies stay high-level and it is easier to understand their application with a clear use case at hand.

Let us take such a concrete example within the privacy by design domain that illustrates the interaction between law and computer science and the application of the mentioned strategies: A toy robot. Especially for smart devices that are distributed globally, it would be useful to have the ability to adapt their behavior (i.e., their program code) depending on the jurisdiction in which the device operates. In this way, producers could avoid having to create several product variants that are shipped in different parts of the world. For instance, under the GDPR, different thresholds exist within different member states of the EU for when parental consent is needed. While for some member states the threshold is 16 years of age (e.g., Germany), for others it is 13 (e.g., Belgium). Thus, if a product is sold in Germany, parents must consent for their children until they are 16 years old, but not when the product is sold in Belgium, a neighboring country. With geo-location data that is trivial to obtain, the toy robot could know where it is located and which threshold applies and could alter its behavior depending on the specific local thresholds (García, Zihlmann, Mayer, Tamò-Larrieux, & Hooss, 2021). Thresholds are represented in if-then clauses and are thus simple to model: “If the toy robot is located in Germany, the threshold for parental consent is 16”. However, the robot’s program code could capture also more complicated relationships, such as the specific verification requirements that must be fulfilled to ensure that if parental consent is needed, then such consent was adequately provided; and these complex circumstances can, likewise, be executed conditionally on the robot’s prevailing legal context.

To model relationships that the law sets out in a more machine-friendly way, it is, however, useful to develop *structured vocabularies* (such as the Data Protection Vocabulary, DPV) and use *ontologies*, which provide a formal conceptualization of the properties and relations among different elements, e.g., in a given legal norm. In this chapter, we explore those

foundations in which law and computer science interact in practice. Understanding these interactions and the ways to translate legislation into program code is important to understand how automatically processed regulation has evolved (Chapter 3: Automatically Processable Regulation) and the challenges it raises (Chapter 4: Challenges and Controversies).

2.1 IF-THEN: LAWS AS CODE

In the second wave of legal automation (see Chapter 1: Automation of Law), the automation of law early-on focused on rule-based representation of legal rules that was implemented manually. Nationality acts are a good legal domain to do so, with the famous example of the *British Nationality Act* modeled into code (Sergot et al., 1986): The researchers involved in this endeavor took the norms of the *British Nationality Act* and expressed the contained provisions in conditional statements (i.e., rules) using *predicate* logic (also called first-order logic). While *propositional* logic only deals with logical objects (i.e., with propositions that can be true or false) and logical operators (such as conjunctions—“AND”—and negations—“NOT”), *predicate* logic uses *quantified* variables over non-logical objects, such as a person (see Box 2.1).

BOX 2.1 EXAMPLES OF PROPOSITIONAL AND PREDICATE LOGIC

The sun is shining.

- This is a proposition; it is true or false depending on the current context.

What time is it?

- This is not a proposition because it does not express a definite truth value, and it cannot be assigned a definite truth value.

There is a value x that makes this statement true: “ x is green”.

- This is a predicate-logic statement that expresses that a proposition holds for some values (i.e., at least one), but not for others, of a quantified variable.

All lawyers are people.

- This is a predicate-logic statement that expresses that a proposition holds for all values of a quantified variable.

“ $2+2=9$ ”

- This is a proposition; it is false.

Either it is sunny today, or it is not sunny today.

- This is a proposition; it is true. Specifically, it is a tautological statement that is always true.

That is, while in propositional logic, the statement “A child whose mother is Swiss is a Swiss citizen” is, *itself*, the object of study and can only be true or false, predicate logic permits introducing predicates and variables, such as $\text{isSwiss}(x)$ to denote that the predicate *isSwiss* applies to the variable x . Quantification is achieved through two types of quantifiers:

First quantifier: The *universal* quantifier, \forall (to read as: “for all”), is used to express that a clause holds for all values of the quantified variable; for instance, the statement:

$$\forall x(\text{isLawyer}(x) \rightarrow \text{isPerson}(x))$$

expresses that, for all x , the predicate *isLawyer* applying to that x implies that also the predicate *isPerson* applies.

Second quantifier: The *existential* quantifier, \exists (to read as “there exists”), is used to express that a clause holds for some (i.e., at least one) value of the quantified variable, such as in the statement:

$$\exists x(\text{isLawyer}(x) \wedge \text{isSwiss}(x)).$$

This expresses that there exists an “ x ” such that the predicates *isLawyer*(x) and *isSwiss*(x) are both true; “ \wedge ” is read as “and”.

Through these facilities, predicate logic has, in principle, the ability to express statements that are found in legal documents; to exemplify this, we take a provision from the current *Swiss Citizenship Act* already mentioned in the prior chapter (see Box 2.2).

Expressed in first-order logic, this norm could be represented as shown in Box 2.3:

This clause makes the *conditional* statement (through the implication \rightarrow in the third line) that the predicate *isSwissFromBirth* applies globally to entities x (i.e., $\forall x$) given that:

BOX 2.2 ART. 1(A) SWISS CITIZENSHIP ACT, ACQUISITION BY DESCENT

1. The following persons are Swiss citizens from birth:
 - a. a child whose parents are married to each other and whose father or mother is a Swiss citizen;

BOX 2.3 FIRST-ORDER LOGIC REPRESENTATION OF ART 1(A) SCA

$$\forall x (\\ \quad (isChild(x) \wedge hasParents(x, y, z) \wedge married(y, z) \wedge (isSwiss(y) \vee isSwiss(z))) \\ \quad \rightarrow \\ \quad isSwissFromBirth(x) \\)$$

- i. the predicate *isChild* holds for the entity *x* and (signified by the first conjunction, \wedge)
- ii. the predicate *hasParents* holds for that entity together with two other entities *y* and *z*, and (signified by the second conjunction, \wedge)
- iii. the predicate *married* holds for the entities *y* and *z*, and (signified by the third conjunction, \wedge)
- iv. the predicate *isSwiss* applies to the entity *y* or to the entity *z* or to both (signified by the disjunction, \vee).

We can hence use *predicate logic* (see Box 2.4) to express statements about non-logical individuals (i.e., the entities *x*, *y*, and *z*) by using logical operations (\wedge , \vee , \rightarrow , etc.) that are also part of *propositional logic*. A rule that is expressed in this way, as shown in the above statement from the Swiss Citizenship Act, can in this way be explicitly represented in a computer

BOX 2.4 EXAMPLES OF PREDICATE LOGIC

We identify the predicates in the following statements, and express the statement in predicate logic using conjunctions (\wedge), disjunctions (\vee), and implications (\rightarrow).

"There are students who passed the exam"

- Predicate: "being a student" and "having passed the exam" are predicates
- $\exists x (isStudent(x) \wedge passedExam(x))$

"For any number, if the number is even, it can be divided by two without remainder"

- Predicate: "being even" and "can be divided by two without remainder" are predicates
- $\forall x (isEven(x) \rightarrow canBeDividedByTwo(x))$

system (or another logic-based automaton). This implies that what we introduce here is one of the ways how it could be *automatically computed* whether the predicate *isSwiss* applies to an individual if the system is supplied with enough information about other predicates of this individual (i.e., *isChild*, etc.) and about other relevant entities—in this case, the child’s parents. The automation of the British Nationality Act followed this same idea: Here, *Prolog*, a logic programming language (Clocksin & Mellish, 2003), was used to represent the act’s content. Then, the system took factual information, such as the marital status of the parents of a child and the citizenship of the parents of the child, to determine the nationality by, effectively, *automatically applying* the legal norm to the given facts. Such systems—which take a set of rules and facts that are stored in a knowledge base and draw logical conclusions to yield new facts (or rules)—are referred to as *inference engines* or *rule interpreters* (see Contissa & Sartor, 2022, with reference to Turban, Aronson, & Liang, 2005). Inference engines are a central component in AI, and specifically feature in *expert systems*: Computer systems that aim to emulate human decision-making through rules-based reasoning on facts from a given field, such as (part of) the legal domain.

Do you accept that our proposed logic statement about Swiss citizenship is correct? Indeed, we expect that many readers will *readily* accept the predicate logic statement we provided in Box 2.3. This makes it unlikely that we would face fundamental opposition if we proposed an automation system that utilized our proposed predicate logic to automatically determine citizenship. Box 2.5 shows one way how our logic statement could be turned into actual runnable program code, in this case, in the Python programming language. The code introduces an object of type `Person` and defines the function `determine_swiss_citizen_from_birth_art1a(Person)` that computes whether an object of that type is a Swiss citizen from birth according to Article 1a—we say that this function hence is an attempt to *express* or *encode* Article 1a of the Swiss Citizenship Act.

However, already the seemingly trivial conversion from human-readable regulation to predicate logic that we illustrate in this example demonstrates some of the caveats that such formalization brings with it in principle: First, in and of itself, the way we implemented the function *is_child* is tied to a threshold of 18 years of age and does not contextually adapt (see Line 25 of the code snippet in Box 2.5). This might be subject to dispute, and the underlying regulation would need to be qualified further:

BOX 2.5 AN ATTEMPT TO CONVERT ARTICLE 1A OF THE SWISS CITIZENSHIP ACT FROM FIRST-ORDER LOGIC FORM TO EXECUTABLE CODE IN THE PYTHON PROGRAMMING LANGUAGE

```

1     from datetime import date
2
3     # The structure for objects of type Person is created.
4     class Person:
5         def __init__(self, first_name: str, last_name: str,
6           birth_date: date, is_swiss = False, parent1 = None, parent2 =
7           None, spouse = None):
8             """
9             A Person has a first and last name, birth date, and
10            a boolean condition is_swiss.
11            A Person can be linked to up to three other persons
12            through parent and spouse relationships.
13            """
14            self.first_name = first_name
15            self.last_name = last_name
16            self.birth_date = birth_date
17            self.parent1 = parent1
18            self.parent2 = parent2
19            self.spouse = spouse
20            self.is_swiss = is_swiss
21
22        def is_child() -> int:
23            """
24            A function to determine whether the person is a
25            child by checking whether their age is under 18
26            """
27            today = date.today()
28            age = today.year - self.birth_date.year - ((today.
29            month, today.day) < (self.birth_date.month, self.birth_date.day))
30
31            return (age < 18)
32
33        def __str__(self):
34            return f"{self.first_name} {self.last_name}, born on
35            {self.birth_date}, Swiss: {self.is_swiss}"
36
37    # As an example, create three persons where two persons (one
38    # of which is Swiss) are the parents of another person, and are
39    # married
40    jane = Person(first_name="Jane", last_name="Doe", birth_
41    date=date(1982, 2, 2), is_swiss=True)
42    john = Person(first_name="John", last_name="Doe", birth_
43    date=date(1980, 1, 1), is_swiss=False, spouse=jane)

```

```

34     jane.spouse = john
35
36     joane = Person(first_name="Joane", last_name="Doe", birth_
date=date(2008, 3, 3), parent1=jane, parent2=john)
37
38     # Evaluate whether a Person is Swiss from birth according to
Art 1a
39     def determine_swiss_citizen_from_birth_art1a(person:
Person):
40         """
41         An attempt to implement Article 1a of the Swiss
Citizenship Act (see text in this Chapter)
42         """
43
44         # Evaluate whether the person is a child
45         person_is_child = person.is_child
46
47         # Evaluate whether the person's parents are married to
each other
48         person_has_married_parents = (person.parent1.spouse ==
person.parent2)
49
50         # Evaluate whether at least one parent is Swiss
51         at_least_one_parent_is_swiss = (person.parent1.is_swiss or
person.parent2.is_swiss)
52
53         # Return the logical conjunction of these three conditions
according to Article 1a
54         return person_is_child and person_has_married_parents and
at_least_one_parent_is_swiss
55
56         # Print whether the Person "Joane Doe" satisfies Article 1a
of the Swiss Citizenship Act
57     print(determine_swiss_citizen_from_birth_art1a(joane))

```

Until which age is a person a child, and are there factors other than age that determine this? As illustrated in the introduction to this chapter for consent, different jurisdictions might apply different thresholds. Here we thus need to add the item “Child” to a controlled vocabulary that gives a clear definition of what a function that determines childhood should return, and this could be done through natural language, leading to the creation of an ontology (see Subchapter 2.2: Encoding Legal Knowledge for Machines), or it can be accomplished through additional predicate logic statements and based on another piece of regulation that *is_child* refers to.

To reveal further, more abstract and broader issues, we translate back *strictly* from our proposed predicate logic statement to a human-readable sentence, yielding the following version of Article 1a, which *differs from the original* (Box 2.6).

BOX 2.6 VERSION OF THE SCA (1) IN CONTROLLED LANGUAGE

1. An entity is a Swiss citizen from birth if it holds that:
 - a. the entity is a child, and it has parents y and z, and y and z are married to each other, and (y or z) is a Swiss citizen.

By comparing to the original text, we note that the notions of *mother* and *father* have been lost (this is also true for the program code shown above). However, what if the original text was *meant* (or even *required*) to be interpreted more strictly? This would imply that the child in question does not only need to have *two parents* but that it specifically needs to have a *father* and a *mother*. This highlights vividly that the conversion of legal clauses into automatically processable regulation is a sensitive task: It requires at least legal experts and (logic) programmers to sit together to ensure that the legal text is represented in a form that is not only automatically processable (i.e., computable) but that also (still) represents the intended meaning of the law, especially if the interpretation of the law might evolve. Such evolution is the norm rather than the exception: New legal sources might clarify (or redefine) the meaning of a piece of regulation, or societal mores might change.

This connects directly to issues about *freezing the law in code* (see Chapter 4: Challenges and Controversies); it furthermore gives rise to another family of issues that are frequently encountered when aiming to express legal norms in a form that is processable by computers: That of *non-monotonic* logic. When considering logic statements in mathematics or computer science, we typically consider monotonicity, i.e., that the adding of further premises cannot make a valid argument invalid (nor vice versa): In the above example, if *isSwissFromBirth(x)* holds for a given x based on a set of premises, the adding of further clauses to express “unless someone has a criminal background” for instance (*isCriminal(x)*) cannot in a monotonic logic system defeat (or annul) that conclusion. However, many legal norms are not meant as monotonic logic statements—they are rather intended to

provide *tentative default conclusions* that may be retracted based on further evidence. This is particularly visible when considering that appeals may overturn—defeat—legal conclusions that have been drawn by a judge. In the same sense, the capturing of legal knowledge in computer systems should be considered fundamentally *defeasible*.

Zooming out another step illustrates another highly relevant issue: While the original legal text, i.e., Art. 1 of the Swiss Citizenship Act, is supposedly understandable by a large enough part of the population to ensure democratic oversight (see Chapter 6: Educational Shifts Induced by Automatically Processable Regulation), the predicate logic, and certainly the resulting program code in the Python programming language are not so easily approachable. This is not only a problem of *access* to the automatically processable form which could be solved through open-sourcing the legally relevant code, i.e., the publication of the program code in a readily accessible repository; rather, it is a much more fundamental issue that *even with access* to this code, laypeople would not be able to scrutinize it. In the paper on the British Nationality Act mentioned above, the researchers represented the legal knowledge in a computable format and also mapped out the underlying structure of the law into a decision tree that visually enables someone not familiar with the legal text to determine, based on the required facts, if an individual is British or not (Sergot et al., 1986). While such additional human-friendly representations address the underlying problem of widespread inability to scrutinize automatically processable regulation, they cannot solve the problem in principle: When a representation of law is executed that cannot be understood by a large part of the population, society is asked to trust the individual or system that created this representation and to trust the execution environment. Naturally, there is also an element of trust in the application of law even without automatically processable regulation: Residents of a country need to trust the judges' ability to correctly interpret the law and that the judges' environment will allow for this correct interpretation; in other words, judges should not be corrupted and should be free of coercion. When we analyze problems with automatically processable regulation, we make the assumption that they are implemented in a democratic country with a separation of powers; the impact of automatically processable regulation on the law in authoritarian, non-democratic, or corrupt regimes is outside the scope of this book (Ginsburg & Moustafa, 2008).

Turning existing pieces of legislation into code is one way of creating automatically processable regulations. Another one is to flip the approach on its head and to have regulations issued already in a digitally ready format when they are issued by governments, administrations, or legislators (see Chapter 5: Needed (Public) Debates). Such an approach, from the outset, could solve certain issues, the most blatant one being issues related to ambiguity, vagueness, or more generally, what we term open-texture (see Chapter 3: Automatically Processable Regulation), and it can address the issue of public oversight introduced above.

Legislation in a digitally ready format would ideally come as a remedy to what can be considered bad legal drafting. Two legal scholars did not mince their words when they discussed the topic in what has become a seminal paper on the subject, writing that “the disorderly syntax is one of the legal profession’s most visible embarrassments”, and that, more generally, legal drafting “is badly done and needs to be improved” (both from Allen & Engholm, 1978). A few may counter-argue that lawyers and legislators merely try to be accurate and that, therefore, legal norms must result in complicated sentences, or in what is dubbed “legalese”. Yet, this argument does not hold under scrutiny. Researchers have asked lawyers which is more understandable between two semantically equal versions of a contract, that written in legalese or that written in non-convoluted English (Martínez, Mollicab, & Gibsona, 2023). The answer is unsurprisingly that those written in non-convoluted English are easier to understand. The same researchers found that lawyers rated contracts written in more simple English as equally enforceable as their counterparts in legalese, making the use of legalese then non-justifiable. And so, if even trained lawyers who should be at ease with legalese are, in fact, not, this raises the question of why legalese should be used in the first place.

The hope would, therefore, be that with digitally ready legislation, it could be possible to improve the quality of writing, understandability, and automatic processability at the same time. A key question is then however: How could legislators *negotiate* bills that are in a digitally ready format? It should be apparent from the above that reading software code would require specific skills that elected officials would not necessarily have. And here comes something in between, something that is not natural language and not code, but which is still easily readable without prior training and which can then be converted directly into code: Controlled language (Fuchs, Höfler, Kaljurand, Rinaldi, & Schneider, 2005).

There exist several different controlled languages (see for an overview Kuhn, 2012), but they share a combination of characteristics: They restrict the use of lexicon, the use of grammar, and the structure of text. By constraining the use of *lexicon*, it is possible to remove *vague* terms (see Box 2.7), for instance by forbidding the use of many adverbs and adjectives such as “periodically”, “reasonable”, “easy”, etc. By constraining the use of *grammar*, it is possible to limit syntactic ambiguity. “Arthur goes to see Carl; his t-shirt is blue”—does “his” refer to Arthur’s or to Carl’s? In controlled language, the sentence would require the author to reuse either the name Arthur or Carl and forbid the use of “his”. By constraining the *structure* of the text, it delineates where to find a definition. A period, as in *periodically*, should be, for instance, specified at most in the preceding sentence but not earlier. With the introduction of all these restrictions, language becomes less rich (say if you want to write a poem or literature), but it also allows direct translation into computer code.

BOX 2.7 EXAMPLES OF AMBIGUITY (FROM SENNET, 2021).

- Lexical: "duck" is both a noun and a verb
- Pragmatic: The classical "can I go to the bathroom?", mixing capability and authorization
- Under-specification: "sanction" can mean to approve or to penalize
- Reference transfer: "he is parked" refers to his car rather than to the person

The rules on how to write in controlled language might seem abstract. However, a step-by-step translation from natural language to controlled language can help illustrate the process. To do so, we take the sentence in natural language:

“All meetings with unvaccinated people are prohibited unless they are excused”.

We now transform this sentence into *Logical English* (Kowalski, Dávila, Sartor, & Calejo, 2022; Kowalski, 1982). Logical English uses predicate and propositional logic statements to translate legal texts into statements close to the if-then types (but allowing for more flexibility). In our case, the sentence would read as follows (adapted from the reference above):

“A meeting is prohibited if a person attends the meeting and the person is unvaccinated and it is not the case that the person is excused”.

Logical English has possibilities when it comes to integrating different forms of logic but from the short excerpt, this also comes at the cost of being possibly less intuitive to read. A controlled language that is more intuitive to read is *Attempto Controlled English* (ACE; Fuchs & Schwitter, 1996) which caters to first-order logic and sentences will therefore be mostly of the form if-then and contain many logic operations (simple sentences such as “The dog is brown”. are also accepted). In ACE, the sentence translates to:

“If a person is not vaccinated and the person is not excused then a meeting with the person is prohibited”.

ACE allows the introduction of new words in the lexicon (simply by using quotation marks) and forces the use of “the” to refer to a previous object (as in “the person” after the “and” in the example above). Other rules include: (1) The sentence must finish with a period; (2) There cannot be personal pronouns he/she/her/their/etc. (and therefore the pronoun “they” in “they are excused” needs to be defined); (3) While sentences may have an attached relative sentence in principle, there are many restrictions. For example, “A brother who I have lives”. passes the test, but “A brother with whom I share a mother lives”. does not (remember: “My” is not allowed if you want to say “My brother”).

While Logical English would then be translated into Prolog directly, ACE can be mapped onto *Discourse Representation Structures* (DRS), a language representing *Discourse Representation Theory* (DRT), which is a “formal account for representing the meaning of natural language discourse”, see Figure 2.1 for an example of DRS (Abzianidze, Bos, & Oepen, 2020, p. 23). And so, DRS models DRT, but more importantly, DRS can be interpreted by machines, and therefore, when this stage is reached, then the next step into automatically processable regulation is immediate (Table 2.1).

We have tested how humans would view and work with ACE and whether it could be a good candidate for the creation of digitally ready legislation (Guitton et al., 2025). When translating legal texts into ACE, we first noticed that this translation increased the texts’ score on “ease of readability” by an average of 14 points (Flesch-Kincaid score). The Flesch-Kincaid score is computed from the syntax and length of words in

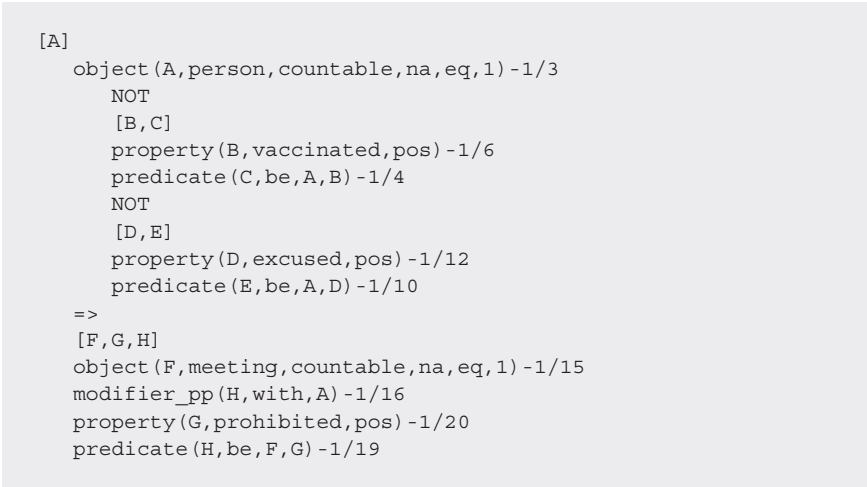


FIGURE 2.1 DRS version of ACE sentences.

TABLE 2.1 Overview of Natural, Logical, Controlled, and Programming Languages

Natural English	All meetings with unvaccinated people are prohibited unless they are excused
Logical English	A meeting is prohibited if a person attends the meeting and the person is unvaccinated and it is not the case that the person is excused
Attempto Controlled English	If a person is not vaccinated and the person is not excused then a meeting with the person is prohibited
Prolog	is_prohibited(A) :- attends(B, A), is_unvaccinated(B), not is_excused(B)

a sentence; and because ACE forces to break down sentences and to use simpler (and hence, on average, shorter) words, this led to such a change in score. Going further, we found out that the use of ACE also improved the comprehension of laws. And so, while ACE could be used in applications for automatically processable regulation geared towards accessibility to increase understanding of specific aspects of laws, it can *at the same time* increase the understanding of everyone reading the law, from laypeople to legislators discussing it when still at the level of being a bill. This finding interestingly contrasted with one of our expectations, however: We also found out that participants perceived *more* open-texture (i.e., ambiguous, vague) terms in statutes in the form of ACE than in their original form. This

is a bit paradoxical but may be attributable to readers reading ACE for the first time and therefore not being used to the forced wordy detours caused by its restrictions. But the main conclusion remains that ACE could be a good candidate to implement digitally ready legislation as it preserves (or even improves) the human-led process of negotiating law while enabling a direct mapping of legislation that is created in this manner to logic programs. Currently, discussions around digitally ready legislation are still early and only emerging; it may be that no country implements it, or if they do, that they choose a different path than via controlled language. But if countries do consider implementation, ACE—or another controlled language—should be part of the toolkit under consideration.

2.2 ENCODING LEGAL KNOWLEDGE FOR MACHINES

In the above, we have introduced approaches to how legal language might be (manually) turned into a form that is processable by machines using logic primitives. We have shown how this could be accomplished in a way that remains approachable to the general population—specifically, through controlled language. However, the alert reader might have already noticed that this discussion sidelined a very relevant issue: If we consider the example from above that

*All *meetings* with *unvaccinated* *people* are *prohibited*
unless they are *excused**

then we notice that the items in asterisks (a meeting, unvaccinated, etc.) are used directly as named literals in the text and in the resulting logic program. To the logic program, they hence do not carry any further meaning, and replacing **unvaccinated** with **happy** (or any other concept) does *not change any aspect of the execution of the program*. This is akin to how the *renaming of a variable* in a computer program code does not change the execution of the computer program. To illustrate this, consider the following two statements in the C programming language:

```
[S1] int taxAmount = 0; taxAmount = taxAmount + 200;
```

and

```
[S2] int numPrimates = 0; numPrimates = numPrimates + 200;
```

S1 and S2 both lead to the *very same low-level instructions* being executed on a computer—specifically, in this case, both these programs yield the

very same instructions for a processor. For instance, when this code is compiled for an Apple M2 processor that features the ARMv8 instruction set architecture, the very same instructions are executed by the processor hardware, namely:

```
mov    r4, #0                ARMv8 Machine Code: e3a04000
add    r4, r4, #200          ARMv8 Machine Code: e28440c8
```

Seemingly part of a convoluted detour to electrical engineering and processor design, we put so much emphasis on this aspect since it is important for readers to understand that the human-readable program code of a program contains semantic information that is *not* preserved when the program is executed. Variable names in the above C program (*taxAmount*, *numPrimates*) and also in the logic programs that we introduced in the previous section *carry meaning for programmers, but not for processors*. In the above example, S1 might be part of a tax calculator while S2 might be part of a Zoo simulation game—however, for a processor, the two statements are *exactly* the same.

Computer Science has, however, come up with an approach to giving *meaning* to variables that are handled by computer programs, and specifically a way that can ensure that such meaning remains compatible across programs that have been created by different entities (e.g., different individuals or different organizations) and that may operate in different domains. With such *semantic technologies*, we encode meaning separately from data and application code and create explicit links of data and code with that meaning, which itself is structured according to curated vocabularies or ontologies. Applying semantic technologies to the example above, the variable *numPrimates* might, for instance, be connected to the well-defined concept of *Primate* in the *BioTop* ontology,¹ which is a “top-domain ontology that provides definitions for the foundational entities of biomedicine as a basic vocabulary to unambiguously describe facts in this domain”. In *BioTop*, the concept of *Primate* is unambiguously identified using the *Internationalized Resource Identifier* (IRI)² <http://purl.org/biotop/biotop.owl#Primate> and can thereby be linked to other concepts, for instance, to express that the concept “Human” (specified in <http://purl.org/biotop/biotop.owl#Human>) is a subclass of the concept of “Primate” (which puts the Zoo simulation game in perspective)—we will see how such association is performed in greater detail later.

The linking of variables that are handled by a computer program to agreed-upon, published, and shared semantic concepts permits programs to hold a shared understanding of the underlying information, and it also allows them to reason on top of this information at run time. The overarching goal—implicitly or explicitly (Calbimonte et al., 2023)—is to permit the programming of computer systems on the *knowledge* level, which complements or replaces today’s common hard-wiring of meaning and relationships into program code at the time the program is designed (like in S1 and S2). As an illustrative example of this concept, consider instructing a child about how they may get to their Judo class after school (this example is extended from Calbimonte et al., 2023). It is perfectly feasible that such instruction happens on the level of the *specific environmental features* that the child will encounter, and hence to tell the child to enter a *specific* means of public transport at a *specific* time and location to reach a *specific* destination: “You need to take bus #2 at 5:37pm and leave the bus when it reaches the station Singenberg”. For a child who is instructed in this way, there is no immediate need to *understand* many of the underlying concepts. The child may, for instance, remain oblivious to the fact that other bus lines exist, or that other buses run on the same line but at different times. This makes this way of instruction particularly straightforward: It is simple to convey and requires relatively little knowledge on the side of the child to be followed; this is akin to the lines of code in S1 and S2. However, parents who emphasize their child’s autonomy will instead opt for instructing their child more at the *knowledge* level. This includes how a public transport system functions in general, which involves knowledge about ticketing, routes, schedules, plans, and possibly even how to use digital tools for navigation. While this way of instruction is not as direct as the alternative shown above, it provides the child with a higher level of autonomy with respect to navigating its environment—and this knowledge stays relevant even when the environment is dynamic, for instance when the child arrives late at the bus station, when bus #2 is canceled, or when the destination station is moved.

Programming on the knowledge level hence increases the autonomy of the programmed systems. This is specifically interesting in the field of automatically processable regulation: While we may program an unmanned aerial vehicle (UAV) to always follow a given path (e.g., a yellow indicator) in its environment—thereby hard-coding it to an environmental

feature—another option is to equip the UAV with knowledge about its navigation domain, such as about where it may go (e.g., by limiting the altitude or defining no-fly zones), and when (e.g., night flying restrictions). The specific navigation decision is then taken by the UAV at run time. In practice, this means that the UAV will execute code that takes into account the specified regulation, e.g., whether it is “night”. This however means that this whole approach depends on the *alignment* of the UAV’s understanding of the underlying concepts with the respective concepts of the regulation’s originator: If one defines “night” as starting when the sun reaches 18 degrees below the horizon (this is how the “astronomical night” is defined) but the other defines it as starting at 8 pm, conflict is programmed, literally. Alternatively, both might refer to a shared concept and ensure that they can (independently and consistently) evaluate this concept at run time.

2.2.1 Types of Machine Knowledge

Before we consider in greater depth how computer systems may be equipped with knowledge in the illustrated sense, we investigate what different types of knowledge might be relevant for machines in the first place. Some knowledge is *ontological* in nature—it describes how the world works by describing concepts and categories in a given subject area as well as relationships between them. *Ontological knowledge* may be expressed as propositions (i.e., knowledge about facts, such as in our example on citizenship or permitted meetings, above) and can be combined with non-propositional, *procedural*, knowledge (or know-how) that describes *how* a given task may be achieved by an agent. For instance, to appropriately use a coffee machine, a human requires ontological knowledge about coffee, hot water, and drinks, and they require procedural knowledge—for instance, acquired from a user manual that describes how the machine is operated. Specifically interesting in the scope of this book is knowledge that is *normative* in nature—this defines notions such as obligation, permission, prohibition, or dispensation—where the field “follows the legal (Hohfeld, 1919) and deontic logic (Von Wright, 1963) traditions in understanding a (social) norm to include laws and other prescriptions or proscriptions on social behavior” (Singh & Singh, 2023). Such knowledge includes—in an organizational context—a definition of the *mission* of a *group*, the different *roles* that members of the organization might adopt, and the *obligations*, *permissions*, *prohibitions*, and *dispensations* that correspond to these roles

(Chopra, Torre, & Verhagen, 2018). They might also define the *sanctions* that are imposed if an agent violates a rule.

After motivating the equipping of machines with shared knowledge and exploring different types of knowledge, we now turn to the question of *how* we may contextualize (or “semantically lift”) concepts in a computer program. Today, one of the most prevalent ways to represent such knowledge is in the form of *semantic networks*, i.e., directed or undirected graphs where the vertices represent concepts, and the edges represent relations between the concepts they connect. This way of representing knowledge dates back to the third century AD, and specifically to a representation that the philosopher and logician Porphyry used to illustrate a “scale of being” and that became known as the *Tree of Porphyry*. This provides a classification of Substance, where, for instance, *Animal* has the subcategories *Human* and *Beast*—the categories are represented as an (implicitly) directed acyclic graph, i.e., a tree (see Figure 2.2).

Today, across many research areas and applications, such knowledge is expressed in *knowledge graphs* (Hogan et al., 2021), and a large variety of tools and techniques are available to support the creation of knowledge graphs as well as for querying them and for reasoning on top of the contained information, at large scale. Today, the most widespread approach to representing knowledge graphs is in the form of concepts and relations between concepts, instantiated in the *Resource Description Framework* (RDF). RDF is a standard by the World Wide Web Consortium (W3C) that can be used for representing graph data. RDF can hence, specifically, be used to represent knowledge graphs triples of the form *Subject-Predicate-Object* (S-P-O), where the *Subject* and the *Object* are nodes and the *Predicate* is a relation between these nodes. While, in RDF, only *Objects* may be literal values (such as numbers or strings), *Subjects*, *Predicates*, and *Objects* can be identified through IRIs (Internationalized Resource Identifiers) — this is the very mechanism that permits the usage of concepts that carry shared meaning, such as BioTop:Primate (<http://purl.org/biotop/biotop.owl#Primate>) from our example above. Since individual RDF triples statements represent individual relationships between concepts, a collection of RDF triples induces a *labeled, directed multigraph*: It is a graph that is *labeled* (since its nodes and edges carry labels) and *directed* (since an S-P-O triple induces a *directional* relationship—*Zurich (S) is located in (P) Switzerland (O)* holds, but *Switzerland (S) is located in (P) Zurich (O)* does not; and it is a *multigraph* because there may be multiple vertices between

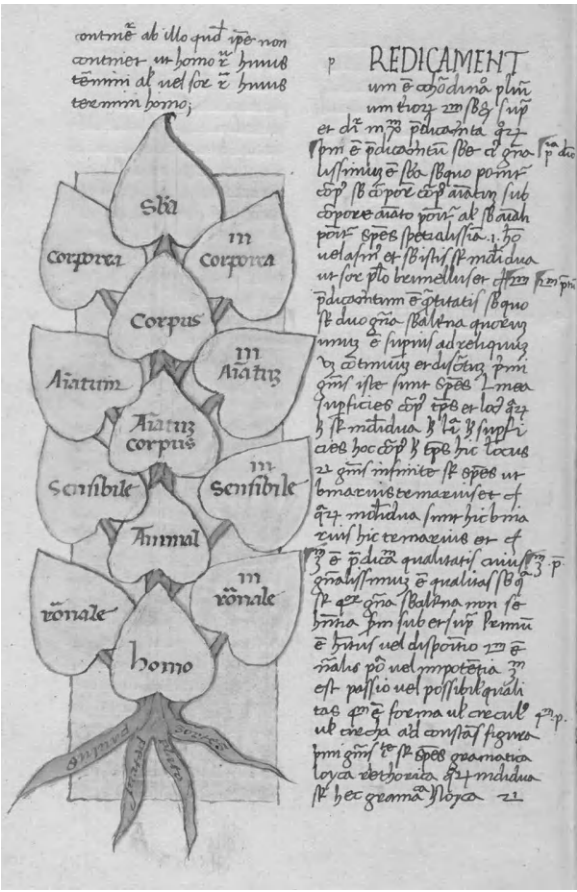


FIGURE 2.2 Tree of Porphyry (Perugia, 1475; reproduced from the Lawrence J. Schoenberg Collection; Creative Commons Public Domain). (https://open.library.upenn.edu/Data/0001/ljs457/data/master/0067_0017.tif)

the same two nodes (e.g., both, Zurich (S) is located in (P) Switzerland (O), and Zurich (S) is largest city of (P) Switzerland (O) hold).

To grasp more details and the value of such knowledge graphs for representing shared ontological, procedural, and normative knowledge among machines, consider a sentence and example that we took from Wikipedia: “Paul Schuster was born in St. Gallen” (see Figure 2.3). While to a human (who can read English), this sentence carries meaning, to a computer program it is merely a sequence of letters. However, we may annotate the individual tokens, where we use IRIs as introduced above to refer to

```

<div vocab="http://dbpedia.org/ontology/" typeof="Person">
  <span property="name">Paul Schuster</span> was born in
  <span property="birthPlace" typeof="Place" href="https://www.wikidata.org/wiki/Q25607">
    <span property="name">St.Gallen</span>.
  </span>
</div>

```

FIGURE 2.3 Example compatible with semantic Web.

concepts that are shared knowledge—in the following, we use the W3C HTML+RDFa standard to apply such annotation:

This is the very sentence “Paul Schuster was born in St. Gallen”, but with additional metadata (in angle brackets) that specify annotations for tokens in the statement in compliance with the RDFa standard. Specifically:

1. We use the property “*vocab*” to set the specific structured vocabulary that we apply across the example—think of this as setting a context for the rest of the example. The vocabulary is *http://dbpedia.org/ontology*, and this vocabulary contains the other concepts we use in our example and within the scope of the `<div>` tag (i.e., between `<div ...>` and `</div>`). This includes the property *name* (*https://dbpedia.org/ontology/name*), the property *birthPlace* (*https://dbpedia.org/ontology/birthPlace*), and the classes *Place* (*https://dbpedia.org/ontology/Place*) and *Person* (*https://dbpedia.org/ontology/Person*).
2. We use the property “*typeOf*” to define an instance of type *https://dbpedia.org/ontology/Person* that is further specified in the statement. Concretely, we apply the property *https://dbpedia.org/ontology/name* to the string “Paul Schuster” to specify what the name of the person is.
3. Similarly, we apply the property *https://dbpedia.org/ontology/birthPlace* to link our instance of type *Person* (i.e., the *Person* named “Paul Schuster”) to an instance of type *https://dbpedia.org/ontology/Place*. For this instance, we, on the one hand, specify its name—“St. Gallen”—again through the *https://dbpedia.org/ontology/name* property. In addition, we use a hypermedia reference (*href*) to link the place to the IRI *https://www.wikidata.org/wiki/Q25607*. If this seems cryptic, keep reading.

Through these annotations, and since we respected the HTML+RDFa standard, a suitable computer program that reads (i.e., parses) this statement will derive exactly five RDF triples:

1. There is an instance of type *Person*. Formally, as an RDF triple:
`_:a <http://www.w3.org/1999/02/22-rdf-syntax-ns#type>`
`<http://dbpedia.org/ontology/Person>.`
2. This same instance (of type *Person*) has the *name* “Paul Schuster”. Formally, as an RDF triple:
`_:a <http://dbpedia.org/ontology/name>`
`“Paul Schuster”.`
3. This same instance has the *birthplace* <http://www.wikidata.org/entity/Q25607>. Formally, as an RDF triple:
`_:a <http://dbpedia.org/ontology/birthPlace>`
`<http://www.wikidata.org/entity/Q25607>.`
4. <http://www.wikidata.org/entity/Q25607> is of type *Place*. Formally, as an RDF triple:
`<http://www.wikidata.org/entity/Q25607>`
`<http://dbpedia.org/ontology/type>`
`<http://dbpedia.org/ontology/Place>.`
5. <http://www.wikidata.org/entity/Q25607> has the *name* “St. Gallen”. Formally, as an RDF triple:
`<http://www.wikidata.org/entity/Q25607>`
`<http://dbpedia.org/ontology/name>“St. Gallen”.`

The given annotations hence *semantically contextualize* the sequence of letters “Paul Schuster was born in St. Gallen”. The below Figure 2.4 illustrates how these statements can be arranged in a knowledge graph, where the node labeled `_:a` corresponds to the entity “Paul Schuster”.

Now, the fascination and power of the introduced way of using IRIs to refer to concepts is that a computer system can *dereference* many of these IRIs—including all of the IRIs we use above—and thereby discover *additional* semantic relationships.

We encourage the reader to try this: When using a Web browser to follow (i.e., dereference) the IRI <http://www.wikidata.org/entity/Q25607>, the client is redirected to the publicly reachable webpage *St. Gallen—Wikidata*

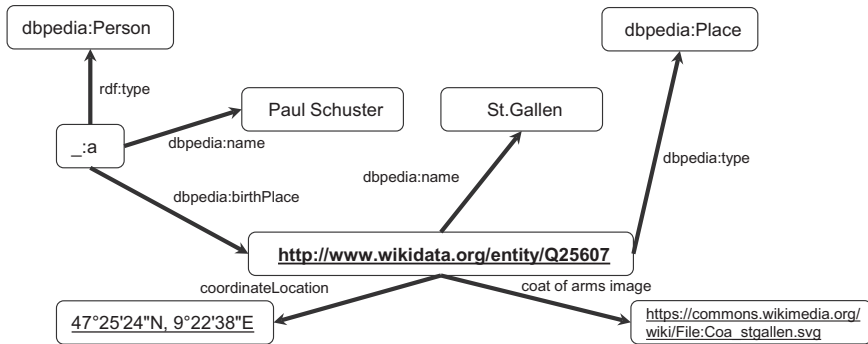


FIGURE 2.4 Example of a knowledge graph based on the sequence “Paul Schuster was born in St. Gallen”.

(<https://www.wikidata.org/wiki/Q25607>) that on first sight resembles a structured Wikipedia entry. However, upon closer inspection, a (human or machine) visitor will discover additional hyperlinks on this page that are *also semantically grounded*. One of these is the IRI <http://www.wikidata.org/entity/P17>—this is a property that is named “country” and, upon dereferencing it (and being forwarded to the human-readable webpage <https://www.wikidata.org/wiki/Property:P17>), we discover that it is a property that is defined as “sovereign state that this item is in (not to be used for human beings)”. Using this property, our concept of the city of St. Gallen, i.e., <http://www.wikidata.org/entity/Q25607> is linked to the concept of Switzerland, <http://www.wikidata.org/entity/Q39>—and machines can follow these links just like humans can. That is, the resources *together* state the triple (in RDF) that (<http://www.wikidata.org/entity/Q25607>, <http://www.wikidata.org/entity/P17>, <http://www.wikidata.org/entity/Q39>) which, in a form that is simpler to understand for humans, states that “St. Gallen” holds the “country” relation towards “Switzerland” (colloquially, “St. Gallen is in Switzerland”). A computer system that consumes our triples above and can browse the given IRIs in this way hence is automatically able to (unambiguously) deduce that “Paul Schuster was born in Switzerland”. This demonstrates how knowledge graphs and RDF can be used to make statements about circumstances where these statements are then automatically and interoperably (we sometimes say “seamlessly”) linked to concepts that have been defined by others. Such definition can happen in standards and can be represented formally—using RDF—as

an ontology. In computer science, this is an “explicit, formal, and general specification of a conceptualization of the properties of and relations between objects in a given domain” (Wyner, 2008); ontologies hence provide explicit specifications of domain concepts and of their relations so that these concepts can be used within computer programs in a uniform way.

Before we apply this approach to the legal domain, we introduce several of the largest reusable-controlled vocabularies and ontologies that are today published and used all around the World Wide Web. To cover different fields of application, we mention (socially oriented) *Friend of a Friend*, (media-oriented) *Dublin Core*, (search-oriented) *schema.org*, and (device-oriented) *Semantic Sensor Network*. In addition to these (and many other) domain vocabularies, there are valuable cross-domain ontologies such as the *Quantity, Unit, Dimension, and Type* (QUDT) ontology. Finally, so-called *upper ontologies* and *bridging ontologies* exist and are used to link equivalent concepts across domain-specific spaces.

