WHITE PAPER

Demystifying Legal Automation

Paving the route towards compliance automation for analytics and data science

2020

IMMUTA™

Introduction

Legal automation is usually seen with suspicion by lawyers and others[1]. This is because substituting the law enforcement process with a 'mere' technological or self-executing enforcement process, is inherently problematic as opacity inevitably creeps in. It thus becomes much more difficult to abide by the Rule of Law[2] and revert to the traditional legal triade: responsibility, accountability and liability: Who is in charge? Who is answerable? Who could be held liable if something goes wrong? When substituting the law enforcement process with a 'mere' technological process, the coder is de facto in charge, whereas he is not answerable and most likely cannot be held liable.[3]

Not all forms of legal automation are substitutive, however. Some should be described as supportive rather than substitutive. Legal automation is supportive when it aims to reduce human intervention with a view to ensure that a human with relevant legal expertise is properly informed or equipped for decision-making, and remains both in charge and answerable. This tracing facilitates allocation of liability, which is likely to be shifted to the organization which the human represents.

Such an approach is particularly useful in the field of compliance. Compliance is a key driver of organizational processes and policies but tends to produce friction between roles and functions leading to a minima or overly-complex compliance strategies. The main cause of this friction comes from repeated re-interpretations, translations and specifications of high-level legal norms, in order to produce executable rules to meet the needs of a variety of business-sponsored use cases.

We thus introduce the concept of ***compliance automation*** and define it as an approach that aims to reduce human intervention with a view to ensure that a human with relevant legal expertise is in charge, and is able to make executable, traceable and contestable decisions.

We then show how ***compliance automation*** **can make the difference** to effectively regulate analytics and data science environments. This whitepaper offers the beginning of a framework to effectively operationalize compliance automation, through the creation of fully-integrated automated policies. It thus clears the path towards effective and scalable compliance for data analytics.

The whitepaper is organized in three main sections. We start the analysis by unpacking the concepts of substitutive and supportive legal automation. We then explain what problems compliance automation can solve, analyzing a couple of common compliance anti-patterns. We finally make some recommendations for the building of fully-integrated automated policies and the enabling of organizations seeking to scale compliance across their data analytics use cases.

1 See e.g., Frank A. Pasquale, A Rule of Persons, Not Machines: The Limits of Legal Automation 87 George Washington Law Review 1 (2019); Sylvie Delacroix, Computer systems fit for the legal profession?, 21(2) Legal Ethics 119 (2018). But Richard and Daniel Susskind are more optimistic. See Richard Susskind and Daniel Susskind, The Future of the Professions – How Technology Will Transform the Work of Human Experts, (OUP 2015). See also J Harper, The Lawyer Bubble : A Profession in Crisis (Basic Books 2013). The roots of the crisis the legal profession is facing are probably more profound, see JC Smith, Machine Intelligence and Legal Reasoning – The Charles Green Lecture in Law and Technology 73 Chicago-Kent Law Review 277 (1998) . At the same time, large-scale adoption of machine-learning technologies, even if not adopted by lawyers specifically, threaten privacy, identity, autonomy, non-discrimination, due process and the presumption of innocence. See M. Hildebrandt, Smart Technologies and the End(s) of Law: Novel Entanglements of Law and Technology (Elgar 2015); M Hildebrandt, 'The Force of Law and the Force of Technology' in M McGuire and T Holt (eds), The Routledge Handbook for Technology, Crime and Justice (Routledge 2017).

2 "[T]he mechanism, process, institution, practice, or norm that supports the equality of all citizens before the law, secures a nonarbitrary form of government, and more generally prevents the arbitrary use of power." Naomi Choi, The Rule of Law in Encyclopedia Britannica, available at https://www.britannica.com/topic/rule-of-law. See also Jeremy Waldron, The Rule of Law in Stanford Encyclopedia of Philosophy (2016), available at https://plato.stanford.edu/entries/rule-of-law/.

3 See e.g., Lawrence Lessig, Code 2.0:Code and Other Laws of Cyberspace (2006), available at http://codev2.cc/download+remix/Lessig-Codev2.pdf.

What the different forms of legal automation are

Substitutive legal automation is generally understood as the transformation of legal norms into code or machine-readable instructions, so that these norms can be made self-executing. This could be seen, at first glance, as being much more effective than traditional legal processes as once standard-setting is done, the need to monitor compliance and trigger the enforcement of process in case of breach disappears: user behavior is shaped ex ante through the Code.

However, such an approach is problematic for at least two reasons.