- Friend of a Friend (FOAF)³: FOAF is an ontology that can be used to express information about social agents (e.g., their name and address), their activities and interests, and their relations to other people and objects. It is supported by several popular online content management systems such as WordPress.
- Dublin Core (DC)⁴ is a set of metadata terms for digital and physical resources (e.g., videos, webpages, books, works of art) that is standardized internationally through ISO 15836 (the Dublin Core Metadata Element Set, or DCMES).
- Schema.org⁵ is an initiative that was created by Bing, Google, and Yahoo! and whose goal is standardized structured data markup for webpages. The schema.org vocabulary today consists of 806 concepts that range from “Places” to “Reviews” to “Legislation” and “Organization”, and is used by over 45 million sites (according to schema.org in the year 2024).
- The Semantic Sensor Network Ontology⁶ is used to describe “sensors and their observations, the involved procedures, the studied features of interest, the samples used to do so, and the observed properties, as well as actuators”. It can, for instance, be used to specify requirements on sensors and actuators in an automation system, enabling

automatic reasoning about whether one sensor can be replaced by another (e.g., from a different manufacturer).

- QUDT⁷ (Quantities, Units, Dimensions, Datatypes) is comprised of several linked ontologies that together represent a large variety of quantity and unit standards (e.g., meters, volts), and can be used for conversions and dimensional analysis of equations, and, in general, for enabling interoperability of input and output data of technical systems.

In the area of semantic technologies—knowledge graphs, RDF, ontologies, and controlled vocabularies—our readers might also encounter other standards, specifically the OWL Web Ontology Language and the LKIF Legal Knowledge Interchange Format—these are both frameworks that are designed to facilitate the management and sharing of structured knowledge. While OWL permits the specification of more detailed relationships between objects and their attributes across various domains in RDF, LKIF is designed to model legal reasoning and to enable the interchange of legal knowledge that is modeled differently across heterogeneous systems, by providing a basic ontology of legal concepts. To give an example of such interchange, we briefly discuss the European Law Identifier (ELI) which represents an initiative towards harmonizing legislation across Europe (Filtz, Kirrane, & Polleres, 2021). Together with the European Case Law Identifier (ECLI), ELI proposes a technical specification for the identification of legal documents (using IRIs within a controlled vocabulary) and suggestions for vocabularies to be used to describe them in a machine-readable way. To accomplish this, the standards, for instance, reuse the Dublin Core vocabulary as mandatory properties to provide metadata about legal documents.

After providing this contextualization regarding how knowledge might be represented for machines, we next apply these techniques to express a legal circumstance in a way that is processable by a computer; if this machine is programmed in terms of the shared ontologies and vocabularies that are employed, we say that the machine “understands” the represented situation. Consider the potential of a system where this endeavor succeeded: The ability to express a regulation (say, in tax law, or in privacy law) in a way that can be understood by machines in this sense would, for instance, permit automatic compliance monitoring. A company’s systems could then access the machine-readable representation of the regulation

that is provided by the government on a secure server and could fuse this information (e.g., about obligations or permissions) with its own program code. In the privacy field, this would, for instance, permit the mapping of legal obligations to the program code of a device that records personal data, which would permit enforcing the processing of this data in a compliant way (García et al., 2021). While this program code needs to reuse the same shared vocabularies and ontologies that are used in the representation of the regulation, importantly, it does not have to be written in a way that is specific to the regulation. We are hence *programming in terms of regulatory knowledge* rather than in terms of *specific regulatory artifacts*: Similar to the child who holds ontological and procedural knowledge about the public transport system rather than being “hard-coded” to taking a specific bus at a specific time, the program code of this legal system is not hard-coded to the specific regulation, but rather is programmed in terms of general information about regulation (e.g., what is an “obligation”).

In other words, we have reached a situation where the system that consumes information about regulations is *decoupled* from the system that shares this information. Speaking from a software-engineering standpoint, such decoupling brings the monumental advantage that it permits *independent evolution of consumers and producers*—in our specific example, this means that the law (and its automatically processable version) may change at run time (i.e., *after* the consuming system has been deployed) without breaking any of the systems that consume it! The reader is familiar with this very concept (and its value) already from their daily experience of browsing the World Wide Web: While webpages are frequently updated by their publishers (some even change every second), users are not required to frequently update their Web browsers (e.g., Firefox). Beyond the possibility of achieving this decoupling, and thereby creating systems that stay up-to-date with current regulations even after they are deployed, we see further efficiency and accessibility advantages of regulation that is made machine-understandable in this way:⁸

- Pieces of machine-understandable representation of regulation could be *automatically checked for inconsistencies or loopholes* (both by the legislative as well as by private entities).
- Machine-understandable regulations could be *automatically and verifiably transformed into further representations*, for instance in

different languages or in a form of language that caters to larger parts of the population (e.g., Simple English).

- Systems could be *automatically certified* with respect to their compliance, and certificates could be automatically issued (or withdrawn).
- When used in the context of specific benefits and taxes, the system could *check the eligibility of a person* to the ever-evolving and complex requirements that an update to the legislation brought about (see Mes Aides for instance, detailed in Chapter 7: Exercises).
- Products (e.g., self-driving vehicles or other cyber-physical systems) could *automatically be made to conform* to the local regulations with little human intermediation (Bhuiyan et al., 2020).
- Based on the possibility to automatically enforce regulation in program code, individuals could make use of systems that *extend regulation*; e.g., an individual could issue a directive that, within their private home, audio recording is banned between 8pm and 8am (see Tamò-Larrieux, Mayer, & Zihlmann, 2021).

While several of these potentials (especially those aiming at improving access to legal knowledge) might be readily accepted by society, they also bring with them considerable issues. We discuss these issues in detail in Chapter 4: Challenges and Controversies.

2.2.2 Exemplifying the Translation of a Legal Norm into Its Machine-Processable Version

We next give a concrete example of how a piece of regulation might be expressed in machine-understandable form while reusing public shared vocabularies and ontologies. Concretely, based on an example we described in Guitton et al. (2023), we re-use here the introduced semantic tools to express a specific article from GDPR, namely Art. 7(1). This article expresses requirements on data controllers, i.e., legal persons that “decide the how and why of a data processing operation” (European Data Protection Board, 2024), regarding *demonstration of consent*. In the following, we describe aspects of this article in a automatically processable way and also show concrete run-time instances that may be validated with respect to GDPR Art.7(1), which reads as follows:

Where processing is based on consent, the controller shall be able to demonstrate that the data subject has consented to processing of his or her personal data.

To create a machine-processable version of this article, we bring together several domain-specific and cross-domain shared vocabularies. The first vocabulary that we make use of is the Data Privacy Vocabulary (DPV; version 2.0 from August 2024), whose objective is to capture the usage and processing of personal data considering different legislative requirements.⁹ DPV is being developed by a W3C Community Group of interdisciplinary scholars and interested industry stakeholders and can be used to specify common rules (namely representing permissions, prohibitions, and obligations) that are associated with the handling of personal data in the context of GDPR. An example of this is the DPV concept “Legal Basis” (specifically, <https://w3id.org/dpv#LegalBasis>), which formalizes the concept of legal bases of processing in GDPR, and is described in DPV as “Legal basis used to justify processing of data or use of technology in accordance with a law”. In DPV, “Legal Basis” is further specified, using subclasses, according to the legal bases that are present in GDPR—Legal Obligation (<https://w3id.org/dpv#LegalObligation>), Legitimate Interest (<https://w3id.org/dpv#LegitimateInterest>), Official Authority of Controller (<https://w3id.org/dpv#OfficialAuthorityOfController>), Public Interest (<https://w3id.org/dpv#PublicInterest>), Vital Interest (<https://w3id.org/dpv#VitalInterest>), and Consent (<https://w3id.org/dpv#Consent>). Note that DPV only defines these (and other) terms, such as Permission (<https://w3id.org/dpv#Permission>) as “A rule describing a permission to performing an activity”; however, DPV does not indicate formal semantics (e.g., using formal logic) of what it *means* that a Permission exists. For this purpose, the authors of DPV refer to other proposals that permit expressing these formal semantics, such as the Open Digital Rights Language (ODRL) and RuleML (Boley, Paschke, & Shafiq, 2010) whose goal is to represent rules in a machine-understandable and executable form.

In our proposed formalization of GDPR Art. 7(1), we combine DPV with FOAF. We further base our formalization on the assumption that any processing of personal data that is not explicitly permitted within the scope of GDPR Art. 7(1) is prohibited. With this assumption, we propose to formalize GDPR Art. 7(1) as an obligation that the data controller must be able to demonstrate that it has valid permission to process certain personal data

owned by a data subject. According to GDPR Art. 7(1), such an obligation becomes active when (1) the action of processing those personal data in a specific process is executed and when (2) the process is based on consent. Our proposed formalization is given here in Box 2.8 and explained below, referring to the individual lines of this formalization as appropriate.

BOX 2.8 FORMALIZATION OF GDPR ART. 7(1)

```

1 @prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
2 @prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
3 @prefix foaf: <http://xmlns.com/foaf/0.1/> .
4 @prefix dpv: <https://w3id.org/dpv#> .
5 @prefix dpv-loc: <https://w3id.org/dpv/loc#> .
6 @prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
7 @prefix ex: <http://example.org/> .
8
9 ex:Alice rdf:type foaf:Person, dpv:DataSubject ;
10   foaf:firstName "Alice" .
11
12 <https://alice.pod/voiceData-2023> rdf:type dpv:PersonalData ;
13   dpv:hasDataSubject ex:Alice .
14
15 ex:ACME rdf:type foaf:Organization, dpv:DataController ;
16   foaf:name "A Company that Makes Everything" ;
17   dpv:hasProcess ex:AnalyzeSpeech .18
19 ex:AnalyzeSpeech rdf:type dpv:PersonalDataProcess ;
20   dpv:hasProcessing dpv:Analyse ;
21   dpv:hasPersonalData <https://alice.pod/voiceData-2023> ;
22   dpv:hasLegalBasis dpv:Consent .
23
24 ex:SpeechAnalysisConsentRecord rdf:type dpv:ConsentRecord ;
25   dpv:hasDataSubject ex:Alice ;
26   dpv:hasPersonalDataProcess ex:AnalyzeSpeech ;
27   dpv:hasConsentStatus dpv:ConsentGiven ;
28   dpv:hasJurisdiction dpv-loc:EU ;
29   dpv:hasIdentifier "63ded36f-4acd-4f3c-991e-6cb636698521" ;
30   dpv:isIndicatedAtTime "2024-10-31T4:44:44"^^xsd:dateTime ;
31   dpv:hasNotice "https://acme.org/data-processing-policy" .

```

This representation first specifies the required linked vocabularies (such as DPV and FOAF) through @prefix directives in Lines 1–7; in this way, we may in the rest of the file use the given prefixes (such as dpv, dpv-loc, foaf) to refer to concepts. As a concrete instance of the processing of personal data, we consider the case where a data controller processes the voice data of a data subject. In our example, the respective voice recordings of the

data subject, `ex:Alice`, are referred to using the IRI `https://alice.pod/voiceData-2023` where we assume this data to be accessible by authorized software programs; specifically, the data can be accessed by the systems of a company called “ACME”. This software then analyses the voice data, for instance, to manage the data subject’s personal calendar (“Create a new meeting tomorrow at 4pm!”) or to actuate devices in their smart home (“Switch on the living room lights!”). In our automatically processable representation, we define `ex:Alice` (as a data subject, Lines 9–10) and link her to her voice data (Lines 12–13). We further define `ex:ACME` as a data controller (Lines 15–17). Note how we make use of FOAF to specify in machine-understandable form that `ex:Alice` is a person (Line 9) and ACME an organization (Line 15), and of the DPV vocabulary to declare that `ex:Alice` is a data subject (Line 9), that `ex:ACME` is a data controller (Line 15), and that the voice data is personal data (Line 12) of Alice (Line 13). We furthermore specify that Alice’s personal data is being handled by ACME (Line 17). Then, Lines 19–22 explain (to a machine) that there is a `dpv:PersonalDataProcess` of type `dpv:Analyse` that is carried out by `ex:ACME` (all according to DPV, Lines 19–20), that this concerns the voice data (Line 21), and that the handling of the data has consent as legal basis (Line 22).

These conditions are sufficient to fulfill the requirements set in GDPR Art. 7(1) in the specific case of Alice’s voice data processing because they describe in the respective vocabularies that the processing (`ex:AnalyzeSpeech`) of personal data about a Data Subject (`ex:Alice`) by a Data Controller (`ex:ACME`) is based on consent (`ex:AnalyzeSpeech dpv:hasLegalBasis dpv:Consent`). This is, hence, sufficient to satisfy the activation clause of GDPR Art. 7(1) for this instance: “Where processing [of personal data] is based on consent [...]”.

To complete the example, we require a representation of the data subject’s permission to process the personal data. This permission to process is commonly represented as a record of the given consent, and we interpret GDPR Art. 7(1) as an obligation to be able to demonstrate this record. To explain this in an automatically processable form, we next express this record using DPV and link it to the data controller, data subject, and data processing instances from above. This is shown in Lines 24–31 of the above representation: We specify a specific `dpv:ConsentRecord` (according to DPV, Line 24) that is linked through DPV to `ex:Alice` (Line 25) and to the processing of Alice’s personal data (and thereby to `ex:ACME`, Line 26). We furthermore use DPV to state that this record expresses that `ex:Alice`

has consented (Line 27) and that the scope of the consent is the EU (Line 28). Finally, we again use DPV in Lines 29–31 to specify practically relevant information: The identity of Alice’s specific consent record, the time when this consent was given, and a link to ACME’s data processing policy.

An automatically processable representation of GDPR Art. 7(1), as introduced above, allows software components to process this regulation in a uniform and compatible manner. For instance, this would permit the automation of compliance checking with GDPR Art. 7(1) in a straightforward way: To achieve this, software would routinely scan all of an entity’s processes of type `dpv:PersonalDataProcess` and ensure that, for each such process that has consent as legal basis, there is a valid and currently active consent record. If it does not find a valid record, it could furthermore automatically contact the data subject to obtain consent or alternatively signal that the non-consensual data processing needs to be stopped (or stop it outright). The same system could also become active when consent is withdrawn—for instance, it could be configured to delete the underlying personal data.

2.3 ENTERS MACHINE LEARNING

The approaches to legal automation that we have discussed up to this point—the formulation of regulation in a way that encodes a logic program and that ties the regulation’s vocabulary to an agreed-upon ontology—rely on symbols, i.e., rules, relationships, events, that are interpretable and often hand-crafted. The focus of these approaches rests on generating interpretable legal knowledge to “understand” symbols in legal documents. All examples that we discussed in the previous subsections fall under this category. Due to their focus on (human-interpretable) symbols, these methods are commonly referred to as “symbol-based methods” or “symbolic AI methods” (Sheth, Roy, & Gaur, 2023; Zhong et al., 2020). Such symbolic methods were also the first approaches to natural language processing, i.e., the field of computer science that aims at enabling computers to process, manipulate, and generate human language.

A distinctly different approach to what we introduced in the above sections has been gaining a lot of momentum over the decade that precedes the writing of this book: Machine learning—the study and development of algorithms that learn *statistical* relationships in their input data and then generalize these relationships to classify or predict given new data—has also been applied to regulatory processes, such as extracting legal norms (e.g., information retrieval systems), linking relevant cases, analyzing legal texts (e.g., finding open-textured terms, see Chapter 3: Automatically

Processable Regulation), and to create question-answering systems for entire bodies of regulation (Ashley, 2017). To do so, different approaches can be taken, and with the rise of large language models, a dominant research focus has been on large, data-driven, embedding-based methods, i.e., methods that enable converting raw data (e.g., text) into numerical data, which a machine learning model can process more efficiently. Distinct from symbolic AI methods, these approaches are referred to as “sub-symbolic”, since the underlying algorithms are not grounded with specific predefined symbols. They rest on extracting legally relevant features to predict rules, relationships, and events from large datasets; today, this approach is dominant in the natural language processing field.

To make this difference more tangible, we again turn to our example with the Swiss Nationality Act from Section 2.1. There, we specified part of the nationality act as a first-order logic statement (Box 2.3). As discussed, with a *symbolic* approach one could take this statement and semantically ground each of the predicates, e.g., by linking them to an ontology or other logic statements. For instance, the definition of $\text{hasParents}(x, y, z)$ would require an agreed-upon understanding of what a “parent” is, what “having parents” means, and about how many such parents are needed. In some cases, it could be simple to semantically ground the predicate, for instance, the property $\text{isChild}(x)$ could be tied (again through a logic rule) to the date of birth of the entity “ x ” (notwithstanding the caveats mentioned above). After these definitions are agreed upon, programs could be written that reuse the logic statement together with the definitions to automatically determine citizenship—concretely, to evaluate the specified implication. These programs would be compatible with one another due to the standardized grounding of the statement, and their decisions could be scrutinized by tracing them to this underlying grounding.

A program that evaluates whether a person has Swiss citizenship from birth could also be created *without* explicitly encoding this logic. Specifically, with a simple sub-symbolic approach, a statistical discriminator (e.g., a linear regressor) could be used that takes, for a few thousand individuals, their birth certificate data (e.g., date of birth and parents’ names) along with data from their parents’ passports as input. If this input is coupled with information about whether each of these few thousand individuals should be classified as a Swiss citizen from birth, this data can be used to *train* the discriminator, i.e., to adjust its parameters so that it, when given only the birth certificate and passport data, reliably

discriminates individuals into the categories “Swiss from birth” and “not Swiss from birth”.

To illustrate how machine learning methods work in practice, we give examples of each of the main abstract training methodologies for this type of AI: *Supervised learning*, *unsupervised learning*, and *reinforcement learning* (see Figure 2.5 as well). Without going into depth, we would like to provide here some fundamentals as machine learning algorithms can be trained in these different ways. *Supervised learning* means that an algorithm is trained on a labeled dataset, which contains the pairing of input data with desired/correct output data (like in the example with Swiss citizenship assignment above). From this labeled dataset, the *training* data, the algorithm learns to classify new inputs. While learning, the algorithm adjusts its internal parameters to minimize the difference between what it predicts and the actual correct outcomes. Often, supervised learning is used in classification tasks like detecting spam or labeling images and in regression tasks to create a closed formula that matches the data points as closely as possible. An example of using supervised training could be with a robot-judge learning from case law that is provided in an automatically processable format. The robot-judge would be adjusting (learning) how it should rule on cases on the basis of past cases, and apply this to new cases. An obvious issue with supervised learning concerns *overfitting*: The algorithm matches very well with the dataset it was trained on, but it performs much less accurately on new data.

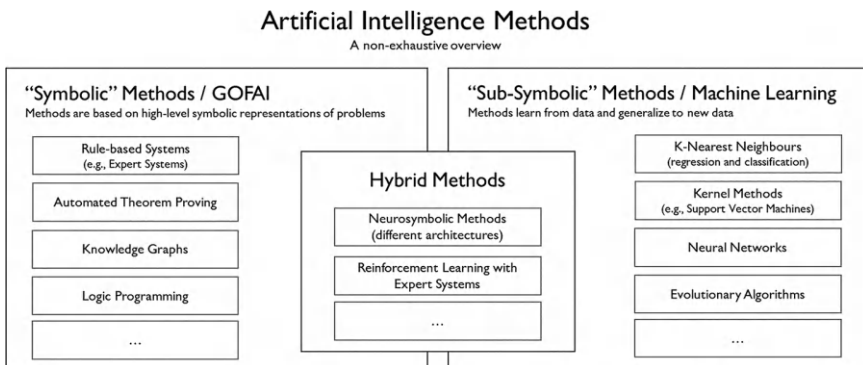


FIGURE 2.5 Breaking down AI methods. GOFAI in the figure stands for “Good old-fashioned AI”.

Unsupervised learning involves training on *unlabeled* data, meaning that the algorithm must find patterns without predefined labels. Within that process, the algorithm identifies inherent patterns and groups similar data points together. Such algorithms are used when clustering datasets such as shopping data of customers. Often, unsupervised learning is used for more generative tasks compared to discriminatory tasks. To reuse the example of a robot-judge just provided, unsupervised learning would be the equivalent of feeding the robot-judge with case laws *not* in an automatically processable regulation format but in their raw format. The robot-judge would have to figure out which part of the case law corresponds to the statement of the facts, to the applicable laws, to the legal conclusion, and so on. The application boundary of supervised and unsupervised learning is blurred today, which is examined in the field of weak supervision (or semi-supervised learning), where supervised and unsupervised methods are combined, e.g., by first performing unsupervised clustering and then using (less) labeled training data to label the clusters. This attempts to overcome one of the main drawbacks of supervised learning—the requirement of labeling the training data.

The last basic approach to machine learning that we introduce is *reinforcement learning*. Here, an agent (e.g., a human) interacts with the algorithm by giving feedback about actions conducted by the algorithm in the form of rewards or penalties; the algorithm then adapts in response to this feedback by updating its policies for action selection. In doing so, the agent guides the model toward better decision-making, which means maximizing cumulative rewards over time. Reinforcement learning is often used in tasks where an agent needs to make sequential decisions, such as robotic control (Lee, Hwangbo, Wellhausen, Koltun, & Hutter, 2020). This has tremendous benefits, e.g., to correct on the fly for false categorization and accordingly guide the algorithm back onto the “correct” track, but also downsides, as actions that the algorithm tries out might lead to irreversible results (e.g., if it breaks a physical object); because of this, and since the algorithm needs to wait to receive its rewards or penalties, reinforcement learning is regarded as a particularly resource-intensive approach. To stick with the robot-judge example, if reinforcement learning was applied in this context this would amount to having the system make decisions and then rewarding or penalizing it on these decisions. This would lead the reward-maximizing system to learn, over time, a conviction policy that maximizes rewards.

At the time of writing, the field of AI—with its roots in Alan Turing’s “machine intelligence” in the 1950s to the introduction of the Transformer architecture that led to the 2020s’ boom in AI technologies (and specifically in generative pre-trained transformers; GPT)—can, from a high level that distinguishes between symbolic and sub-symbolic methods, be structured as shown in Figure 2.5. This illustrates some of the relationships between current AI approaches and also includes hybrid systems that integrate symbolic and sub-symbolic components in different architectures (see Sheth et al., 2023). Sub-symbolic and symbolic methods each have advantages and disadvantages that are especially relevant within the field of automatically interpreting law—for instance, as described above, that the working of a symbol-based method is simpler to scrutinize than with complex sub-symbolic methods.

When applying symbolic, sub-symbolic, or hybrid methods to the legal field, several of their trade-offs become very visible, and we conclude this chapter with a discussion of some of these. Even if two systems—one based on machine learning (e.g., a neural network) and the other based on a symbolic approach (e.g., a rule-based system)—produce the same (functional) outcomes for a problem at the same resource requirements (e.g., time or compute power), the methods still exhibit different non-functional properties.

One relevant non-functional property concerns the *interpretability* of the system and its outcomes—this is relevant for explaining *why* the system reached a specific conclusion and, hence, for scrutinizing the system’s functioning, for instance with respect to biases. Since their symbols are readily understandable by human users, typically, symbolic methods are preferable if a system needs to be highly interpretable, but this assessment needs a bit more scrutiny itself. For instance, a rule-based system exhibits the very rules it uses to reach its decisions, meaning that these rules can be traced to find out exactly why an outcome was reached. This ability suffers with increasing scale, as systems that are made of millions of rules are also not *practically* interpretable anymore. A similar situation exists in machine learning: A simple linear regressor can be readily interpreted by people who understand its symbolism (i.e., what the parameters mean); however, this is not true anymore for more complex methods, such as Support Vector Machines or deep learning systems. The reason for this lack of interpretability is not that these models’ parameters cannot be

observed—this is possible—however, there are too many of these parameters (e.g., billions in current Transformer architectures), and the parameters are not tied to human-understandable symbols. Humans are, in this sense, disintermediated to a greater extent in machine learning methods since they are unable to assign meaning to the learned parameters that these methods use during inference; due to its high societal relevance, this aspect is highly relevant when considering the automation of legal processes.

Another relevant non-functional property is the *data provenance, data intensity, and training automation aspect*: One of the two main reasons for the current boom in AI methods is the wide availability of readily accessible training data today. Many of the underlying primitives for today’s AI systems (e.g., perceptrons or the concept of backpropagation) have been known since the 1960s—but their use remained limited to niche applications until the availability of swathes of training data (e.g., pictures on the Web) in the early 2010s. Symbolic systems, on the other hand, were traditionally “trained” by hand—with individuals entering rules or assigning meaning to concepts through the definition of ontologies (the concept of ontology *learning* was introduced in the late 2000s). This, hence, describes an elegant tradeoff when deciding on a method: If large amounts of reliable (and possibly already labeled) training data are available, a machine learning approach might be preferable since the training can be automated to a higher degree in this case.

The decision of which method to select when attempting to automate a (legal) system hence needs to consider functional as well as non-functional factors: What data is available, and is it available in automatically processable form? What quality level (accuracy/fidelity) is desired? What performance (delay/resource intensity) is desired? What level of interpretability is required? What level of human disintermediation is acceptable? And how many resources (time/money) are available? The determination of the “right” technology to apply in a specific case is far beyond the scope of this book and requires technical as well as domain expertise; however, in the subsequent chapters we will demonstrate how several of the introduced approaches are used in a variety of projects, and we also distill the specific challenges and issues that the application of AI methods to the legal field brings.

NOTES

- 1 <http://purl.org/biotop>
- 2 Internationalized Resource Identifiers expand the set of permitted characters in Uniform Resource Identifiers (URIs) which are well-known from our everyday usage of Web browsers. They are an Internet standard: *RFC 3987 - Internationalized Resource Identifiers (IRIs)* (ietf.org)
- 3 <http://xmlns.com/foaf/spec/>
- 4 <https://www.dublincore.org/>
- 5 <https://schema.org/>
- 6 <https://www.w3.org/TR/vocab-ssn/>
- 7 <https://qudt.org/>
- 8 This example and some text in the following paragraphs are taken, in some cases verbatim, from the article: Guitton, Mayer, Tamò-Larrieux, Garcia, & Fornara (2024)
- 9 <https://w3c.github.io/dpv/dpv/>

Automatically Processable Regulation

AFTER HAVING SURVEYED TECHNOLOGIES and approaches that in principle enable to automate regulation and briefly discussed their trade-offs, we now turn to the concrete application of these in a variety of scenarios, with the aim of deriving a typology of automatically processable regulation and, in Chapter 4: Challenges and Controversies, addressing the scandals that have emerged with automatically processable regulation projects. To get a sense of different projects in the field and the deliberation on the dimensions of the typology that follows below, we introduce two automatically processable regulation projects: A robot judge project in Estonia, and a social benefits project in France.

Estonia, the most northern of the three Baltic states, is often considered a trailblazer when it comes to digital government. Over the past 15 years, many operations and services that involve the state have been made available online. One exception so far has been the automation of court rulings, but in 2019, journalists of the *Wired* reported on a new robot judge project in their article “Can AI Be a Fair Judge in Court? Estonia Thinks So” (Niiler, 2019). The article stated that the robot judge would be offered to Estonian residents seeking to settle smaller disputes of less than € 6,400. The name “robot judge” might have been hyping up the project, as in reality, this appeared to have been a software program that delivered an authoritative state-approved response for cases of claims that were relatively uncontested. Within the initial pilot phase, the software delivered a

response for 65,000 “disputes” (for a country with a population of 1.3 million). After the initial pilot, Estonia discontinued the project—although it is unclear why. One of the assumptions was that two Ministries, namely Finance and Justice, came head-to-head in their different mandates. While the Ministry of Finance was interested in speeding up resolutions of cases and, hence, supportive of the project, the Ministry of Justice had concerns surrounding legal concepts. It remains that the project, despite rumors that it would soon re-start, still hasn’t re-emerged as of 2024.

Another automatically processable regulation project comes from France, where an administrative unit wanted to tackle the complex system of eligibility for social benefits (Alauzen, 2021; Merigoux, Alauzen, & Slimani, 2024). There is a wide range of such benefits in France, with many rules on income and wealth that make eligibility determination hard for laypeople. With the incoming of a left-wing president in 2012, François Hollande, the agenda of access to social benefits as prescribed by the law gained momentum and gave impetus to start the project called “Mes Aides” (“my social benefits”). Very quickly, problems started. The state office in charge of administering the actual decisions for social benefits refused to share the software code that they were using, as they feared that creating Mes Aides would support social benefit fraud. Their refusal even ran against the law mandating that the software code for decision-making be made public. The development team for Mes Aides had hence to start again from scratch without considering the actual implementation for deciding on granting or not social benefits; they had, therefore, to bring together the logic behind the law covering 30 plus social benefits. The launch of Mes Aides, despite these early challenges, proved very popular, with 2.3 million connections to the site in 2019—so popular that even social workers advising families were using the site. But in 2020, the project nonetheless was shut down for a while, leading to anger amongst citizens, and leading an NGO to duplicate the website. In 2021, the state took the project over again and expanded its scope with a new website on Mes Droits Sociaux.¹

3.1 TERMINOLOGIES AND TYPOLOGIES

With these two examples in mind, let us now consider how to classify them. When we first introduced the term “automatically processable regulation”, the reasoning behind it was to provide a typology to the expanding field of the automation of law that had produced *multiple terminologies* used by researchers and policymakers, such as computational or computable law

(Deakin & Markou, 2020b), code-driven law (Hildebrandt, 2018), rule as code (Mohun & Roberts, 2020), legal informatics (Ashley, 2017), or legal AI (Cobbe, 2020) to name just a few. These terms often overlap with one another, yet at times with slightly different scopes: For instance, the OECD report titled “Cracking the code: Rulemaking for humans and machines” defines rule as code as “an official version of rules (e.g. laws and regulations) in a machine-consumable form, which allows rules to be understood and actioned by computer systems in a consistent way” (Mohun & Roberts, 2020, p. 3). Hildebrandt (2020, p. 67) defines code-driven law as “legal norms or policies that have been articulated in computer code, either by a contracting party, law enforcement authorities, public administration or by a legislator. Such code can be self-executing or not, and it can be informed by machine learning systems or not”. Or, as another example, Katz, Dolin, and Bommarito (2021, p. 3) do not provide a strict definition of legal informatics but state that it is the “academic discipline that underlies transformational technologies” in “document review in litigation, to compliance, case prediction, billing, negotiation and settlement, contracting, patent management, due diligence, legal research”.

We argue that having a clear understanding of the terminologies employed in the field is central to enable comparisons among different applications of the automation of legal processes. The term automatically processable regulation, therefore, comes with *three dimensions* to form a comprehensive typology (Guitton et al., 2022b). The first dimension centers around the aims that are being pursued through the automation, the second dimension analyzes the potential of divergence of interests in a given project, and finally, the third dimension looks at the degree of mediation by computers. Before analyzing two concrete instances of automatically processable regulation and their classification, it is important to take the time to familiarize ourselves with these three dimensions as they will be recurring explicitly or implicitly throughout the book.

The first dimension is probably the most intuitive one and classifies the answer to the question: What is the *primary type of benefits sought after* by the automation process? While often we think of automation as a process to increase efficiency, automatically processable regulation can be geared towards a more efficient legal process or a more accessible one. While aiming for efficiency and accessibility may go hand in hand, we argue that a binary decision can, in most cases, be taken following the decision tree we provide in Figure 3.1 and analyzing who benefits from the automatically

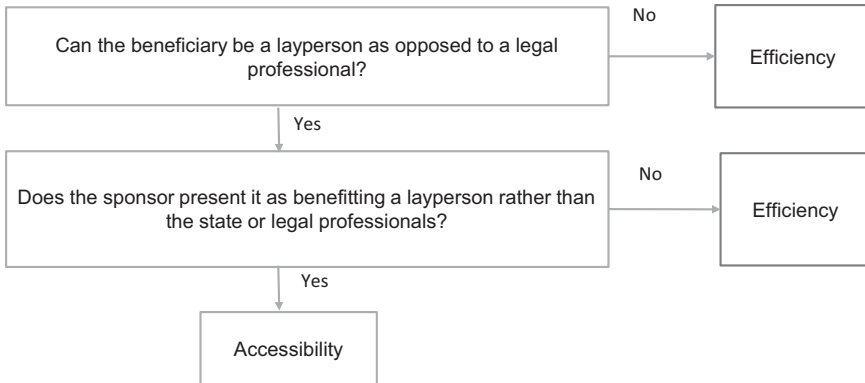


FIGURE 3.1 Decision tree to determine the aim of an automatically processable regulation implementation.

processable regulation implementation and how this benefit is communicated to individuals.

The second dimension distinguishes between different stakeholders involved in the implementation of automatically processable regulation (e.g., when a government agency determines to automate access to rebates and involves an external party to design and deploy the system), and the *different interests each of the stakeholders* have within the implementation (e.g., government agency wants to streamline tedious work, civil servants within the agency want to remain in charge of the decision-making, and the developers want to be remunerated for their software). Typically, one can distinguish between sponsors (i.e., an institution commissioning or sponsoring an automatically processable regulation implementation), implementers (i.e., institutions or entities in charge of developing and implementing the automatically processable regulation), beneficiaries (i.e., the target group who should benefit from the automatically processable regulation project), and users (i.e., the actual person in the end using the automatically processable regulation). While disentangling the different stakeholders in a given project is often straightforward, the alignment of goals among them is often more concealed and, thus, less accessible. For instance, an automatically processable regulation project might only involve two different groups of stakeholders, but these two groups might have highly diverging interests with respect to the project's outcome. In light of this, we proposed a conservative measurement of the total potential divergence of interests by taking the maximum class between two values:

TABLE 3.1 Classification of the Divergence of Interests

Whichever Returns the Higher Class	f1: Number of Distinct Actors	f2: Degree of Observed Divergence of Aims across and within Non-overlapping Actors
Class 1	1	Non-existent or small
Class 2	2	Medium-small
Class 3	3	Medium-high
Class 4	4	High

The first value being the number of stakeholders involved and the second value being the perceived alignment of goals (as illustrated in Table 3.1).

The last dimension looks at the degree of automation in terms of the *involvement of humans* in guiding the specification and execution of a computer program. A strong focus rests on the underlying data and code base used to transform regulation into its automatically processable form. For instance, with respect to the data, we can distinguish between curated and uncurated data, with the curated one being explicitly mediated by a human, and an uncurated one only implicitly being mediated by a human. On the one side, we understand by a curated data and code base that the logic of such data and programs analyzing it can be changed by editing the content of the data structures as well as the statements within the code. For example, this is the closest to the manual encoding of expert systems (notably with the Swiss Citizenship Act) seen in Chapter 2: Law and Computer Science Interactions. More data-driven automatically processable regulation applications, on the other side, using machine learning approaches are typically less curated due to the sheer size of the underlying data sources. We refer the reader to the previous chapter for examples of this, too, notably on what concerned sub-symbolic applications.

The degree of human mediation is, in turn, determined by the domain of law and its complexity. For instance, while an automatically processable regulation project applying a certain threshold to determine a tax benefit can easily be programmed by a human, a complex analysis and pattern recognition of case law will likely be implemented using a machine learning approach, hence disintermediating humans. Even if information on those factors—data, code, and domain—will not always be accessible, we can estimate the likely degree of mediation by humans by assigning weights to each factor. We score each of these from 0 to 4 and take the Euclidean distance from the point of origin: This is the distance of a vector in a space from the point (0,0,0), where here the space is three-dimensional for a project P.

In our case, the Euclidean distance is expressed as:

$$\text{mediation degree} = \sqrt{\text{data}^2 + \text{code}^2 + \text{domain}^2}$$

As we round it to the nearest 0.5, this gives a maximum mediation degree of 7 and a minimum of 0.

The three dimensions—aims, stakeholders and their divergence of interests, and degree of human disintermediation—provide a good basis to classify and discuss applications of automatically processable regulation (see Chapter 4: Challenges and Controversies). Of course, other approaches to classify the myriad of automatically processable regulations emerging in the literature or in practice can be and have been proposed. For example, a typology that emerged from a research team working on an EU-funded project on “Counting as a Human Being in the Era of Computational Law” divides automatically processable regulations into four dimensions, including the type of system analyzed (app, dataset, paper), the functionality (listing different aspects of automatically processable regulations such as ones focused on the drafting of legislations, automated compliance, and prediction of litigations), the user targeted by an automatically processable regulation project (e.g., an individual, a lawyer, judge), and whether the automatically processable regulation project is code or data-driven. Compared to our typology, there is a certain alignment: The aspect of targeted user is also part of our typology, but it is used to determine the aim of the project as efficiency or accessibility; the aspect of code or data-driven is captured by the degree of mediation by computers. Both typologies will require to make judgments when assessing a project, but the typologies will target different parts of projects. Our focus on divergence of interest is not reflected in the EU-funded typology while the functionality aspect from the EU-funded typology is not in ours. Divergence of interest came as a dimension out of interviews which we carried out and where we noticed that some projects were more benign than others, less prone to politicization too, and because we were, independently, about to embark on another study seeking links between projects and the issues that they could trigger. On the other hand, categorizing by functionality can be helpful for other (research) purposes.

Yet other classifications exist. For instance, Whalen (2022) takes a broader view by including technology outside of the automatically processable regulation sphere per se, making a division between deep/shallow and between legal/generic technologies. This has the advantage of

being able to map technologies that are not necessarily only developed for the legal spheres but still find their usage there. Another one comes from McMaster (2019), who, when looking at the very specific question of copyright issues when commercializing legal AI, breaks down between suppliers, incubators, service providers, legal counsels, clients, and investors. And so, specific purposes will drive the shaping of creating typology to categorize projects.

3.1.1 Applying the Typology to Our Initial Examples

Let us retake the two examples outlined at the very beginning of this chapter which we have used to elaborate on the typology in Guitton et al. (2022b), the robot judge and Mes Aides. These two projects point to clear differences between automatically processable regulation projects. For starters, we see very different *aims*. While the robot judge project is about enforcing the law and reducing backlogs in courts, the social benefit project wants to give French residents better access and understanding of what they are legally entitled to. Another difference illustrated is related to the complexity of the projects, not because of human-to-human interactions but because of the underlying technical requirements. From the onset, it appears that the ambitions of both projects are widely different. The robot judge aims to tackle only non-contentious claims, while the social benefit project has to make its way through the intricacies, exceptions, and counter-rules of many different laws. But they would also repose on different technical bases: One of them requires data concerning past cases to inform how to decide on current cases, whereas the other is not based on case law but on the mere interpretation of statutes. As explained above, we can combine these different factors into one, which we call the degree of *disintermediation of humans*. In other words, how much are humans removed from the overall process and being replaced by an automated process? This question can be answered by analyzing three sub-factors: Domain, code, and data. In the case of Mes Aides, many different social benefit laws had to be integrated, and there is a certain complexity to each of them, as well as considering the impact that they have when considered as a whole and not only separately—hence this justifying a 3 over 4 for the complexity of the legal domain. We assume that this would also translate to a certain complexity in the code—hence score it with a 2 of 4, and as no past data is required to be used, a zero for underlying data feeding into the software is appropriate. On the other hand, for the robot judge, the extent of the domain of law is also extensive (3/4), but we assume a slightly higher code

complexity to integrate a machine learning approach from the 65,000 on which it ruled, and so with code and data both scoring a 3 over 4.

The two automatically processable regulation projects have also similarities, notably the controversies among different stakeholders. Many entities are involved in both projects—from legal advisers, implementers, testers, those paying for the project, those benefiting from it, government entities, and social workers—and there are difficulties in aligning different stakeholders’ interests. This is where the dimension of the *potential for divergence of interests* between all these actors comes in. Smaller projects, or those that do not involve the state or where stakes are less high (e.g., because it is “only” about disputing a parking fine) may trigger less divergence. The potential for divergence can be informative though of the potential for mishaps too: The more different stakeholders try to pull the project in different directions, the more large mishaps could occur (see Chapter 4: Challenges and Controversies). This is why to measure this potential for divergence, we use two sub-metrics: One focuses on the number of distinctive actors (e.g., between beneficiaries, sponsors, implementers, users), and the other looks more via human judgment at the degree of observed divergence across and within these non-overlapping actors. Each sub-metric goes from one to four, and the resulting aggregate number is the maximum between the two. In our mentioned automatically processable regulation projects, the robot judge and Mes Aides, the involvement of many stakeholders and publicly stated divergence of interests translates into a high score of 4 for both. Overall, Table 3.2 provides an overview of the resulting classification of these two projects.

3.1.2 Challenges When Classifying Applications of Automatically Processable Regulation

For all of these three metrics—divergence of interest, human disintermediation, and project aims—which we have introduced to characterize projects for comparability, there are notable difficulties. Often, divergences of interests among different units are not made public, as management fears a bad reputation. This makes it difficult to measure the exact degree of divergence of interests. For the disintermediation of humans, while the law is public (and hence the *domain* sub-factor is publicly known), the implemented code rarely is, and it is typically also not published how the curation of the underlying data takes place. This means that many assumptions have to be made. Lastly, the distinction between efficiency/accessibility can also be difficult to make. For instance, and similarly to

TABLE 3.2 Classification of Mes Aides and Robot Judge According to the Typology

Name	Mes Aides	Robot-judge
Sponsors	Public service*	Ministry of Economy (MoE)
Implementers	Public service*	Unknown
Beneficiaries	Citizens	Ministry of Justice (MoJ)
Users	Citizens	Citizens
Aim	Accessibility	Efficiency
f_1 : # Distinct actors	2	3
f_2 : Degree of observed divergence	High: Those in charge of innovation and those administrating the benefits	High: MoE concerned with cuts; MoJ with upholding the rule of law
Potential for Divergence of Interests: $\max(f_1, f_2)$	4	4
Domain factor	3	3
Code factor	2	3
Data factor	0	3
Degree of mediation by computers (rounded to the nearest 0.5)	3.5	5

* denotes the same unit across roles. In other words, the sponsors and implementers were in this case the same.

divergence and disintermediation, how sponsors view the primary goal of a project might not be public, although inferring it is often a possibility. More problematic is that the frontier between efficiency and accessibility can be blurry. Arguably, Estonian residents who go through the robot judge also learn (again by inference) about case law from how the output of their own case looks like. And also arguably, those using Mes Aides could contribute to a sort of pre-screening of applications, whereby those who apply genuinely have grounds to believe that they should be granted social benefits. This could consequently result in the official social benefits office having to consider fewer “superfluous” applications with an outcome of no allowance at all. Furthermore, those utilizing the online tool do not have to consult a legal professional, once more contributing to the efficiency of the legal system. But we consider that these are only secondary effects, and that is why the categorization reposes on the *primary* aim of the project. More generally, it would even appear that any type of accessibility project will result in efficiency gains as a by-product, and the two types of aims are, therefore, by far not mutually exclusive.

When presenting such a typology to help compare projects, fair criticisms of the project aims can be brought up. For instance, the extent to which one has to be a layperson for the accessibility to be valid can be questioned. Those with legal training learn where to find and how to navigate between the different legal sources and materials (statutes, case law, commentaries, legal scholarship), as well as how to interpret them. Often such knowledge is specialized in one legal domain (tax law, family law, data protection law). A data privacy lawyer will struggle to give an opinion on a company's debt issuance program. And thus, even legal professionals could benefit from tools for access to the law. Therefore, defining the goal of accessibility to the law only to laypeople is somewhat restrictive and maybe not nuanced enough. Yet, even if legal professionals outside their area of expertise might benefit from automatically processable regulation tools, they still have, importantly, knowledge and practice on where to find legal knowledge, how to put it together for interpretation, and can even rely on a network of legal professionals without upfront payment to exchange or validate their findings. A layperson will not, or definitely not to the same extent, have such a knowledge base to build upon. This is why defining access to the law to laypeople, even if restrictive and narrow, is what we can use to effectively distinguish projects (Guitton et al., 2022b).

3.2 EFFICIENCY YOU SAID?

As mentioned above, often the push for automation in any domain is linked to efficiency gains. This is also true for the legal domain, where tools to automate legal practitioners' work are promoted with the high-efficiency gains for the industry and where smart contracts are sold on the premise of making business transactions more seamless. Also, within the domain of automated driving and autonomous vehicles, cars that are able to respect traffic regulations are marketed as a way of relieving passengers from having to care about traffic (and traffic law) and hence may benefit more from the time spent traveling from say Lausanne to Geneva, e.g. for relaxation or work.

The following quote by Max Tegmark illustrates this quest for efficiency and the premise that automatically processable regulation will lead to more efficient as well as just legal systems:

Since the legal process can be abstractly viewed as computation, inputting information about evidence and laws and outputting a

decision, some scholars dream of fully automating it with robo-judges: AI systems that tirelessly apply the same high legal standards to every judgment without succumbing to human errors such as bias, fatigue or lack of the latest knowledge.

Tegmark (2017, p. 105)

The quote shows that efficiency-driven implementations of automatically processable regulation are often accompanied by other ideals and promises that go hand in hand with efficiency, such as the promise of reducing costly legal uncertainty, (human) errors, and biases in the application of the law. The first set of promises circles around *legal uncertainty* and resulting human errors in the application of the law. These are costly because they lead to mistakes in how the law is applied, which in turn creates individual and societal costs to overcome the errors. Such errors often occur because the law needs to be interpreted and because, in the legislative process, compromises are struck on wordings that lead to syntactic and semantic ambiguities with respect to the concrete application of the law. In particular, syntactic ambiguities arising from the clarity of the structure of a sentence could be avoided and would need to be overcome when turning legal text into legal code. Such syntactic ambiguities can be specific wording or even punctuation marks. To illustrate this, an interesting case arose in New England, USA, where truck drivers brought forward a case for the compensation of overtime by their employer, Oakhurst Dairy. The case revolved around a local law in the State of Maine which stated that employers were *not* to pay overtime in the case of:

the canning, processing, preserving, freezing, drying, marketing, storing, packing for shipment or distribution of: 1) agricultural produce; 2) meat and fish products; and 3) perishable foods.

The question thus was whether the part of the norm refers to packing alone or includes also the distribution as a separate step. The debate in front of the courts was whether the part “packing for shipment or distribution” meant “packing for shipment or packing for distribution” or just “packing for shipment” or “distribution”. For truck drivers, this had a serious consequence, as they were not involved in the packing (for which there is an exemption to overtime pay), but were involved in the distribution. The court sided with the truck drivers as it argued that the lack of a comma in the critical passage of the norm indicated that the passage referred only to

the packing. If the passage had read “packing for shipment, or distribution” the law would have been applied differently (Victor, 2018). In the end, the company settled with the drivers.

Such *syntactic ambiguities* as the one that was resolved in a dispute between Oakhurst and its drivers would likely be discovered when turning the legal text through an interdisciplinary team into an automatically processable regulation (or by using controlled language, see Chapter 2: Law and Computer Science Interactions). At least that is one of the conclusions we learn from New Zealand’s initiative to transform their Rates Rebate Act into automatically processable regulation. In their final report they write that “[t]he process of developing rules statements identified gaps in the logic of legislation that had not been previously identified, which range from technical gaps that could be resolved by legal drafters to gaps that require further analysis” (Stevenson, 2019, p. 6). Aside from getting rid of legal uncertainty in these cases, it would also minimize errors in the application of the law by humans and thereby reduce the costs of implementing such laws by the state (e.g., when issuing rebates) or companies (e.g., when calculating allowances based on labor law). Both government bodies and companies would implement the law via digital interfaces that enable citizens or employees to access legally guaranteed benefits they are allowed under the law.

Another set of promises, linked to the ones discussed about *reducing human errors* (see the aforementioned quote from Max Tegmark), is the promise that automating legal processes would reduce biases in legally relevant decision-making. This reasoning is often brought forward when discussing judicial processes with the mentioning of a study conducted in Israel that showed that rulings of judges might be impacted by hunger (Danziger, Levav, & Avnaim-Pesso, 2011a,b). The study looked at 1,112 decisions of 8 judges and noted that judges ruled roughly 65% in favor of defendants but that this rate gradually dropped to 0% before lunch more picking up again right after lunch. The study was heavily disputed, with critics arguing that the order of the cases was scheduled according to their likely outcome (Weinshall-Margel & Shapard, 2011) (the authors replying that this was incorrect (Danziger, Levav, & Avnaim-Pesso, 2011b)), and others positing via separate experiments that the effect was overestimated (Glöckner, 2016).

These studies highlight an old debate on how (flawed) our judicial system is and whether it is at all possible to apply the law impartially. Automation,

of course, such as the ideal of a robot judge, elevates this promise even more, yet, as multiple studies on algorithmic fairness have shown, creates a new set of problems for society (Barocas, Hardt, & Narayanan, 2023). Machine learning models trained to issue new court decisions are necessarily based on data from past cases. However, relying on past data—especially with little or no human mediation—is problematic, as older court decisions may reflect the biases and prejudices that were prevalent at the time the decision was taken, and training a model on this data is hence likely to perpetuate these biases. This is especially problematic as past legal systems are known to discriminate against certain groups based on race, gender, religion, and other characteristics—and supposedly to a greater extent than current systems (Zatz, 2013). In addition, laws and legal standards evolve over time, and a model trained on old decisions will not be aware of recent changes in the law, which might lead to incorrect or outdated rulings. And even if the law has not, in fact, changed, the way that courts interpret it might have shifted; such evolution of judicial thinking cannot be reflected in a model that has been trained on data from before the time the shift took place. Similarly, models that are trained on past data are also likely to lack the specific societal context of the current case. This is very visible with respect to changing societal values and norms—changes that are not adequately reflected in a statistical model of past data (see Chapter 4: Challenges and Controversies). Technological advancements can also change the context in which legal issues arise, but such advancements are again not reflected in past data. Legal reasoning itself sometimes requires creativity and the ability to leave established precedent. However, a model trained on past decisions might be overly reliant on precedent and lack the ability to think creatively (Hildebrandt, 2020). These aspects, hence, point to a “freezing of the law” (see Chapter 4: Challenges and Controversies) if past data are used to train decision models for current and future cases. However, the law needs to remain adaptable to handle new and unforeseen situations. Finally, in case of a wrong ruling, the root cause of this could be due to a variety of issues: That the learning algorithm is faulty, that incomplete or erroneous data was supplied to the training system, that the system’s output was misinterpreted by a human; etc. Due to the likely black-box nature of the trained model, it is then, however, on the one hand hard to find the root cause in such cases, and it is hard to assign accountability for mistakes. Going full circle, another problem then becomes evident when considering that, in many legal systems, decisions serve as precedents for future cases. However, if a machine learning model issues

decisions, it is important to ensure that these decisions are well-reasoned (rather than just made) and that they can hence serve as sound precedents.

While machine learning may have the potential to increase the efficiency of court rulings, it is thus evident that relying on models that are trained on old data carries significant risks and challenges that need to be carefully considered and addressed, including accountability, contextualization, legal evolution, and biases. Indeed, many implementations of automatically processable regulation projects have shown the biased results that can occur. For instance, as recently as 2023, a Dutch state agency with the acronym DUO, the one handling requests for student financing, came under investigation because its algorithm disproportionately pointed to students with an immigrant background as subjects of being controlled for abuses of loan or grant fraud (DutchNews, 2024). A very similar case occurred in France too concerning the detection of fraud for family benefits by using algorithms which organisations alleged embedded biases, with the case brought to the highest administrative court (La Quadrature du Net, 2024). Many more similar examples exist worldwide, beyond only the Netherlands and France; we see this pattern from the United States of America (Egan & Roberts, 2021), to India (Tapasya, Sambhav, & Joshi, 2024), to many other countries (see notably Chapter 4: Challenges and Controversies for a deep-dive on a few selected cases).

The question of using automatically processable regulation to increase the efficiency of the legal system is thus clearly a double-edged sword: Yes, processes can be automated and made more efficient for everyone, but not everyone benefits from this automation, and automation is likely to create the issues discussed above. This also becomes apparent when thinking back at the structural changes that occur when law becomes automated. Typically, text-based norms need to be applied in a case-by-case manner, which requires a lot of time and resources. However, with automatically processable regulation, the buffer between the circumstances triggering an application of a legal norm and its actual application collapses (Diver, 2020). Hence, automatically processable regulation changes the pace of application of the law, from human-paced to *machine-paced* (Guitton, Tamò-Larrieux, & Mayer, 2022a). This can be particularly problematic in instances where human emotions and empathy can play a role in certain judicial proceedings (Bandes, 2000; Hulst, van den Bos, Akkermans, & Lind, 2017). We here do not refer to these emotions swaying the court decision, which is a separate concern since also this type of human emotional mediation is suppressed in an automatically processable regulation-based

system; rather, we refer to these emotions and perceived empathy supporting victims to psychologically process what has happened, and possibly improve their ability to overcome traumatizing past events. However, with automation and a fully digital administration, empathy gets threatened (Ranchordàs, 2022). We have yet to see and understand how this temporal collapse of the letter of the law and its enforcement impacts human beings. Research on the subject matter has looked at how outcomes from a robot judge compared to a human judge are perceived (Chen, Stremitzer, & Tobia, 2021) and found that the perception to favor humans is driven by “belief about the accuracy of the outcome and thoroughness of consideration” (p.127), while there are still circumstances where there is a perception of humans and robots being equally fair (e.g. when both offer a hearing or when both can show to clear a specific accuracy threshold). Other researchers have analyzed different AI tools within the judiciary and classified them according to the different phases in which they are applied: Information acquisition, information analysis, decision selection, and decision implementation (Barysé & Sarel, 2024). Their empirical research indicates that automation of information acquisition is perceived as fairer by individuals, especially those who have a legal profession, than automation of information analysis, decision selection, and decision implementation. Legal professionals specifically deem automated decision implementation as being an unfair practice within the judiciary context. These study results complement research on the fairness of automated judicial decision-making: Such studies on the perception of fairness within the context of AI have shown so far, that often algorithmic decision-making is perceived as less fair than human decision-making (e.g., in the context of automated human resource decisions, Newman, Fast, & Harmon, 2020), but also that the task which is being automated matters in terms of the perceived fairness Lee (2018). This latter study by Lee (2018) tested how individuals perceive the automation of different managerial decisions (e.g., work assignment, work scheduling, hiring, work evaluation), which either required mechanical (e.g., objective measure, quantitative data analysis) or human skills (e.g., subjective judgment, emotional intelligence). The study showed that when decisions requiring human skills were automated, the decision-making was perceived as less fair and trustworthy and evoking negative emotions, while decisions that require mechanical skills are perceived as equally fair if made by an algorithm or human. While these studies reveal interesting insights into the complexity of human perception of automated decision-making, still lacking today is research that

investigates the impact of automatically processable regulation projects, specifically with the angle on how the *speed* at which a legal decision is rendered impacts individuals. In other words, empirical research on the impact of machine-paced legal decisions rather than human-paced decisions is needed to understand how individuals perceive this collapse of the law and its enforcement. Such studies should be conducted in court decision settings but also outside of courts, since the premise of automatically processable regulation is to be applied widely for all interactions that are guided by law—that is, most of our everyday interactions as societal beings. As argued in the literature, it is likely that machine-paced decision-making in the legal context leads to increased anxiety, in particular for elderly individuals who are not used to this change (Ranchordàs & Scarcella, 2021). In fact, many individuals appreciate human contact and conversations prior to obtaining a legal recommendation or order: Human mediation of legal decisions is hence not only technically relevant but also psychologically. Automation removes this human touch and might also mediate away a central element of an administration: Being able to show empathy for individual circumstances and forgiving certain (legal) mistakes (Ranchordàs, 2022).

For these many reasons, efficiency is often not the right angle to approach automatically processable regulation implementations. While we will discuss in the next chapter the challenges of automatically processable regulation and controversies that have been raised, it is key to first sketch out the characteristics of those legal norms that lend themselves to automation better than others, that is: To characterize those norms that we are (technically) *able to* automate, and where we also (ethically) *should consider* automating them.

3.3 WHEN AND WHEN NOT TO

Some legal norms will be more easily transformed into automatically processable regulations and raise fewer challenges in the process. The question is, how to identify those norms? The Service Innovation Lab (LabPlus) of the government of New Zealand² which ran the experiment of bringing different stakeholders to the table for 3 weeks to turn legal text into legal code noted in their final report (LabPlus, 2018) that *five characteristics* of legislations (or part thereof) facilitate the transformation into automatically processable regulation. These characteristics are (1) legislation that involves calculation, (2) legislation that prescribes a process that is used repeatedly, (3) legislation that prescribes a compliance process or

Part of the regulation involves:

Examples:

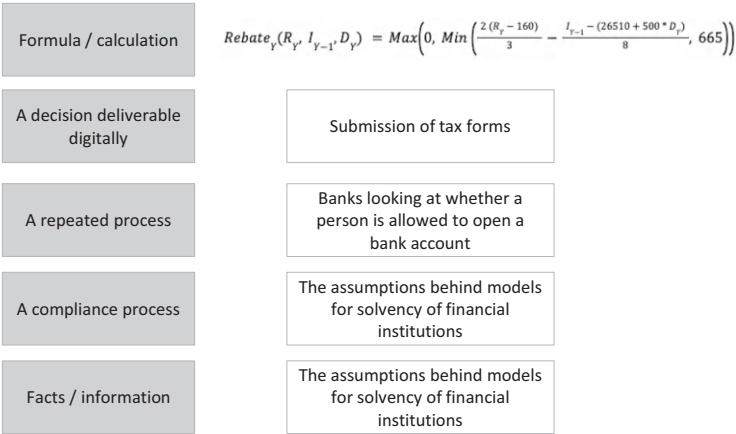


FIGURE 3.2 Excerpts of what to factor in when gauging whether parts of the regulation are well-suited for automatically processable regulation (this is, however, not a decision tree nor as binary as it appears!).

obligation, (4) legislation that prescribes a process or a system that can be delivered digitally, and (5) legislation that involves a process that requires factual information to determine application, eligibility, entitlements, or coverage. The already mentioned OECD “Cracking the Code” report that used the terminology of “rule as code” (Mohun & Roberts, 2020) builds upon these criteria and provides us with factors to decide whether automatically processable regulation should be considered or not, irrespective of social issues. We build on this framework and suggest further additions, notably when it comes to considering open-texture terms.

In the following, we discuss each of these categories shown in Figure 3.2 and provide examples of what can be done and how so. While the OECD report seems to argue that these factors can provide a decision tree of when to and when not to automate, the description that follows shows the nuances that such a decision has to take into account and the lack of an easy protocol or answer to the question of “when and when not to” automate the law.

3.3.1 “Involves Calculations”

As an example of a regulation that involves calculation, we take the Rates Rebate Act from New Zealand and illustrate how this might be turned

into automatically processable regulation and what could be done with the resulting automatically processable regulation. In the context of the Rates Rebate Act, a ratepayer is an individual, household, or business that pays fees or charges for a specific public service or utility provided by a government or a utility company. These fees or charges are typically assessed based on the rate or usage of the service but may be subject to rebates. Such rebates, hence, are a type of financial assistance program that is offered by (typically, local) governments to provide relief to individuals or households who may be struggling to afford their property tax bills, which may represent a significant financial burden. To qualify for a rates rebate, applicants usually need to meet certain income thresholds or other criteria. The exact process for applying for a rates rebate and the amount of assistance provided depends on the policies and regulations of the specific municipality or local government offering the program. Figure 3.3 shows a specific piece of regulation that specifies the amount of rates rebate a ratepayer of a residential property is entitled to.

In our concrete example, the rebate an individual (residential property) ratepayer, P , is entitled to in a rating year, Y , according to this regulation, depends solely on the rates payable for the rating year, R_Y , the ratepayer's income for the preceding tax year, I_{Y-1} , and the number of dependants of the ratepayer at the beginning of Y , D_Y . It is, hence, possible to efficiently turn this piece of regulation into executable code. Specifically, an interested layperson might produce the following closed-form mathematical equation based on the legal text:

$$\text{Rebate}_Y(R_Y, I_{Y-1}, D_Y) \\ = \text{Max} \left(0, \text{Min} \left(\frac{2(R_Y - 160)}{3} - \frac{I_{Y-1} - (26,510 + 500 * D_Y)}{8}, 665 \right) \right)$$

3 Rates rebate

- (1) A ratepayer who, at the commencement of a rating year, was the ratepayer of a residential property is entitled, on application in that year, to a rebate of—
- (a) so much of the rates payable for that rating year in respect of the property as represents—
 - (i) two-thirds of the amount by which those rates exceed \$160, reduced by—
 - (ii) \$1 for each \$8 by which the ratepayer's income for the preceding tax year exceeded \$26,510, that last-mentioned amount being increased by \$500 in respect of each person who was a dependant of the ratepayer at the commencement of the rating year in respect of which the application is made; or
 - (b) \$665,—
- whichever amount is smaller.

FIGURE 3.3 Art 3 of the Rates Rebate Act, as at 28 October 2021.

This mathematical equation can be readily implemented using any high-level programming language, for instance the Python programming language (Box 3.1):

BOX 3.1 RATES REBATE ACT (ART. 3) IN PYTHON

```
def calculate_rebate(R, I, D):
    # Make sure to handle potential division by zero if necessary and
    # validate inputs as needed
    try:
        R = float(R)
        I = float(I)
        D = int(D)

        # Calculate the rebate
        income_penalty = (I - (26510 + 500 * D)) / 8
        rebate = max(0, min(2 * (R - 160) / 3 - income_penalty, 665))
        return rebate
    except (ValueError, TypeError):
        # Return None for invalid inputs
        return None
```

When comparing to the actual implementation, it becomes evident that the Innovation Lab simplified the formula to quite an extent, writing: $(\text{household_income} - \text{base_rate}) / (\text{dependants_number} * 0.1 * \text{rate})$. Our inquiries as to why have remained without an answer, but here would be a putative one: According to their own minutes, they spent a great deal of resources to map out the process and seem to have in the making neglected the actual formula, still at the core of it. That one author of this book was able to see the error is comforting in a way: It means that their transparency allowed for it. But it is also a reminder of the important role of researchers and civil society in keeping the state in check, despite no formal institutional mandate to do so. Had the mistake been implemented in an actual functioning process, correcting it would certainly have been extremely difficult with an important power imbalance: On one side, a whole bureaucratic machinery potentially unwilling to own up to their mistake, on the other, one attentive scholar.

3.3.2 “Can Be Delivered Digitally”

From an accessibility point of view, we believe that many individuals would see it as beneficial if they had access to this calculation with few barriers, for instance, to evaluate the effect that taking on a new job would have on their rates rebate. Such a service could clearly be delivered digitally by a

government, for instance through the creation of a Web application that takes user inputs (rates, income, number of dependents) and returns the correct rates rebate for a user. Based on the above code, it is readily possible to create a Web service that calculates this function and returns the result to users of the service. To illustrate the simplicity of this step once the law has been turned into automatically processable regulation, we have tasked the GPT-4 large language model (which was the most advanced language model at the time of writing) to create, in the Python programming language, a Web application that makes use of the Flask framework and realizes the automatically processable regulation version of the given Rates Rebate Act. The prompt used was simply: “Create a simple python/flask web server that takes as input an income I, rate R, and number of dependents D and returns as output the result of this mathematical equation: $\text{Rebate}(R, I, D) = \text{Max}(0, \text{Min}(2(R - 160)/3 - (I - (26510 + 500 * D))/8, 665))$ ”. We have subsequently verified that the generated code indeed performs the requested function and makes the implementation of this Web service available here (part of the generated code is shown in Box 3.1). Please note that we explicitly do not advocate using automatically generated code without human verification, and certainly not for an application on which ratepayers might base financial and life decisions.

3.3.3 “Repeated Process”

Many processes in bureaucracies—be they within the state or private institutions—have a certain repetitive element, the extent of which will vary. In the above example of a calculation of a rate rebate, not only the calculation is a repetitive element, but further aspects of the process will be: Selecting relevant documents, submitting them to an office, verifying the validity and relevance of the document, extracting figures relevant to the said computation, communication of the decision to the applicant, and enforcement of the decision (e.g., transfer of funds). Such a description of the process is rather generic, and as such, describes not only a process to apply for rates rebate but more generally (from how much overtime a worker is entitled, how much taxes a resident has to pay, how much social benefits can someone obtain, and so on). Even processes *without any calculation* would be well-suited: There are e-banking solutions that allow users to open up a bank account via the submission and analysis of documents without the involvement of a human operator. This is similarly possible only because (part of) the regulation on requirements that a bank has to perform has been turned into an automatically processable regulation form.

The movement under the name “rule as code” therefore sees it as reductionist to only look at the encoding of the regulation. They see it as an opportunity to revamp processes that were likely created to cater to the regulation, but with the technology and needs of the time, both of which may now be outdated. They, therefore, advocate to carefully map out the full process in which an automatically processable regulation would be taking place and identify the potential for improvement, which would lead to better policies. Those advocates have been living to their words: In the aforementioned example of the Rates Rebate Act, the team, living up to the transparency ideal, put out their drafts of how they look at the process and how it integrates. The resulting graph for such a small piece of regulation is so extensive that it is nearly impossible to read without zooming in on specific parts.³ But they basically break down the process of a user first becoming aware of the law, receiving the form, filling out the details, finding more information, submitting the form, remedying errors and/or providing more information, potential payment issues, and finally, receiving payments. The whole process can take roughly 6 months.

For such processes, it is unlikely that there is only one secluded part of the regulation to encode, as was the case when calculations were involved, for instance, with the dozen of lines above turned into a nice closed-formula. Instead, many regulations will impact the code, and conversely, many lines of code will be scattered around to represent just one part of a regulation, making it difficult to assess equivalence. In order to simplify the reading of code and whether there is indeed equivalence, scholars have developed Catala, “a programming language for law”, as the title of one of their publications announces it (Merigoux, Chataing, & Protzenko, 2021). The language aims at highlighting where in the software code part of a regulation has been encoded. Their code base provides examples (see Figure 3.4) of mixing calculations and processes: Family benefits, taxes, or inheritance taxes.

3.3.4 “Compliance Process”

Legal compliance is costly, and thus it is not surprising that many companies have looked into means to automate certain compliance practices. For instance, in the financial sector, regulators have to approve how financial institutions model their assets and liabilities to ensure that they have sufficient funds to cover a wide range of different scenarios. Models have to make a many assumptions, and regulators have to check these assumptions, be convinced of their standing, and approve the models. The

Article L521-3 Chacun des enfants à charge, à l'exception du plus âgé, ouvre droit à partir d'un âge minimum à une majoration des allocations familiales.

```

489 champ d'application CalculAllocationsFamiliales :
490   règle majorations_allocations_familiales.droits_ouverts de enfant
491   sous condition
492     (enfant dans ménage.enfants) et
493     (enfant != ménage.enfant_plus_âgé) et
494     (enfant.âge >= 1521_3.âge_limite_alinéa_1 de enfant)
495   conséquence rempli

```

Toutefois, les personnes ayant un nombre déterminé d'enfants à charge bénéficient de ladite majoration pour chaque enfant à charge à partir de l'âge mentionné au premier alinéa.

```

500 champ d'application CalculAllocationsFamiliales :
501   règle majorations_allocations_familiales.droits_ouverts de enfant
502   sous condition
503     (enfant dans ménage.enfants) et
504     (nombre de ménage.enfants >= 1521_3.minimum_alinéa_2) et
505     (enfant.âge >= 1521_3.âge_limite_alinéa_1 de enfant)
506   conséquence rempli

```

FIGURE 3.4 An example from Catala reproduced from <https://github.com/CatalaLang/catala>.

back-and-forth between regulators and companies is resource-intensive, and can also make the application of similar regulation to a wide range of companies difficult with so much judgment involved. This shows, that legal compliance processes are, due to their repetitiveness (see Subchapter 3.3.3: “Repeated Process”) and the financial incentives to automate it, often prone to be the first processes to be turned into automatically processable regulation.

Also, in data protection law the possibilities of automation to comply with the law have been showcased. For instance, the DAPRECO knowledge base (Robaldo et al., 2020) aims to facilitate GDPR compliance by providing a structured and machine-readable representation of the GDPR’s legal provisions and constraints. It allows for the automated assessment of whether specific actions or data processing activities comply with GDPR requirements, which can be a valuable resource for anyone who seeks to adhere to European data protection regulations. DAPRECO is a knowledge base that contains rules and constraints related to GDPR. These rules are written in LegalRuleML, which is a formalism designed to represent the logical structure and content of legal documents, including laws and regulations (see Chapter 2: Law and Computer Science Interactions). To give readers a rough idea of what LegalRuleML representations look like, Box 3.2 transcribes “every man is obliged to run”:

BOX 3.2 EXAMPLE OF A LEGALRULEML REPRESENTATION FROM ROBALDO ET AL. (2020)

```
<lrml:PrescriptiveStatement key="someuniquekey">
  <ruleml:Rule closure="universal">
    <ruleml:if>
      <ruleml:Atom>
        <ruleml:Rel i ri="man" />
        <ruleml:Var key=":x">x</ ruleml:Var>
      </ ruleml:Atom>
    </ ruleml:if>
    <ruleml:then>
      <lrml:Obligation>
        <ruleml:Atom>
          <ruleml:Rel i ri="run" />
          <ruleml:Var keyref=":x" />
        </ ruleml:Atom>
      </ lrml:Obligation>
    </ ruleml:then>
  </ ruleml:Rule>
</ lrml:PrescriptiveStatement>
```

Agarwal, Steyskal, Antunovic, and Kirrane (2018) proposed a GDPR compliance assessment tool that extends *Open Digital Rights Language (ORDL, an ontology to express rules in an automatically processable manner)* in a way that not only digital rights but also legislative obligations can be represented. When parsing the text, the authors used ODRL to begin by extracting the text that specifies obligations from the legal documents. Subsequently, they pinpointed and defined the connections between these extracted obligations, aligning them with the legislative framework. This requires identifying which text parts are related to different components (like duty and dispensation) and understanding the links made within legislation (e.g., when articles point to other articles within the GDPR or to other legal sources). Afterward, they converted these modeled obligations into a compatible format that can be understood by the compliance system, such as an RDF format (see Chapter 2: Law and Computer Science Interactions). Similarly, the tool called SPECIAL was developed to comply with consent and transparency obligations under the GDPR and it uses RDF to express data processing activities (Kirrane et al., 2018). Within their framework, the authors showed how data processing activities can be

modeled in order to automatically verify if the data processing and sharing of data comply with the relevant usage policies of a company.

3.3.5 “Factual Information”

The final characteristic that LabPlus mentions as facilitating the expression of legal norms as automatically processable regulation is when a rule is applied based on factual information. We extend this notion to the characteristic that the regulation in question should be constituted of unambiguous factual bits of information that are clearly composed. For example, a norm that states “the speed limit is 50 km/h” is based on clear facts that can be measured to establish a rule; and such a rule can also be monitored and enforced automatically, as the presence of speeding cameras in many countries demonstrates (although here as well, different way to interpret the law, encode it, and issue fines can easily ensue, see notably Shay, Hartzog, Nelson & Conti, 2016). On the other hand, though, in a norm that states “drivers must drive safely”, the term “safely” needs to be contextually determined under the given circumstances. Thus, the norm “drivers must drive safely” requires an ex-post analysis of the situation, taking into account the factual circumstances (e.g., visibility on the road, traffic), balancing different societal interests (e.g., safety on the road, certainty to be able to drive a certain speed on highways). While it is also possible to turn such a piece of regulation into automatically processable regulation and to monitor it automatically, this will require a great deal of interpretation. To illustrate this, imagine creating a “safety camera” that permits to automatically monitor whether drivers are indeed driving safely and that would, in the circumstance that they are not, issue *unsafe driving* tickets—fully automatically. Towards the construction of such a system, and early in the process, the required sensors will need to be determined along with the algorithms that process the sensed data, which necessitates a clarification of the term “safety” in this circumstance. In the context of driving, safety might encompass following traffic laws such as driving within speed limits, maintaining control of the vehicle, not driving under the influence of substances, staying attentive, and responding appropriately to ambient conditions. Suppose that “visibility” is determined to constitute a relevant component of these conditions and that it refers to the driver’s ability to see and be seen by others. This is affected by weather, time of day, and obstructions as well as vehicle lighting and signals. Sensors that could be used to measure visibility include light sensors (to determine ambient light levels), cameras (to capture videos of the road and driver’s surroundings), LIDAR

or radar (to measure the distance to objects and other vehicles in adverse weather conditions), and weather sensors (to detect adverse weather conditions such as fog, rain, or snow, which can severely reduce visibility). The data collected by these sensors may then be processed to analyze ambient conditions, to assess vehicle lighting (e.g., by checking if the vehicle's headlights and taillights are functioning and being used appropriately), and to evaluate the driver's response (e.g., to monitor if the driver is adjusting their speed and following distance). Suppose that these combined sensing and analysis systems are effective in measuring the constituting factors of "safe driving"; next, thresholds would need to be derived to permit the automatic assignment of fines to different unsafe driving conditions, based on these measurements, which constitute a large challenge by itself. Especially when contrasted with regulation that states "the speed limit is 50 km/h" instead of "drivers must drive safely", the issues when automating the latter become apparent.

These problems are rooted in the use of *open-textured terms*, in this case around "driving safely": Unlike norms which unambiguously determine a legal result based on one or more objectively measurable facts (Sullivan, 1992), open-texture terms are ones that must be interpreted based on established norms or underlying principles and from which different legal conclusions and results can be drawn depending on the current context in which the open-textured norm is interpreted—including the political, societal, and cultural context (Buchholtz, 2019; Sullivan, 1992). Open-textured norms thus refer to norms with semantic indeterminacy, i.e., terms that are prone to trigger different understandings among individuals (Vecht, 2020). This is also why legal (philosophical) scholars have argued that turning open-textured norms into automatically processable regulation is especially problematic as it requires encoding one definition and thus not representing the whole scope of possible interpretations (Cobbe, 2020; Diver, 2020; Hildebrandt, 2020). To illustrate this open-textured terms like "good faith" or "reasonableness" are often evoked, where a clear tradeoff between legal certainty and flexibility of interpretation is apparent.

For lawyers, the question will come up of "what is *not* open-textured"? The question is legitimate as multiple court cases and our own research have shown that lawyers may question nearly every word of a legal norm. This is in line with Hart's famous argument that the law is *inherently* open-textured (Hart, 1994; Schauer, 2013). Court cases have illustrated this phenomenon, such as for instance the Swiss Federal Supreme Court that had to determine whether a man who stopped at the red light,

unbuckled his seat belt, and buckled it back up before continuing the journey must be considered “driving” and thus could be fined for not being buckled in “while driving”. In other words, the question brought forward was, is being stopped at a red light still considered “while driving” and thus unbuckling the seatbelt a violation of the traffic law norm that states that one needs to be buckled in a while driving? The court decided that yes, even if one is stopped at a red light this needs to be considered “while driving” and thus that the fine was legitimate (Swiss Supreme Court, Decision 6B_5/2011). The Swiss Federal Supreme Court case shows, that individual words challenge the very application of the law itself, enabling different interpretations which are also changing with time.

At times open-textured norms are called ambiguous or vague norms. Ambiguity is a concept that has been studied in linguistics and defined as containing different facets, such as lexical ambiguity (e.g., when words can be both nouns or verbs) and syntactic ambiguity (e.g., as the one referred to above in the Oakhurst case) (Sennet, 2021). In addition to ambiguity, there are terms that are called “essentially contested concepts” (Gallie, 1956), meaning terms that are value-laden to the amount that there will never be one agreed-upon definition of the term. When it comes to the analysis of turning legal text into automatically processable regulation, the nuances between ambiguous terms, contested concepts, under-specifications, or vagueness do not matter that much, and we, therefore, refer to them under the umbrella term “open-textured terms”.

Research on open-textured terms and how to deal with it when automating the law has been proposed: Previous research has tackled the challenge of open-texture by introducing a “value-ontology” that features two axes, one in which security appears to conflict with freedom, and the other in which equality seems to oppose utility (Benzmüller, Fuenmayor, & Lomfeld, 2020). While such a duality can be questioned (Solove, 2011), there is an ongoing discourse that underscores the critical need for a more comprehensive understanding of how to approach open-textured terms. Existing studies have evaluated the complexity of legislation through computational analysis, examining aspects such as structure (e.g., paragraphs, sections, articles), language (e.g., token count, word length), interdependence (e.g., citations), indeterminacy, vocabulary diversity, and readability (Bourcier & Mazzega, 2007; Katz & Bommarito, 2014; Palmirani & Cervone, 2013; Walzl & Matthes, 2014). Of these aspects, indeterminacy bears the closest connection to open-texture, and Walzl & Matthes (2014) identified 62 legal terms falling under indeterminacy. Notably, their work

does not specify the methodology used to create the list (or provide the list itself). In our own research, we show that it is relevant to understand the number of norms that are open-textured within a given law (Guitton, Mayer, Tamò-Larrieux, Garcia, & Fornara, 2024). This is key to determine also what parts of the law can and cannot (or should not) be turned into automatically processable regulation. Preferable approaches to crafting automatically processable regulation involve clear, unambiguous rules, as interpretation reintroduces transaction costs that automatically processable regulation seeks to reduce. When open-textured terms emerge, it necessitates human input or consultation with appropriate case law to ascertain the provision's final outcome. Nevertheless, from a legislative standpoint, open-textured terms can serve a purpose and prove cost-effective in specific situations (Huber & Shipan, 2012), as allowing for open-texture can lower the expense of drafting legal documents. Consequently, assessing the extent of open-textured elements within a law becomes instrumental in determining the feasibility and value of employing automatically processable regulation for that law. To address this, we proposed a framework and applied it to different legislation. We showed that a surprising amount of norms contain open-textured terms; yet also point to the further research needs in this field as our first attempt has many limitations (Guitton, Tamò-Larrieux, Mayer, & Djick, 2024). Most notably, we have only been looking at open-texture at the “text level”, meaning whether the terms were vague, ambiguous, or under-defined when only considering the legal text in which they are written. But law is not a compartmentalized endeavor: Within a legal field, other texts, be they statute or case law, could refine the meaning of what was originally thought to be an open-texture. To be able to know whether such a refined definition exists, however, requires in-depth knowledge of the legal field—hence being more expensive to ask annotators for it—or a breakthrough in mapping the network of related laws, a breakthrough which has not happened yet. When it comes to merely the “text level” though, our research has shown that there were good reasons to be hopeful in automating this flagging. We asked the famous large language model gtp-3.5-turbo (we also asked others, but this one came on top) to flag open-texture and asked reviewers to evaluate the output (Guitton, Gubelmann, Karray, Mayer & Tamò-Larrieux, Forthcoming). The reviewers looked for false negatives, false positives, and true positives, and when combined into a single F_1 score, we obtained that GPT reached 84%, something relatively high. But also high was the number of clauses within the GDPR as containing at

least one open-texture at roughly 80% (or higher in the case of using large language models). This is an interesting finding in so far that there is an old debate about the preponderance of open-texture as to whether it is everywhere or not. Our results show that, while it is not literally everywhere, it is very well present. In a sense, this may not come as a surprise as after all, legal professionals and students are trained in considering opinions and arguments from someone opposing them. Still, as a consequence of this high number of open-texture, all clauses containing at least one open-texture term become much more difficult to be turned into an automatically processable clause.

Furthermore and lastly, throughout this research on open-texture, we have uncovered another illuminating finding: That there is very little agreement between annotators whether terms are open-texture at all. The statistics on the preponderance of open-texture are based on at least two annotators agreeing but that happened rather rarely. This means that flagging open-texture is actually even more difficult than what we have presented so far, and that its preponderance could be even higher if we considered as open-texture any terms where at least one person thinks that it is.

Both of these findings—on the preponderance and on the lack of agreement between annotators when it comes to open-texture—has real world consequences. If we think back of the projects mentioned at the opening of this chapter, namely the robot judge in Estonia and Mes Aides in France, both will have had to have been encoded with only one interpretation of the text, even though it is fair to assume that many more possible interpretations of many clauses would have been possible. This lack of reflection for the plurality of interpretations is an issue: Who is to say that the interpretation in the automatically processable regulation is the most legitimate and the most appropriate, simply because it stems from a branch of government? How to dispute it, especially if one disagrees with it when considering the letter of the law? These questions point towards the different challenges that arise when turning law into its automatically processable form which will be the subject of the next chapter.

NOTES

-
- 1 See <https://www.mesdroitsociaux.gouv.fr>.
 - 2 <https://serviceinnovationlab.github.io/>
 - 3 <https://docs.google.com/drawings/d/1G3cs66o7u-xwJWr2FT46cd5UivJ-HoipcS3I-Q5bSOY/edit>

Challenges and Controversies

CHANGING THE WAY THE law is being applied brings its own set of challenges and has downsides, especially when *not* done in a responsible manner. Many of these challenges are not new but have been discussed for many years, especially within the literature on AI or other technologies and the ethical, social, and legal challenges that arise from their deployment. The discussion on legal automation mirrors to a great extent discussions held on responsible AI, yet with some notable differences that we cover within this chapter. Nonetheless, the overarching themes such as accuracy, reliability, interpretability, non-discrimination and fairness, economic and human impacts (e.g., with respect to job replacement), or trustworthiness are all ones that also arise when thinking about AI and law more specifically. This is useful for the field, as we can rely on frameworks, best practices, approaches, and even tools that have been developed in other contexts and apply them to the legal automation sphere. While we cannot apply these frameworks blindly, they can provide a starting point to discuss the issues arising when turning law into its automatically processable form and even help us make sense of controversies that have emerged in real life because of the implementation of automated decision-making systems that triggered a legal effect.

A famous framework in this field comes from the High-Level Expert Group of the EU (HLEG) which established ethics guidelines for trustworthy AI in 2019. These guidelines have centered on accountability, human

empowerment and agency, oversight, technical robustness and safety, privacy and data governance measures, transparency, non-discrimination and fairness, and societal and environmental well-being. Other organizations and researchers have relied on similar categories to frame the discussions on the issues that arise with the increased use of AI in different domains (Fjeld, Achten, Hilligoss, Nagy, & Srikumar, 2020; Floridi et al., 2018; Loi, 2020). Also, governments have discussed guidelines to ensure that automated decision-making systems deployed within government agencies and often also with a direct impact on citizens follow responsible practices. In the Netherlands, for instance, the Ministry of Interior developed in collaboration with a Dutch university an instrument to assess the impacts of such automated decision-making tools on human rights (Utrecht University, 2021). These guidelines are open-ended questions that should foster an informed and inclusive debate among different stakeholders. Such initiatives have been seen in other countries as well (e.g., UK Government, 2021). While such initiatives and discussions are needed (see also Chapter 5: Needed (Public) Debates), we nonetheless often see a gap between guidelines, frameworks, metrics, or checklists and the real-world implementation of those documents. A meta-review (Prem, 2023) reports 106 AI frameworks, criteria, metrics, or checklists of issues and concludes that the surveyed constructs do not provide much input as to translating their insights into practical and implementable recommendations—this is particularly problematic since we argue that the target audience of such frameworks is not confined to lawyers, compliance officers, legal scholars, and regulators; rather, they should enable sponsors and implementers of automatically processable regulation projects to understand and follow the (valuable) knowledge therein in concrete projects. The author of the meta-review furthermore notes that most of these frameworks “do not consider the practice of AI system development such as the typical trial-and-error approach” (Prem, 2023, p. 702). These results are not surprising, as the real world is messy, and the implementation of guidelines requires constant negotiations and decisions about how to trade off and balance conflicting interests. Yet, even if this is a hard problem, it is one we need to tackle and address as incremental steps are needed to enable long-term changes towards responsible automatically processable regulation.

4.1 REPRESENTATIVE AND BALANCED AUTOMATICALLY PROCESSABLE REGULATION

As already explained, choices need often to be made between several possible interpretations when turning law into its automatically processable form. Many laws include terms that are vague and ambiguous. At times, this is on purpose in order to achieve political consensus or so that laws apply to a wide range of scenarios. Moreover, the *interpretation of the words and norms evolves* throughout time too. Social mores evolve, and law either captures or is a reflection of this evolution. For instance, some laws may exist on paper but are not enforced anymore, such as in Switzerland, where there was a ban in place in many cantons to live in concubinage during most of the twentieth century. The bans were only lifted formally in the mid-1990s but were already not enforced in practice for years as an increasing number of people lived in such arrangements, and accepting it corresponded to the newly formed consensus (Forstmoser & Vogt, 2012). The issue is how automatically processable regulation should reflect this evolution of the interpretation and enforcement of the law. It may be particularly difficult sometimes to say when society has reached a turning point. The US provides a vivid example with the Supreme Court overturning the right of abortion granted under *Roe v. Wade* in June 2022. Three months later, police gave a woman a fine for driving alone in a lane reserved for those driving with at least two passengers (so-called high-occupancy lanes). The woman disputed the fine arguing that the overturning of *Roe v. Wade* meant that the US state recognized her child as another human being. Her argument worked, and the ticket was dismissed (Romo, 2022). But is this case of a fine in a high-occupancy lane sufficient to say that this interpretation can be applied across the board to other contexts as there has been a cultural change in society? It is not clear, and with a lot of push from both sides of the political spectrum—liberals and conservatives in the USA—which have their own moral views, one recognizes the difficulties of implementing within an automatically processable regulation project the “correct” interpretation of the law.

Another concern is the actual legislative evolution of statutes (rather than their mere interpretation and enforcement). Because policymakers update laws, automatically processable regulation versions of a given law should ensure that their design allows them to easily identify which parts of their code relate to which part of the law and when there has been an update to these parts. This should ensure that the automatically processable

regulation implementation is in line with the most current version of a statute. Without falling into techno-determinism, this could appear to be a problem with technical solutions—much more so at least than the issues raised about the evolution of the interpretation of law within society.