First, transforming legal norms into code or machine-readable instructions is not a straightforward exercise. In fact, legal rules are generally seen as malleable and open-ended. This is because legal concepts are tools used by lawyers to achieve an end and their content varies depending upon the definition adopted, the theory underpinning the approach, the interpretation method chosen and sometimes simply convenience. Legal concepts can hardly be fixed into a closed list of necessary and/or sufficient properties.[4] In fact, they are often considered to be providing an open list, ending with an etc clause.[5]

In the field of data privacy and data protection, a good example of an open list would be the security principle as formulated in the EU General Data Protection Regulation (GDPR)[6].

> 1. Taking into account the state of the art, the costs of implementation and the nature, scope, context and purposes of 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 organisational 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 organisational measures for ensuring the security of the processing.[7]

4 H. L. A. Hart, The Abscription of Responsability and Rights, and Definition and Theory in Law, Proceedings of the Aristotelian Society, New Series, Vol. 49 (1948 – 1949), pp. 171–194, available at http://legacydirs.umiacs.umd.edu/~horty/courses/readings/hart-1948-ascription.pdf.

5 H. L. A. Hart, The Abscription of Responsability and Rights, and Definition and Theory in Law, Proceedings of the Aristotelian Society, New Series, Vol. 49 (1948 – 1949), pp. 171–194, available at http://legacydirs.umiacs.umd.edu/~horty/courses/readings/hart-1948-ascription.pdf.

6 Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance), OJ L 119, 4.5.2016, p. 1–88.

7 GDPR Article 32.

Second, even if legal concepts were to be found to be self-contained or closed, legal determination remains a heavily-regulated social process built upon fundamental principles, such as the Rule of Law and due process,[8] which aim to protect – directly or indirectly – individual rights and liberties. Yet, by delegating the transformation or translation of legal norms into Code to coders, standard-setting is de facto delegated to coders rather than lawyers, and by producing self-executing instructions, the protection of individual rights and liberties – such as the rights to non-discrimination, due process or the presumption of innocence – is then potentially undermined, in particular when the Code is only an approximation of substantive provisions.

approach to legal automation that aims to reduce human intervention with a view to ensure that a human with relevant legal expertise is properly informed prior to decision-making, and is both in charge of the decision-making and answerable, thereby facilitating allocation of liability, which is likely to be shifted to the organization which the human represents.

Take the example of the much debated field of online content regulation. When the automated matching of copyright works and user files through the process of fingerprinting prevents users from uploading their files on platforms, even in cases in which no copyright infringement has occurred, the right to freedom of expression is undermined.[9]

In order to protect individual rights and liberties, it is essential to be able to precisely identify who is making the legal determination and the reasons for the determination, so that ungrounded or poorly-grounded decisions can be contested and challenged. 'Who is in charge?', 'Who is accountable?', and 'Who could be held liable?' are key questions, to which a clear answer should be given before initiating a process of automation.

Although substitutive legal automation has changed name over time, its inherent flaws remain. 'Robotic law enforcement' is a form of substitutive legal automation and as such is inherently problematic, as it hides the obvious: *"Until some "master algorithm" can code its own progeny, human beings will always be responsible for legal determinations,"*[10] explains Prof. Frank Pasquale, who rightly insists that "[w]*ithout attributing algorithmic judgments and interpretations to particular persons and holding them responsible for explaining those judgments, legal automation will undermine basic principles of accountability."*[11]

The example that Frank Pasquale uses to explain the concept of robotic law enforcement is that of red-light cameras and parking tickets. He writes:

> "Red-light cameras are one version of robotic law enforcement. All that is necessary for the robot to enforce traffic law is a simple set of rules declaring that any person who owns a car that passes under a light when it is red shall be fined a certain amount and possibly lose his or her license to operate the car."[12]

8 "Due process, a course of legal proceedings according to rules and principles that have been established in a system of jurisprudence for the enforcement and protection of private rights. In each case, due process contemplates an exercise of the powers of government as the law permits and sanctions, under recognized safeguards for the protection of individual rights." Encyclopedia Britannica, Due Process, available at https://www.britannica.com/topic/due-process.

9 See Daphne Keller, Problems with filters in the European Commission's platform proposal (2017), available at http://cyberlaw.stanford.edu/blog/2017/10/problems-filters-european-commissions-platforms-proposal.

10 Frank A. Pasquale, A Rule of Persons, Not Machines: The Limits of Legal Automation 87 George Washington Law Review 1 (2019).

11 Frank A. Pasquale, A Rule of Persons, Not Machines: The Limits of Legal Automation 87 George Washington Law Review 1 (2019).

12 Frank A. Pasquale, A Rule of Persons, Not Machines: The Limits of Legal Automation 87 George Washington Law Review 1, 12 (2019).