The evolution of interpretations and legislations themselves has given rise to the fear that, notwithstanding any technical implementation, automatically processable regulation is prone to “*freezing the law*” (Hildebrandt, 2020). Freezing the law means that interpretations of the law would be hardcoded without the possibility of updating them. This fear is especially stark when it comes to automating judicial processes based on case law, as such algorithms would have severe repercussions: A system trained only on old cases with old interpretations could never give a new interpretation to the law, hence freezing its interpretation. Relatedly, there is a certain risk when seeking to solve the problem of representativeness and evolution to fall into the trap known under the term “*legal singularity*” (Deakin & Markou, 2020a). This term refers to achieving certain (or at least predictable) outcomes of the law. There would, as a consequence, hardly be any need for lawyers to argue and counter-argue cases. Many find this to be an unattainable scenario and/or one we should not strive for (Cobbe, 2020; Pasquale, 2022), while others that it is attainable and an ideal we should strive for (Aidid & Alarie, 2023). The risk is, however, rather clear: By seeking to weed out uncertainty in the law, which potentially is a defining feature of the law itself, we flatten out nuances, push out different interpretations, and, in so doing, not only weaken the law. Most likely, we do so at the cost of impacting minorities rather than majorities, with majorities invoking democratic principles that the most numerous are “right”. Take the definition of the word “marriage” as the union between a woman and a man. Heterosexuals couples are more numerous than same-sex ones. Just based on the word, and provided that there are no new statutes clarifying the meaning of the word, it may be tempting for a software engineering team to side with the definition of the most populous, hence minimizing the number of people negatively affected by the definition (following moral principles of consequentialism rather than deontology, or in other words, a standard application of the trolley dilemma). By doing so, the minority, here same-sex couples, would be left unable to enter into a software program that they are married.

4.2 COMPUTER SAYS NO

Imagine you have a brand-new car with your local traffic regulation embedded in the car as an automatically processable regulation. For the vehicle to be accepted on the road, the regulator requested that—as per the law—any driver should always and under any circumstances comply with the law. The car, hence, does not allow you to digress the law. For instance, if the speed limit is 50 km/h in the city, then your car will not respond to you accelerating above this limit. One day, your partner has a medical emergency—a pregnant woman is about to deliver, a heart attack, a life-threatening wound— and you decide that the quickest and best way is to take your car to drive to the hospital. It may be night time, and no one is at the red light, but the car will refuse to let you burn it and will let you wait until it turns green. On the highway, despite being nearly empty, you will not be able to go anywhere above the speed limit. In short, you, the driver, lose the ability to act on your decisions, regardless of whether there was a legitimate reason for you to do so. The driver loses their agency—freedom of action—and there could be a serious underlying assumption in this hypothetical example: That the machine’s decisions are better than those of humans’.

Such an assumption goes by the name of “*algocracy*”, meaning that the algorithm can or should overrule human decisions at any time (Danaher, 2016). As previous examples have made clear, there are many reasons to be fearful of algorithms, and to embed such a decision hardcoded in the machine is one additional one. It is probably a fair assumption to posit that we cannot foresee every single situation when designing rules, on paper or as automatically processable regulation. Human judgment will have to come into play, especially in situations that are imbued with unclarity. Emergencies could be one such situation, but the point even goes beyond that and has a political dimension to it, one of being able to conscientiously disobey either the law or what someone might deem is the wrong interpretation of the law as encoded into the automatically processable regulation. To reuse the example of the high-lane occupancy mentioned earlier where a pregnant woman drew attention to the status of the unborn child: If the car had overruled her intention of using the lane, she would not have been able to make her point. But law, social, and political life thrive on challenges being brought to it via disobedience notably, as many social movements, from the Yellow Umbrella movement in Hong Kong to the Black Lives Matter movement, illustrate so well. In other words, disobedience

is a vital tool in democracies to challenge any institution with power, be it the executive governments or courts deciding on how to enforce laws or legislators deciding on how to craft them. Those making the decisions on how to interpret and consequently implement automatically processable regulation have similarly certain powers against which, within this tradition of democratic institutions, staging protest and disobedience should be possible. And this can only be possible if humans are left with a certain agency to overrule decisions made by machines. A key distinction to make is that inherently automatically processable regulation or any automatic decision-making does not take away agency or the capacity to disobey.

4.3 TRANSPARENT AND CONTESTABLE AUTOMATICALLY PROCESSABLE REGULATION

There are but too many examples that showcase the severe consequences when automatically processable regulation projects are not transparent, contestable, or implemented responsibly. A case in point is the United Kingdom Post Office scandal which became public in 2009 but has a long history prior to this event (Wallis, 2021). The scandal involved the United Kingdom Post Office, which pressed charges against their sub-postmasters. Sub-postmasters are those investing in a Post Office and, at the same time, contracted out to run a branch of the post (a bit similar to a franchisee opening a McDonald's). At the center of the scandal was a software keeping tabs on cash flows called Horizon, which was used by the United Kingdom Post Office to bring forward charges against specific postmasters who allegedly committed fraud within their branches. The creators of Horizon insisted that it worked completely error free (a dubious claim that should already have raised red flags); however, as court transcript shows, Horizon was far away from working perfectly, and many obvious computing errors had occurred. Ironically, when Horizon sought to update their systems much later, namely in 2016, to correct known bugs, someone asked whether they could provide the list of such bugs—a request which Horizon's mother firm denied but confirmed the existence of known error logs. The major incident that brought those errors to light occurred in Falkirk, a town of 32,000 inhabitants located between Edinburgh and Glasgow, where the postmaster had noticed a strange behavior: "I have one screen that says I have a £4 gain, and the screen next to it says I have a £13,000 loss on the same stock unit" (Thomson, 2009). In this very instance, a "system error" had caused the mismatch. Upon finding the error, the United Kingdom Post Office made, however, no efforts to see

whether the error could have happened in any of the other systems scattered across the country and used by “78,000 people to process six million transactions every working day”, according to the Office’s own statistics. This example illustrates not only the lack of transparency over the existing automated decision-making but also a lack of responsible implementation, especially once serious errors were made apparent. The United Kingdom Post Office however, continued to rely on Horizon as a system to determine against which postmaster to press charges, which had serious consequences for individuals involved in these trials: The Standard, a newspaper in the United Kingdom, claimed that the consequences of the trials led to bankruptcy and personal hardship, even contributing to at least one person committing suicide after being wrongly accused of stealing £60,000 (Sinclair, 2021). Just to name one public example: Seema Misra, a sub-postmaster, collapsed upon hearing her conviction of theft, then went on to serve 4 months of a 15-month sentence while pregnant. The cases went to a court of appeal, during which a long list of bugs appeared (Flinders, 2019). The court consequently concluded that “bugs, errors and defects did in fact exist in the Horizon system and on numerous occasions had caused financial discrepancies in Subpostmasters branch accounts” (Justice for the Subpostmasters Alliance, 2019). In yet another case, sentences were very much out of line with the alleged theft: One judge awarded the Post Office £321,000 in costs and fees for alleged theft (later exonerated) of £26,000. Convictions were eventually overturned and with strong words that the prosecution amounted to an “abuse of process” and an “affront to justice”, while the *Financial Times* called it “one of the greatest miscarriages of justice in modern British legal history” (Uddin, 2023). But the overturn of the convictions also only started occurring in 2021. In other words, it took more than 10 years to overturn prosecutions, and many are still waiting for this to take place: As of December 2023, only 93 of 700 had been overturned (Uddin, 2023). One reason for this delay has been that the United Kingdom Post Office was actively working against this, “orchestrat[ing] a cover-up hiding crucial information from MPs and campaigners”, as a journalist covering the story put it in a book (Wallis, 2021). But the pace for accountability, prominence on the political agenda, and monetary compensation for those affected picked up notably after the diffusion by the channel ITV of a 4-part mini-series entitled “*Mr Bates vs. the Post Office*” after Alan Bates, one of the subpostmasters who organized and represented the aggrieved group. The diffusion took place at the turn of the new year 2024 and it unleashed media articles in nearly

all British newspapers, statements from the prime minister, Rishi Sunak, with promises, and scrutiny of the government's link with Fujitsu, the parent company that had created the Horizon software. The story is hence a vivid reminder: Any IT system can go wrong, and automatically processable regulation would not be any different from this rule, and overturning this wrong can be extremely tricky and lengthy, bearing a high toll for the victims, notwithstanding that not all victims of a wrongly implemented automatically processable regulation could benefit from a TV show popularizing their fate.

The United Kingdom Post Office scandal highlights therefore the need for *transparency* over automatically processable regulation implementations. We can break down the transparency requirement into aspects that are key even before an automatically processable regulation tool is launched (*ex ante* aspects) and ones that come into play once the tool is used to generate (legally binding) decisions (*ex post* aspects). With respect to the *ex ante aspects*, transparency of the computer code is needed for experts to assess it (also needed is the justification behind its implementation in the first place). There could have been outright errors in the transposition, and showing the code would help expose these. But also, how did the state (or any other implementers if in the case of private companies) make decisions around transposing the law into computer code? They have had to make choices; some of these could be debatable. As laws can be extremely complex, understanding the transposition can also be no easy feat, with its defeasibility, exceptions to exceptions, and references to other articles, sometimes implicitly just by re-using a specific terminology. One of the ways for automatically processable regulation to shine a light on this transposition is, for instance, by juxtaposing the code and the legal text so that by hovering over one, it would highlight where the correspondence is located. Catala, a language developed in France and already mentioned in the previous chapter, aims at doing so (Merigoux & Huttner, 2020).

The *ex post transparency aspects* are related to the explainability of the output generated by automatically processable regulation. An example of a tool that violated the requirements of explainability was the original implementation of the COMPAS tool which stands for *Correctional Offender Management Profiling for Alternative Sanctions*.¹ In short, COMPAS supports judges in the US in their decision of whether a person could be a recidivist. Such risk assessment tools are common around the world, with the United Kingdom using *Offender Assessment System* for probation, the *Oxford Risk of Recidivism Tool (OxRec)* for predicting violent re-offending,

or smaller states using similar risk assessment still in an analog format. In the COMPAS case, a judge has to give certain inputs and receive an output (Dijck, 2022). Which input goes into the software is not public, neither is the algorithm, and neither the defendant nor anyone else will know what the output of COMPAS has been—and so, gauging to which extent it has affected any judge’s decision is close to impossible. This makes, by default overseeing and auditing the outcomes of the tool difficult to impossible as well. Many cases illustrate the inadequate oversight over automated decision-making tools that impact individuals’ life.

For instance, in the Netherlands, another infamous case is the Dutch Child Care Benefit scandal, which the government ignored for too many years (House of Representatives of the States General, 2020). Similar to the student loan case mentioned in Chapter 3: Automatically Processable Regulation, in the Dutch Child Care Benefit scandal Dutch tax authorities accused parents of fraudulently claiming child allowances on the basis of an algorithm for fraud detection, the “risk classification model” (Henley, 2021). Small administrative mistakes led to the “risk classification model” classifying parents as fraudsters, without recourse for parents and caregivers (and later journalists, politicians, and oversight bodies) to understand how this classification happened (Amnesty International, 2021). As a consequence, the tax authorities forced more than 10,000 families to re-pay tens of thousands of euros, leading many to extreme personal situations with bankruptcy and divorces following. Finally in 2020, the tax authorities admitted that they singled out those families because of their ethnic origins and/or dual nationalities; however, such an admission should have come at least 3 years earlier, as already in 2017, the National Ombudsman published a report strongly criticizing their approach and recommending that parents be compensated (European Commission for Democracy through Law (Venice Commission), 2020). In addition, already in 2019, the Central Government Audit Services and Dutch Data Protection Authority started an investigation of the risk classification model (Heikkila, 2022). In short, the government should have been aware that something was amiss and should have taken action consequently. A parliamentary committee publishing their investigation in their report titled “Unprecedented Unjustice” came to the same conclusion. Investigations showed that the “risk classification model” used individuals’ nationality prominently as a driver to flag potential fraudsters (see for a summary and further sources in Amnesty International, 2021). One month following the publication of the parliamentarians’ report, the entire Dutch government resigned

in acknowledgment of their wrongdoing. But in a twist of events, Dutch political parties could not find another agreement for the coalition, and the Prime Minister, Mark Rutte, was back at the helm 3 months later, in 2022 without political accountability being served. He remained Prime Minister until July 2024 when he then became NATO General Secretary in October 2024.

Important lessons learned not only from the Dutch Child Care Benefit scandal but also from the United Kingdom Post Office scandal are that one audit alone will not bring the government and especially the agency in question to change their policies; it took several investigations, and different actors had to apply pressure (not only the victims, but civil society groups, and politicians too). Only having one audit trail as a guarantor to catch errors is unlikely sufficient to catch them, and more importantly, it won't be sufficient to redress any wrong-doing. Furthermore, in the two previous examples, audit mechanisms rested on already established processes within (democratic) institutions, but in light of the technical complexity of automatically processable regulation projects, much more technical expertise will be required, and debates will equivalently need to take place on whether we will need to create a separate instance to investigate such implementations. Another lesson learned is that oversight and audits need to occur in a timely manner: Time that it took for redress to occur in both cases was roughly a decade. Within this very long timeframe, many tragedies followed. A quicker resolution timeframe should seek to avoid such dragged-out processes. Hence, by making audit processes clear, an examination of the possible misdeeds should also become quicker.

Ensuring that automatically processable regulation projects are transparent is linked to enabling contesting certain implementations. *Contestability*, however, is only possible when the baseline assumption for every implementation is that the algorithm can, in fact, be wrong at times. Contesting an algorithmic decision can involve high stakes; for instance, if one contests an automatically processable regulation implemented by a state that led to an official decision (e.g., bail decision, fine). Other times, contesting an automatically processable regulation decision will involve claims against a company that developed and implemented a tool (e.g., legal technologies applied in law firms to increase the efficiency of lawyers). Here, contesting a decision will be more geared towards assigning liability (e.g., if a lawyer realizes the legal technology is erroneous and requests the company developing the tool to amend it or if an individual who sought legal advice and obtained an erroneous advice based on the tool requests

compensation by the lawyer). The contestability of a decision overall will be important in order to ensure that individuals trust that if things go wrong, someone will be held responsible for it. Contestability is thus linked with responsibility as it is about identifying deficiencies to improve and make sure that similar cases do not emerge again. In the case of automatically processable regulation or automatic decision-making mishaps, such processes become often protracted because of the amalgam between the two and because of the personal stakes: Individuals fear for what it means for them personally. As a consequence, a blame game ensues that can mean that there are little or no incentives to owe up to one's mistakes.

4.4 REPLACEMENT OF THE HUMAN TOUCH AND WORKFORCE

Being fearful of change, even regardless of whether technology plays a role, is, if not natural, at least very commonplace. Change introduces uncertainty and challenges our ability to plan ahead. It has been, hence, a near-constant of human-machine interactions that humans have feared technology and feared notably being replaced by it. A case in point is the arrival of the automated teller machines (ATM) in the mid-to end of the 1960s. At the time of their introduction, anyone who wanted to withdraw cash from their accounts had to go to one of the bank's branches, speak to a clerk, and confirm their identity before they could successfully go out with cash. Upon the introduction of the ATM, clerks hence feared that their profession would slowly disappear. But the effects were not those really expected. As the number of clerks required to operate a branch decreased, banks started opening more branches, bringing back up the overall demand for clerks. Over the 20-year span of 1970–1990, the number of clerks working at bank branches then roughly doubled (from 250,000 to 500,000) (Bessen, 2015). While doing so, the tasks and job requirements of a clerk also changed: They went from the mere handling of cash to a role closer to a client/customer relation manager where they were marketing the bank's services (Bátiz-Lazo, 2015). It hence went from a rather low skill type of job to one more demanding in terms of being able to build personal contacts.

Similar fears have started emerging regarding the use of automatically processable regulations and how much they will replace legal professionals (Susskind & Susskind, 2015). It is still much too early to say whether the fears are misplaced and will turn out like the ATM story, or whether there could be more to it. Notably for automatically processable regulation, the

number of jobs impacted is really varied: Legislators and their staff may have to write digitally ready legislation (see Chapter 5: Needed (Public) Debates); judges may have to consider appealed cases not from a lower court but from a robot judge and thus may have to understand how the robot came to certain conclusions; public servants may have to implement, run, and maintain automatically processable regulation version of regulations; and lawyers will probably have to use legal automation to facilitate their crafting of arguments and to find similar cases on which to draw inspiration, or use automatically processable regulation to zero in on cases with higher chances of success. Those outside of the traditional legal profession are likely to use automatically processable regulations to get a better understanding of the law, making it easier for companies and citizens alike to ensure compliance.

With these changes, the pace of legal processes (e.g., obtaining an official decision, appealing it, challenging a contractual term) will be higher—as with ATMs increasing the pace of transactions, legal processes could become quicker. It has a certain appeal to policymakers when considering the very high number of backlogs that judges in certain parts of the world may have. But it also has one large drawback. Certain legal processes, by their nature of being drawn out, offer certain upsides too. Divorce proceedings give the chance for parties to reconsider when they may have decided in the heat of a moment. It can give them time to go through therapy and work out an alternative solution. For victims of violent crime, court hearings can be part of a closure process, which would otherwise not exist if the interaction took place automatized. For certain particularly lonely people, talking to a human clerk, even amidst an administrative procedure, offers a chance for a little human contact, which, again, an automated decision-making system would deprive them of. By reducing decision time and removing many of these options, it is uncertain how widespread the negative aspects could be.

4.5 IS THIS FOR REAL?

Many current implementations of automatically processable regulation are to bring certain services of public administration to users. When doing so, they seek to replicate the tools used within public administration to make decisions. But, in contrast to actual decisions of a public administration, a few automatically processable regulation projects only simulate a decision but do not offer a legally binding one (e.g., the Rates Rebate Act calculator or Mes Aides, both mentioned in Chapter 3: Automatically

Processable Regulation). In such situations, citizens might be led to believe that the tool that they are using is the real one, that is, that it leads to a binding decision, when this is not the case actually. If the tool used by a citizen in the end leads to a different decision than the one made by an official public administration, it is easy to imagine unpleasant reactions. For instance, if an online tool to check the eligibility of social benefits indicates to a family their right to receive from the state certain financial support, but when officially applying for it, they are denied the financial support. The situation is not hypothetical: France's tool Mes Aides (see Chapter 3: Automatically Processable Regulation) was made public after some debates; Mes Aides had been developed using a tool called OpenFisca, a platform that facilitates writing code implementing law (rule as code, or, in our terminology, automated processable regulation), and many countries have used the platform to create similar tools to Mes Aides and *simulate* official government response. Despite Mes Aides becoming very popular, this had not replaced the potential for diverging answers between the tool and the public administration issuing the social benefits. And even if the distinction for many is clear, it does not alleviate the difficulty for citizens to understand why the systems used within the public administration cannot be shared with them. While public administrations often bring forward potential fraud as a reason, others have argued that there should be only one implementation of the social benefit laws which is public (as documented by Alauzen, 2021) as that implementation of automatically processable regulation could genuinely help citizens understand the current regulation better, true to the dictum that "everyone should know the law" (see Chapter 5: Needed (Public) Debates).

Even if an automatically processable regulation implementation is real in the sense of having a real-world impact on the actual decision-making process, it is only "real" if it is affordable and usable for individuals. The question is, thus, to what use are automatically processable regulations if they are out of reach for potential users? Automation of law does bear the possibility to increase the accessibility of the law to many people, possibly a tremendously positive turn on how the law has operated so far in our societies (for more on this, see Chapter 5: Needed (Public) Debates). But this will only be possible if the tools are in and of themselves also accessible to laypeople. One hope is for instance, that prices for obtaining legal advice will be lowered. Legal professionals are currently and almost anywhere on earth very pricey (Hadfield, 2000). If automatically processable regulation solutions do not bring viable alternatives, or if they do not contribute

to lower legal professionals' prices for clients, then this would represent an important missed opportunity. In addition, legal questions and especially answers are often complicated as the law involves many complex and, at times, difficult-to-understand moving parts (Blank & Osofsky, 2020). When providing answers to laypeople, systems should strive at not over-simplifying (and hence distorting by doing so what would be the law), but they should also ensure that laypeople are able to use and navigate their solutions. The concern is somewhat less prevalent when the users of automatically processable regulation are not laypeople.

4.6 TOWARDS RESPONSIBLE AUTOMATICALLY PROCESSABLE REGULATION

Similarly to transparency which contains *ex ante* and *ex post* elements, responsibility can relate to distinct areas, namely the *encoding itself*, the *data sourcing* (i.e., obtaining and processing of data), the *overall implementation* (i.e., its final deployment), and its *outcome* (i.e., the decision the tool takes). The encoding of automatically processable regulation relates to responsibility for errors in the code, which relies on underlying input data that can either contain certain biases from its sampling and/or lead the algorithm to a specific bias. We will talk about biases in Chapter 6: How Education Should Shift, but just to name a few already, well-studied biases are: Label biases, where data is assigned a wrong or misleading tag, selection bias, where the training data is not representative of the population, and within this, more specifically, demographic and population bias where the project leads and/or implementers systematically introduce bias according to their own characteristics, and evaluation bias where the model could perform well on trained data but much less so when generalized.

The overall implementation refers to the larger framework within which the automatically processable regulation is deployed, such as the server on which the automatically processable regulation runs, whether on cloud or in-house, processes to ensure the cyber security of the servers and data, handling of troubleshooting and escalation, and so on—basically, anything which any implementation of an IT service would require. Depending on the exact implementation, different failures can arise. For instance, servers need to have regular updates installed (“patches”), most notably to correct newly found vulnerabilities that a hacker could exploit to siphon out sensitive data (e.g., who has claimed for a social benefit). But more benign problems can also arise: The server can freeze and need rebooting, or the application can start having bugs for certain users on

a specific Web browser, and the issue would have to be investigated with potentially new code deployed to fix the issue. And so, basically, once development has more or less ended, maintenance of the code base and of the infrastructure will have to continue for the whole duration that the platform will be online.

Finally, with respect to the outcome, an automatically processable regulation could perform according to specifications, but there could still be an undesired decision resulting from it—for instance, because the outcome is not morally or socially acceptable. All scandals mentioned in this book had possibly good rationales behind their implementations; nonetheless, undesired results emerged, and uproars followed as a sign that society was not comfortable with the tradeoffs that such implementations involved. When code performs according to specifications but with undesired results, the blame could be on many units: Those who wrote the specification, those who approved it, and those looking at mechanisms to review and catch such outcomes. From the onset, it is foreseeable that the complexity of transposing legal texts into computer code necessitates the involvement of many different stakeholders with diverging interests.

As it emerges from this chapter, the challenges in creating automatically processable regulations are numerous. So numerous that anyone seeking to develop such a solution might feel at a loss when reading the list. The aforementioned concerns have appeared in the past in concrete cases. Taking each issue on its own and advocating for a solution can appear simplistic though: Tradeoffs will have to be struck in many different areas. For instance, making an automatically processable regulation tool more user-friendly might require additional costs in development, which in turn can drive up the price and thus its affordability. Another example can be for an automatically processable regulation to implement several possible interpretations of the law and let users choose their own interpretations. This way, there is less an issue in trying to determine the “best” interpretation possible of open-texture as this would acknowledge that different interpretations are possible, but this would also lead to less guidance for the layperson who will not know how to make a choice. More generally, in order to guide the development of automatically processable regulation on how to navigate this range of challenges and avoid controversies, a responsible automatically processable regulation framework is needed (Guitton et al., 2024).

Many frameworks for responsible use, say of automated decision-making within government, already exist, but none really for automatically

processable regulation specifically. As the Netherlands has been embroiled in the scandal on racial profiling presented earlier, its Ministry of the Interior jointly with a local university developed an instrument (abbreviated IAMA) to assess the impacts of automated decision-making algorithms from a human-rights perspective, as the press release states (Utrecht University, 2019): “Step by step, discussion points are described that must be addressed before the algorithm is implemented. By shedding light on the course of a careful decision making and implementation process for algorithms, IAMA [the very instrument in question] can help prevent situations”, such as the scandal that shook the country. The framework distinguishes between three aspects: The preparation (e.g., the goal of the algorithm), the input data, and the throughput. The format of the framework is clear, the questions it puts forward do not seek to push towards a prescriptive answer (they are mostly open-ended), and the overall impression is much less overwhelming for users than for other frameworks.

The United Kingdom government, although not officially as a response to a specific scandal (the developments with the United Kingdom Post Office appeared, *a priori* at least, to have not had any bearings), also developed its own framework, justifying it as “current guidance can be lengthy, complex and sometimes overly abstract” (UK Government, 2021), thereby echoing the conclusion from Prem (2023) mentioned above that frameworks are often not operationalized. This framework, geared towards public servants, starts with a warning, offering a rebuke to the determinism of technical tools to solve problems (as exemplified in Kroll (2015)): “Algorithms are not the solution to every policy problem”, states the British government’s report (UK Government, 2021). The seven-point framework then provides a case study on how to apply it, is as easy to follow as the one from the Netherlands, and several of its main points explicitly take systems’ users into account: (1) Test to avoid any unintended outcomes or consequences; (2) deliver fair services for all of our users and citizens; (3) be clear who is responsible; (4) handle data safely and protect citizens’ interests; (5) help users and citizens understand how it impacts them; (6) ensure that you are compliant with the law; and (7) build something that is future proof.

Furthermore, and more broadly, several frameworks have already emerged to attempt to provide remedies to the many issues of AI. One such framework was developed by the Alan Turing Institute and focuses on AI systems in the public sector (Leslie, 2019). The framework merges three parts together, each three constituted of several concepts:

Respect-connect-care-protect, fairness-accountability-sustainability-transparency, and process-based governance. The report then focuses on providing consistent definitions of the used concepts, but it is difficult to extract concrete operationalizations, which undermines the utility of guiding the implementation of individual projects.

Building upon these frameworks, we want to offer a more pragmatic and tailored way of evaluating the creation of automatically processable regulation. In any democratic country, any act of creating new legislation is time-consuming. Mostly, it is to allow for consultations to heed several points of view, for negotiations, and for trade-offs, and it bears the upside of being able to flesh out the many consequences that a new piece of legislation could have. Creating automatically processable regulation can be in many ways similar to creating a new piece of legislation, especially—depending on who creates it and whether it is a state institution—, as it can have profound repercussions. Taking the time to consider how the many different issues highlighted here could come up and what ways to possibly hedge them appears hence sensible. What we want to very much avoid is that implementers take what they assume are mere technical decisions when they, in fact, have much deeper societal and political aspects.

In the following (see Table 4.1), we propose a framework to kickstart the debates on how to implement automatically processable regulation responsibly which we have published in the *AI & Society* journal. The framework should not (cannot) reasonably be a box-ticking exercise; it is a stimulant for debates between several stakeholders, including (but not exclusively): Sponsors of projects, different designers and implementers, legal advisers, end users, and affected groups. And it should lead to finding mitigation strategies adapted to the context at hand. The full framework can be consulted under Guitton, Mayer, Tamò-Larrieux, Fosch-Villaronga, Kamara & Pałka (2024).

Ideally, any implementation of an automatically processable regulation would come with a publication of an assessment of the choices made along the different dimensions of the framework (or others), and where tradeoffs occurred. This would allow for the public to be well-informed and take a political stance on it. How likely is this to happen and is there support for this? Currently, this is hard to say. A public debate would need first to take place and not only on this point but on a variety of others; in other words, this needs to be part of a larger discussion on the investments, implementation, and use of automatically processable regulation. What other topics of discussion should entail is the topic of the next chapter.

TABLE 4.1 Framework to Guide the Debates Among Arising Issues in Automatically Processable Regulation Projects taken from Guitton, Mayer, Tamò-Larrieux, Fosch-Villaronga, Kamara & Palka (2024)

Issue Type	Lead Questions to Evaluate the Issue
Vagueness and balancing of interests	<p>Were several possible and valid interpretations of the law implemented?</p> <p>Is there a technique in place to annotate elements of vagueness (e.g., to express probabilistic certainty about the interpretation)?</p> <p>Are vague terms implemented in a way that clearly differentiates them from the rest in the technical implementation of the automatically processable regulation system (to allow flexible modification or configuration)?</p> <p>Does the system evaluate the different outcomes that different interpretations may have, and raise an issue when these different outcomes are fundamentally divergent?</p> <p>Is the human enactment of the same regulation (e.g., in precedents) in line with the automatically processable regulation implementation?</p> <p>Was there sufficient AB testing, verification, and quality assurance before deployment of the automatically processable regulation implementation?</p>
Evolution of norms and statutes	<p>Can the system be adjusted to new interpretations of the law?</p> <p>Can the system enact such adjustments automatically?</p> <p>Is, and how is, the evolution of morale and social norms reflected in the implementation?</p> <p>What is the mechanism of how the automatically processable regulation keeps track of changes, both in terms of new interpretations from evolving social norms and from new statutes and case law?</p>
Lack of interdisciplinarity	<p>What is the demography and professional training of the individuals involved in an automatically processable regulation project? Are the different viewpoints sufficiently represented (business, society, citizen)?</p> <p>Was ethical validation performed and ethical approval sought?</p> <p>How were the points of view of those from a professional/academic minority among the involved taken into consideration?</p> <p>Did reviews between fundamentally different expert groups take place?</p> <p>Who gave instructions to developers and how did developers seek advice when in doubt?</p>

(Continued)

TABLE 4.1 (Continued) Framework to Guide the Debates Among Arising Issues in Automatically Processable Regulation Projects taken from Guitton, Mayer, Tamò-Larrieux, Fosch-Villaronga, Kamara & Palka (2024)

Issue Type	Lead Questions to Evaluate the Issue
Agency	<p>If a human judgment conflicts with the output from the algorithm, is there a process in place to ensure that the human can overrule the algorithm’s decision?</p> <p>Is this appropriately recorded, e.g., in a logfile?</p> <p>Is the implementation sufficiently transparent to show the rationale behind the decision, allowing humans to weigh the arguments against and for breaking the law in an exceptional situation?</p>
Natural pace	<p>Are users of the system made aware that the system delivers decisions at a “non-human” speed?</p> <p>Does the project team anticipate that such a faster processing time would lead to any negative psychological effects on certain users of your system?</p> <p>If applicable, has psychological counsel been sought to verify that this is not an issue?</p>
Workforce replacement	<p>Who is impacted by work replacement?</p> <p>Does it replace work that some people enjoy doing? Does it replace work which was an essential life-support for those doing it?</p> <p>Does it replace work that offered compensation (financial, status-wise, etc.) that some regarded as either fair or even attractive?</p>
Implementation transparency	<p>Can individuals, public authorities, and interested stakeholders have access to the implementation code, training datasets, trained models, and information on the automatically processable regulation implementation and deployment?</p> <p>How can different technical implementations enable more transparent and malleable approaches? Does the public know what they should know?</p> <p>Are the code, training datasets, and trained models easily accessible, for example, without burdensome procedures or intermediaries?</p> <p>Does the state play a role in educating its citizens in reading and understanding automatic processable regulation?</p>
Process transparency	<p>Who verifies that processes are in place to catch errors and correct any wrongs?</p> <p>Is this process communicated publicly, and clearly?</p> <p>Who verifies that data is collected, retained, and managed appropriately?</p> <p>How does the audit process take place?</p> <p>Should private companies that turn public regulation into automatically processable regulation or leverage automatically processable regulation come under an auditing process, and to what extent?</p>

(Continued)

TABLE 4.1 (*Continued*) Framework to Guide the Debates Among Arising Issues in Automatically Processable Regulation Projects taken from Guitton, Mayer, Tamò-Larrieux, Fosch-Villaronga, Kamara & Palka (2024)

Issue Type	Lead Questions to Evaluate the Issue
Affordability	<p>Are the costs to the end user more manageable than through professional support?</p> <p>Should the state support, through subsidies or other means, the development of tools to make the law more accessible, hence fostering the rule of law?</p>
Usability	<p>To which extent is the development user-centric?</p> <p>Are there aspects of the projects (micro or macro) that are unclear as to whether there has been a public debate around, and whether or how to bring this debate about?</p> <p>How can we ensure that there is a public debate if the implementation comes from the state?</p>
Responsibility	<p>In case of mistakes in automatically processable regulation, will it be possible to ascribe responsibility to one organizational unit (or a person within that unit), hence guaranteeing clear ownership and associated responsibility which in turn incentivizes developers to take precautions?</p> <p>Is the division of responsibility between encoding, inputting data, project management, and the resulting output clear?</p> <p>Is the division of tasks clear, or is it part of a complex organizational setup prone to hiding a lack of ownership?</p> <p>Is the hierarchy also well-established when it comes to decision-making? Or is the culture aimed towards group leadership, with groups loosely defined?</p>
Reality	<p>Is it clear to users whether the tool is a simulation or whether it is exactly the same tool that will be used for the official decision-making process?</p> <p>Why can the simulation and the actual decision-making system (not) be the same? Is the justification strong enough?</p> <p>Or is this just a showcase of the public service's inefficiency?</p> <p>Are the messages displayed to users specific enough on when and how simulation can differ from real usage of decision-making systems?</p>
Contestability	<p>Can individuals technically and legally reverse the process by contesting the outcome? How cumbersome is it to appeal to the encoding or the outcome?</p> <p>What cost, if any, to the users does such appeal generate?</p> <p>How can arguments about fairness be brought in during the contestation of decision-making?</p>

NOTE

- 1 In their updated version, they claim that it is the most “streamlined, transparent, user-friendly, and rigorous RNA [Risk and Needs Assessment] available today”. Needless to say that we couldn’t verify this claim but that caution in light of marketing parlance is warranted: <https://www.equivant.com/the-making-of-the-compass-r-core/>

Needed (Public) Debates

DEBATES ON AI HAVE emerged all over the world. One driver of such discussion is a rooted distrust and fear of the risks that new machine learning-based systems may bring (Guitton, Tamò-Larrieux, & Mayer, 2024). While governments around the globe have taken different stances on how to approach this topic, a lot of attention has rested on ensuring that we humans can rely upon and trust the AI-powered systems that surround us. The quest for trustworthy AI and determining ways to ensure that regulation can promote trust in these systems (see Tamò-Larrieux, Guitton, Mayer, & Lutz, 2023) has led to the first regulations. Notable is the AI Act in the European Union, which has received lots of public attention and led to heated public debates especially on the question of how wide its scope should reach (Veale & Borgesius, 2021). Regulators in Europe have chosen to classify AI systems according to the risks that they pose to society and outright ban certain practices. For instance, AI systems that manipulate or deceive individuals' decision-making processes are prohibited. Likewise, AI systems used to predict the risk of an individual to commit a criminal offense, which is based solely on the profiling and the assessment of personality traits of that individual, are prohibited. Yet, this prohibition does not apply (think of defeasible logic described in Chapter 2: Law and Computer Science Interactions) when a human is kept in the loop and objective and verifiable facts are available to confirm the automatic risk assessment. Aside from such prohibitions, the AI Act contains a long list of requirements that AI systems that pose a high-risk need to fulfill. Under the AI Act, different AI systems may be qualified as high-risk, in particular when they are used as a safety component of a product or when they are

integrated into a product and need to undergo a third-party conformity assessment based on another regulation. Obligations that occur to the providers of high-risk AI systems range from documentation and transparency duties to human oversight and risk mitigation measures that must be implemented. What the AI Act reveals is that with respect to AI more generally, a political debate on different scopes and risks emerged and led to a final regulation. While the AI Act does cover different domains such as law enforcement, to date, a clear sense of how the AI Act applies to different legal automation tasks developed by academics, private companies, or public state institutions, does not exist. In fact, a more general debate on how legal processes should be automated (through law as code or with machine learning approaches) has yet to occur. We take this as a starting point to contrast different topics picked up within the course of this book and kickstart a more general debate that we argue is needed within the field. Indeed, the examples elaborated in the previous chapters showcase that different implementations of automatically processable regulation bring with them opportunities for citizens (e.g., understanding their rights to social benefits) and challenges (e.g., biased systems). Within the administration, automation of legal processes can enable more proactive administration, such as, for example, what the DINUM (Direction Interministérielle du Numérique) is promoting in France, where the rights of residents to social benefits should not rely on applications. Yet such efforts also need to take into account the needs and wishes of citizens that they, in the end, should serve. Chapter 4: Challenges and Controversies has notably made clear how real cases have gone wrong and affected individuals. As a consequence of this, politicians, journalists, and, more generally, civil society have asked many questions. While it is regrettable that scandals are needed for such debates to take place, it is still a silver lining if public debates emerge. Due to the pervasiveness of the role of law, its constant evolution, and an unavoidable role that the state plays as a basic guarantor of the law, it is key that such debates involve various stakeholders.

5.1 LAW FOR ALL: IS THERE A MANDATE TO MAKE LAW ACCESSIBLE?

It is one of the greatest anomalies of modern times that the law, which exists as a public guide to conduct, has become such a recondite mystery that is incomprehensible to the public and scarcely intelligible to its own votaries.

Lee Loevinger quoted in (Genesereth, 2015)

One of the oft-quoted principles of law is “ignorance of the law is no excuse” (“ignorantia juris non excusat”). A basic and first requirement for such a principle to be perceived as fair is if the law is accessible to everyone—even before asking about its understandability. However, history shows that only in more recent times, accessibility has become a cornerstone of how the law should be. Starting in Roman times, guardians of the law were called pontifices who “followed oral traditions that captured unwritten customs that were both religious and secular in orientation” (Herzog, 2018, p. 33). The oral tradition meant that the law was not only dependent on pontifices but was also rather inaccessible. Only around 500 BC, with the 12 Tables, laws started becoming publicly accessible in the Roman Republic. This trend of publishing the law continued with the judges in ancient Rome regularly posting “formulas” which included the cases they would agree to head, how the process for specific legal questions would go on, and which arguments the parties could bring about. These publications became known as “edits” and were legally binding, and around 100 AD, they were collected in an “official compilation” (Herzog, 2018, p. 40). As the form of law continued to change throughout the next 1,000 years, a major point impeding accessibility remained: Namely that in Europe, and this until at least the reformation years, only roughly 5%–10% of the population was literate (Harris, 1989). Knowledge of the law in the meanwhile occurred for the many illiterates via more gruesome and physical means. For instance, by publicly displaying tortured and dismembered bodies for all to see and rumor about, lay citizens could get short insights into what certain deeds could lead to and, in this sense, how they were against the norm. The lack of literacy was even leveraged much later during the Middle Ages by kings focusing on written text as a basis for the law, through which kings sought “to control the local normative order” (Herzog, 2018, p. 246), arguably to minimize the impression of arbitrariness.

With the French Revolution came a new understanding of what accessing the law meant. Before, judges and jurists had sufficient power so that their (possibly arbitrary) interpretation of the law could become the de facto law. Following the revolution, the creation and empowerment of legislative chambers meant this role moved upstream to the legislators, and importantly, it moved to a centralized institution. Printing legal texts coming out of the chambers would ensure dissemination; and higher literacy rates throughout Europe could, in theory, help increase accessibility and interpretability. Yet, judges retained (and still do nowadays, although more in common law countries than in civil law ones) the power to shape

the interpretation of the law through case law. With case laws not widely published, accessibility to the law remained hampered.

In light of the difficulty of getting access to the law, it might be relatively fair to think that there are cases where it could be excusable not to know it. Jumping across the ocean to the United States we find that the question did come up with a case that reached the US Supreme Court and which had consequently to rule when ignorance of the law can be an excuse. In 1955, Virginia Lambert was arrested and convicted because Los Angeles, where she lived, had a law requesting convicted felons to stay in the city for more than five days to register with the police (Brooke, 1992). And Lambert had been convicted four years earlier of forgery. Interestingly, Lambert worked for a local attorney, so within the legal professions, but claimed not to know of the requirement, and that the law which referred to “punishable as a felony” was not clear to any layperson what it involved and what not. A local court sentenced her to 3-year probation and a \$250 fine, and an appellate court confirmed the sentencing—at which point it reached the Supreme Court. The Supreme Court argued that she should have been given notice as part of an “essential due process” and defined two criteria when it is excusable not to know the law: (1) the offense is “*malum prohibitum*” (a prohibition due to a statute as opposed to an evil in and of itself, otherwise known as *malum in se*), (2) and that the offense “is purely passive” (Brooke, 1992, p. 289). By doing so, the court reversed Virginia Lambert’s conviction and the case became a hallmark, a “revolutionary opinion” in the words of a legal scholar (Brooke, 1992, p. 289), as it showed that “ignorance of the law” could be a valid argumentation in specific cases, most notably regarding due process considerations in criminal proceedings. Over the next few years, the Supreme Court further refined what it meant to make it clear that it applies only to a handful of cases (e.g., anyone in possession of dangerous material would be expected to think that there is a high chance of regulations existing that they would have to follow).

However, the dictum “ignorance of the law is no excuse” also makes sense as a legal fiction as it is necessary to the functioning of the law to prevent people from evading sanctions for merely not knowing the law (warranted that they could prove it). Such a legal fiction appears fairer though, if access to the law is easy for laypeople. The question becomes then how to interpret the word “accessibility” to the law. In France for instance, the constitutional council published in December 1999 (99–421 DC) a decision that a goal for the law must be, in order to be compliant

with the French constitution, accessible and understandable (“*accessibilité et intelligibilité de la loi*”, the context was regarding the limit of the ability for the government to rule by decrees). Certain commentators interpreted this as a principle of a need to avoid legal ambiguity. As different people, depending on their background and ability, could differ in their skills to being able to understand a text, avoiding ambiguity might be a hard goal to set (Benezech, 2020), also given the prevalence of open-texture in law. Operationalizing accessible and understandable law might involve turning regulations into automatically processable regulations, yet also this is not without its challenges as already discussed in prior chapters.

While from the point of view of citizens, the benefits of having access to the law are obvious, there are domains in which restrictions often occur under the heading of national security. To illustrate this, we can think back to the debates that emerged after the Edward Snowden revelations in 2013, especially the ones surrounding the FISA court—the Foreign Intelligence Surveillance Court. The court reviews demands for wire-tapping of foreign targets, but they meet in secret; their cases are secrets, and hence, how they interpret and apply the law is not accessible to the public (Hogle & Abdo, 2021). There are certain legitimate grounds for this secrecy, but critics voiced two concerns. First, by not knowing the interpretation, it was impossible to know whether certain general bulk wiretapping was in scope and, hence, legal following both a court’s (secret) ruling and the court’s own interpretation of its competence (Hogle & Abdo, 2021). A second concern was that 99% of the submissions appeared to be accepted, making, in appearance, the court a mere rubber-stamping instrument at the hands of the intelligence services (Abramson, 2013). Counter-arguing, FISA mentioned that there is a lot of back-and-forth on each submission with requests for complementary information for instance, and that the high percentage only reflected “final” submissions (Kris, 2018). Since the Snowden scandal, the FISA court has also published yearly statistics on the number of applications, the number granted, modified, denied in part, and denied completely. Different appellate courts since also ruled that bulk data collection was not authorized by Congress and hence illegal (Greenberg, 2015; Satter, 2020).

While in the case of the FISA court, the inaccessibility to the law is on purpose, other times, inaccessibility can be the result of technical hurdles: In the United Kingdom, for instance, access to all case law through centralized systems was for a long time non-existent. Until 2022, there was no systematic publication of court cases, with no centralized system for

all court judgments (Ministry of Justice, 2022). The technical hurdles that had to be overcome to connect all the different courts stemmed from their loosely independent acting, without seeing the necessity to bring handled cases within a single repository (Flood, 2023).

From the point of view of a country's residents, the benefits of better access to the law are almost obvious: Importantly, it reduces the need to consult legal professionals and allows for more independent thinking. The benefits extend overboard to other parts of society. From the legislator's point of view, one major benefit of rendering the law more accessible is that it is more likely to achieve its intended aim. For instance, in the case of any legislation on giving social benefit rights to residents, the laws will have carved out exceptions and rules; legislators would want to see those applied as they voted them to be. Yet, when accessibility to the law is equated with providing citizens with a system that lets them apply the law in their context, like issuing social benefits, we have seen that the picture is a lot more nuanced: With *Mes Aides*, in France, for instance, the ministry faced opposition from social worker employees arguing that their role was to process applications, and not that of a tool. This pointed to different tensions in approaching the role of the state and of the public service. Automatically accessible regulations do not change only the gatekeepers to the law (as much as they did in the past), but they also go to the core of the social contract between the state and its residents. A clear definition of this social contract—whether it includes this view of the state as championing access to legal education as a basic necessity for the good citizenry—should come as a result of a democratic debate, in line with the political system in which this debate would be inscribed.

5.2 TO WHICH EXTENT SHOULD THE STATE STRIVE TO MAKE THE LAW DIGITALLY READY?

Governments have launched initiatives to prepare their legal systems for even more automation: Denmark, for example, has put in place a systematic approach, overhauling its legislative processes. All future legislation will have to undergo a digital impact assessment to chart its effects on existing and digital systems, and its potential for automation (Danish Agency for Digitisation, 2021). The United Kingdom and Switzerland use a different approach, having established regulatory sandboxes to experiment with digitalization. Other countries, such as Finland, adapt existing legislation on a case-by-case basis, e.g., when rules hinder the use of emerging technologies.

Different *terminologies* have come up along these initiatives to prepare the legal system for the increased digitalization of our environment, but two preponderant ones are digitally ready policymaking (European Union, 2024) and digitally ready legislation (Plesner & Justesen, 2022). On a European Union level, for instance, *digitally ready policymaking* refers to the process of formulating digitally ready policies and legislation by considering digital aspects from the start of the policy cycle to ensure that they are ready for the digital age, future-proof, and interoperable. It also encompasses the use of innovative methodologies and tools in the policy design, analysis, and implementation process. The goal is to enable a smooth transition from a policy into its digital implementation. The European Union sets out different components and enablers needed to achieve such digitally ready policies and provides policymakers, service managers, and IT professionals working on policymaking or for the government with short courses and toolboxes to ensure that the bridge between the drafting of legal texts and its implementation is created.

In Denmark, the term *digitally ready, or digital-ready legislation*, emerged, which seems to have inspired many other initiatives, such as the ones of the European Commission mentioned, as well as ones in other European member states. For instance, in Austria, the Austrian Ministry for Digital and Economic Affairs hosted in 2019 the Danish Agency for Digitisation to share learnings from their digitally ready legislation project to gain insights for their Austrian project “Das Digitale Amt” which explores how to identify and remove legal barriers to digitalization. In the Czech Republic, the Ministry of Industry and Trade issued the “Digital Czech Republic” strategy paper that includes the development of a digitally friendly legislation environment that is conducive to digital technology. Yet again similarly, in Germany, the federal government developed a “digital check” to verify if a new law is suitable to be digitally ready. Rules are not only about the exact formulation of the law, but also about the data and data privacy it would involve, IT infrastructure, and how to simplify its communication via visual help. But let’s go back to Denmark. In 2018, the Danish Parliament reached a political agreement that requires all legislation proposed from July 2018 onwards to be digitally ready. The goals of this agreement were to ensure that legislation is written clearly and in simple terms, that it lends itself to be implemented within the administration in a digital form, that it enhances the public services by making them more coherent and efficient, and that it leads to better access, transparency, and trust of citizens to public services. In order to achieve these goals, the

parliament agreed on seven principles used when drafting new regulations that encourage the implementation of digital solutions and new technologies. These seven principles are listed in Table 5.1.

Compliance with the principles is monitored by a Secretariat, established in 2018 by the Agency for Digitisation (which later, in 2022, changed its name to Agency for Digital Government). The Secretariat for Digital Ready Legislation works in close dialogue with the ministries, screens draft legislation, assists ministries on how to develop digital-ready legislation, and provides consultation responses to legislation submitted to the parliament (Agency for Digitisation, 2021). The Secretariat has provided within 2018 and 2020 over 240 responses to specific legislative proposals since it became mandatory for Danish ministries to assess their implementation. These responses look as shown in Box 5.1.

TABLE 5.1 Seven Principles from the Agency for Digitisation (2018)

Principle 1: Simple and clear rules	New legislation has to be clear in its use of terms and has to have unambiguous rules if possible. This makes it more clear and more simple.
Principle 2: Digital communication	All communication between citizens and companies and the public authorities has to be digital, so it has to go through digital means. Every time new legislation makes it mandatory to use digital communication between citizens and the authorities, there always has to be a thought on the people that cannot use digital means of communicating and they have to have a possibility to communicate in another way.
Principle 3: Enable the automated processing of the case	If possible, decisions made by authorities have to be made on objective criteria, which means that the decision by time can be automated.
Principle 4: Coherence across government – homogeneous concepts and reuse of data	Different parts of legislation across different sectors in the public sector use the same terms and refer to the same registers with the same information so the authorities can share data instead of picking up the same data twice from the citizens. This makes a more effective public sector and it makes better service for the citizens and the companies.
Principle 5: Safe and secure data handling	Digital legislation is not all about effective service and administration, it's also about making it transparent and safe for citizens to share their data with the public sector.
Principle 6: Use of public infrastructure	Use the already made public infrastructure for communication, instead of inventing something new.
Principle 7: Prevent fraud and error	Using digital systems to check if data is right and that the information given by citizens and companies are true.

BOX 5.1 VERBATIM EXAMPLES FROM DANISH AGENCY FOR DIGITISATION (2021, PP. 43–44)

IMPACTS ON PEOPLE

“The secretariat notes the ministry’s assessment that the bill does not have significant public implementation impacts. However, the secretariat recommends that, in accordance with the requirement to assess and describe the impact of the bill on citizens rights the ministry may consider explaining in more detail whether people whose application is declared invalid will be informed accordingly. Such information could, for example, make use of Digital Post and serve to ensure that people do not inadvertently miss the opportunity to exercise their right to reapply within the period on the basis of an erroneous perception that their application remains valid”.

DATA PROTECTION

“The secretariat recommends that the ministry may consider elaborating further on the technical measures envisaged to ensure that any personal data, including statements made under the proposed Section X, are processed safely and securely. This is particularly in the context of the Board’s examination of the grounds for the annulment of the declarations made”.

In addition to monitoring the compliance with the principles, impact assessments on the digital consequences of the implementation of new legislative drafts are conducted. These impact assessments look at the impact of new legislations on IT systems, organizations, data protection, and their influence on civil life. The impact assessment which is conducted at an early stage in the pre-legislative work and policymaking process, is guided by different questions that cover these four elements, such as:

- “What are the consequences with regard to existing information systems, is the development of new information systems required, and does this entail any significant risks, for example in relation to when the legislation is supposed to enter into force?”,
- “Is the law and the proposed administration of the law in compliance with data protection legislation?”, and
- “Does the legislation contribute to greater transparency, better accessibility for citizens and businesses, and a more consistent approach? Does it ensure that digitalisation respects the citizen’s rights under national law?” (Agency for Digitisation, 2021).

The impact assessment should be done by an interdisciplinary team with a main focus on experts in law and IT, which leads to an ongoing exchange of expertise.

Two years after the principles and assessments were put into place in Denmark, an evaluation of the concept of digitally ready legislation occurred. In their report, the Agency for Digitisation of Denmark provided an overview of their main findings as well as examples. Their main findings were that the efforts had indeed led to safer, more secure data management regulations and better support of citizens in daily life activities. In addition, they noted an increased awareness and knowledge in the ministries about digitally ready legislation, which also occurred thanks to the dialogues with the Secretariat. They noted that fewer barriers to digitally ready legislation were experienced and that legislation prior to their initiative was overall less digitally ready. Their analysis ended with some forward-looking recommendations: To systematically map out the potential to revise existing legislation according to the principles set for digitally-ready legislation, to strengthen the focus of digitally ready legislation at earlier stages in the political decision-making process, and to target more training activities to policymakers and individuals involved in the policymaking process. With respect to the first recommendation, the European Commission granted additional resources to the Danish Agency for Digitisation in 2022 to develop a method together with the OECD to find and prioritize legislation that needs a “digital review” with the aim to “develop best practice method in the form of a ‘guide’ that can be used to promote digital ready legislation” (Agency for Digital Government, 2022).

Finally, it is important to note that digital-ready legislation has many overlaps with overall efforts within the field of e-government or digital government as well as with rule as code initiatives (see Chapter 3: Automatically Processable Regulation). Digital government or e-government stands more generally for the use of digital technologies within public administration to modernize its operations (e.g., moving from paper-based forms to digital ones). While efforts to make legislation digitally ready also fall within this field, digital government or e-government is broader and does not only concern the process of rulemaking and implementation of such rules. At times digitally ready legislation initiatives also fall under the term *rule as code* (that we introduced in Chapter 3: Automatically Processable Regulation), which seeks to make the rulemaking and application more efficient by side-stepping the current process of having one entity that makes the rules (lawmakers), having entities that

interpret and formalize the rules (lawyers, judges), and having entities that have to adapt to these interpretations (organizations) and often translate them into machine-readable formats for a more efficient adoption (organizations, third parties providing automatically processable regulation).

5.3 HOW TO PROMOTE LEGAL DESIGN THINKING?

Throughout the discussions of how to make the law more digitally ready and processes more ready for automation emerges the theme of legal design. Legal design is a term that combines various *elements* that must be disentangled. Importantly, legal design wants to be *human-centered or user-centered*: Legal design puts the user at the center of the (automated) legal solutions that are being developed and focuses on how to better represent and communicate legal information to the user (Bazzi, 2021). Human-centered design means here that the focus has to rest on understanding the “target audiences’ deep-seated and compelling needs, in order to craft interventions that will improve these people’s experience and deliver them value” (Hagan, 2015). In addition, legal design is about being *proactive and preventive*, meaning actively promoting the understanding of the as well as preventing mistakes arising from a misunderstanding of the law. Thus, automatically processable regulation and legal design overlap on many points.

The reasoning behind the push for legal design as a proactive and preventive approach to law is that common sense and studies have shown that large parts of the population are not able to correctly identify what the law states. Studies have shown that especially in areas of the law that are not salient in one person’s life, citizens do not know what the law states (Pleasence, Balmer, & Denvir, 2015). More generally, research has shown that citizens do not know their rights in various fields of the law: A study looking into 1760 legal problems showed that for over one-third of these problems people did not know at all what their rights were (Denvir, Balmer, & Pleasence, 2013). There are probably multiple reasons for this lack of legal-savviness, and the abundance of legalese jargon and open-texture terms in legal norms does definitively not mitigate the challenge of understanding legal processes and gaining access to one’s rights (see Chapter 3: Automatically Processable Regulation). Legal design attempts to address these challenges by taking a more human-oriented mindset and aiming to reshape legal processes in a manner that empowers individuals.

The earliest attempts within the field now called “legal design” build upon visualization techniques and aim to visualize legal rights in a

more intuitive manner. Still today, this is an important aspect of legal design and multiple initiatives have relied upon such an approach. One example that can be mentioned here is the United Nations Development Programme that worked closely with a Legal Design office in Turkey to publish visualizations that provide a roadmap for women who are victims of violent attacks and who need protection (UNDP, 2023). Visualizations are, of course, a core aspect of legal communication. As seen throughout this book though, they are not the only measures that can be employed to promote access to legal processes. For instance, the Stanford Legal Design Lab, one of the earliest (if not the earliest) formal centers for legal design and kickstarted by Margaret Hagan, contributes to legal design projects that combine visual design elements with technical tools. While those technology-based approaches can range from websites to automatic text message alerts, software code that automatically implements certain regulations can also be used to create more human-centered legal processes. Think, for instance, back at the Rates Rebate Act in New Zealand or the Mes Aides application in France. Both these forms of automatically processable regulation have in common a more citizen-centered application of the law. The governments created websites or calculator applications that can be used by laypeople in order to gain access to the law's benefits. Within the research community, within government bodies, and also in practice, we see this change of mindset that aligns with the vision of legal design to combine different expertise (including those of designers, lawyers, and computer scientists) to craft means to enforce rights, gain benefits, and ensure protection. A case in point is the Laboratory for Numerical Innovation (LINC: Laboratoire d'Innovation Numérique) at the French data protection authority referred to as CNIL (the Commission Nationale de l'Informatique et des Libertés). The LINC enables the CNIL to dive into specific topics of interest (e.g., such as online tracking via cookies) analyze with an interdisciplinary team the effects of different technical artifacts (e.g., cookie consent interfaces), and propose concrete tools to enhance understanding for European citizens.

The projects described by legal design thinking are ones that work with already established legislation, but the concepts can, of course, already be applied prior to a law being enacted. In fact, the movement towards making law digitally ready, as described above, also fits within the broader debates on how to improve the law itself and ensure that law is simple and can be structured digitally. In addition, we see approaches that employ foresight methodologies to create realistic futures and debate with different

stakeholders how, within those future scenarios, current regulation falls short to protect citizens. On this basis, new ideas for new legislation or protection measures can arise. Testing approaches for new regulations or omissions of regulation is a field also referred to as experimental regulation which includes the use of sandboxes.

Experimental regulation has three distinct features: It has a temporary character, it follows a trial-and-error approach, and it is collaborative as it involves different stakeholders (Ranchordas, 2021). To successfully assess the effectiveness of a regulation, such experimental approaches must follow clear objectives to know how to evaluate the success or failure of an experimental regulation. Such measurement criteria will depend on the context in which the experimental regulation applies. For instance, analyzing the impact of an experimental regulation that allows the use of AI to recommend travel vaccination to patients can be done within a rather short period of time. In contrast, to analyze the impact of the system on reducing infection rates overall, a much longer period of time would be needed. In addition, experimental regulation must be applied to a representative sample of individuals to be able to have inclusive findings. Different countries have used experimental regulation to test the effectiveness of proposed regulations in specific sectors (e.g., urban planning, traffic safety, education). In France, the Constitution even explicitly states that “statutes and regulations may contain provisions enacted on an experimental basis for limited purposes and duration” (Art. 37-1).

One approach to experiment with regulation, or rather the lack thereof, is the use of *regulatory sandboxes*. The term comes from computer security studies where an environment is isolated to test new code. The idea is simple: Within the enclave that is the sandbox, new ideas and code can be tested without impacting the rest of the environment. In the legislative context, a sandbox provides the entities that are part of the sandbox with the opportunity to test new ideas, products, and services without the fear of regulatory sanctions. For example, financial technology companies, often referred to as FinTech, that want to provide new products and services may benefit from sandboxes that enable them to forgo large parts of financial regulation. Sandboxes, however, are highly supervised, and the regulatory authorities that allow for the sandbox are constantly kept in the loop to ensure that no negative effects occur to a larger part of the market or population. This supervision is important as sandboxes challenge, in theory, certain legal principles such as legal certainty and equal treatment—making it thus necessary to ensure that through adequate objectives and

guardrails, these principles are adhered to (Ranchordas, 2021). Especially today with the European AI Act and the provisions within that regulation on enabling AI sandboxes, more policy and academic attention will be put on such experimental approaches. The hope of policymakers is that through experimentation, closer cooperation among AI service providers and authorities will arise leading to best practices for AI development and deployments. Importantly AI sandboxes contribute to generating a needed evidence base for regulatory learning. Such an evidence base is key to enhance our understanding of how AI impacts different domains of our society and how regulation can ensure that individual needs and societal concerns are best addressed.

We have seen in this chapter that there are many questions on which the public should be consulted to give their opinion in order for automatically processable regulation to remain within the democratic foundations of our societies. The list of topics mentioned here should by no means be taken as exhaustive. With new scandals coming about, we can also expect other topics to emerge. But we consider the larger issue thematized in this chapter rather central on the role of the state to promote an understanding of the law, be it via automatically processable regulations or via legal design. Similarly, the state plays a role in shaping not only how educated citizens are in law, but also how educated they are on a range of other topics which would play a central role if citizens are to understand the wider context of how automatically processable regulation works: Beyond law, this would include some basics of computer science, of data privacy, and of how even geopolitics impact its development. How to weave in these topics with existing curricula is what we aim to show in the next chapter.

How Education Should Shift

AS DEMONSTRATED IN THIS book with many practical examples as well as theoretical arguments, we are only at the beginning of an era where we expect technology and law, together, to reshape many assumptions that have been a cornerstone of society for many decades, from how decisions are taken, to accountability handled, law shaped, and more. The *responsible* integration of algorithms, analytics, and regulation might yield a future with more efficient, more accessible, and more transparent legal processes, as discussed in this book—and where the challenges and issues we raised have been overcome. Next to the responsible implementation of, and oversight over, automatically processable regulation, another societal aspect stands out as highly relevant: The necessity of educating citizens, young and old, about how automation is changing the regulatory realm that affects them.

Different *types of literacy*, from functional literacy that pertains to reading and writing, to numerical, media, and civic literacy, are required to permit functioning democratic societies (Milner, 2002). In fact, the lacking ability to consume and understand politically relevant publications undermines individuals' possibility to participate in societal and political discourse. Hence, just as reading and writing became foundational skills to communicate with the advent of printing, a similar shift occurred to include in most secondary education curricula the fundamentals of physics, chemistry, or geography as the fields evolved and the broad impact on society to have individuals understand the primary concepts of these fields

became clearer. Given the proliferation of computer-based automation in our today's everyday environments, education systems worldwide have been striving to similarly enrich secondary education with an overview of the principles of computer technology. In the early 2000s, these efforts were concentrated on the use of computer programs, such as word processors and spreadsheet programs. In many countries, we have since observed a (very welcome) shift of this focus to the principles of computing, including algorithmic thinking, first computer programming in high-level programming languages, and—in some secondary schools—even basics of automaton theory. Thus, more students at a young age are exploring the increasingly societally relevant question: What is computable in principle?

The field of computer science studies the nature of computation and its uses, where “[c]omputability refers to the possibility of solving a mathematical problem by means of a computer, which can either be a technological device or a human being” (Edelkamp & Schrödl, 2012). The most common examples are computations performed by computers; the architecture of modern processors determines which computations they can carry out in principle, and this is mapped to facilities provided by high-level programming languages that are usable by human programmers. Thus, computability means that, in theory, there is an algorithm out there that could solve a given problem at hand—an algorithm, in turn, is a finite sequence of well-defined instructions that together realize a computation.

The expected expansion of automation to the automatically processable regulation field further should emphasize curricular contents that focus on automation. In addition, this expansion complicates the educational landscape: Students in *specialized tertiary and further education*—specifically in study programs in the fields of Law, Computer Science, and programs related to Public Governance—will require a firm grasp of the interactions between computer science and regulation. But also, the general population, through *primary and secondary education*, will need to be equipped with a fundamental understanding of these interactions to ensure that they are able to engage in informed dialogues, safeguard their rights, and evaluate the legitimacy of specific innovations—from microtargeting on social media to autonomous driving and its legal implications.

This calls for an extension of education programs on all levels, and the question to answer is where such educational materials are weaved in and what modules get shortened to enable such discussions. These are difficult questions, and we rather posit that many of the elements we describe below can be integrated into current curricula. For instance, discussions on how

we ended up with our digitalized environments can be part of (technology) history courses; discussions on data analytics could be integrated into mathematics courses where they naturally fit as application-oriented statistics. In fact, such changes might soon be required by law: We already today see first signs of this development in strongly related fields. Specifically, the European Union's Data Act (2023/2854) demands that Member States "ensure that the tasks and powers of competent [national] authorities [...] include: promoting data literacy and awareness among users and entities falling within the scope of this Regulation of the rights and obligations under this Regulation" (Article 37 (5) Data Act). Recital 19 of the Data Act further specifies that:

Data literacy refers to the skills, knowledge and understanding that allows users, consumers and businesses, in particular SMEs falling within the scope of this Regulation, to gain awareness of the potential value of the data they generate, produce and share and that they are motivated to offer and provide access to in accordance with relevant legal rules. *Data literacy should go beyond learning about tools and technologies and aim to equip and empower citizens and businesses with the ability to benefit from an inclusive and fair data market.* The spread of data literacy measures and the introduction of appropriate follow-up actions could contribute to improving working conditions and ultimately sustain the consolidation, and innovation path of, the data economy in the Union. Competent authorities should promote tools and adopt measures to advance data literacy among users and entities falling within the scope of this Regulation and an awareness of their rights and obligations thereunder. (emphasis added)

While this final version of the text is weaker compared to earlier proposals of including a dedicated article defining data literacy within the European Parliament's first reading of the Data Act (in Article 3a PA_TA (2023) 0069), it nonetheless shows that policymakers are aware of the changes brought forward by the digitalization of our lives and are pushing towards more comprehensive strategies to educate European citizens. We see the same developments within the AI Act of the EU, where AI literacy is regulated within Article 4 stating (albeit also weaker as the first introduction of AI literacy by the Parliament):

Providers and deployers of AI systems shall take measures to ensure, to their best extent, a sufficient level of AI literacy of their

staff and other persons dealing with the operation and use of AI systems on their behalf, taking into account their technical knowledge, experience, education and training and the context the AI systems are to be used in, and considering the persons or groups of persons on whom the AI systems are to be used.

With the move towards more automatically processable regulation, curricular elements will be required that cover regulation from an automation viewpoint, casting light on the benefits and issues of the automation of law that we discuss in this book. In the following, we propose such elements: We first introduce a strategy that targets the preparation of children in primary and secondary education and introduce concrete curricular elements that illustrate this strategy. We then move to specialized education and propose curricular elements that should be added to study programs in law and public governance, and elements that we recommend being introduced in computer science study programs; finally, we discuss how cross-disciplinary courses might be constructed. In the next and final chapter of the book, we provide selected exercises that target specialized (and/or continuous) education curricula based on this book's content.

6.1 PRIMARY AND SECONDARY EDUCATION

In the following, we propose three crucial building blocks that can serve as pointers for primary and secondary education, starting with understanding how more digitalized environments are driven by data and algorithms, towards a better understanding of what data trails users leave and their privacy implications. Lastly, data analytics courses introducing basic statistics and bias training are needed. Each of these modules already represents an educational challenge by itself, but all of them are required in a curriculum that prepares citizens to understand the implications of progressing legal automation.

6.1.1 Ubiquitous Computers and Ubiquitous Computing

We first go back to the roots and look at the question of how society ended up in a situation where we are discussing the automation of law. The answer, in short: *Ubiquitous computing*—computational processes, and machines that execute them, and which today permeate our everyday lives. And with this ubiquity comes the ability of our everyday environments to sense, transmit the sensed data, and calculate using this data. In fact, the very word computer, as the noun of “to compute”, means “the one who calculates”.

While the term originally, and for several hundred years, referred to people, today, when we say *computer*, we typically refer to a programmable electronic device that performs arithmetic or logic computations. We are continuing a journey of computing that started with mainframe computers, access to which was time-shared among dozens or hundreds of operators, and continued to the personal computer and to today's situation where one person has access to several computers—in fact, several hundred computers if we consider that a smartphone today contains dozens to hundreds of individual computing units (i.e., Central Processing Units; CPUs and Graphics Processing Units; GPUs), and that we carry an embedded computer with every banking card. Computers are hence not only disappearing physically—through miniaturization, low-power operation, and wireless communication—but they also disappear mentally: We do not perceive computing as such anymore, but as an essential, implicit component of the objects we use, from adaptive lighting systems to driving assistance. This phenomenon is similar to the mental disappearance of the electric power system, but its reach is broader since the distribution of electric power is still bound to wires in today's everyday environments while ubiquitous computers operate wirelessly, thus avoiding conspicuous cabling.

Looking back at the history of computers, we see that humans had a fascination with computers for a long time. From early computing mechanisms, such as the Antikythera Mechanism (possibly to predict stellar constellations), increasingly complex mechanical calculators emerged (such as Babbage's proposed Analytical Engine in the year 1837). Important contributions to today's computing came from mathematics, such as Leibniz' introduction of binary arithmetic (in the year 1689) where arithmetic and logic functions can be mapped directly: In binary, arithmetic addition corresponds to a logic either/or, i.e., an exclusive OR-function (" $1 + 0 = 1$ " while " $1 + 1 = 0$ ", with overflow), and multiplication corresponds to a logic AND-function (" $1 \times 0 = 0$ ", " $1 \times 1 = 1$ "). Remember from Chapter 1: Automation of Law that these Boolean operations (OR, AND, etc.) were systematized by George Boole (in the year 1847): Their inputs and results are either true or false (accordingly 1 or 0). Finally, Shannon (in the year 1937) integrated electronic circuits and binary code—the idea that electrical (or, at the time, electromechanical such as in Zuse's Z3 computer) switches could be used to solve all logical problems is the unifying concept behind all electronic digital computers today. Further important developments towards our modern computers were made during World War II (WWII) in the United Kingdom, for instance,

to crack the Enigma machine, which is a cipher device that was developed to protect communication in the German military during WWII. A prominent role in this operation was played by Alan Turing who, before joining the allied naval cryptanalysis division, was working on mathematical models of computation—Turing machines—that describe machines that operate on symbols on an infinite strip of memory tape according to a (finite) table of rules. These rules prescribe the machine's next action (replace the symbol on memory; move left in memory; move right in memory; halt) while considering its current state and the currently read symbol on the memory tape (see Box 6.1). A Turing machine hence provides an abstract model of a computer (or of computation in general) and most programming languages that exist today can simulate a Turing machine (we say that such languages are “Turing complete”); hence, these languages can be used to express all task that are (in principle) accomplishable by computers. The ability to argue about automation in such abstract terms has been crucial for many basic insights into computer science; one of these important results is that it is not decidable in general whether a given algorithm will halt or continue to run forever (the *Entscheidungsproblem*) and that it is, hence, not decidable whether a given computer program will eventually produce a given output, or not. This represents the foundation to think about what we described in Chapter 2: Law and Computer Science Interactions and why automatically processable regulation is possible in the first place; it also provides the foundation for symbolic and sub-symbolic AI: Remember that GPT (Generative Pretrained Transformers) models run on Turing machines.

BOX 6.1 TURING MACHINE: WHAT DOES IT DO?

Think of it as a tape that can have values 1 or 0 or be blank
What you need:

- You need to tell the machine what state it is in
- You need rules associated with the state
- You need a start and an end



Let's say the goal is to report back if there is an even number of 1 in this tape. The tape ends with a 0 for simplicity so that we know this is the end. If we had to do it, we would just count and cross the pairs and report back, but for a machine, you need to go step by step with the rules:

- Let's say we start in the state even (just assume there was a long list of numbers prior): We say: If you see a 1, then set the state to ODD and move one to the right. If the next was now a 0, the state would be odd, so the answer of the machine needs to be 0.
- But it is a 1 so we move back the state to Even. If now the next was a zero, it would indicate that there is an even state.

You can run this as long as you want, you can change the actions, etc., but the premise is the same, you set a state and a rule and move left or right depending on what you want the machine to do.

"Turing-complete"

Read a string of 1s ending with a 0 and tell me if there even number 1s? (true = 1; false = 0)



Move one
right or left

States and Rules
States: Even or
Odd | Start and
Halt
Rules: (see table)

State	Symbol	Effect
Even (start)	1	Set state Odd Move head to the right
Even	0	Write a 1 Set stage Halt
Odd	1	Set state Even Move head right
Odd	0	Write a 0 Set state Halt

Understanding this and similar historical contexts help to situate the different interplay of drivers that shape our digital environments, from technical foundations to social and economic factors that propelled certain developments (e.g., the quest for data, the proliferation of big platforms). In the context of a school curriculum and targeting the preparation of young learners for a world in which even the law is automated, we argue that the basics of automation should, hence, have a place in primary school. Already today, schools take up automation principles in the form of curricular elements that focus on *computational thinking* (or *algorithmic thinking*), where a wide range of teaching approaches and materials exist: These range from the translation of cooking recipes into formal algorithms with basic *operations* (e.g., `add(ingredient)`) as well as *do-while loops* (e.g., for mixing) to remote-control games where a (blindfolded) child receives formal instructions from the group (e.g., “2 steps back”; “turn around”; “walk straight until you hit the wall”). The introduction of computational thinking elements develops children’s structured problem-solving

abilities and logical reasoning, and might even also highlight certain limits of polysemical words (or, said differently, open-texture) with children playfully and knowingly trying to deviate from the intended goal.

Further ideas about automation (and the modeling of automatons) can be introduced already early in a school curriculum as well. For instance, while it might not be obvious (due to its abstract nature) how children could be taught about the Church-Turing Thesis (a central result in computer science), we argue that such concepts can be simplified and transferred to align with a child's level of understanding (see Box 6.2). We also argue that this is valuable, since it enables children to understand—on a basic level—the technology that surrounds them, ideally leading to an emancipated state where they are more than just passive users of technology.

BOX 6.2 INTRODUCING THE CHURCH-TURING THESIS IN PRIMARY SCHOOL

We start by discussing problems that are readily understood by young learners, like simple math problems or puzzles. We explain that, just like they can already solve these problems in a structured way (e.g., by starting with the edge pieces of a puzzle), engineers and scientists create structured methods to solve much more complex problems. We introduce the idea that machines might help solve problems and use examples (e.g., calculators or robots). We discuss how these machines follow specific sets of instructions (or programs) to perform tasks and solve problems and link this to any curricular element that the children have already seen about computational thinking and computational problem-solving. Next, we introduce computers as a type of machine that can be programmed to carry out a wide variety of tasks, and we introduce that anything that can be solved by a step-by-step method (algorithm), also can be solved by a computer; this represents a simplified version of the Church-Turing thesis (a note that not everything can be solved by a step-by-step method should also be made—as an example, the problem of enumerating all Real numbers starting from 0.0 can be used). As practical examples and activities, the children might follow a simple set of instructions to create a solution, which permits subsequent elements that explore what types of problems are not solvable in this way. Together, this permits an approach to understand what types of tasks can be solved algorithmically and by computers.

6.1.2 Connected Computers

Technological changes and automatically processable regulation do not come separated from other large socio-political changes underway in

society. Concerns about sustainability, as much as high geopolitical tensions around AI have already shown their mark: Large language models consume a large amount of energy and hence could contribute to the growing problem of climate change depending on how this energy is generated (Pitron, 2021; 2023).

Yet another sociopolitical challenge in which automatically processable regulation is embedded concerns geopolitical tensions. The United States of America, China, as much as the European Union have framed the hardware and software technologies around AI as a “race” which they need to win, effectively pitting themselves against everyone else. Accordingly, they have engaged in actions distorting the free market, from restricting via trade sanctions access to companies in their country or the sales of component towards other countries, heavily subsidizing the industry or their own local champions, and turning up the rhetoric on protectionism in general (Miller, 2022). The effects to be expected are less cooperation (even in other non-related policy areas), a more volatile environment, and a contribution to tensions all certainly not conducive to innovation and the spreading of change.

The transition to green energy, as much as trade restrictions to “win” the technology race, are complex enough that they offer different ways to approach them and to teach about them. And there is no doubt that both will have an impact on the degree of uptake in the field of automatically processable regulation. It would also be naïve to pretend that the current dominating market power of technology firms could not extend to automatically processable regulation. Depending on your vision of the world, taken all together—climate change, geopolitics, and antitrust—will represent either an optimistic or a pessimistic way to look at changes coming alongside automatically processable regulation. It is not the place to argue either way—we merely want to bring out here two related points: First, that the backdrop for technical development has a rich history that can make for interesting teaching material; and second, that education on automatically processable regulation can be embedded in many other existing school subjects, such as history.

Next to the creation of computer systems themselves, another important ingredient to the proliferation of computers and computing in our today’s everyday experience has been the development of the Internet, i.e., mankind’s steps towards a worldwide computer network. Notably, this was set against an interesting political backdrop, the Cold War and its nuclear threat—as much as today’s developments are similarly *not* taking place in a political vacuum.

One way to start the story of the creation of the Internet is to go back to the year 1960, when the United States Department of Defense wanted to create a resistant communication network, meaning a survivable communication network, as they feared that any first nuclear strike on the country would destroy means of communication and probably cause collateral damage (Brand, 2001). This led to the suggestion of a novel networking paradigm where data would be divided into packets that would find their own routes to the destination—and if one pathway was blocked (or destroyed), the packets would themselves select a different route. The researcher Paul Baran started working on such a “distributed adaptive message block switching” communication network in 1960 at the Rand Corporation, a research entity largely funded by the United States Air Force that gave researchers a lot of flexibility in choosing research topics (Naughton, 2000). In 1964, he published the results of his research in a book entitled “On Distributed Communications”, making public 4 years of state-funded research linked with the US nuclear strategy. His design included digital switches to comply with the Department of Defense’s requirement for large bandwidth. However, experts at AT&T, which at the time held the monopoly on telephone lines in the United States and would have had to implement the project, were dubious about digital technology. The project was, therefore, dropped and not implemented.

In the United Kingdom and the same year that Paul Baran published his results, the Labour Party won the elections after promising to act in order to bridge the “technology gap” that many Britons felt they were falling into (Abbate, 2000). The British Government gave funding to the National Physical Laboratory to develop computing. Time-sharing, where scientists or businesses could use computers and pay in function of the Central Processing Unit (CPU) time they used, was an emerging business at the time. However, a major obstacle was that workers had to physically come to the location of the computing mainframes to deliver the data to be processed and the programs to process it. Donald W. Davies thought of the potential of developing a data communication link between time-sharing computers to reduce their unutilized processing time and to reduce the need for scientists to physically travel. Unaware of Paul Baran’s work, Donald W. Davies started working on a packet switching network that he presented in 1966 and started implementing at the National Physical Laboratory under the name “Mark I” in 1967. That year, the United States Advanced Research Projects Agency (ARPA) was also working on the same issue but had not heard of prior research on the

subject. They organized a conference at the end of 1967 to expose their problem. Roger Scantlebury, from the National Physical Laboratory, who was in the audience, approached them after their presentation and told ARPA's representatives about Paul Baran's and their own work on decentralized networks; they were applying for funding at the time to deploy their network. However, the National Physical Laboratory did not obtain funding to deploy its network nationwide, and the rapid implementation of ARPA's network (known as ARPANET) that followed in 1969, now with knowledge of Paul Baran's earlier work, quickly eclipsed the work from the National Physical Laboratory (Abbate, 2000).

The ARPANET followed the requirements of the United States Department of Defense. Apart from being decentralized, these included having a reliable standard, "especially in the presence of communication unreliability", as well as providing "availability in the presence of congestion" (DeNardis, 2007, p. 682)—this requirement is still visible in today's Internet and in our everyday use experience: While, under load, circuit-switched networks become "full" (i.e., unavailable), packet-switched networks like the Internet become "slow" (but remain available). In 1969, a computer, referred to as a *host*, was in charge of creating a packet and was linked directly to another computer charged with routing the packets (Naughton, 2000). This mini-computer—in today's terms, a hybrid of a network interface and a router—was known as the *Interface Message Processor* (IMP). The packet formed by the host contained the destination host address and the link identification number that was entered manually into each host. Each IMP held a list containing all other IMPs' addresses and the number of the link they were connected at. Adding a new host with its IMP was, therefore, cumbersome, as all lists needed to be updated with the new host addresses, and hindered the scalability of the network. On the other hand, users had the capacity to check information about the device (and potentially its user) they were communicating with. On top of the IMP, the Network Control Program was introduced in the year 1970, to allow hosts to communicate with each other via a socket number in order to run concurrent processes. Traffic was still directed via the IMP, but this was now invisible to the host which was just concerned with the higher abstraction layer at the Network Control Program. In July 1978, a new standard then separated the addressing protocol (the *Inter-network Protocol*, *IP*) and the communication protocol (the *Transmission Control Protocol*, *TCP*), making the network structure more scalable, and removing the need for all hosts to keep a local list of all other hosts on the network. This marks the origin of the today-ubiquitous IP address.

IP addresses contained (and still contain today) network identifiers that allow exchanges between *different* networks via the Gateway-to-Gateway Protocol (GGP). This change permitted different networks to communicate with each other and was required for the following proliferation of the ARPANET, including outside of the United States. The first expansion of the network outside of the United States occurred in 1973 with a connection to Norway's NORSAR, an agency still existing in charge of monitoring seismic data, and at the time, the US Department of Defense was interested in investigating seismic data in the context of nuclear events (Lukasik, 2010). Back then, the allocation of addresses was in the hands of a few individuals: Jon Postel, a researcher at the Information Sciences Institute at the University of Southern California, started allocating addresses on the ARPANET in 1972 and continued with IP addresses in 1978. Jon Postel, along with another pioneer of the ARPANET, Vinton Cerf, requested in March 1972 that all entities submit their socket number to discard unused services but also to help find resources. The listing was first published in March 1973, with the *ARPANET News*. To be able to access the ARPANET, an entity had to have a research contract with ARPA and afford between \$55,000 and \$107,000 in hardware and software. Due to the high costs involved, the Information Sciences Institute could, therefore, manage the list in this centralized way for the time being. However, due to the funders being the United States Department of Defense, this was not without its challenges. The involvement of universities from the very start of ARPANET came at a time when students and university staff were reluctant to be associated with military affairs, as the involvement of the United States in the Vietnam war was widely controversial. ARPANET's managers had, therefore, to "de-politicize" the research project when addressing universities, but had to emphasize its contribution to military affairs when addressing the Congress. On 28 February 1990, ARPANET was formally decommissioned, but the allocation of addresses remained within the supervision of the Department of Defense until 2000. Still, in 1990, Jon Postel created the *Internet Assigned Numbers Authority* (IANA) under a contract with the Department of Defense and with the University of Southern California. IANA allocated specific ranges of addresses to networks that assigned them to smaller and sub-networks. In 1998, the *Internet Corporation for Assigned Names and Numbers* (ICANN) was created under the Department of Commerce, although not without resistance from various international stakeholders, and in 2000, IANA became part of it. Hence, the Department of Defense, and especially Jon Postel, kept an

almost monopolist centralized capacity over the decentralized network for around 30 years.

Following the development of packet switching, the TCP and IP protocols, and the institutions to govern the Internet, another central step towards today's connected society was the creation of the World Wide Web. Many, albeit wrongly, use "Web" as a synonym for "Internet" (and vice versa). However, while the Internet is the above-introduced system of connected host computers, the Web is an information system that is built on top of that infrastructure (an example of another system that is built on top of the Internet is e-mail). In 1989, Sir Tim Berners-Lee created the first Web server, Web browser, and also the today-still prevalent way of representing Web resources (e.g., websites) for human users—the Hypertext Markup Language (HTML)—we can use this infrastructure to access information that is stored on Internet-connected Web servers. And, as indicated by the term "hypertext", these resources are not merely linear collections of information (like a book that is read page-by-page) but are rather interlinked using hyperlinks. This approach proved extremely successful in making distributed information readily accessible to human users (and machines as well) and marked the beginning of the widespread use of computer networks by everyday users in the 1990s. Thanks to the decreased cost of personal computers and the increased speed of Internet access, thanks to broadband Internet, which permitted transferring more and more interesting resources, including pictures and videos, more individuals started to use the Web for daily work and entertainment. In the year 2024, the International Telecommunication Union estimated that around two thirds of the population—i.e., around 5.5 billion people—have access to the Internet and the variety of applications it supports.

Looking at these two aspects—the development of computers and networks—already provides rich opportunities for different perspectives, either from an algorithmic perspective, historical, or political one. They all help situate generally ubiquitous computing better in its overall context and, hence, more specifically, automatically processable regulation as well. A further building block and consequence of ubiquitous computing and automatically processable regulation has been the large amount of data required, produced, and flowing, creating a host of privacy-related issues that nearly everyone can easily relate to. Ubiquitous computing, boosted by the data hunger of modern machine learning systems, hence gave rise to a broad societal challenge that is highly relevant for everyone today: Privacy implications of data trails online.

6.1.3 Data Trails and Their Privacy Implications

Today, data is being tracked whether we are online or offline. Connected hardware sensors (think surveillance cameras or ubiquitous microphones) collect more and more data in the physical space, and in online environments, behavior tracking has become widespread today. Behavior tracking refers to the collection, storage, and analysis of data about users' actions on webpages or in other online platforms such as mobile apps. The actions include pages visited, links clicked, time spent on a page, and products viewed or purchased; they are sometimes referred to as *click trails*. Click trails are marks we leave while we click ourselves through the World Wide Web or similar virtual environments. The purposes of online behavior tracking are manifold and can have little to many implications in the real world. For instance, companies use tracking to understand how individuals navigate websites to ensure a smooth user interface. This branch of research is often part of user experience (UX) research, which is a central component of creating user-centered products and services (Hassenzahl, 2013). While different techniques can be employed, we often encounter *A/B testing* of specific features of a website; this refers to the display of different variants of a feature to different users, to find out which variant maximizes an objective function (e.g., high click-through rates or low delays). A real-life example of A/B testing—and one that brought awareness on the topic—came in 2014 with revelations that Facebook had been hiding certain emotional words from the feed of nearly 700,000 people and measuring which effect this had on their “likes” and statuses (Gibbs, 2014). The results were even published in an article in the prestigious Proceedings of the National Academy of Sciences, with the researchers concluding that “emotional states can be transferred to others via emotional contagion, leading people to experience the same emotions without their awareness” (Kramer, Guillory, & Hancock, 2014). In this very case, an uproar ensued with the journal publishing an “editorial expression of concern” in which they wrote that “[q]uestions have been raised about the principles of informed consent and opportunity to opt out in connection with the research in this paper” (PNAS, 2014).