Frank Pasquale does not reject all forms of legal automation, however. He distinguishes between substitutive legal automation and decision-support tools and explains that the latter *"are not a replacement of the human with the algorithmic, but rather another step toward improving a sociotechnical system of human decision makers and machine-aided decision analysis."*[13]

Building on this distinction, we introduce the concept of supportive legal automation and define it as an *approach to legal automation that aims to reduce human intervention with a view to ensure that a human with relevant legal expertise is properly informed prior to decision-making, and is both in charge of the decision-making and answerable, thereby facilitating allocation of liability, which is likely to be shifted to the organization which the human represents*.

This approach, we argue, is particularly useful in the field of compliance, which is distinct from judicial and legislative law-making. Compliance is frequently described as a set of rules, processes and tools used by organizations to comply with legal and/or businesses requirements. Within corporations, these processes are initiated by legal or compliance departments.

Contrary to judicial and law-making processes, compliance processes are not entirely owned by lawyers or legal roles. In fact, as we will explain in the following sections, the complexity of roles and functions within organizations often has the effect of duplicating decision-making and giving non-lawyers a wide margin of interpretation when translating and making legal norms more specific to guide action on the ground. This is true in particular in analytics and data science environments, in which legal input is often perceived as being too generic and unhelpful.

We refer to the term *compliance automation* as a subclass of legal automation, and define it in the following way: *compliance automation is an approach that aims to reduce human intervention with a view to ensure that a human with relevant legal expertise is properly informed prior to decision-making, is in charge of decision-making, and can make executable, traceable and contestable decisions*. We therefore go beyond the common understanding of the term, which defines it as a tool used to simplify compliance procedures. Our approach is fundamentally human-centric and seeks to improve both compliance procedures and compliance outputs by eliminating compliance anti-patterns.

Compliance anti-patterns are behaviors triggered to respond to a recurring problem, but which are usually ineffective and, in some cases, undermine the *raison d'etre* of the requirement to be complied with. In the world of data science, compliance anti-patterns occur frequently.

13 Frank A. Pasquale, A Rule of Persons, Not Machines: The Limits of Legal Automation 87 George Washington Law Review 1, 54 (2019).

What compliance automation solves

A compliance journey is usually made of multiple decisions involving different actors with different expertise. It necessitates translating high-level, open-ended legal norms into a series of more specific instructions and mitigating measures, which are usually called rules and controls. By way of example, the requirement of data minimization, which obliges organizations to tailor the amount of data to their processing purposes so that they only process what is adequate and necessary to pursue their purposes, can lead to a wide variety of specifications.[14]

For organizations engaged in data analytics activities and wanting to develop these activities further, the compliance journey can be particularly complex.

Diagram 1 illustrates what usually happens within analytics and data science environments, i.e., a three-step process starting with the drafting of policies, followed by the formulation of standards, and then the creation of rules and controls. This is what we call the PSR (Policy-Standard-Rule) process. Policies are usually set by the legal or compliance department most likely under the supervision of the Director of Data Privacy or her equivalent. Generic recommendations such as *"data should be anonymised as often as possible"* will usually be produced at this stage. More specific standards will then be expressed at a second stage, such as a list of available anonymization techniques.

These high-level recommendations are then translated by an intermediate team in order to generate executable rules and controls for each data science use case or project. This intermediate team is often led by the "Data Governance Team," acting under the supervision of the Head of Data Governance. These rules and controls will then have to be enforced.

Each project owner will be subject to these rules and controls, usually with some checks to perform and forms to fill in before she can obtain the green light to initiate a data analytics project. These checks could entail contacting the legal department to better understand which decision should be made at the project level (e.g., in order to anonymize the data). Most of the time this will mean producing a very detailed description of the entire project including references to the data sharing infrastructure, so that the legal department is put in a position to produce meaningful advice. As a result, the workflow will be significantly extended and delays will occur.