While such practices seem at their core not that invasive or impactful for individuals navigating the Web, similar practices are employed to deliver more *personalized content*. Personalized content may range from a music streaming service suggesting songs based on what a user has listened to in the past, to news or video outlets suggesting content an individual user might like more. Most of us will be familiar with such suggestions

and the rabbit-hole phenomenon (Roose, 2020), a term that indicates the mechanisms at play utilized by large online platforms like YouTube to continuously suggest personalized content that is attractive to the user and ensures a constant stream of entertainment that keeps the user engaged, and consuming. In fact, and maybe paradoxically, Netflix who also is in the market of personalizing recommendations, produced a documentary on the subject matter called *The Social Dilemma* which can be used in class to highlight the problems of recommender systems, yet must be contextualized as the documentary has also been heavily criticized for its sensationalism (Newnham, 2021). Combined with the echo chamber or filter bubble effect (Pariser, 2011), meaning that users do not realize that the content that is shown to them is not representative but rather originates from their assignment to a specific (maybe very small) bubble, such personalized content can create severe disconnections among what an individual believes and what more general parts of the population believe. This disconnect may, in times of crisis like the COVID-19 pandemic, have a severe impact on how society reacts to scientific information (Salvi et al., 2021). It is thus not surprising that large online platforms have come under media attention for such practices and have even been sued for being responsible for the personalized content that they provide. In a Supreme Court case in the United States of America, *Gonzales vs. Google* (Gonzalez v. Google LLC, 598 U.S. 617 (2023)), the court was asked whether online platforms should be held responsible for the personalized content they provide, especially for content that is terrorism-related. The facts of the case at hand were tragic: A family, the Gonzales, lost their daughter in a terrorist attack that occurred in Paris in 2015. The shooters were terrorists who had been indoctrinated, among others also on YouTube by watching content that was promoted by the Islamic State or its supporters. The question at hand was whether Google, the parent company of YouTube, was responsible for its recommendation system that further recommended terrorist content. In the end, the Supreme Court ruled in favor of the online platforms, i.e., them not being responsible for the recommendation system in place, and thus upholding a provision that is known as the “26 words that created the Internet” (Young, 2023), namely Section 230 of the United States Communication Decency Act. The 26 words of this provision are: “No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider”. The idea behind this rule is that we need to distinguish between the distributor (who is passive and is not liable for the content it

distributes) and the publisher (who is active and exercises editorial control over the content it publishes). A platform is seen as a distributor; even if they are screening the materials for obscene, illegal, terror-related, and objectionable content, this act of screening the content does not make them publishers. Legal provisions (known as Good Samaritan provisions) protect the good faith removal or moderation of third-party material that a platform provider deems “obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected” (see Section 230(c)(2)). A similar provision can also be found in the European Union, under Article 7 of the newly enacted Digital Services Act.

To even have the ability to implement such recommender systems, we need data and tools able to track behaviors online. In fact, a variety of technical means are available for such behavior tracking and citizens should have basic knowledge of these methods and of how they can be tracked using these techniques, which we propose should form an essential element already in primary school curricula. The most prominent behavior-tracking approaches today—that we propose to include in educational programs about online privacy—include the tracking of IP addresses, Web beacons, cookies, fingerprinting, and session recording tools; and we believe that these can be readily explained to school children by referring to analogs that support the sensemaking of this audience.

Since every device that is connected to the Internet receives a distinct *IP address* that is used whenever the device accesses online content, this address can be used to track visitors to websites. Through geolocation services, IP addresses can be used to find a user’s location—depending on the setup, down to the country, community, or (rarely) street address. This is commonly used by websites to detect actions like multiple voting and can also influence the display of local currency prices, product availability, special offers, or even block access from certain regions. When the same monitoring infrastructure is shared, visitors’ IP addresses can also be tracked across sites. This can be explained to children through its similarity to offline monitoring systems that track individuals while they walk on the street, visit different shops, and talk to different people, and can, in this way, be prepared for young audiences. As prize-winning teaching material on privacy available under www.learnprivacy.ch put it in one of their lessons: If you do something (consume information, buy a product, access restricted content) on the Internet, then this is more similar to doing this very action on the town’s market square than to doing it in your

living room. Individuals may bypass such restrictions by using virtual private networks (VPNs), which has become a widespread way to circumvent IP-based blocking or redirection by (state) entities.

Another tool that especially many European Web users are superficially familiar with—due to the mandatory banners and consent notices—are *HTTP Cookies*. Cookies are files that are stored on a user's computer by a website and are recalled by subsequent visits to this site or to other sites; this may be done by the entity that controls the visited site or by third parties that supply content (e.g., advertisements) to a site. Cookies, hence, represent a means of tracking behavior across sessions—from a user experience perspective, this permits pages to retrieve elements of the previous browsing state of the user (e.g., the last time a user visited a website or the state of their shopping cart) and thereby allows to greatly enhance the browsing experience. However, this very feature, from a privacy perspective, amounts to the tracking of user behavior within and across pages and the creation of (cross-page) user profiles.

A technique that is less visible than cookies is *Web beacons*. These refer to content that is embedded in a webpage or an email and that is loaded, typically from a third-party server, when the page or email is accessed by a user. The content that is loaded may, for instance, be a tiny picture—possibly only a single pixel that cannot be visually detected by users, e.g., because it has the same color as the background of a page. Beyond such rendered (but invisible) content, any element that a website loads may be used as a beacon, including embedded scripts, styling information, or fonts that are loaded from a server. Since this loading action can be detected by the server that serves the beacon, and since the user's IP address will hence be shared with that server, Web beacons are a prominent technique to track online behavior that is less visible than tracking based on cookies. Again, through simple analogies, the core ideas and implications of cookies and Web beacons can be readily communicated to primary school children (see Box 6.3).

BOX 6.3 INTRODUCING COOKIES AND WEB BEACONS IN PRIMARY SCHOOL

Imagine you go to a bakery to buy a breakfast croissant, and they give you a special sticker with your name on it. Every time you go back to that bakery, you show them your sticker, and they remember your favorite croissant and give it to you. To even better know what you like, the bakery has tiny,

invisible goblins that watch you when you come in; they keep watching you, and report to the baker every time you look at a certain cookie or stand by a particular shelf. In this way, the baker finds out a lot of information about you, not only information that you tell them—like that you love chocolate cookies—but also information that you would have liked to keep to yourself, for instance, that you would really like to try out the bakery’s new type of bread that everyone else in your class hates but you don’t dare buying.

When you visit a website on the Web, this site similarly gives your computer a tiny file called a “cookie” that is similar to the bakery’s sticker—and every time you visit that website, or other websites that belong to the same organization, they remember who you are and what you like, like your favorite games or where you left off in a story. These cookies, therefore, collect information about what you do online. And the bakery’s goblins also exist on the Web: They are tiny, invisible pictures or codes on a website that tell the website owner that you’ve looked at a certain page—these are called “Web beacons”, and they are like tiny spies that report what you’re doing.

Web beacons and cookies help website owners understand what you like to see on their site. But this means that someone is keeping track of your actions, which can feel like being watched. So, remember that when you are online, you are typically leaving a trail of crumbs about wherever you go and what you do there, and others can see and use that trail. Therefore, understand that your online activities can be seen and shared by others. Tell your parents and teachers about what you’ve learned about cookies and Web beacons. And, if you are unsure about a page, ask someone who knows better than you do and discuss with them whether it is appropriate that you are tracked.

A less obtrusive method of online tracking involves the creation of a *fingerprint* of a user’s device and device configuration: When browsers access websites, they (for technical reasons) share a variety of information about their own configuration and the configuration of the host device. This includes the specific operating system and browser version as well as information about the locale of a user, their system language and keyboard layout, and even the installed fonts. Together, this information permits creating a fingerprint of a browser and, by extension, often of a user¹ that can be used to create user profiles even in the absence of HTTP cookies or Web beacons. Finally, websites may use sophisticated *session recording tools* that capture not only your site-to-site browsing behavior but also track user mouse movements as well as click and scroll events, thereby permitting detailed analysis of browsing behavior. Originally introduced to

optimize the user browsing experience, for instance, to understand what design elements reduce the time needed by users to identify and click a button, session recording tools are today used for user profiling and tracking, with the aim to adapt website appearance and content to individuals.

The deployment of all these techniques has today become trivial for website owners: Data aggregators and data brokers provide tools for the embedding of user profiling software on websites, and a wide range of session recording tools is available for simple embedding as well. Data aggregators also combine several methods, even across websites that do not share tracking mandates, thereby creating comprehensive profiles of a user's online browsing behavior.

Of course, the law counters such developments. Especially data protection and privacy regulations set limits to data collection and impose obligations to data controllers, i.e., the entities that determine the means and purposes of the data collection. Understanding the scope of such regulations, the tools that it provides for individuals (and specifically for minors), and its limitations, are an important component of primary and secondary education. In particular, data protection and privacy regulations that have importantly shaped the global debate, such as the European General Data Protection Regulation (GDPR) or the Californian consumer privacy laws which include the California Consumer Privacy Act (CCPA) and California Privacy Rights Act (CPRA) should be mentioned. These regulations establish a baseline of how data processing should occur and what rules entities processing data must adhere to. Importantly, these regulations have set out rights that were created to empower individuals with respect to data controllers. Key rights under the GDPR are the right to access, modify, and delete data that is being processed about oneself, the right to data portability (i.e., porting data from one service provider to another), and the right not to be subject to automated decision-making. These rights act as safeguards, giving individuals some sense of control over their personal information and the ability to navigate the intricate web of data usage. Education on these tools becomes an essential aspect of both primary and secondary education, shaping individuals who are not only proficient in using technology but also well-versed in exercising their rights to data privacy. Yet these rights also come with limitations. A major one is that it relies on individuals to take up action (Hagendorff, 2018). To do so, not only does one need resources, but also know-how about what is being tracked and why taking up action matters. The matter of the fact is, however, that today many feel resigned from taking action, a phenomenon

that has been referred to as privacy cynicism (Hoffmann, Lutz, & Ranzini, 2016). Privacy cynicism provides a potential partial explanation of the *privacy paradox*, which describes the consistent high self-reported privacy concerns of online users together with low actual privacy protection behavior: Users resign in the face of seemingly overwhelming privacy threats. To change this, regulation, education, and technical tools must work together to provide a greater sense of empowerment.

6.1.4 Data Analytics: Understanding Statistics and Biases

For instance (but not limited to) the tools that we introduced above, in today's age of big data, vast amounts of (personal) information are automatically collected and analyzed; this collection forms a basis for decision-making and data-driven profiling: The automatic analysis of data to identify patterns or tendencies among groups of individual. It is, hence, becoming increasingly important for citizens to have a *basic understanding of statistics and bias*. Data-driven profiling can be an influential tool for various sectors, from business to governance. However, when misinterpreted or misused, these insights can perpetuate stereotypes, create discriminatory policies, or foster a misguided public perception. Statistical biases may enter at various stages: During the initial collection of data (e.g., label and selection bias), the design of algorithms that analyze the data (e.g., sampling and sorting bias), and the interpretation and usage of results (e.g., evaluation bias). In fact, with increased focus on big data processing and AI techniques that rely on data-driven machine learning, we have seen an increase in research on statistical biases as also showcased by increased demand for conferences on the subject matter of fairness (e.g., the ACM Fairness, Accountability, and Transparency (FAccT) conference). Some of the most prominent unintended bias challenges have been label bias, selection bias, demographic and population bias, sampling bias, overfitting and underfitting, evaluation bias, and user-interaction bias as documented in Table 6.1.

These biases sneak into the design of automatically processable regulation and decision-making systems. Many examples that have showcased discriminatory and unfair decisions by computers exist, and many are likely to stay hidden. Throughout this book, different examples have been mentioned, such as in France and in the Netherlands. (Chapter 4: Challenges and Controversies). Other examples could be mentioned, such as infrastructures that were put in place (often in a rush, see Newlands et al., 2021) during the COVID-19 pandemic. In 2020/2021 many companies

TABLE 6.1 Biases to be aware of

Bias	Description
Label bias	Data used for data analytics might be labeled wrongly or poorly. This will, in turn, impact the correctness of the attributions.
Selection bias	Data in a dataset might be selected in a way that does not reflect accurately the population.
Demographic and population bias	Data used has too many data points on one demographic group (e.g., male), leading to an overrepresentation of that group compared to another group (e.g., female).
Sampling bias	Data does not adequately represent the entire population of interest.
Overfitting and underfitting	This happens when a model learns the training data too well, including noise and irrelevant patterns, making it perform poorly on new, unseen data. When the opposite is the case, i.e. the model is too simplistic and cannot capture the underlying structure of the data.
Evaluation bias	This bias occurs when the criteria used to evaluate a model's performance are flawed or incomplete.
User-interaction bias	When humans are involved with the system, this can lead to affecting a model's behavior. E.g., the personalized content mentioned above may skew a recommender system to continue to promote similar content yet the content was chosen not because of its match for the user but only because of its prominent display.

and schools were forced to put in place new mechanisms to allow business as usual. In the United Kingdom, for instance, where grades during a final exam determine a student's ability to go to certain universities, the Office of Qualifications and Examinations Regulation (Ofqual) put in place a system that calculated a student's final score based off (weighted) grades obtained during the year, with the weights reflecting a school's standing. However, it turned out that the system in place put too much weight on the school and its geographical location, leading to many individual students in schools with lower averages (and often located in less wealthy neighborhoods) being assigned a grade which they felt was too low and not representative of their own potential. This led to massive protests with banners stating "ditch the algorithm" (in much less nice albeit straightforward language). What this shows is how demographic and population biases are often ingrained into historical datasets. In fact, many researchers have pointed to this problem: Data is always from the past so a machine can only learn to infer something that lies in the past (Wachter, Mittelstadt, & Russell, 2018). This is important to remember when more data-driven approaches are taken. The plurality of possible biases makes it crucial that children, as well as adults, understand how and when such biases can occur. Most importantly, citizens must learn— from an early age

(remember that individuals in many countries are allowed to vote from the age of 16)—to understand the difference between population-level statistics and individual-level application of statistical findings: While profiling may be a valid technique when considering statistics on population-level, it has a high likelihood to lead to unfair bias when applied to an individual. For the informed citizen, recognizing the potential pitfalls of relying too heavily on data without questioning its origins, methods, and interpretation is essential. This becomes even more important when legal decisions are based on such data, thereby importing and perpetuating biases. As we navigate this data-saturated era, a critical mindset—paired with a foundational understanding of statistics and bias—will ensure that we promote fairness, inclusivity, and accuracy, while challenging and refining the narratives presented to us both within and outside of the automation of legal processes.

While a detailed mathematical understanding of statistics (and, more broadly, machine learning) builds on top of secondary school mathematics, the foundations of the introduced issues with statistical bias and the principled problem of the application of statistical results to individuals, are accessible already in primary school: Children already have experience with (unfair) stereotypes, and this can be readily generalized (see Box 6.4).

6.1.5 Pioneer?

There are already concrete suggestions for schools seeking to integrate a few of the suggestions made above about the different aspects of digital lives. For example, Harvard’s Berkman Klein Center’s materials on *Safety, Privacy, and Digital Citizenship* have been aligned with the relevant standards of the International Society of Technology Education.² Other examples would also include the recently adopted reform *Lehrplan 21* (“Curriculum 21”) in the German-speaking parts of Switzerland that all public schools need to emphasize digital and media competencies and has led to the development of curricular materials for children in the first 2 years of primary school. Within this context, the teaching material *Secrets are Allowed*,³ a joint project of the Data Protection Authority of the Canton of Zurich and the Zurich University of Teacher Education, has been awarded with the 2019 Global Privacy and Data Protection Award. And the German Federal Commissioner for Data Protection and Freedom of Information in the year 2022 assigned the creation of children’s books of the well-known Pixi series—on e.g., Transparency (“But Why?”), Freedom of Information, and Personal Data (“This is Personal!”)—for a pre-school

BOX 6.4 GAME ON BIASES: BIAS BAMBOOZLE!

As a concrete method to teach children about the implications of data analytics, profiling, and (algorithmic) bias, we propose a game that we refer to as “Bias Bamboozle”. This game may be played in a classroom or workshop environment, and it may also be adapted to an adult audience.

The goal of Bias Bamboozle is to teach children about the inherent biases that can exist in data-based profiling and the importance of critical thinking and fairness. In addition to tokens for tracking scores, the game consists of a deck of cards, each showing:

- **Profile cards:** Each profile contains the name, picture, hobbies, and occupation of an individual.
- **Scenario cards:** Each scenario contains a situation or task that needs to be solved.
- **Data cards:** Each data card shows information about a bias that connects to the profile cards and the scenarios (e.g., that people with certain hobbies are more likely to be good at particular tasks). The biased data sheets contain a variety of biases—some might be based on hobbies, others on appearance, and others on job titles, etc. This will help children understand that biases can come in various forms.

A round of Bias Bamboozle progresses as follows:

1. A scenario card is drawn; this card outlines a particular job or task that needs to be done, for example, “Find a person to lead the school’s science project”.
2. Players then draw 3 profile cards each from the deck. They need to pick one person from their cards that they believe is best suited for the job based on the information provided.
3. The data cards are introduced. These contain misleading information, for example, suggesting that “people who enjoy painting are 50% less likely to be good at science”. Players can choose to use the data sheets or rely on their intuition.
4. Each player explains their choice. They can cite the biased data sheet, the actual profile, and scenario, or their intuition.
5. After everyone has made a selection, the scenario card is flipped to reveal a better choice based on a more complete and less biased set of data. In many scenarios, it will not be possible to make the best choice purely based on statistical information, and the children and teacher should discuss which additional factors (that are not contained in the statistical information) need to be considered in these cases.
6. Discussion Phase: After the round ends, the children and teacher discuss the biases that influenced each player’s decisions and the information on the data sheets. Children should be encouraged to

talk about why they made certain choices and how they felt when they realized their decisions might have been influenced by bias. They may write down their own biases on additional data sheets, and as homework, they might verify whether their bias is based on evidence or not.

As players progress, introduce more nuanced biases or even conflicting biases to make decisions more challenging and discussions more in-depth. The learning outcomes of Bias Bamboozle include an understanding that data can be misleading, a recognition of the importance of questioning sources, and a confrontation with one's own biases and verification of how far these are substantiated. Overall, children should learn the importance of fairness and treating individuals as unique, not just based on stereotypes. Through its interactive engagement with bias and the teacher-led discussions on the topic, Bias Bamboozle should help children engage with the concept of bias and promote critical thinking and discussion.

target group, and on Data Protection for school children aged 10–14 years of age. There is, therefore, evidence that it is feasible and desired to integrate more content about one of the many major societal aspects of ubiquitous computing, and specifically the foundations of automatically processable regulation, into early-school curricula. We expect more such curricular elements to be introduced that highlight the societal importance of sensing systems, communication infrastructures, and data analysis methods—and their consequences. The relevance of these aspects and of the invisible computers that enable them will keep increasing in the foreseeable future, and this can be catered to by integrating them across existing subjects from mathematics to history, rather than by teaching them linearly within a dedicated “Computer Science” class.

6.2 SPECIALIZED EDUCATION

We argued that already early schooling should include elements on the foundations of the automation of societal processes and its implications to ensure the functioning of democratic societies. On this basis, upcoming students at universities or applied colleges also need to be aware of how automation is changing their fields, in order to be well-equipped for the future of legal practice, policymaking, and public service. This book gathers diverse information on how automation and AI are altering the legal field and aims to make it accessible to a wider audience beyond interdisciplinary researchers in the field. Through this book, we have put an accent

on making sure, through real-world examples, that the materials within this book can be used to educate the next generation of students interested in research situated at the intersection of law, computer science, and political science.

Understanding the changes and how their challenges can be mitigated is the key for being able to leverage automation for the benefit of individuals and society. Targeted at tertiary education environments and based on the content elaborated in this book, we propose a syllabus (Table 6.2 below) that can be used and amended to specific educational needs. This syllabus includes not only a structure for the individual classes but can be combined with the exercises we propose in Chapter 7: Exercises. Together, the content of the book and the individual assignments enable students to learn through a problem-based approach that is grounded in current research in the area of AI and law.

TABLE 6.2 Syllabus for Specialized Education on AI and Law Based on the Content of this Book

General course description: In this course, we delve into the dynamic intersection of law and technology, exploring through real-world projects how automation and AI are reshaping the legal landscape. By examining examples of rule as code and executable versions of the law, students develop a deep understanding of how technology is changing legal practice and policymaking. This forms the basis for students developing the ability to devise strategies to mitigate the individual and societal challenges that are caused by these changes. Together, this class sets the stage for a comprehensive exploration of the implications of automation in law and equips students with the tools to navigate this evolving field effectively.	
Session 1: Introduction	We start this course by understanding how the concepts known as computational law, legal AI, legal tech, rule as code, code-driven law, executable version of the law, automatically processable regulation, etc. evolved. We introduce several real-world projects from both public and private institutions to contextualize the content of this course.
Session 2: Automation of the Law I	We delve into the basics of logic, ontologies, and controlled language, laying the foundation for understanding how technology is reshaping legal practice. By examining these concepts, students develop insights into the intersection of law and technology, preparing them to navigate the evolving landscape of automation in legal fields effectively.
Session 3: Automation of the Law II	We explore how to leverage machine learning techniques both for the formalization of law and, more generally, for the creation of automatically processable regulation. Furthermore, we delve into assessing the feasibility frontier for encoding laws, offering insights into the complex interplay between legal norms and computational methods.

(Continued)

TABLE 6.2 (Continued) Syllabus for Specialized Education on AI and Law Based on the Content of this Book

Session 4: Open-Texture	We define the term open-texture and conceptualize how to measure it in practice. We ask students to experience for themselves the prevalence of open-texture in law, and the low inter-rater agreement of flagging open-texture. We discuss the role of open-texture from the viewpoint of legal philosophy and discuss practical issues that arise when open-texture is present in (legal) documents, including the automation barriers this implies. Upon this basis, we analyze ways to encode different interpretations of legal norms into artificial agents and show the limitations of current research on the topic.
Session 5: Typologies of Legal AI Projects	We look back at all the examples encountered so far and classify them according to different typologies provided in the literature. We compare different real-world projects and their impact on citizens, policymaking, law enforcement, etc. This session serves to ground the course contents to date and forms a necessary basis for the discussions in the subsequent sessions.
Session 6: Responsible Automatically Processable Law I	We examine a host of societal issues surrounding the encoding of laws, including what possible potential remedies can be considered. We look at real past cases and discuss frameworks to foster the development of responsible automatically processable regulations. We apply those frameworks to understand how they impact the development of automatically processable regulation. Lastly, we examine the relation between the different dimensions of a typology and the types of issues triggered by automatically processable regulation.
Session 7: Responsible Automatically Processable Law II	We explore and reflect on current regulations of automation, AI, and computational law. This includes understanding how newer regulation on AI, such as the European AI Act, influences the field of AI and law. We notably bridge with the previous section by examining which areas of automatically processable regulation creation and usage are the most problematic and how, if at all, regulation could act as a further remedy.
Session 8: AI, Law, and Democracy I	We delve into the crucial topic of social acceptability of automatically processable regulation, examining how emerging technologies intersect with societal norms and values. Additionally, we explore outstanding debates in AI and law, addressing key questions and controversies surrounding its implementation and impact on legal systems worldwide.
Session 9: AI, Law, and Democracy II	We look at how existing democratic processes could account for the emergence of automatically processable regulation, notably which, how, and where debates should take place. We organize such debates between participants to showcase the range of possible opinions on topics that could, at the surface, appear technical but which have critical political, moral, and ethical underpinnings.
Session 10: Outlook	Abstracting from the domain of legal automation, we focus on the importance of critical thinking and emotional intelligence in navigating the evolving landscape of law and technology. We ask ourselves: In an ever-increasing automated world, how can we ensure that critical thinking and emotional intelligence prevail to ensure a well-functioning society?

It is great to see that freely available syllabi on the subject matter have started to emerge, such as from the ERC-funded “Counting as a Human Being in the Era of Computational Law” project mentioned in Chapter 3: Automatically Processable Regulation. Such openly accessible materials are a great resource for interdisciplinary researchers and educators who aim to integrate new directions and viewpoints with traditional fields that are deeply intertwined with societal processes, such as law. Yet, not only legal education should be updated to reflect these newer developments, but computer science disciplines also need a fundamental understanding of the concepts discussed in this book. In fact, courses on ethics and legal foundational courses are often part of the education of computer scientists (e.g., at EPFL with mandatory social and humanities science classes, at the University of St. Gallen with its Contextual Studies program). In such courses, the introduction of a foundational understanding of how automation is changing the legal discipline as well is needed (i.e., course elements on law and computation). Aspects include understanding how key rights (e.g., right to privacy) can be challenged and improved through technology and how we ensure that developments of social structures typically formalized via the law, e.g., through social movements and civil disobedience, are still guaranteed in a highly automated society (Custers, 2023; Tamò-Larrieux, Mayer, & Zihlmann, 2021).

An example showcasing how project-based learning approaches with interdisciplinary classes full of students in law and computer science can lead to interesting results for the field of AI and law comes from the United States: In a study, Escher et al. (2022) investigated the translation of legal text into computer code, analyzing the societal implications and development processes of legal algorithms. They explored how software developers or computer science students encoded a particular section of the law depending on the constellation they were faced with programming it alone, programming it as a pair of computer scientists, or programming it together with a student in law. After creating the automatically processable regulation, they were then surveyed on their confidence in their program and its implications. Their analysis leads to interesting discussions, such as discussing differences in opinions on whether such tools should be used in courts (with legal students being more reluctant to endorse such a step, and computer science confident of their translation of law into an algorithm) as well as discussing the sources that need to be relied upon to determine how to encode the law. Such discussions showcase that, in addition to disciplinary tertiary offerings in the fields of law and computer science, *cross-disciplinary courses* are required, combining expertise from social sciences and computational sciences.

At the University of St. Gallen, in 2022, two of the authors together with a professor of technology studies, co-created such a cross-disciplinary course covering social, legal, and technical aspects of quantified health tools; the course was offered to students of Business Administration, Finance, Economics, International Affairs, Law, and Computer Science. In it, the complex entanglement of technology, legislation, and societal change was investigated using real-world case studies, academic theories, and legal norms that shape the technology at hand. In the course, we picked three overarching concerns—Privacy, Openness, and Data Lives (see Box 6.5)—and treated each of these concerns from each disciplinary viewpoint, before integrating these viewpoints.

BOX 6.5 CROSS-DISCIPLINARY COURSE OUTLINE ON PRIVACY-INVASIVE TECHNOLOGIES COVERING TOPICS OF PRIVACY, OPENNESS, AND DATA LIVES AND ITS IMPLICATIONS

Privacy: Concerns over privacy implications of new technologies are not new. In fact, throughout history, the introduction of new technology—from photographic film to the telegraph and to facial recognition technology—has been accompanied by heated debates over what limits need to be set to ensure that their use respects other people’s privacy needs. The first session on privacy focuses on the technical foundations of security and privacy-preserving technologies. It is important to us that you see that these provide tradeoffs rather than “silver bullets”, and that you are aware of these tradeoffs. Technical developments have triggered regulatory reactions around the globe. You will get familiarized with one regulatory reaction that has been actively propagated in Europe: Data protection law. We will discuss how technical measures elaborated upon in the previous session are also part of the solution to addressing arising privacy challenges, by discussing the developments in privacy-enhancing technologies and the codification of privacy by design. You will get familiarized with the concept of datafication and we will discuss how datafication challenges privacy. You will be introduced to several conceptual approaches to protect privacy in a datafied society. Moreover, we will contrast legal understandings and contextualization with theories of data assemblages and data shadows found in social sciences.

Openness: Before elaborating upon the technical foundation and possibilities for openness, we will discuss the legal rights individuals have with respect to their own data. This includes the right to know and obtain information about personal data that is being processed as well as the right to data portability. We will put these rights into the broader context: Who makes use of such rights? What hurdles are individuals encountering? How useful have these rights proven to be so far? Who should access data from your wearable device? Who needs to process the data for a service to

function and make (economic) sense to a manufacturer? How can products and services around the quantified self-movement and health technologies be designed to empower individual consumers? All these questions point to the bigger issue of how open and accessible technology should be, with many legal battles that follow. Within this “openness” block, we will investigate the technical foundations of (open) interfaces between services and will discuss hurdles to interoperability—not only from a technical but also from a broader economic perspective. We will discuss how more open software ecosystems might be created through the decentralization of data and of control.

Data lives and implications: In this block of the course, we will explore the intricacies of data creation and discuss how data-driven technologies have become essential to how society, government, and the economy work. How can we begin to grasp the scope and scale of our new data-rich world, and can we truly comprehend what is at stake? We reflect upon what we have learned throughout the course and contextualize the findings. We discuss not only how health technologies have altered social values but how technology overall has impacted every corner of our daily lives

From Schneider, Tamò-Larrieux, & Mayer (2022).

We argue that such integrating courses are needed because they gear students toward having a holistic understanding of complex issues. This makes them more prepared for the real world, as decisions on whether to launch a product or service or how to regulate technologies do not occur in a vacuum. The ability to not only establish but to critically review connections between knowledge from different fields will become important especially in a world that is populated by artificially intelligent systems. With AI tools such as ChatGPT on our side, the future of specialized education will more and more need to move towards ensuring that students comprehend the underlying principles, connect knowledge across domains, and think critically about the principles and the connections from a human viewpoint. We need to stay curious beyond our specific fields of expertise and to contextualize knowledge with what we as humans care about, and are willing to commit to.

NOTES

- 1 Test it under <https://coveryourtracks.eff.org/>
- 2 See under <https://dcrp.berkman.harvard.edu/tool/safety-privacy-and-digital-citizenship-introductory-materials>
- 3 See under https://learnprivacy.ch/English/OEBPS/MB0501_split_001.xhtml

Exercises

THE GOAL WE PURSUED when writing this book was to combine the diverse scholarship on how automation and AI are altering the legal field and make it accessible to a wider audience. To do so, we relied extensively on legal and computer science scholarship, government initiatives, and case studies. Yet, in order to be well-equipped for future legal practice, policymaking, and public service, it is useful to not only read about current and future changes in the domain of law but also, more hands-on, go through exercises to critically reflect on those changes. As we have seen in the previous chapters, more specialized education on the subject matter is emerging and it will be key to establish good exercises that spark interest and ideas on how to further debate about “AI and Law”. Within this chapter, we provide some possible exercises and provide at the end of the book some guiding approaches for the solutions. The exercises below are divided according to the book’s chapters. However, some of the exercises draw across the themes discussed in the book so it is advisable to first read through all the chapters before tackling these exercises. Happy problem solving!

7.1 DESIGNING A LEGAL DECISION TREE

Different tools have been developed, such as GraphDoc, an open-source application that provides an intuitive and easy-to-use interface to visually display constructed flowcharts: <https://maastrichtlawtech.github.io/graphdoc/>. GraphDoc is available at no cost on the Maastricht Law and Tech Lab Github page and provides an easy interface to generate decision trees (Figures 7.1 and 7.2):

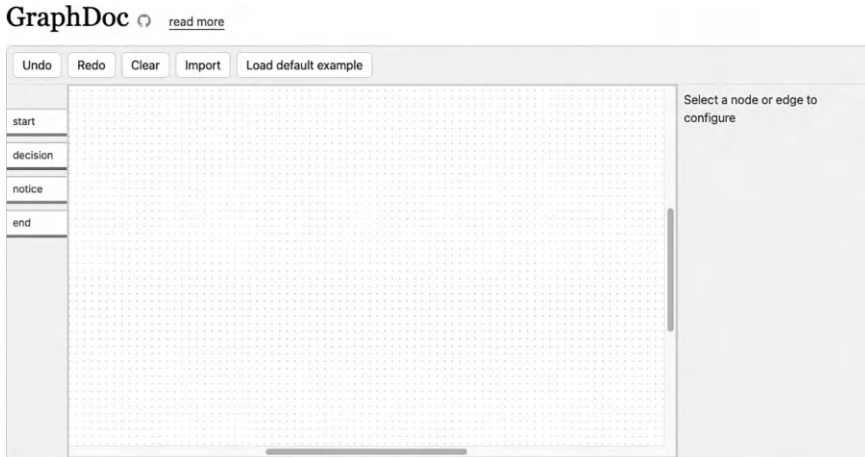


FIGURE 7.1 GraphDoc's user interface.

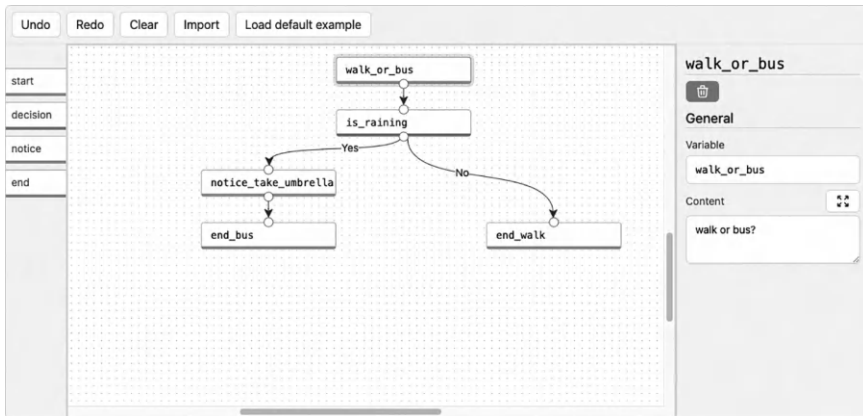


FIGURE 7.2 A default example with GraphDoc.

With it, simple decision trees can be built. The instructions on GraphDoc explain how to use the tool: “Start by dragging nodes from the left sidebar. Connect nodes by dragging edges from and to the ports of the nodes. Configure nodes and edges by clicking on their cells and filling in the details in the config sidebar on the right. Nodes and edges can also be removed from this sidebar”. In this exercise have a look at the material scope of the General Data Protection Regulation and sketch a decision tree to determine whether some data processing falls under the scope of the

law or not. To do that, open the GDPR and familiarize yourself with the material, personal, and territorial scope. You can start at a high level and refine it along the way.

7.2 TURNING NATURAL LANGUAGE INTO CONTROLLED LANGUAGE

As Chapter 2: Law and Computer Science Interactions previously presented, controlled language is situated in between natural language and computer code: Once something is expressed in controlled language, it is directly possible to obtain a logic program from it. In this exercise, you try this yourself while making use of the Attempto Controlled English Web service. Go to <http://attempto.ifi.uzh.ch/ape/>. There, type in the following (including the period at the end):

If the light is red then the car must stop.

When you click on the button “Analyse”, you will see the resulting DRS, but you will also see help in case you entered something incorrectly. For instance, try as well:

If the light is red then my car must stop.

You should now see the following with the description that the use of the pronouns “my” is not allowed with Attempto (Figure 7.3)—this is because the parser is unable to ground subjective statements.

A full description of the syntax rules can be found here:

http://attempto.ifi.uzh.ch/site/docs/syntax_report.html

And you will also notice that certain words are unknown—the lexicon can also be found here:

http://attempto.ifi.uzh.ch/site/docs/ace_lexicon.html

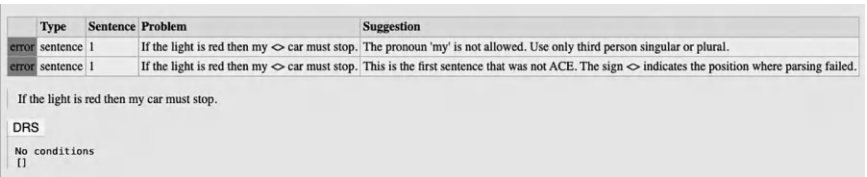


FIGURE 7.3 An example with Attempto where a personal pronoun causes a parsing error.

Basically, just to give you a few starting rules:

- The sentence needs to end with a period;
- The sentence can start with “If”, but then it needs to have a subsequent (without a comma) “then”;
- New words may be introduced by simply putting quotation marks around them;
- Relative sentences using “which” and “that” (but without commas) are allowed;
- No pronouns are allowed.

The goal of the exercise is to transform (parts of) an existing piece of legislation. To start off on the easy side, you could pick a law related to becoming a national of a country, as such laws very often entail rules already closely formulated in an “if-then” fashion. Once you have a clause in controlled language, note it down and go through the following:

1. Which compromises did you have to make?
2. In which way *may* the meaning deviate from the originally intended meaning?
3. Via another iteration, can you find remedies to your answers in (1) and (2)?

Once you have managed with a simple law and clause, move on to a more difficult one. You will also increasingly encounter the effects of open-texture, thereby helping to ground our treatise of this subject in the book. We encourage you to discuss your controlled-language version of a law with that of a peer!

7.3 MODELING A RULE

This exercise is designed to create a better understanding of the semantic structures needed to model legal clauses. Relying on an example that Jason Morris provides (see for more information the guiding approaches for solutions once you are done with this exercise), we will use the game *Rock, Paper, Scissors* in this exercise and ask you to think through the listed steps below. Before we do so, for everyone not familiar with the

game, here are its rules: *Rock, Paper, Scissors* involves two players, where, after an articulated count to three (typically while uttering the three words of the game's name), the players simultaneously form one of three shapes with their outstretched hand: *Rock* is represented by a closed fist; *Paper* by a flat hand; and *Scissors* by having two fingers form a V. After each round, a winner is determined following three rules: *Rock beats Scissors*, *Paper beats Rock*, and *Scissors beats Paper*. If the two players have the same sign (*Rock*, *Paper*, or *Scissor*) then this equals a draw.

In this exercise, you should create a formalization of this game that might be represented as an ontology. To do this, think about the different concepts that occur in the *Rock, Paper, Scissors* game, their properties, and the relationships between the concepts. To guide you through these processes we suggest you tackle the following questions:

- What are the core concepts in the game? Define these concepts, and label them using individual or composite nouns.
- What are the properties of each concept? Add them to the concepts.
- What are the relationships between the identified concepts in the game? Define these relationships and label them using individual or composite verbs that relate the concepts to one another. Not all concepts need to be related.
- How can the rules of the game be formalized while utilizing the defined concepts and relationships?
- Finally, are you able to express your concepts, properties, relationships, and rules using a language or framework that exists already?

7.4 CLASSIFYING AUTOMATICALLY PROCESSABLE REGULATION PROJECTS

In the following, we provide you with descriptions of different automatically processable regulation projects. These project descriptions offer a chance to discuss the typology dimensions presented in Chapter 3: Automatically Processable Regulation. Please note that projects described in 7.4.1 and 7.4.2 should be well familiar to the careful reader as we have mentioned them throughout the book, but they are presented as summaries which can be used as a standalone by instructors (and others). For all cases, we leave it to the readers to see whether they can identify where the difficulty lies for each case. Lastly, readers can fill out the following Table 7.1 for each

TABLE 7.1 Applying the Typology to Any Case Study by Determining the Different Dimensions

Typology	Answer
Primary aim: (efficiency/accessibility)	
Potential for divergence of interest (maximum between the two)	Number of distinct actors (beneficiaries, sponsors, implementers, users) Degree of observed divergence of aims across and within non-overlapping actors (1 to 4)
Degree of mediation by computers (each from 0 to 4, aggregated using Euclidean distance)	Domain factor Code factor Data factor

case to determine the location of the projects on a two-axis plan (potential for divergence/degree of computer mediation), with different coloring for the two primary aims.

7.4.1 Rates Rebate

New Zealand’s Service Innovation Lab (Higgison, 2020), under its mission to create “Better Rules” (Fraser, 2021) implemented in 2018 benefits calculators¹ for two acts, the Rates Rebate Act and the Holidays Act. In this case study, we focus only on the Rates Rebate Calculator (RRC), which calculates partial refunds for individuals who paid rates to the council in order to help individuals pay for their housing (it thus targets low-income families). The calculator is for everyone to use. The Rates Rebate Act codifies the condition under which an individual or a family is able to apply for a rebate; it is a rather short law with more straightforward criteria that are applied to determine eligibility and which documents must be submitted to prove specific conditions of eligibility (e.g., income tax forms).

The RRC was developed by a multidisciplinary team over the course of a 3-week workshop in 2018. The calculator was the result of the “Discovery Sprint” initiative launched by the Service Innovation Lab (LabPlus) of the New Zealand Government which aimed at “exploring the challenges and opportunities of developing human and machine consumable legislation for effective and efficient service delivery” (Digital Government NZ, 2018). The full report of the Discovery Sprint describes, among others, the key questions that were addressed during the workshops, the approaches and insights generated, as well as the team members present during this

initiative. The core team included designers, legislative drafters, policy analysts, software developers, and strategic advisors.

The team developing the RRC took different steps to translate the legislation into what they refer to as pseudocode and, finally, software code. The first step was the creation of decision models and flow models, which are necessary for creating the structure/process of the RRC. The second step was the extraction of key terms and their definition in a formal language. To that end, they used OpenFisca²: OpenFisca enables the representation of rules in code by translating specific expected input/output into a formal logic. But OpenFisca is geared towards simulation, not towards supporting actual decision making.

Overall, the results of the Discovery Sprint were that while it is difficult to produce machine-readable rules the best way to go about it is to create multidisciplinary teams and take a user-centric approach when designing new systems. While not all legislation will be suitable for being translated into code, “common frameworks, reference points, and data points (like concept and decision models and ontologies) will assist multi-disciplinary teams to co-design policy and legislation and, once developed, can be used as blueprints for the development of human- and machine-consumable rules without the need for further translation of the intent and logic (which, in turn, reduces the time and resources required and the chances of errors)” (Webster, 2018).

7.4.2 Mes Aides

More than 10 years ago, the French state conducted an extensive review of people’s knowledge about their eligibility for social benefits. They concluded that many people could, in fact, be eligible for benefits but were also struggling to assess how certain actions would affect their benefits, for instance by taking up a job, due to the complexity of the rules. Furthermore, many either did not know about their eligibility or did not know how to check it, leading to them not applying for social benefits that they were entitled to. A logical next step was to attempt to remedy these issues, and one of the ideas was to do so through a new legal tech solution called Mes Aides. The original goal of Mes Aides was that anyone in France could use the platform to evaluate their eligibility to up to 30 benefit schemes.

The project started under the left-wing government of President François Hollande, who launched it with much noise on October 30, 2014 (Alauzen,

2021; Merigoux et al., 2024). While Mes Aides was welcomed by many, the case study also exemplifies the diverse interests at play: What may sound like a desirable goal, namely to make the law more accessible, is not without opposition. At the core has been the issue that Mes Aides is only a simulation tool, based on OpenFisca, a platform also put at the disposal of other governments but which can only also support such simulations. And a simulation is not the same as the automated decision-making implementation used by the genuine offices handling benefits; it is vastly less potent, albeit still useful in certain ways. It lacks officialdom. Moreover, the simulator does not cover all cases, crucially not when such cases would be too complex. It was only a simulation tool because the exact code was closely kept as a secret by another ministry which did not want to share it (either by fear of showing how deprecated and unnecessarily complex the code was or on dubious charges of state secret—something that similarly happened with the French tax software too). The team behind Mes Aides had, therefore, to re-encode the benefits law from scratch using a modern language (Python), which was, however, too difficult to integrate with legacy systems and too slow to handle the millions of queries performed by the state officials. From the onset, it was clear that this would hence remain a simulation tool.

Opposing forces to the project put forward three main arguments: First, it would turn the benefit-paying institutions into a “service provider”; second, it would overload them with queries; and third, it would heighten fraud. On top of that came the worry that people who would obtain a positive answer on the simulator but a negative one from the real office would have difficulties understanding the difference, despite the online simulator stating as clearly as it could its nature.

In the end, Mes Aides turned out very popular: In 2019 alone, there were 10 million connections to the site, and 1,970 single visitors came on it more than a hundred times, hinting that professional social workers must have used the platform too on behalf of citizens. And yet, despite this relative success, investments of €1.25 million across 5 years, the ministry scrapped it in 2020, replacing it first with a site merely providing textual information (see Figures 7.4 and 7.5). Another competitor website, developed by yet another ministry but with ten times less traffic (MesDroitsSociaux) was supposed to pick up the slack, along with another simulator for young people under 30, counting roughly 2 million visitors a year. In light of the



FIGURE 7.4 Screenshot on Mes-Aides.gov.fr on December 13, 2019.



FIGURE 7.5 Screenshot from Mes-aides.org on May 10, 2022.

popularity of Mes-aides.gouv.fr, the state has now kept the website but redirects users to the two other alternatives it has set up.

In a twist of the event, a self-branded group as a “citizen community” reused the source code and copied it onto a new website, meaning that it is still accessible to the larger public, although no longer on a state-powered website but under a registered NGO. In meanwhile, beta.gov.fr, the “state start-up”, released the code publicly on GitHub: <https://github.com/betagouv/aides-jeunes>.

Regardless of how desirable the goal of making the law more accessible to the public is in order to ensure that as many people understand it as they should, the French case for *Mes Aides* highlights a crucial point that is probably unavoidable: Power struggles within large bureaucracies and between ministries—notably those acting as gatekeepers to the state secret and those seeking automated transparency.

7.4.3 Victor

Judges within the Brazilian Supreme Court use Victor to determine the admissibility of cases, with Victor notably testing whether a case can have sufficient repercussions for the society for the court to take it on, a principle known as “general repercussions”. The test is for the court to work on socially relevant cases rather than to be a last-instance type of recourse. To administer such a test manually takes time, an average of 45 minutes, and has to be carried out for roughly 50 thousand applications per year.

The University of Brasilia co-developed it along with the Supreme Court in 2019 (Becker & Ferrari, 2020). Controversies have been mostly on two fronts: Regarding the unknown extent to which judges fall under biases by following Victor’s recommendation (especially for rejected cases); and, following the introduction in 2021 of the Brazilian Data Protection Law, automated decision-making needs to be “fair, transparent, and informed”, a standard Victor possibly does not meet.

Lastly, a machine learning algorithm trained on past cases meeting the standard for “general repercussions” constitutes the technical underpinning. More specifically regarding data, training occurred on more than 118 thousand appeals filed between 2017 and 2019 and 3 million case dockets, and the court further feeds to Victor new cases it has been receiving since then (Conselho Nacional de Justiça, 2019).

7.4.4 DoNotPay

Receiving a parking fine can be vexing, but sometimes, contesting the fine is also possible—provided that one knows the law, how it is applied, and through which processes. One individual originally developed DoNotPay in 2015 for the United Kingdom market to assist defendants in their contestation quests. But the developer later extended it to other regions and to broader scopes than merely disputing parking offenses, for instance, for filling Freedom of Information access requests. Originally,

the “robot-lawyer” seems to be using mostly a type of expert system, but the company’s CEO, Joshua Browder, then announced on X.com in January 2023 plans to have an AI chatbot, writing: “On February 22nd at 1.30pm, history will be made. For the first time ever, a robot will represent someone in a US courtroom. DoNotPay A.I will whisper in someone’s ear exactly what to say. We will release the results and share more after it happens. Wish us luck!”. Two days later, however, Browder had to retract his announcement following threats by state prosecutors of a jail time of 6 months (Cerullo, 2023). While Browder did not disclose on which basis, it is most likely that the legal threat was on the basis that all parties must consent to be recorded in the courtroom, a consent they did not have, and the tool would have needed recording to be able to process what was being said to then suggest an appropriate answer.

This already showed that there have been mixed reactions, especially to the further iterations of the DoNotPay application. In the United States of America, more cases and challenges emerged. A law firm appears to have perceived DoNotPay as a new competitor, and hence as a threat to their own market. The law firm filed a lawsuit against DoNotPay for practicing law without a proper license, effectively “pitting real lawyers against a robot lawyer”, according to the judge (Merken, 2023b). Yet, the judge dismissed the lawsuit as lacking real damage in Autumn 2023. In another lawsuit, still ongoing at the time of writing, the same claim was made of practicing law without a license, but the claimant also alleged that they obtained “substandard and poorly done” results (Merken, 2023a). It would notably appear that DoNotPay has repeatedly missed deadlines and mis-handled client cases: For instance, one customer had their plea changed from “not at fault” leading to the customer having to pay, while another stated that they would not have purchased the services had they known that no real lawyer was behind it—a claim showing a clear lack of due diligence (or good faith)—and that many documents delivered which were simply unusable or contained mistakes related to misspelling names (Merken, 2023a). In relation to misleading claims that the companies made, the Federal Trade Commission fined the company a modest USD 193,000 in September 2024 following a complaint that “the company did not conduct testing to determine whether its AI chatbot’s output was equal to the level of a human lawyer, and that the company itself did not hire or retain any attorneys” (FTC, 2024).

7.4.5 Overtime Regulation

The original aims of Canadian civil servants were two-fold: To update the piece of legislation and, in so doing, to provide for an automatically processable regulation version of it. The full title of the legislation—for anyone seeking to consult it—is the Motor Vehicle Operators Hours of Work Regulations. What made it a good candidate for such an exercise in the eyes of the public servants was that it's only two pages long, and so, rather short. But the two goals quickly proved conflicting: In order to create an automatically processable regulation version of the law, external developers and (technical) consultants had to be on board; in order to be involved in the legislative process, even at the draft level, a security clearance is required as the work is sensitive and can have repercussions on specific markets. However, the external staff could not obtain a security clearance, or at least not so rapidly. The original aim, hence, had to be truncated to “only” creating an automatically processable regulation version.

The piece of legislation determines when truckers can take (and be paid) for overtime. An automatically processable regulation version of it can hence be useful to truckers and their company so that they can better plan, and for regulators, so that it facilitates both their oversight (controlling) of the trucking companies and the clarity of the regulation, thereby raising the chance of correct compliance. In order to turn it correctly into an automatically processable regulation, the process had to involve the kickstarter of the project, the Canada School of Public Policy, policymakers, regulators, truck companies, and engineers. Apart from the figures stated in the statute relating to how overtime works, there was no requirement to include any other type of data. Lastly, implementation occurred using OpenFisca (a tool easily available that facilitates turning pieces of regulation into automatically processable regulations). A particularity of OpenFisca, though, is that any projects using it are only a form of “simulation”: In other words, no projects currently running around the world and based on OpenFisca use it for the actual automated decision-making mechanism within a state office.

7.5 IDENTIFYING OPEN-TEXTURED TERMS

Note: This exercise is best conducted in pairs of two people.

In past experiments which we have conducted, to identify open-texture, four questions were found to be particularly valuable. These are:

1. Is there more than one quantity/quantification associated with the word? e.g., periodically, monthly.
2. Does the word include a broad spectrum of meanings? e.g., secured, appropriate behaviors.
3. Is there no agreement on a single version of the definition or standard? e.g., freedom of expression, terrorism.
4. Does the word represent a value, or is it value-laden, for example, because it presupposes the acceptance of specific moral principles or beliefs? e.g., economic welfare, liberty.

These questions are not supposed to be mutually exclusive and certain terms could be identified as open-texture by answering “yes” to more than one question. We illustrate how the questions can be utilized in practice with the GDPR Art. 32(1) titled “Security of Processing” (see Box 7.1).

BOX 7.1 ART. 32(1) GDPR

Taking into account the state of the art, the costs of implementation, and the nature, scope, context, and purposes of the processing, as well as the risk of varying likelihood and severity for the rights and freedoms of natural persons, the controller, and the processor shall implement appropriate technical and organizational measures to ensure a level of security appropriate to the risk, including inter alia as appropriate:

- a. the pseudonymisation and encryption of personal data;
- b. the ability to ensure the ongoing confidentiality, integrity, availability, and resilience of processing systems and services;
- c. the ability to restore the availability and access to personal data in a timely manner in the event of a physical or technical incident;
- d. a process for regularly testing, assessing, and evaluating the effectiveness of technical and organizational measures for ensuring the security of the processing.

For this exercise, follow this process:

1. Start by highlighting all terms which you identify as open-texture;
2. Separately, ask someone else to do the same, without any sharing of results at this stage.

3. Calculate your inter-annotator agreement via the Cohen's Kappa measure (Table 7.2). To do this, fill out Table 7.2 by:

i. Computing how many open-texture do you and your pair agree on? This is the value a.

ii. Computing the three remaining cells. Person A will look at the table vertically so that, from Person A's perspective,

c = Value a minus the number of OT terms marked down by Person A

and from Person B's perspective,

b = Value a minus the number of OT terms marked down by Person B

Finally, the value d is calculated by subtracting values a and b (or alternatively values a and c) from 135, which is the number of words in the article.

iii. Next, calculate:

$$p_0 = (a + d) / (a + b + c + d)$$

$$p_{OT} = (a + b) * (a + c) / (a + b + c + d)^2$$

$$p_{non-OT} = (c + d) * (b + d) / (a + b + c + d)^2$$

$$p_e = p_{OT} + p_{non-OT}$$

$$Kappa = (p_0 - p_e) / (1 - p_e)$$

Is your agreement with your partner higher than 60%? Then your partner and you already rank exceptionally high with respect to your alignment of interpreting legal texts!

4. See whether you can increase your agreement by going through the list and discussing it.

TABLE 7.2 Aid to Calculate Cohen's Kappa

	OT _A	Non-OT _A
OT _B	a = OT _{A&B}	b = OT _B but Non-OT _A
Non-OT _B	c = OT _A but Non-OT _B	d = Non-OT _{A&B}

Finally, keeping in mind automatically processable regulation, try to think of how you would go about encoding the term “*state of the art*” in a way that remains flexible enough to keep up with the evolution of the term’s meaning. As a starting point in your reflection, consider using a taxonomy for cyber-security (see, for instance ENISA (2022)—the practical grounding of this example is that several laws mandate that data be secured with “state of the art” cyber-security mechanisms) and how you could query the “state-of-the-art” of its different elements (e.g., by looking at vendors, associations, etc.). Be as concrete as possible.

7.6 DEBATING ABOUT ISSUES OF AUTOMATING LEGAL PROCESSES

Consider the following statement from Bench-Capon and Sergot (1985):

Nevertheless, it is totally unacceptable in general that legal decisions should be taken by machine, whether the machine can explain its conclusions or not. The nature of the law itself limits the usefulness of computer programs that are intended primarily to *take* legal decisions. Computer programs in law become more widely applicable if they are regarded not as decision takers, but as legal decision-taking *aids*. Used in this way, they are tools for the analysis and solution of legal problems. The construction of proofs, with a view to identifying possible lines of reasoning, is the principal *aim* of consulting such a program [emphasis from original]

Bench-Capon & Sergot (1985)

Do you agree with the authors’ take on the role of computers being limited to aiding and not taking decisions? If yes, why? If not, why not, and do you have counterarguments?

To help in answering those two questions, consider the following: Take one of the several cases presented throughout the book. Take as well the framework from Chapter 4: Challenges and Controversies. Go through each row of potential issue and write down/discuss them:

1. To what extent does the issue arise in this context?
2. How can a remedy be implemented? Be as concrete as possible.

3. Which areas would require, if at all, intervention from state institutions—be it in the form of a legislative debate, priority of (legal) principles, interpretation of the legal text, etc.?

In a separate step, consider pairing up with someone or debating the options. Alternatively, consider the effects of a specific socio-politico orientation on the debate: Create two or more camps (possibly involving several people), assign to them a specific socio-politico position (e.g., ranging from conservative to liberal, or from lawyers vs. engineers, again, real or *imaginary*), and ask them to debate the issues/remedies from the point of view of their assigned position. This should bring out differences in approaches based on *perceived* priorities by the camp. Assign at least one person as an observer who summarizes and reports at the end on the different arguments brought about, and on whether consensus could be reached.

NOTES

-
- 1 <https://www.govt.nz/browse/housing-and-property/getting-help-with-housing/getting-a-rates-rebate/>
 - 2 <https://openfisca.org/en/>

Epilogue

AFTER LAYING OUT THE nuts and bolts of the interactions of automation and AI with law through technical explanations, stories, and debates, we would like to leave readers with a few more thoughts on where this might lead to. We would like notably to explore two last areas: Personalized law, and regulating automatically processable regulations.

PERSONALIZING LAW?

How can law be personal or personalized? One of the most fundamental tenets of law is that it applies to everyone. Research on personalized law (Ben-Shahar & Porat, 2021; Busch & Franceschi, 2020) argues, that instead of writing the law generally to cater to the widest possible number of scenarios, as is currently the case, we should consider ways to break it down to the many instances possible to enable a more personalized application of the law. Therefore, a possibly better term than “personalized law” is the term “micro directives” (Casey & Niblett, 2019). Automatically processable regulation is one fundamental way through which this vision for personalizing law could come to materialize. For instance, automatically processable regulation could run down many simulations and suggest directly the exact appropriate wording catering to each scenario. Automatically processable regulation could also ensure that the law remains accessible and understandable, despite it being lengthier, notably as automatically processable regulation could ensure that citizens can query the law and obtain answers to clarify their doubts.

Expressing statutes in less broad terms would have the advantage notably of reducing, potentially, the use of open-texture of the law by making the application of the law much more concrete. A consequence of this, however, would be to decrease the power given to the judiciary by shifting the interpretation of the law away from the court and back to the legislatures. In one way, this would make it more democratic in the sense that

members of parliament are elected officials representing the voice of the people and their interpretation of their intent of the law would be given more weight. On the other hand, courts act as a counter-weight in the balance of power and also a pillar of democratic societies in order not to lead to excesses. And so, at the moment, it is not very well understood either how personalization of legal norms would look concretely in practice or what consequences for our societies this would entail. And yet, in the meanwhile, humbler projects have emerged which would offer more direct help to people. Take the field of privacy law. Who reads privacy notices—or for this matter, who has ever read one privacy notice at all? One seminal paper called it “the biggest lies on the internet” (Obar & Oeldorf-Hirsch, 2018). Yet, users do express concerns about privacy, even if their behaviors don’t reflect it, and by that, we mean their behavior of overly sharing rather than their behavior of not reading privacy policies. It follows that it may be possible to present privacy policies in a way that caters to users’ concerns, as much as to their abilities to read and comprehend texts. It could be possible to modulate the presentation of the text according to personal preferences, ranging from which parts of a privacy policy are displayed first, to tuning the sentence structure and words used to a person’s skills: More legalese for those with a background in law, more controlled language for the rest. In this case of application, the law (aka the information duty referred to in the law) itself does not change but only its presentation.

The research into automatically processable regulation will still contribute to making this form of personalization of legal norms possible. In the very example from above, research into comprehension of the law when written in various forms will be essential, as much as how to extract personal features and link them to preferences. Furthermore, an automatically processable regulation version of the law will allow more flexibility when modulating the presentation of each part of the law, without having to do it manually but simply by specifying general rules. Again, we note that personalized law can happen without automatically processable regulation, but automatically processable regulation does come in handy in the endeavor. Just to name a few other applications of personalized law to show the potential breadth of applications: Speed limits in specific zones could be lowered when in the presence of driver convicted of driving under the influence; or default in inheritance law, instead of being based on average preferences, could reflect differing preferences between gender, aged, and wealth groups. Applications of personalized law could hence reach into

many different facets of life, and possibly contribute to increasing citizens' well-being while keeping down the negative side effects.

REGULATING AUTOMATICALLY PROCESSABLE REGULATION

A separate point that we have so far not broached is about regulating the use of automatically processable regulation. What we have presented so far is that the choices made in implementing automatically processable regulations matter and can have far-reaching negative consequences. And so, as is the case often with such technologies—GMOs, airplanes, or social media—, the question arises whether it can make sense to try to have a tighter legal framework around the development, deployment, or use of automatically processable regulations. At the moment, most automatically processable regulations don't operate in a legal void either, though. Those leveraging AI will have to comply with AI regulations (e.g., the AI Act in Europe), and others will still have to navigate product liability, data protection, and cyber security rules, just to name a few. But as various scandals have shown redress of wrongs can still be very tricky, and relying on the already established processes is not enough. To resort to an analogy: Planting a nail down in a tilted fashion is a lot easier than pulling it out and nailing it down correctly again. Unlike with the Post Office scandal in the United Kingdom, not everyone will be able to have the chance to get a film done to show a wider audience how wronged they have been, so that, eventually, courts, government, and private companies recognize their mistakes—nor should it be the preferred way for redress to occur. And so, it may be necessary, as automatically processable regulation becomes more widespread, to resort to statutes. There are currently either early discussions or existing legal frameworks when it comes to the use of automated decision-making tools by governments. This is less so for their use by private companies even though the damage can also be important—think of automatic denials for loans or insurance coverage. Also, privately developed and operated systems on which individuals rely to obtain an answer to a legal question could also matter greatly, and existing regulations could prove insufficient to cover for their numerous applications equally as to the numerous items on the list of things that can go wrong. While we do not necessarily advocate for more regulation, we'd simply like to point out to the readers that this is an area in which we could see in the future more coming to the fore. And how regulations emerge on automatically processable regulation could very much vary from country to country.

From the examples highlighted in this book, a couple of points regarding cultural differences to automatically processable regulation should have come out. Best practice in New Zealand includes considering the process in its entirety, involving many different stakeholders, and being transparent about the advancement of the discussion, even in its draft form, can be embraced. In the very case of New Zealand, no backlash could be observed from the over-sharing of documents on progress, even in its very raw forms. In France, a different bureaucratic culture has been palpable, with early push-back against the solution and difficult access to information (e.g., access to the source code despite a law mandating it—also a case in point in how more regulations does not necessarily solves the issue especially if not combined with a strong enforcement mechanism). In the case of *Mes Aides*, the online tool at the disposal of French residents to check how their eligibility to various social schemes could be impacted by them taking up job opportunities or even how their eligibility to a scheme is impacted by them receiving another scheme, clerks expressed concerns that they would have to deal with a higher number of applications while the number of clerks would remain the same. Clerks also feared that they would have a higher workload to explain rejections, as the online tool was merely a simulation and did not reflect the actual encoding of the decision-making tool. Therefore, the best practice for France will look different than the one for New Zealand, especially as the gap between ideals and *de facto* implementation will vary. While the points raised by New Zealand may therefore look sensible, it would be futile to expect that they are applied similarly over the board across states and private institutions alike. Seizing the potential for automatically processable regulation will necessarily mean being sensible to many nuances in the approach to law—and only time will tell whether we as a society have been able to heed these nuances to create a society that is not only more efficient but also more responsible.

Acknowledgments

IN EMBARKING ON THIS exploration of the intersection between AI, law, and the automation of legal processes, we extend our gratitude to the institutions and individuals that have provided the ground for our intellectual pursuits. Our thinking has been shaped by a great many interactions across the years, with great pioneers in the field of AI and law, such as our co-authors of the short piece *Pervasive Computational Law* published in *IEEE Pervasive Computing* Prof. Dr. Kevin Ashley from the University of Pittsburgh, Prof. Dr. Giovanni Sartor from the University of Bologna, Prof. Dr. Mattias Grabmaier from the Technical University of Munich, Prof. Dr. Gijs van Dijck from Maastricht University: Thank you for those valuable exchanges!

During the course of our journey, we had the honor to interview many people, especially ones involved in state-sponsored automatically processable regulation projects. We want to thank again everyone who took the time for advancing research from their busy schedules and talking to us, at times due to the time zones at very odd hours. These semi-structured interviews were an ideal starting point to understand the complexity of automating governmental processes and legal rules on the ground and have shaped our way of thinking about the subject. Moreover, we have asked many students to work on annotations so that we could run experiments, and we have exchanged with many scholars on the topic throughout a series of workshops organized from 2021 to 2023. We are very much indebted to everyone who helped us along this interesting journey. We are also extremely thankful for the time that they took to share with us their various valuable insights, as much as we are to our respective institutions. This joint research would not have been possible without the support and open-mindedness of the School of Computer Science at the University of St. Gallen, the Law and Tech Lab at Maastricht University, and the Faculty of Law of the University of Lausanne. We would also like to thank the

Hasler Foundation for financing research into exploring the frontier of what is currently possible with turning law into automatically processable regulation, which branched out into many joint publications, including on open-texture, digitally ready legislation, and trust.

Part of the funding allowed us to host workshops, which strongly contributed to enrich our collaborative journey through multiple discussions. The workshops were held in Switzerland, each time with a different twist: The first one was general on many different automatically processable regulation directions with the discussions leading to the establishment of the responsible automatically processable regulation framework reproduced in this book too. Here we thank our workshop participants and co-authors Prof. Dr. Dimitri van Landuyt, Prof. Dr. Eduard Fosch-Villaronga, Prof. Dr. Irene Kamara, and Prof. Dr. Przemyslaw Palka for establishing together a framework that can be applied on the ground to test and mitigate possible challenges of automatically processable regulation. The second workshop focused on the current limitations of turning law into automatically processable regulation. Finally, within our third workshop, which compared to the other workshops was not international, we sought to bring together the different strains of research within Switzerland.

Once again, we would like to stress our deepest thanks to the “silent” workers who made many of this research possible: The students of Maastricht University who poured over texts for hours, going through repetitive tasks, for us in order to be able to test our different hypotheses. Participants in our study on digitally ready legislation also deserve accolades for taking time out of their busy schedules to allow us to see how they were understanding legalese vs. controlled language differently—to members of parliament, public servants, and students, thank you.

This book is therefore a collective effort to merge our different perspectives to provide knowhow at the intersection of various disciplines that are interacting. We therefore wish to thank everyone who has enriched our endeavors and propelled us forward in our quest to understand and shape the evolving landscape of artificial intelligence and the law. Our dearest thanks therefore, go to our mentors who have helped us throughout our academic journeys, as well as our families who have enabled us to stay curious and pursue these challenging yet rewarding research avenues.

Guiding Approaches for Solutions

EXERCISE 7.1: DESIGNING A LEGAL DECISION TREE

We created our own documentation of the decision tree on the material scope of the GDPR. This work was conducted over an internship program 2022 at the Law and Tech Lab at Maastricht University and refined later on. Here our attempt:

This tool starts with an announcement that it will help the end-user determine whether the GDPR is applicable to them. A hyperlink to the GDPR is provided accordingly.

The tool first tests the **material scope**.

It asks whether one performs on the data any of the operations that constitute processing (node **processing**). Examples of activities that can constitute processing according to Art. 4(2) GDPR are mentioned. If the answer is “No”, one is led to the node **notice_no_material_scope**, that states that the GDPR only applies when personal data is being processed as defined in Article 2(1) and 4(1) GDPR and that the material scope is not fulfilled. Then the user is directed to node **gdpr_not_applies**, and told that the GDPR is not applicable to their case.

If “Yes”, one is directed to a notice that explains what personal data is, using Art. 1; Art. 4(1) Recitals 14, 27, and 30 of the GDPR.

Then, one is directed to the node **identified_or_identifiable** that asks whether the information that is processed relates to an identified or identifiable natural person. The end-user is given two options—“Yes”, “No”. If the answer is “No”, one is directed to the node **notice_no_material_scope** that states that the material scope has not been fulfilled in accordance with Art. 2(1) GDPR, and then to node **gdpr_not_applies** and informed that the GDPR is not applicable to their case.

If the end-user clicks “Yes”, they arrive at a node **sensitive_data**, as this is a sub-category of personal data. Here, they are asked if they are any of the special categories of data, as defined in Art. 9(1) GDPR. If the end-user clicks “Yes”, they are directed to node **notice_sensitive_data** where they are informed that they indeed process sensitive personal data, and the processing of which is prohibited unless the exceptions apply. The end-user is told the explicit consent exception, and then referred to Article 9(2) GDPR for further exceptions. From there the end-user is directed straight to the node **controller**.

If the answer is “No”, one is directed to node **non_sensitive_data** where one is asked if one processes any of the following: Name, address, phone number, e-mail address, online identifiers (such as IP address, cookie identifiers or RFID tags), identification numbers (such as social security-, health insurance-, tax identification- or citizen service numbers), identification documents (such as identity card, driver’s license or passport), photos or video recordings based on Article 4(1), Recital 30 GDPR and guidance by the European Commission (hyperlink given). If one presses “Yes”, one is directed to the node **notice_material_scope** that informs them that the material scope was fulfilled. If they press “No”, they are directed to the node **potentially_identifiable** where they are asked if they collect, store, or process *at least two* of the following: Year of birth, gender, physical characteristics (such as weight, skin color, hair color, eye color, stature, height, dress size), details of profession, education, information on income and financial status (such as bank details, loans), information on family and marital status, certificates, job references, employee appraisals, curriculum vitae, telecommunications data (such as connection data, contents of the communication), location data or postal code, customer data (such as orders, customer account data, delivery addresses), psychometric data (such as assessments of a person in terms of knowledge, skills and experience), voice recordings or spoken language or languages in line with Article 4(1) and Recital 26 GDPR.

If the end-user presses “Yes”, they are directed to the node **notice_material_scope** and informed that the material scope was fulfilled. To rely on two of the mentioned characteristics to fall under the scope of the law, is a value statement that the developers of the decision tree made as a conservative measure. The goal is to ensure that users who did not state that they were processing personal data would still be flagged, if more than one of the mentioned categories that typically enables identification are being processed. If they press “No”, they are directed to node **notice_no_material_scope**, that states that the GDPR is not applicable to them since the material scope is not fulfilled according to Art. 2(1) and Art. 4(1) GDPR.

Then the tool tests the **personal scope**:

After a notice on fulfillment of material scope, the end-user is directed to the node **controller**, where the end-user is asked whether they determine the purposes and means of the processing of personal data, either alone or jointly with others. If the end-user answers “No”, they are directed to the node **processor**. If they answer “Yes”, they are taken to node **notice_controller**, where they are informed that they qualify as a controller under Art. 4(7) GDPR and that the personal scope of the GDPR is fulfilled. They are then directed to the node **establishment** to determine if they fall under the territorial scope of the GDPR.

The node **processor** asks whether the end-user processes personal data on behalf of someone else who determines the purposes and means of the processing of personal data, either alone or jointly with others. If the end-user answers “No”, they are directed to node **not_personal_scope** and informed that the personal scope of the GDPR is not fulfilled, as they do not qualify as a controller or a processor. However, they are also told that they might still be considered a third party within the meaning of Article 4(10) GDPR, and the Regulation could apply accordingly. If they answer “Yes”, they are taken to the node **notice_processor**, where they are informed that they are a processor under Art. 4(8) GDPR and that the personal scope of the GDPR is fulfilled.

Then the tool tests the **territorial scope**

Next, the user arrives at the node **establishment**, where they are asked if they are established in the European Union, in line with Article 3(1) GDPR. If the end-user clicks “Yes”, they are directed to the node **notice_territorial_scope** that states that according to Art. 3 GDPR, the territorial requirements are fulfilled.

If the end-user presses “No”, they arrive at node **located_in_eu**, which asks whether the data subjects are located in the European Union and the processing activities are related to either the offering of goods or services to them or the monitoring of their behavior as far as their behavior takes place within the EU, as stated in Article 3(2) GDPR. If the end-user clicks “Yes”, they are directed to node **notice_territorial_scope** that states that according to Art. 3 of the GDPR, the territorial requirements are fulfilled. If the end-user presses “No”, they arrive at the node **public_int_law**.

The node **public_int_law** asks the end-user whether they are established outside of the European Union, but in a place where Member State law applies by virtue of public international law, as stated in Article 3(3) GDPR. If the end-user clicks “Yes”, they are directed to the node **notice_territorial_scope** and again told that the territorial requirements are fulfilled. If the end-user presses “No”, they arrive at the node **gdpr_not_applies** that states the GDPR is inapplicable.

After the **notice_territorial_scope**, the node follows that states the GDPR is applicable to their case.

Figure A1 shows the created decision tree in GraphDoc.

Despite a user-friendly interface, challenges remain. During our quest to simplify the analysis of when a EU legislation is applicable, we noted the at times difficult formulations within a legislation. These unclarities within the scopes lead to potentially unwanted uncertainties. For instance, the material scope of the GDPR is already open-textured due to the nature of the term “personal data” and technical developments, such as the transient processing of personal data (George, Reutimann, & Tamò-Larrieux, 2019). This means that while the material scope at a given time can be established with reliance on current case law, a definitive answer as to its extent cannot be given (Purtova, 2018). Overcoming such uncertainties—to the extent that they are unwanted—would require changing the design of regulations. Some have proposed to address this by rewriting (manually and automatically) legal text into Logical English (e.g., Kowalski et al., 2022). Others, such as policymakers within the EU and especially Denmark, have proposed to think about ways to create new digitally ready legislations that are clear and more easily machine-readable. Our decision trees might help to direct the attention of policymakers to specific issues within the scopes of the analyzed legislation to provide further guidance on how to test for the application of a given law. More generally, our observations and experiences beg the question how technology (in our case: GraphDoc) can or should reflect possible uncertainties when visualizing or formalizing the law.

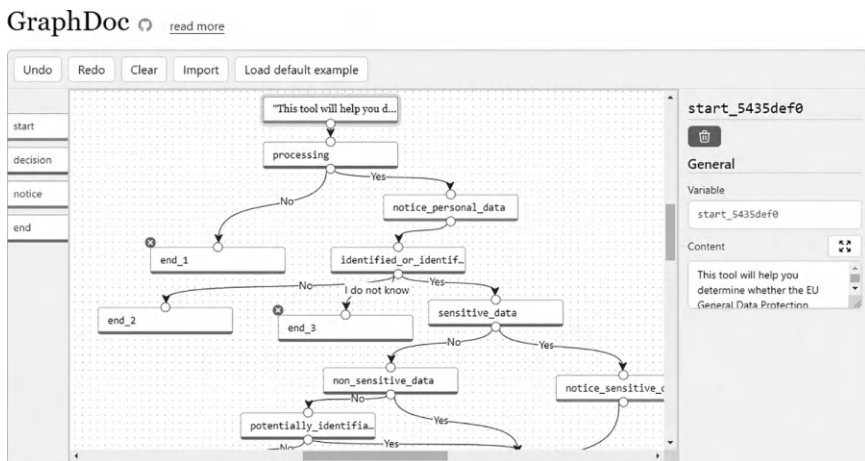


FIGURE A1 Associated GraphDoc to the legal tree.

EXERCISE 7.2: TURNING NATURAL LANGUAGE
INTO CONTROLLED LANGUAGE

The solution will depend on the exact law chosen. Below, we reproduce a table with more examples which might be useful in finding out how to go about the translation. Please note that the table is from Guitton, Tamo-Larrieux, Mayer, & Zumbrunn, et al. (2025) (Table A1).

TABLE A1 Examples of Laws in Normal and Controlled English (ACE)

Law	Article in Natural Language	Article in Attempto Controlled English
Art 1 SR 141.0, Federal Act on Swiss Citizenship	1. The following persons are Swiss citizens from birth: a. child whose parents are married to each other and whose father or mother is a Swiss citizen;	1. a. If the parents of a child are married and the father is Swiss then the child is Swiss. If the parents of a child are married and the mother is Swiss then the child is Swiss.
	b. The child of a female Swiss citizen who is not married to the child's father.	b. If the Swiss mother of a child is not married to the child's father then the child is Swiss.
	2. The minor foreign child of a Swiss father who is not married to the child's mother acquires Swiss citizenship as if at birth on establishing filiation with the father.	2. If a child is minor and is foreign and the Swiss father of the child is not married to the child's mother and the child's filiation with the father is established then the child is Swiss.
Art 19 SR 235.1, Data protection law	3. If a minor child who acquires Swiss citizenship under paragraph 2 has children, they also acquire Swiss citizenship.	3. If a minor is Swiss and has a child then the child is Swiss.
	1. The controller informs the data subject appropriately about the collection of personal data; such duty of information also applies when data is not collected from the data subject.	1. If the controller collects the data and the data is personal then the controller must inform the person about the collection. If the controller does not collect the personal data from the person then the controller must inform the person about the collection.

(Continued)

TABLE A1 (Continued) Examples of Laws in Normal and Controlled English (ACE)

Law	Article in Natural Language	Article in Attempted Controlled English
Art. 12 SR 311.0, Criminal law	<p>2. At the time of collection the controller shall provide to the data subject all information which is required in order for the data subject to assert his rights according to this Act and to ensure transparent processing of data, in particular:</p> <p>a. The controller's identity and contact information;</p> <p>b. The purpose of processing;</p> <p>c. If applicable, the recipients or the categories of recipients to which personal data is disclosed.</p>	<p>2. If the controller starts the collection of the data then the controller must immediately provide the information to the individual.</p> <p>If the individual has the information then the individual can assert the individual's rights under the Act.</p> <p>If the individual has the information then "data processing" is transparent.</p> <p>a. The information must include the controller's identity.</p> <p>b. The information must include the purpose of the n:processing.</p> <p>c. If the circumstances apply then the information must include the recipient of the disclosed data or the information must include the recipient's categories of the disclosed data.</p>
	<p>3. If data is not collected from the data subject, it additionally informs the data subject of the categories of personal data which is processed.</p>	<p>3. If the controller does not collect the personal data from the person then the controller must additionally inform the person about the categories of the personal data.</p>
	<p>1. Unless the law expressly provides otherwise, a person is only liable to prosecution for a felony or misdemeanor if he commits it willfully.</p>	<p>1. If a person commits a felony willfully and the law does not expressly allow it then the person can be prosecuted.</p> <p>If a person commits a misdemeanor willfully and the law does not expressly allow it then the person can be prosecuted.</p>

(Continued)

TABLE A1 (Continued) Examples of Laws in Normal and Controlled English (ACE)

Law	Article in Natural Language	Article in Attempto Controlled English
	2. A person commits a felony or misdemeanor willfully if he carries out the act in the knowledge of what he is doing and in accordance with his will. A person acts willfully as soon as he regards the realization of the act as being possible and accepts this.	2. If a person carries out an act and the person wants to achieve the act then the person acts willfully. If the act is a felony then it is a willful felony. If the act is a misdemeanor then the act is a willful misdemeanor. If a person believes that an act is possible and accepts the possibility of the act then the person acts willfully.
	3. A person commits a felony or misdemeanor through negligence if he fails to consider or disregards the consequences of his conduct due to a culpable lack of care. A lack of care is culpable if the person fails to exercise the care that is incumbent on him in the circumstances and commensurate with his personal capabilities.	3. If a person acts carelessly, and the person fails to consider the consequences of the conduct or disregards the consequences then the person acts negligently. If the act is a felony then the felony is committed negligently. If the act is a misdemeanor then the misdemeanor is committed negligently. If a person does not exercise the care that is incumbent to the circumstances and that is commensurate to the person's capabilities then the lack of the care is culpable.

EXERCISE 7.3: MODELING A RULE

We carry out an Entity Relationship modelling, by starting with the first question about the game's *core concepts* and subsequently defining the relationships among them.

The **core concepts** of the game can be listed as shown here:

Elements that we encounter in a game of Rock, Paper, Scissors: Player A, Player B, Hand Sign, Rock, Paper, Scissors, Game, Outcome, Win/Loss/Draw

We assign Rock, Paper, and Scissors as instances of the class Hand Sign.

We assign can Player A and Player B as instances of the class Player.

A specific game is an instance of the class Game. An instance of a game is assigned an outcome (from the perspective of Player A).

We hence then have three concepts or classes: *Hand Sign*, *Player*, and *Game*

Next, we define the **core relations** among these concepts:

A specific *Game* has exactly two *Players* who participate. This can be formalized as a relation that has the domain *Player* and the range *Game* and that we call “participates-in”. An example member of this relation in a specific game may be (*Player A*, *participates-in*, *Game 1*).

In the game, players are required to show a *Hand Sign* (Rock, Paper, Scissors), we can formalize this relation between the concepts *Player* and *Hand Sign* as a relation that has the domain *Player* and the range *Hand Sign*. We call this relation “shows”. An example member of the relation in a specific game may be (*Player A*, *shows*, *Scissors*).

The rules of the game specify an additional relation between *Hand Signs*, where Rock beats Scissors, Paper beats Rock, and Scissors beats Paper. We formalize this as a relation that has the domain *Hand Sign* and the range *Hand Sign*. We call this relation “beats”. An example member of the relation, according to the game’s rules, is (*Rock*, *beats*, *Scissors*).

Finally, a specific *Game* has a winner. This winner can be determined by (automatically) applying the defined rules (i.e., the relations above). We formalize this as a relation that has the domain *Player* and the range *Game*, and call it “has-won”. An example member of the relation in a specific game is (*Player B*, *has-won*, *Game 1*).

In addition to these relations, further properties of the individual concepts might be specified. For instance, the players might a “name” property, and the game might have a “duration” property.

EXERCISE 7.4: CLASSIFYING AUTOMATICALLY PROCESSABLE REGULATION PROJECTS

The classification of automatically processable regulation requires gathering a lot of information to unveil the aims of projects, the involved parties, and methods used to develop the automatically processable regulation.

Within our article “A Typology of Automatically Processable Regulation” we showcase ten different projects. We gathered information about these projects through analyzing publically available information and interviews. Please check out our paper Guitton, Tamò-Larrieux, & Mayer (2022b) if you want to read more about our classification of these projects.

EXERCISE 7.5: IDENTIFYING OPEN-TEXTURED TERMS

There won't be a correct or incorrect answer for this exercise. As will become apparent to you when conducting the exercise, there is a lot of disagreement on what term is open-texture. We invite readers to consult the original paper (Guitton, Tamò-Larrieux, Mayer, & Djick, 2024) and other research cited within our work that has explored open-texture in law.

EXERCISE 7.6: DEBATING ABOUT ISSUES OF AUTOMATING LEGAL PROCESSES

For potential support, we refer the readers to Guitton, Tamò-Larrieux, and Mayer (2022a).

References

- Abbate, J. (2000). *Inventing the Internet*. Cambridge, MA: The MIT Press.
- Abramson, L. (2013). *FISA Court: We Approve 99 Percent of Wiretap Applications*. Washington, DC: NPR.
- Abzianidze, L., Bos, J., & Open, S. (2020). DRS at MRP 2020: Dressing up discourse representation structures as graphs. In *Proceedings of the CoNLL 2020 Shared Task: Cross-Framework Meaning Representation Parsing* (pp. 23–32), Online. Bangkok, Thailand: Association for Computational Linguistics.
- Agarwal, S., Steyskal, S., Antunovic, F., & Kirrane, S. (2018). *Legislative Compliance Assessment: Framework, Model and GDPR Instantiation*. Cham: Springer.
- Agency for Digitisation (2018). Seven principles for digital-ready legislation. <https://en.digst.dk/digital-transformation/digital-ready-legislation/guidances-and-tools/seven-principles-for-digital-ready-legislation/>.
- Agency for Digital Government. (2022, March 9). *EU Grant Allows Denmark to Develop Method to Further Digital-Ready Legislation*. Denmark: Agency for Digital Government. en.digst.dk.
- Agency for Digitisation. (2021, May). *Evaluation of the Effort to Make Legislation Digital-Ready*. Denmark: Ministry of Finance, Agency for Digitisation.
- Aidid, A., & Alarie, B. (2023). *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better*. Toronto: University of Toronto Press.
- Alauzen, M. (2021). Splendeurs et misères d’une start-up d’Etat: Dispute dans la lutte contre le non-recours aux droits sociaux en France (2013–2020). *Réseaux*, 225, 121–150.
- Allen, L. E., & Engholm, C. R. (1978). Normalized legal drafting and the query method. *Journal of Legal Education*, 29, 380–412.
- Amnesty International. (2021, October 25). *Xenophobic Machines: Discrimination through Unregulated Use of Algorithms in the Dutch Childcare Benefits Scandal*. London: Amnesty International.
- Ash, E., Kesari, A., Naidu, S., Song, L., & Stammbach, D. (2024). Translating legalese: enhancing public understanding of court opinions with legal summarizers. *Paper presented at the Proceedings of the Symposium on Computer Science and Law*, Boston, MA. <https://doi.org/10.1145/3614407.3643700>
- Ashley, K. D. (2017). *Artificial Intelligence and Legal Analytics: New Tools for Law Practice in the Digital Age*. Cambridge, MA: Cambridge University Press.
- Baldwin, R., Cave, M., & Lodge, M. (2011). *Understanding Regulation: Theory, Strategy, and Practice*. Oxford: Oxford University Press.

- Bandes, S. A. (2000). When victims seek closure: forgiveness, vengeance and the role of government. *Fordham Urban Law Journal*, 27(5), 1599–1606.
- Barocas, S., Hardt, M., & Narayanan, A. (2023). *Fairness and Machine Learning: Limitations and Opportunities*. Boston, MA: The MIT Press.
- Barysè, D., & Sarel, R. (2024). Algorithms in the court: does it matter which part of the judicial decision-making is automated? *Artificial Intelligence and Law*, 32(1), 117–146. <https://doi.org/10.1007/s10506-022-09343-6>
- Bátiz-Lazo, B. (2015, March 26). A brief history of the ATM: How automation changed retail banking, an object lesson. *The Atlantic*.
- Bazzi, F. (2021, April 12). Legal design and the future of law. *Columbia Journal of Transnational Law*.
- Becker, D., & Ferrari, I. (2020). VICTOR, the Brazilian Supreme Court's Artificial Intelligence: a beauty or a beast? Retrieved from https://nja.gov.in/Concluded_Programmes/2021-22/SE-09_2021_Speeches_and_Write-ups/3.%20Victor-Beauty-or-the-Beast.pdf [Accessed on March 2024]
- Ben-Shahar, O., & Porat, A. (2021). *Personalized Law: Different Rules for Different People*. New York: Oxford University Press.
- Bench-Capon, T. J. M., & Sergot, M. J. (1985). Towards a rule-based representation of open texture in law. In C. Walter (Ed.), *Computing Power and Legal Language* (pp. 39–60). New York: Greenwood/Quorum Press.
- Benezech, L. (2020). L'exigence d'accessibilité et d'intelligibilité de la loi: retour sur vingt ans d'existence. *Revue française de droit constitutionnel*, 3(123), 541–563.
- Benzmüller, C., Fuenmayor, D., & Lomfeld, B. (2020). encoding legal balancing: automating an abstract ethico-legal value ontology in preference logic. *Freie Universität Empirical Legal Studies Center (FUELS) Working Paper*, 5.
- Bessen, J. (2015). *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*. New Haven, CT: Yale University Press.
- Bhuiyan, H., Governatori, G., Bond, A., Demmel, S., Islam, M. B., & Rakotonirainy, A. (2020). Traffic rules encoding using defeasible deontic logic. *Paper presented at the Legal Knowledge and Information Systems*, Brisbane, Australia.
- Blank, J. D., & Osofsky, L. (2020). Automated legal guidance. *Cornell Law Review*, 106, 179–246.
- Boley, H., Paschke, A., & Shafiq, O. (2010). Rule ML 1.0: The overarching specification of web rules. In M. Dean, J. Hall, A. Rotolo, & S. Tabet (Eds.), *Semantic Web Rules. RuleML 2010* (Vol. 6403, pp. 162–178). Berlin: Springer.
- Bourcier, D., & Mazzega, P. (2007). Toward measures of complexity in legal systems. *Paper presented at the Eleventh International Conference on Artificial Intelligence and Law*, New York.
- Brand, S. (2001). Founding Father. *Wired, March*. <https://www.wired.com/2001/03/baran/>
- Brooke, A. (1992). When ignorance of the law became an excuse: Lambert & its progeny. *American Journal of Criminal Law*, 19(2), 279–312.
- Buchholtz, G. (2019). Artificial intelligence and legal tech: Challenges to the rule of law. In T. Wischmeyer & T. Rademacher (Eds.), *Regulating Artificial Intelligence* (pp. 175–198). Cham: Springer.

- Busch, C., & Franceschi, A. D. (Eds.). (2020). *Algorithmic Regulation and Personalized Law: A Handbook*. Munich: Beck/Hart Publishing.
- Calbimonte, J.-P., Ciortea, A., Kampik, T., Mayer, S., Payne, T. R., Tamma, V., & Zimmermann, A. (2023). Autonomy in the age of knowledge graphs: Vision and challenges. *Transactions on Graph Data & Knowledge*, 1(1), 22.
- Casey, A. J., & Niblett, A. (2019). A framework for the new personalization of law. *The University of Chicago Law Review*, 86(2), 333–358.
- Cavoukian, A. (2009). Privacy by design: The 7 foundational principles. *Information and privacy commissioner of Ontario, Canada*, 5, 12.
- Cerullo, M. (2023, January 26). AI-powered “robot” lawyer won’t argue in court after jail threats. *CBS News*.
- Chen, B., Stremitzer, A., & Tobia, K. (2021). Having your day in robot court. *ETHZ Working Papers Series*, 05/2021.
- Chopra, A., Torre, L. & Verhagen, H. (Eds.). (2018). *Handbook of Normative Multiagent Systems*. London: College Publications.
- Clocksin, W. F., & Mellish, C. S. (2003). *Programming in PROLOG*. Springer Science & Business Media, Berlin, Heidelberg.
- Cobbe, J. (2020). Legal singularity and the reflexivity of law. In S. Deakin & C. Markou (Eds.), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (pp. 1–30). Oxford: Hart Publishing.
- Conselho Nacional de Justiça. (2019). *Inteligência Artificial: No Poder Judiciário Brasileiro*. Retrieved from https://www.cnj.jus.br/wp-content/uploads/2020/05/Inteligencia_artificial_no_poder_judiciario_brasileiro_2019-11-22.pdf
- Contissa, G., & Sartor, G. (2022). How the law has become computable. In G. Contissa, G. Lasagni, M. Caianiello, & G. Sartor (Eds.), *Effective Protection of the Rights of the Accused in the EU Directives: A Computable Approach to Criminal Procedure Law* (pp. 27–41). Leiden: Brill/Nijhoff.
- Custers, B. (2023). The right to break the law? perfect enforcement of the law using technology impedes the development of legal systems. *Ethics and Information Technology*, 25(4), 58. <https://doi.org/10.1007/s10676-023-09737-3>
- Danaher, J. (2016). The threat of algocracy: Reality, resistance and accommodation. *Philosophy & Technology*, 29(3), 245–268.
- Danish Agency for Digitisation. (2021). Seven principles for digital-ready legislation. Retrieved from <https://en.digst.dk/policy-and-strategy/digital-ready-legislation/guidances-and-tools/seven-principles-for-digital-ready-legislation/>
- Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011a). Extraneous factors in judicial decisions. *PNAS*, 108(17), 6889–6892.
- Danziger, S., Levav, J., & Avnaim-Pesso, L. (2011b). Reply to Weinshall-Margel and Shapard: Extraneous factors in judicial decisions persist. *Proceedings of the National Academy of Sciences of the United States of America*, 108(42), E834.
- Deakin, S., & Markou, C. (2020a). *Is Law Computable?* London: Hart Publishing.
- Deakin, S., & Markou, C. (2020b). *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence*. London: Hart Publishing.
- DeNardis, L. (2007). A history of internet security. In D. L. Bergstra (Ed.), *The History of Information Security* (pp. 595–621). Amsterdam: Elsevier Science.

- Denvir, C., Balmer, N. J., & Pleasence, P. (2013). When legal rights are not a reality: do individuals know their rights and how can we tell? *Journal of Social Welfare and Family Law*, 35(1), 139–160. <https://doi.org/10.1080/09649069.2013.774764>
- Digital Government NZ. (2018). *Better Rules for Government Discovery Report*. Retrieved from <https://www.digital.govt.nz/dmsdocument/95-better-zrules-for-government-discovery-report/html>
- Dijck, G. v. (2022). Predicting recidivism risk meets AI act. *European Journal on Criminal Policy and Research*, 28, 407–423.
- Diver, L. (2020). Digisprudence: the design of legitimate code. *Law, Innovation & Technology*, 13(2), 325–354.
- DutchNews. (2024, March 1). “Student finance group Duo did discriminate in fraud probes”. *DutchNews*.
- Edelkamp, S., & Schrödl, S. (2012). *Heuristic Search: Theory and Applications*. Amsterdam: Elsevier.
- Egan, P., & Roberts, A. (2021, March 26). *Judge: Companies can be used Over Michigan Unemployment Fraud Fiasco*. Detroit, MI: Detroit Free Press.
- ENISA. (2022). *ENISA Cybersecurity Market Analysis Framework (ECSMAF)*. Athens, Greece: ENISA:
- Escher, N., Bilik, J., Miller, A., Huseby, J. J., Ramesh, D., Liu, A., & Banovic, N. (2022). Cod(e)ifying The Law. *Paper presented at the Programming Languages and the Law (ProLaLa)*, Philadelphia.
- European Commission for Democracy through Law (Venice Commission). (2020). *Opinion No.1031/2021 Report of the Childcare Allowance Parliamentary Inquiry Committee*. Retrieved from Strasbourg: https://www.nationaleombudsman.nl/system/files/onderzoek/Rapport%202017-095%20Geen%20powerplay%20maar%20fair%20play_0.pdf
- European Data Protection Board. (2024). Data controller or data processor. Retrieved from https://www.edpb.europa.eu/sme-data-protection-guide/data-controller-data-processor_en
- European Union. (2024). Digital-ready policymaking. Retrieved from <https://academy.europa.eu/courses/digital-ready-policymaking>
- Federal Trade Commission. (2012). Protecting consumer privacy in an era of rapid change: Recommendations for business and policymakers. Retrieved from <http://www.ftc.gov/os/2010/12/101201privacyreport.pdf>
- Filtz, E., Kirrane, S., & Polleres, A. (2021). The linked legal data landscape: linking legal data across different countries. *Artificial Intelligence and Law*, 29(4), 485–539. <https://doi.org/10.1007/s10506-021-09282-8>
- Fjeld, J., Achten, N., Hillgoss, H., Nagy, A., & Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to principles for AI. *Berkman Klein Center Research Publication*, 2020-1.
- Flinders, K. (2019, January 15). Thousands of known errors on controversial post office computer system to be revealed. *ComputerWeekly*.
- Flood, G. (2023). National Archives democratizes access to court data with NoSQL and metadata. *diginomica*.

- Floridi, L., Cowsls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28, 689–707.
- Forstmoser, P., & Vogt, H.-U. (2012). *Einführung in das Recht*. Bern: Stämpfli.
- Fraser, H. (2021, December 7). What is Better Rules? *Digital.govt.nz*.
- FTC. (2024, September 25). FTC announces crackdown on deceptive AI claims and schemes. *Federal Trade Commission*.
- Fuchs, N. E., Höfler, S., Kaljurand, K., Rinaldi, F., & Schneider, G. (2005). Attempto controlled English: A knowledge representation language readable by humans and machines. *Paper Presented at the Reasoning Web*, Msida, Malta.
- Fuchs, N. E., & Schwitter, R. (1996, March). Attempto controlled English (ACE). *Paper presented at the CLAW 96, First International Workshop on Controlled Language Applications*, University of Leuven, Belgium.
- Gallie, W. B. (1956). Essentially contested concepts. *Proceedings of the Aristotelian Society*, 56, 167–198.
- García, K., Zihlmann, Z., Mayer, S., Tamò-Larrieux, A., & Hooss, J. (2021). Towards privacy-friendly smart products. *Paper presented at the 18th International Conference on Privacy, Security and Trust (PST)*, Auckland, New Zealand.
- Gasser, U. (2016). Recoding privacy law: Reflections on the future relationship among law, technology, and privacy. *Harvard Law Review*, 130, 61.
- Genesereth, M. (2015). *Computational Law: The Cop in the Backseat*. Retrieved from theconversation.com
- George, D., Reutimann, K., & Tamò-Larrieux, A. (2019). GDPR bypass by design? Transient processing of data under the GDPR. *International Data Privacy Law*, 9(4), 285–298.
- Gibbs, S. (2014, July 2). Facebook apologises for psychological experiments on users. *The Guardian*.
- Ginsburg, T., & Moustafa, T. (2008). *Rule by Law: The Politics of Courts in Authoritarian Regimes*. Cambridge: Cambridge University Press.
- Glöckner, A. (2016). The irrational hungry judge effect revisited: Simulations reveal that the magnitude of the effect is overestimated. *Judgment and Decision Making*, 11(6), 601–610. <https://doi.org/10.1017/S1930297500004812>
- Governatori, G., Rotolo, A., & Sartor, G. (2005). Temporalised normative positions in defeasible logic. *Paper presented at the Proceedings of the 10th International Conference on Artificial Intelligence and Law*, Bologna, Italy. <https://doi.org/10.1145/1165485.1165490>
- Greenberg, A. (2015, May 7). Court rules NSA bulk data collection was never authorized by congress. *Wired*.
- Guitton, C., Gubelmann, R., Karray, Mayer, S., & Tamò-Larrieux, A. (Forthcoming). Identifying Open-Texture in Regulations Using LLMs.
- Guitton, C., Mayer, S., Tamò-Larrieux, A., Fosch-Villaronga, E., Kamara, I., & Palka, P. (2024). Responsible automatically process able regulation. *AI & Society*, 2024, 1–16.

- Guittou, C., Mayer, S., Tamò-Larrieux, A., Garcia, K., & Fornara, N. (2024, Mar 12–13). A proxy for assessing the automatic encodability of regulation. *Paper presented at the 3rd ACM Computer Science and Law Symposium*, Boston, MA.
- Guittou, C., Tamò-Larrieux, A., & Mayer, S. (2022a). Mapping the issues of automated legal systems: Why worry about automatically processable regulation? *Artificial Intelligence and Law*, 31, 571–599.
- Guittou, C., Tamò-Larrieux, A., & Mayer, S. (2022b). A typology of automatically processable regulation. *Law, Innovation & Technology*, 14(2), 267–304.
- Guittou, C., Tamò-Larrieux, A., & Mayer, S. (2024). How distrust is driving artificial intelligence regulation in the European Union. *European Journal of Law and Technology*, 15, 2.
- Guittou, C., Tamò-Larrieux, A., Mayer, S., & Djick, G. V. (2024). The challenge of open-texture in law. *Artificial Intelligence and Law*, 2024, 1–32.
- Guittou, C., Tamò-Larrieux, A., Mayer, S., & Zumbunn, L. (2025). Controlled language increases comprehension of law for people. *ACM Journal on Computing and Sustainable Societies*.
- Hadfield, G. K. (2000). The price of law: How the market for lawyers distorts the justice system. *Michigan Law Review*, 98, 953.
- Hagan, M. (2015). *Law by Design*. <https://lawbydesign.co/>.
- Hage, J. (2003). Law and defeasibility. *Artificial Intelligence and Law*, 11(2), 221–243. <https://doi.org/10.1023/B:ARTI.0000046011.13621.08>
- Hagendorff, T. (2018). Privacy literacy and its problems. *Journal of Information Ethics*, 27(2), 127.
- Harris, W. V. (1989). *Ancient Literacy*. Cambridge, MA: Harvard University Press.
- Hart, H. L. A. (1994). *The Concept of Law*. Oxford: Clarendon Press.
- Hassenzahl, M. (2013). User experience and experience design. In Interaction Design Foundation (ed.), *The Encyclopedia of Human-Computer Interaction*. Aarhus: Interaction-Design.org Foundation.
- Heikkilä, M. (2022, March 29). *Dutch Scandal Serves as a Warning for Europe over Risks of Using Algorithms*. Politico.
- Henley, J. (2021, January 15). Dutch government resigns over child benefits scandal. *The Guardian*.
- Henseler, S., & Tamò-Larrieux, A. (2022). Research handbook on EU data protection law. In E. Kosta, R. Leenes, & I. Kamara (Eds.), *Chapter 12: Reaching Beyond Its Territory - An Analysis of the Extraterritorial Scope of European Data Protection Law*. Cheltenham: Edward Elgar Publishing.
- Herzog, T. (2018). *A Short History of European Law: The Last Two and a Half Millennia*. Cambridge, MA: Harvard University Press.
- Higgison, E. (2020, August 18). The rise of Innovation Labs in government – part 1. *Digital.govt.nz*.
- Hildebrandt, M. (2015). *Smart Technologies and the End(s) of Law*. Cheltenham: Edward Elgar Publishing.
- Hildebrandt, M. (2018). Law as computation in the era of artificial legal intelligence: Speaking law to the power of statistics. *University of Toronto Law Journal*, 68, 12–35.

- Hildebrandt, M. (2020). Code-driven law: Freezing the future and scaling the past. In S. Deakin & C. Markou (Eds.), *Is Law Computable? Critical Perspectives on Law and Artificial Intelligence* (pp. 67–84). Oxford: Hart Publishing.
- Hoepman, J.-H. (2021). *Privacy Is Hard and Seven Other Myths: Achieving Privacy through Careful Design*. Boston, MA: The MIT Press.
- Hoffmann, C. P., Lutz, C., & Ranzini, G. (2016). Privacy cynicism: A new approach to the privacy paradox. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 10(4), 7. <https://doi.org/10.5817/CP2016-4-7>
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutierrez, C., & Zimmermann, A. (2021). *Knowledge Graphs*. Berlin: Springer.
- Hogle, C., & Abdo, A. (2021, April 20). *The Public Should Have Access to the Surveillance Court's Opinions*. New York: ACLU.
- House of Representatives of the States General. (2020). *Childcare Allowance Parliamentary Inquiry*. Retrieved from The Hague: https://www.houseofrepresentatives.nl/sites/default/files/atoms/files/verslag_pok_definitief-en-gb.docx.pdf
- Huber, J. D., & Shipan, C. R. (2012). *Deliberate Discretion? The Institutional Foundations of Bureaucratic Autonomy*. Cambridge: Cambridge University Press.
- Hulst, L., van den Bos, K., Akkermans, A. J., & Lind, E. A. (2017). On why procedural justice matters in court hearings. *Utrecht Law Review*, 13(3), 114–129.
- Justice for the Subpostmasters Alliance. (2019). *An Introduction to the Judgments of the Common Issues and the Horizon Issues trials*. Retrieved from https://www.jfsa.org.uk/uploads/5/4/3/1/54312921/an_introduction_to_the_ci_and_hi_judgments.pdf
- Katz, D., & Bommarito, M. (2014). Measuring the complexity of the law: The United States code. *Artificial Intelligence and Law*, 22, 337–374.
- Katz, D. M. (2023, Mar 15). GPT-4 Passes the Bar Exam. *Illinois Tech*. Retrieved from <https://www.iit.edu/news/gpt-4-passes-bar-exam>
- Katz, D. M., Dolin, R., & Bommarito, M. J. (2021). *Legal Informatics*. Cambridge, MA: Cambridge University Press.
- Kirrane, S., Fernández, J. D., Dullaert, W., Milosevic, U., Polleres, A., Bonatti, P. A., & Raschke, P. (2018). *A Scalable Consent, Transparency and Compliance Architecture*. Cham: Springer.
- Klitou, D. (2014). *A Solution, But Not a Panacea for Defending Privacy: The Challenges, Criticism and Limitations of Privacy by Design*. Berlin, Heidelberg, Springer.
- Kowalski, R., Dávila, J., Sartor, G., & Calejo, M. (2022). Logical English for law. *Paper Presented at the Proceedings of the International Workshop on Methodologies for Translating Legal Norms into Formal Representations (LN2FR 2022)*.
- Kowalski, R. A. (1982). Logic as a computer language for children. *Paper Presented at the Proceedings of the 5th European Conference on Artificial Intelligence*, Spain.

- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(24), 8788–8790. <https://doi.org/https://doi.org/10.1073/pnas.1320040111>
- Kris, D. (2018, September 2). How the FISA court really works. *Lawfare*.
- Kroll, J. A. (2015). *Accountable Algorithms*. Princeton, NJ: Princeton University, ProQuest Dissertations Publishing.
- Kuhn, T. (2012). *Timeline of Controlled English Languages*. Retrieved from Zürich: <https://attempto.ifi.uzh.ch/site/cnl2012/posters/cnltimeline-poster.pdf>
- La Quadrature du Net (2024, 16 Oct). French family welfare scoring algorithm challenged in court by 15 organisations. *La Quadrature du Net*. <https://www.laquadrature.net/en/2024/10/16/french-family-welfare-scoring-algorithm-challenged-in-court-by-15-organisations/>
- LabPlus. (2018). *Better Rules for Government: Discovery Report*.
- Lee, J., Hwangbo, J., Wellhausen, L., Koltun, V., & Hutter, M. (2020). Learning quadrupedal locomotion over challenging terrain. *Science Robotics*, 5(47), eabc5986. <https://doi.org/https://doi.org/10.1126/scirobotics.abc5986>
- Lee, M. K. (2018). Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society*, 5(1), 2053951718756684. <https://doi.org/10.1177/2053951718756684>
- Leslie, D. (2019). Understanding artificial intelligence ethics and safety: A guide for the responsible design and implementation of AI systems in the public sector. *The Alan Turing Institute*. https://www.turing.ac.uk/sites/default/files/2019-08/understanding_artificial_intelligence_ethics_and_safety.pdf
- Lessig, L. (2003). Law regulating code regulating law. *Loyola University Chicago Law Journal*, 35, 1.
- Lippi, M., Pałka, P., Contissa, G., Lagioia, F., Micklitz, H.-W., Sartor, G., & Torroni, P. (2019). CLAUDETTE: an automated detector of potentially unfair clauses in online terms of service. *Artificial Intelligence and Law*, 27(2), 117–139. <https://doi.org/10.1007/s10506-019-09243-2>
- Loi, M. (2020). *People Analytics must benefit the people. An ethical analysis of data-driven algorithmic systems in human resources management*.
- Lukasik, S. (2010). Why the ARPANET was built. *IEEE Annals of the History of Computing*, 33(3), 4–21.
- Martínez, E. (2024). Re-evaluating GPT-4's bar exam performance. *Artificial Intelligence and Law*, 2024, 1–24. <https://doi.org/10.1007/s10506-024-09396-9>
- Martínez, E., Mollicab, F., & Gibsons, E. (2023). Even lawyers do not like legalese. *PNAS*, 120, 23.
- McMaster, C. (2019). Is the sky falling for the Canadian artificial intelligence industry? *Intellectual Property Journal*, 32(1), 77–103. Retrieved from <https://www.proquest.com/scholarly-journals/is-sky-falling-canadian-artificial-intelligence/docview/2321846038/se-2?accountid=28962>
- Merigoux, D., Alauzen, M., & Slimani, L. (2024). Journal of cross-disciplinary research in computational law. *Journal of Cross-disciplinary Research in Computational Law*, 2, 1.

- Merigoux, D., Chataing, N., & Protzenko, J. (2021). Catala: A programming language for the law. *Proceedings of the ACM on Programming Languages (ICFP)*, 5(77), 1–29.
- Merigoux, D., & Huttner, L. (2020). Catala: Moving Towards the Future of Legal Expert Systems. *HAL*, hal-02936606, <https://api.semanticscholar.org/CorpusID:225231457>.
- Merken, S. (2023a, March 9). Lawsuit pits class action firm against ‘robot lawyer’ DoNotPay. *Reuters*.
- Merken, S. (2023b, November 18). ‘Robot lawyer’ DoNotPay beats lawsuit by Illinois law firm. *Reuters*.
- Miller, C. (2022). *Chip War: The Fight for the World’s Most Critical Technology*. Simon and Schuster, New York.
- Milner, H. (2002). *Civic Literacy: How Informed Citizens Make Democracy Work*. Hanover: University Press of New England.
- Ministry of Justice. (2022). Court judgments made accessible to all at The National Archives. *Gov.uk*.
- Mohun, J., & Roberts, A. (2020). *Cracking the Code: Rulemaking for Humans and Machines*. Paris: OECD.
- Naughton, J. (2000). *A Brief History of the Future: The Origins of the Internet*. London: Phoenix.
- Newman, D. T., Fast, N. J., & Harmon, D. J. (2020). When eliminating bias isn’t fair: Algorithmic reductionism and procedural justice in human resource decisions. *Organizational Behavior and Human Decision Processes*, 160, 149–167. <https://doi.org/10.1016/j.obhdp.2020.03.008>
- Newnham, D. (2021). Why the social dilemma is wrong: And could actually cause more harm than good, *Medium*, Aug 27. <https://daniellenewnham.medium.com/why-the-social-dilemma-is-wrong-17d8b2952187>.
- Niiler, E. (2019, Mar 25). Can AI Be a Fair Judge in Court? Estonia Thinks So. *Wired*.
- Newlands, G., Lutz, C., Tamò-Larrieux, A., Villaronga, E. F., Harasgama, R., & Scheitlin, G. (2020). Innovation under pressure: Implications for data privacy during the Covid-19 pandemic. *Big Data & Society*, 7(2), 2053951720976680.
- Obar, J. A., & Oeldorf-Hirsch, A. (2018). The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication & Society*, 23(1), 128–147.
- OECD. (2013). *The OECD Privacy Framework*. Paris: OECD.
- Palmirani, M., & Cervone, L. (2013). Measuring the complexity of law over time. *Paper Presented at the AI Approaches to the Complexity of Legal Systems: AICOL 2013 International Workshops*.
- Pariser, E. (2011). *The Filter Bubble: What the Internet Is Hiding from You*. London: Penguin Books.
- Pasquale, F. (2022). The resilient fragility of law. In S. Deakin & C. Markou (Eds.), *Is Law Computable?: Critical Perspectives on Law and Artificial Intelligence* (pp. v–xvi). London: Hart Publishing.
- Pleasence, P., Balmer, N. J., & Denvir, C. (2015). *How People Understand and Interact with the Law*. Cambridge: The Legal Education Foundation (June).

- Plesner, U., & Justesen, L. (2022). The double darkness of digitalization: shaping digital-ready legislation to reshape the conditions for public-sector digitalization. *Science, Technology, & Human Values*, 47, 1.
- PNAS. (2014). Editorial expression of concern: Experimental evidence of massive scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences*, 111(29), 10779–10779.
- Prem, E. (2023). From ethical AI frameworks to tools: a review of approaches. *AI and Ethics*, 3(3), 699–716. <https://doi.org/10.1007/s43681-023-00258-9>
- Pitron, G. (2021). *The Rare Metal Wars: The Dark Side of Clean Energy and Digital Technologies*. London: Scribe UK.
- Pitron, G. (2023). *The Dark Cloud: The Hidden Costs of the Digital World*. London: Scribe US.
- Purtova, N. (2018). The law of everything. Broad concept of personal data and future of EU data protection law. *Law, Innovation and Technology*, 10(1), 40–81.
- Ranchordas, S. (2021). Experimental regulations for AI: Sandboxes for morals and mores. *Morals+Machines*, 1, 86–100.
- Ranchordàs, S. (2022). Empathy in the digital administrative state. *Duke Law Journal*, 71, 1341.
- Ranchordàs, S., & Scarcella, L. (2021). Automated government for vulnerable citizens: Intermediating rights. *William & Mary Bill of Rights Journal*, 30, 373.
- Robaldo, L., Bartolini, C., Palmirani, M., Rossi, A., Martoni, M., & Lenzini, G. (2020). Formalizing GDPR provisions in reified I/O logic: The DAPRECO knowledge base. *Journal of Logic, Language and Information*, 29, 401–449.
- Romo, V. (2022, Sept 2). Pregnant woman who claimed her fetus was an HOV lane passenger gets another ticket. *NPR*. Retrieved from <https://www.npr.org/2022/09/02/1120628973/pregnant-woman-dallas-fetus-hov-lane-passenger-ticket>
- Roose, K. (2020). *Rabbit Hole*. Retrieved from <https://www.nytimes.com/column/rabbit-hole>
- Salvi, C., Iannello, P., Cancer, A., McClay, M., Rago, S., Dunsmoor, J. E., & Antonietti, A. (2021). Going viral: how fear, socio-cognitive polarization and problem-solving influence fake news detection and proliferation during COVID-19 pandemic. *Frontiers in Communication*, 5, 1–16.
- Satter, R. (2020, September 3). U.S. court: Mass surveillance program exposed by Snowden was illegal. *Reuters*.
- Schafer, B. (2010). ZombAIs: Legal expert systems as representatives “Beyond the Grave”. *SCRIPTed*, 7(2), 384–393.
- Schauer, F. (2013). On the open texture of law. *Grazer Philosophische Studien*, 87, 197–215.
- Schneider, T., Tamò-Larrieux, A., Mayer, S. (2022). *Privacy-Invasive Technologies - Cross-Disciplinary Perspectives, Course at the University of St. Gallen, Spring Semester 2022*
- Sennet, A. (2021). Ambiguity. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/ambiguity/>

- Sergot, M. J., Sadri, F., Kowalski, R. A., Kriwaczek, F., Hammond, P., & Cory, H. T. (1986). The British nationality act as a logic program. *Communications of the ACM*, 29(5), 370–386.
- Shay, L. A., Hartzog, W., Nelson, J., & Conti, G. (2016). Do robots dream of electric laws? An experiment in the law as algorithm. In: Ryan Calo, A. Michael Froomkin, and Ian Kerr (eds.) *Robot Law* (pp. 274–305). Edward Elgar Publishing, Cheltenham, UK.
- Sheth, A., Roy, K., & Gaur, M. (2023). *Neurosymbolic Artificial Intelligence (Why, What, and How)*. 38(03), 56–62. Retrieved from <https://doi.ieeecomputersociety.org/10.1109/MIS.2023.3268724>
- Sinclair, L. (2021, April 24). Post office worker took his own life after being wrongly accused of stealing £60K. *The Standard*.
- Singh, A. M., & Singh, M. P. (2023, August 19-25). *Norm deviation in multi-agent systems: a foundation for responsible autonomy. Paper presented at the Thirty-Second International Joint Conference on Artificial Intelligence*, Macao.
- Solove, D. J. (2011). *Nothing to Hide: The False Tradeoff between Privacy and Security*. New Haven: Yale University Press.
- Stevenson, J. (2019). *Exploring Machine Consumable Accident Compensation Legislation: Lessons for a structural rewrite of the AC Act and opportunities to make it machine consumable*.
- Sullivan, K. M. (1992). The Justices of Rules and Standards. *Harvard Law Review*, 106, 22.
- Susskind, R., & Susskind, D. (2015). *The Future of the Professions: How Technology Will Transform the Work of Human Experts*. Oxford: Oxford University Press.
- Tamò-Larrieux, A. (2018). *Designing for Privacy and its Legal Framework: Data Protection by Design and Default for the Internet of Things* (Vol. 40). Zurich: Springer.
- Tamò-Larrieux, A., Guitton, C., Mayer, S., & Lutz, C. (2023). Regulating for trust: Can law establish trust in artificial intelligence? *Regulation & Governance*.
- Tamò-Larrieux, A., Mayer, S., & Zihlmann, Z. (2021). Not Hardcoding but soft-coding data protection principles. *Technology and Regulation*, 2021(May), 17–34.
- Tapasya, S. K., & Joshi, D. (2024, January 24). How an algorithm denied food to thousands of poor in India's Telangana. *Aljazeera*.
- Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. New York: Alfred A. Knopf.
- The Engine Room. (2019). *Technology for Legal Empowerment*. Retrieved from <https://www.theengineroom.org/wp-content/uploads/2019/01/Tech-for-Legal-Empowerment-The-Engine-Room.pdf>
- Thomson, R. (2009, May 11). Bankruptcy, prosecution and disrupted livelihoods - Postmasters tell their story. *ComputerWeekly*. Retrieved from <https://www.computerweekly.com/news/2240089230/Bankruptcy-prosecution-and-disrupted-livelihoods-Postmasters-tell-their-story>
- Turban, E., Aronson, J., & Liang, T.-P. (2005). *Decision Support Systems and Intelligent Systems*. New Jersey: Pearson Prentice Hall.

- Uddin, R. (2023, December 28). How Post Office lawyers pushed boundaries in UK sub-postmasters scandal. *Financial Times*.
- UK Government. (2021). *Ethics, Transparency and Accountability Framework for Automated Decision-Making*. Retrieved from <https://www.gov.uk/government/publications/ethics-transparency-and-accountability-framework-for-automated-decision-making/ethics-transparency-and-accountability-framework-for-automated-decision-making>
- UNDP. (2023, January 10). *UNDP Publishes New Guide with Legal Design Türkiye to Improve Violence Survivors' Access to Justice*. New York: UNDP.
- Utrecht University. (2021, Nov 9). *IAMA Makes Careful Decision about Deployment of Algorithms Possible*. Netherlands: Utrecht University.
- Vasdani, T. (2020, February). Robot justice: China's use of Internet courts. *The Lawyer's Daily*.
- Veale, M., & Borgesius, F. Z. (2021). Demystifying the draft EU artificial intelligence act. *Computer Law Review International*, 22(4), 97–112.
- Vecht, J. J. (2020). Open texture clarified. *Inquiry*, 2020, 1–21.
- Victor, D. (2018, Feb 9). Oxford Comma Dispute Is Settled as Maine Drivers Get \$5 Million. *The New York Times*.
- Wachter, S., Mittelstadt, B., & Russell, C. (2018). Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2), 841–888.
- Wallis, N. (2021). *The Great Post Office Scandal: The Fight to Expose a Multimillion Pound IT Disaster which Put Innocent People in Jail*. Bath: Bath Publishing.
- Walzl, B., & Matthes, F. (2014). Towards measures of complexity: Applying structural and linguistic metrics to German laws. *Paper presented at the JURIX 2014: The Twenty-seventh Annual Conference*, German.
- Webster, N. (2018). LabPlus: Better rules for government discovery report. *Business Rules Community*, 19, 8.
- Weinshall-Margel, K., & Shapard, J. (2011). Overlooked factors in the analysis of parole decisions. *Proceedings of the National Academy of Sciences*, 108(42), E833–E833. <https://doi.org/https://doi.org/10.1073/pnas.1110910108>
- Whalen, R. (2022). Defining legal technology and its implications. *International Journal of Law and Information Technology*, 30(1), 47–67.
- Wyner, A. (2008). An ontology in owl for legal case-based reasoning. *Artificial Intelligence and Law*, 16(4), 361–387.
- Xu, H., & Ashley, K. (2023). A question-answering approach to evaluating legal summaries. In J. S. Giovanni Sileno, G. van Dijck (Eds.), *Legal Knowledge and Information Systems* (Vol. 379, pp. 293–298). Amsterdam, Netherlands: IOS Press.
- Young, A. R. (2023). *The 26 Words that Created the Internet*. Brown Undergraduate Law Review. Ithaca, NY: Cornell University Press.
- Zatz, M. S. (2013). The convergence of race, ethnicity, gender, and class on court decision making: looking toward the 21st century. In E. K. Charis & D. S. Thomas (Eds.), *Introduction to Criminal Justice* (pp. 203–218). Redwood City: Stanford University Press.

- Zhong, H., Xiao, C., Tu, C., Zhang, T., Liu, Z., & Sun, M. (2020, July 5-10). How does NLP benefit legal system: A summary of legal artificial intelligence. *Paper Presented at the 58th Annual Meeting of the Association for Computational Linguistics*, Bangkok, Thailand.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

Index

Note: **Bold** page numbers refer to tables; *italic* page numbers refer to figures and page numbers followed by “n” denote endnotes.

- A/B testing 130
- Advanced Research Projects Agency (ARPA) 126–128
- Agency for Digital Government 109–110, **110**, 112
- Agency for Digitisation of Denmark 108, 110–112
- AI Act of the EU 16, 103, 104, 116, 119, 164
- algorithmic thinking 118, 123
- ARPANET 127–128
- Attempto Controlled English (ACE) 29, 30, 31, 172, **172–174**
- Austrian Ministry for Digital and Economic Affairs 109
- automated decision-making 68, 82, 83, 88, 90, 93, 96–97, 135, 136, 153, 155, 157, 164
- automated teller machines (ATM) 92, 93
- automatically processable regulation (APR) 1, 3, 9, 12–15, 19, 25–27, 29, 30, 33, 47–48, 50, 54–55, 70, 70, 86–87, 95–97, 101, 117, 118, 120, 122, 124, 125, 129, 136, 140, 143, 162–165
- accessibility 15, 30, 42–43, 56, 59, 61–63, 72, 94, 105, 106, 108, 111
- efficiency 9, 10, 14, 42–43, 56, 59, 61–69, 91
- formalization 7, 22, 44, 150
- implementations 56, 57, 93–95
- issues 13–15, 25–27, 31, 43, 50, 52, 98, **99–101**, 160–161, 176
- legal design 15, 113–116
- ontologies 13, 18, 24, 32, 34, 37–43, 47, 48, 52, 76, 150, 152
- projects 150–151, **151**, 175–176
 - Canadian civil servants 157
 - DoNotPay 155–156
 - Mes Aides 152–155, *154*
 - Rates Rebate 71, 73, 151–152
 - Victor 155
- responsible 95–101
- sub-symbolic AI 48
- symbolic AI 47, 48, 51, 52
- terminologies and typologies 55–61
- workforce replacement 92–93
- behavior tracking 130, 132
- Boole or Boolean algebra 4, 121
- Brazilian Supreme Court 155
- British Nationality Act 7, 19, 22, 26
- Church-Turing Thesis 124
- civic literacy 15, 117
- Claudette *11*, 11–12
- click trails 130
- Commission Nationale de l’Informatique et des Libertés (CINIL) 114
- computational thinking 123, 124
- conditional statement 19, 20
- controlled language 13, 25, 27–31, **30**, 65, 148–149, 163, 167, 172
- cookies 114, 123, 132–134, 169
- Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) 89, 90

- Cracking the Code report 56, 70
- cross-disciplinary courses 120, 143–145
- Danish Agency for Digitisation 108, 109, 111, 112
- Das Digitale Amt 109
- data
 - analytics 119, 120, 136–139
 - intensity 52
 - literacy 15, 119
 - provenance 52
 - trails users 120, 129–136
- data-driven profiling 136
- Data Protection Vocabulary (DPV) 18, 44–47
- decision tree 26, 56, 57, 70, 70, 168–171
- deep learning systems 51
- Department of Defense 126–128
- digital-ready or digitally ready legislation and policymaking 109, 110, 112
- Digital Services Act 132
- Direction Interministerielle du Numerique (DINUM) 104
- Discourse Representation Structures (DRS) 29, 30, 148
- Discourse Representation Theory (DRT) 29, 30
- Discovery Sprint report 151, 152
- divergence of interests 56–59, 58, 61
- DoNotPay 2, 155–156
- Dublin Core 40–41, 106
- Dutch Child Care Benefit 90, 91
- education 13, 15, 108, 115, 117, 118, 120, 146, 169
 - primary education 117, 118, 120–140
 - secondary 117, 118, 120–140
 - specialized 140–145, **141–142**
- efficiency 9, 10, 14, 42–43, 56, 59, 61–69, 91
 - gains 3
- Entity Relationship modelling 174
- European Case Law Identifier (ECLI) 42
- European Commission 90, 109, 112, 169
- European Law Identifier (ELI) 42
- European Union (EU) 15–18, 47, 103, 109, 119, 125, 132, 170–171
- European Union’s Data Act 16, 119
- existential quantifier 20
- experimental approaches 115, 116
- Federal Trade Commission 17, 156
- FinTech 115
- flow models 152
- Foreign Intelligence Surveillance Court (FISA) 107
- formalization 7, 22, 44, 150
 - GDPR of Art. 7(1) 44–45
- Friend of a Friend (FOAF) 40, 41, 44–46
- F₁ score 4–6, 5, 80
- functional literacy 15, 117
- functional outcomes 51, 52
- Gateway-to-Gateway Protocol (GGP) 128
- General Data Protection Regulation (GDPR) 17, 18, 43–45, 75, 76, 80, 135, 147, 168–171
 - Art. 7(1) 45
 - Art. 32(1) 158
- good old-fashioned AI (GOF AI) 49, 49
- GPT 3, 80, 122, 145
- GraphDoc 171, 171
- High-Level Expert Group of the EU (HLEG) 82
- Holidays Act 151
- host computer 127, 129
- HTTP Cookies 133, 134
- human-centered design 113, 114
- human errors reduction 64, 65
- human-paced legal decisions 67
- humans disintermediation 59–61
- Hypertext Markup Language (HTML) 37, 38, 129
- IAMA 97
- if-then statements 7–8
- information retrieval system 5–6
- Interface Message Processor (IMP) 127
- Internationalized Resource Identifiers (IRIs) 32, 35, 36, 38, 39, 42, 53n2
- Internet Assigned Numbers Authority (IANA) 128
- Internet Corporation for Assigned Names and Numbers (ICANN) 128
- Internet Courts 2

- inter-network protocol (IP) 127–129, 132, 133, 169
 - addresses 127, 128, 132, 133
- inverse document frequency (IDF) 6, 7
- knowledge graph 35, 36, 38–39, 39, 41
- Laboratory for Numerical Innovation (LINC) 114
- legal automation 13, 15, 19, 47, 82, 104, 120
 - first wave 4–7
 - second wave 7–10
 - third wave 11–12
- legal decision tree 146–148, 168–171
- legal design 15, 113–116
- legal informatics 1, 56
- LegalRuleML representations 75, 76
- legal singularity 85
- LegalTech 2
- legal uncertainty 64, 65
- leveling function 16
- LKIF 41
- logic
 - defeasible 9–10
 - first-order 19–21, 23–24, 29, 48
 - predicate 8, 13, 19–22, 24–26
 - propositional 8, 19–21, 28
- Logical English 28, 29, 171
- logical languages 29, 30
- logic gates 8, 8
- logic programming language 22
- machine learning 12, 47–52, 103, 129, 136, 138, 155
 - approaches 104
 - legal knowledge 31–43
 - legal norm 43–47
 - models 56, 58, 61, 66, 67
 - techniques 11
- machine-paced legal decisions 67, 69
- machine-understandable regulations 42–46
- Member State 18, 109, 119, 170
- Mes Aides 43, 55, 60–62, **62**, 81, 93–94, 108, 114, 152–155, 154, 165
 - case 60, 61
- meta-review report 83
- micro directives 162
- Ministry of Interior 83
- National Ombudsman 90
- National Physical Laboratory 126–127
- natural language 2, 6, 11, 12, 24–29, **30**, 48, 148–149, 172
- neural network 10, 51
- New Zealand Government 151
- NGO 55, 154
- non-functional property 51, 52
- normative knowledge 34, 36
- NORSAR 128
- Offender Assessment System 89
- Office of Qualifications and Examinations Regulation (Ofqual) 137
- online tracking 114, 134
- ontologies 13, 18, 24, 32, 34, 37–42, 47, 48, 52, 76, 150, 152
 - knowledge 34, 36
 - learning 52
- Open Digital Rights Language (ODRL) 44, 76
- OpenFisca 94, 152–153, 157
- open-texture 14, 27, 30, 47, 70, 78–81, 96, 107, 113, 124, 162, 167, 171, 176
 - identification 157–160
 - norms 78
- Organization for Economic Co-operation and Development (OECD) 17, 56, 70, 112
- OWL Web Ontology Language 41
- Oxford Risk of Recidivism Tool (OxRec) 89
- PageRank algorithm 7
- personal data 18, 42, 44–47, 111, 119, 138, 144, 158, 168–171
- personalized content 130, 131
- personalized law 162–164
- policymakers 1, 3, 12, 13, 15, 55, 84, 93, 109, 112, 116, 119, 157, 171
- Post Office 87–89, 91, 97, 136, 164
- primary education 117, 118, 120–140
- privacy by design strategies 17
- privacy cynicism 136
- privacy-invasive technologies 144–145
- privacy law 42, 135, 163
- privacy policy 18, 163

- procedural knowledge 34, 36, 42
- Prolog 22, 29
- pseudocode 152
- public debates 2, 13, 15, 27, 83, 93, 94, 98, 103–104
 - law access 104–108
 - law digital impact assessment 108–113
 - legal design 113–116
- Python programming language 22–24, 26, 30, 31, 72, 73
- Rates Rebate Act 65, 70–72, 71, 74, 114, 151
 - accessibility point of view 72–73
 - calculations 70–72
 - compliance process 74–77
 - Python programming language 72
 - repeated process 73–74
- Rates Rebate Calculator (RRC) 151–152
- regulatory sandboxes 108, 114–116
- reinforcement learning 49, 50
- Resource Description Framework (RDF) 35–39, 41, 76
- retrieval system 4–7, 47
- risk classification model 90
- robot judge project 54, 60–62, **62**, 66, 68, 81, 93
- robot-lawyers 11, 156
- Rock, Paper, Scissors game rules 149–150, 175
- ROSS Intelligence 11
- rule as code 56, 70, 74, 94, 112
- rule-based systems 9, 10, 51
- RuleML 44
- Schema.org 40, 41
- secondary education 15, 117, 118, 120–140
- Secretariat for Digital Ready Legislation 110
- Semantic Sensor Network Ontology 40, 41
- Service Innovation Lab 69, 151
- social benefits project 54, 55, 60, 62, 73
- stakeholders 14, 44, 57–59, 61, 69, 83, 96, 98, 104, 115, 128, 165
- Stanford Legal Design Lab 114
- statistical biases 136–138, **137**
- statistical relationships 47
- structured query language 4
- Subject-Predicate-Object (S-P-O) 35
- sub-symbolic AI 48, 51, 58, 122
- supervised learning 49–50
- Support Vector Machines 51
- Supreme Court 84, 106, 131, 155
- Swiss Citizenship Act 20, 21, 23–24, 26, 48, 58
- symbolic AI methods 47, 48, 51, 52
- syntactic ambiguities 28, 64, 65, 79
- technology-based approaches 114
- training automation aspect 52
- transmission control protocol (TCP) 127, 129
- transparency 10, 14, 16, 17, 72, 74, 76, 83, 87–92, 95, **100**, 104, 109, 111, 117, 155, 165, **173**
- Tree of Porphyry 35, 36
- trial-and-error approach 83, 115
- Turing machines 122–123
- ubiquitous computing 120–124, 129, 140
- Uniform Resource Identifiers (URIs) 53n2
- United Kingdom (UK) 2, 83, 87, 89, 97, 107, 108, 121, 126, 136, 137, 164
 - DoNotPay 155–156
 - Post Office 87–89, 91, 97
- United States (US) 17, 67, 106, 125–128, 131, 143
 - DoNotPay 155–156
- United States Communication Decency Act 131
- universal quantifier 20
- unmanned aerial vehicle (UAV) 33, 34
- unsupervised learning 49, 50
- upper ontologies 40
- user-centered design 113
- user-friendly interface 171
- Victor case (Brazil) 155
- virtual private networks (VPNs) 133
- visualization techniques 113, 114
- Web beacons, implications 132–134
- Web resources 129
- workforce replacement 92–93
- World Wide Web 7, 35, 40, 42, 129, 130
- World Wide Web Consortium (W3C) 35, 37, 44