14 This requirement can be found in GDPR, Article 5(1)(c) or Cal. Civ. Code §§ 1798.140(d).

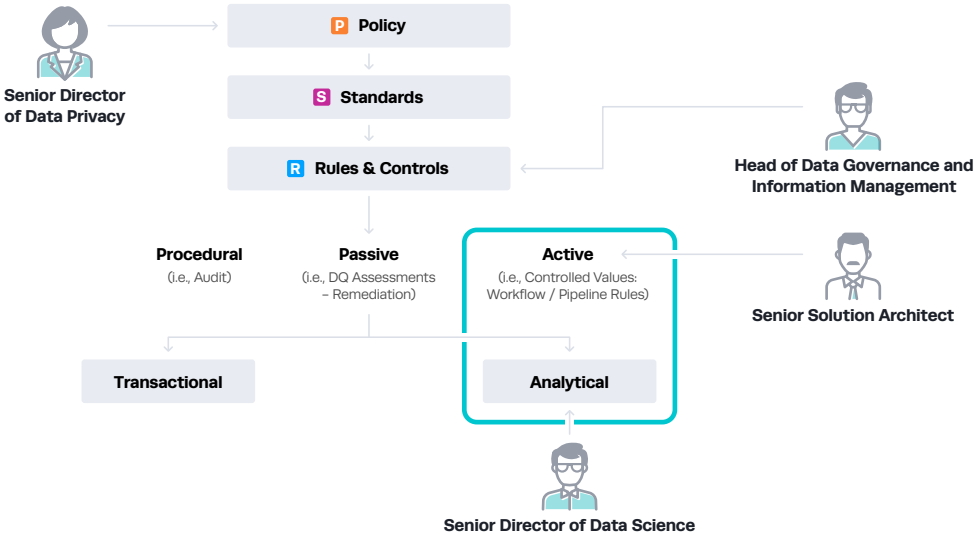

Diagram 1: A typical compliance journey within data-driven organizations

Another way to represent the interaction that is usually happening within organizations is to use a RACI modelisation,[15] common in project management as illustrated in Table 1. RACI is an acronym derived from four key roles used in organizations across sectors: Responsible, Accountable, Consulted, and Informed. Importantly, using RACIs does not necessarily lead to the streamlining and scaling of workflows, although they certainly help clarify nodes and tasks.

	Data Owner	Use Case Sponsor	Data Scientist	Data Governance team	IT team	Legal/Data Compliance/ Data Protection Officer
Completion of use case description	C/I	A	R	C	C	C
Review of use case description						R/A
Legal case-by-case assessment	C			I/C	C	R/A

Table 1: An example of a RACI

Not only can a compliance journey be complex and lengthy, but it can also be undermined by anti-patterns.

Let's take an example. As mentioned above, compliance guidance produced by legal/compliance teams often point to data anonymization as a solution to drastically reduce compliance issues. Most of the time this guidance doesn't explain how anonymization should be achieved. A data scientist wanting to initiate a data analytics project thus usually goes back to the legal/compliance team for further clarification. The legal/compliance team is, however, often ill-equipped to support decision-making at this stage. A back-and-forth therefore begins between the data scientist, the data governance and the legal/compliance teams, until a decision is taken. Unfortunately, in many instances, this back-and-forth is based upon wrong technical assumptions. As a result, the second time the data scientist is faced with a similar question for a new project, he will be tempted to discard the anonymization option from the start.

15 Duncan Haughey, RACI matrix, available at https://www.projectsmart.co.uk/raci-matrix.php.

This is how a compliance anti-pattern emerges: while it is actually best practice to always consider the full range of anonymization techniques in the light of the applicable legal standard before engaging into a processing activity, data scientists or project owners are in fact given an incentive not to make any assessment. Such an anti-pattern has four dimensions:

- **Speed:** Data sharing is slowed or sometimes completely stopped when the sharing of personal data is seen to be too risky.
- **Scale:** When context evolves and for example new attributes become indirectly identifying or utility requires access to new attributes while others could be ignored, organizations have to restart the anonymization process from scratch.
- **Technical soundness:** Legal and compliance professionals, who likely lack technical insight, have no way to understand the effectiveness of the anonymization process.
- **Legal validity:** Legal and compliance guidance for anonymization/pseudonymization are being interpreted in silos, making it more difficult for a sound legal interpretation to emerge.

Other related compliance anti-patterns are likely to appear as well. In fact, the following scenario is quite frequent. When a data user wants to request access to a data set, she submits, most likely through emails, a request to the data set administrator, after having asked for guidance to the legal/compliance and/or data governance teams and filled in a risk assessment form. As hinted above, technical input from the IT teams might be needed, depending upon what access entails: most likely copying and pasting the data. As a result, there are always four or five personas involved in the process: data user, data owners, legal/compliance, data governance and IT.

Risk assessment, which should be undertaken prior to the granting of access, is thus usually highly fragmented. It's typically a long and painful process, initiated from scratch each time a data owner receives a request from a data user. In practice, this means that the data user has to conduct a risk assessment based on guidance produced by the legal/compliance team from scratch each time a new project is envisaged. The data owner is then asked to verify the legitimacy of the request on the basis of another set of guidance, produced once again by the legal/compliance team.

In short, risk profiles for each data analytics project are often generated in their own unique ways. This is another compliance anti-pattern with four dimensions:

- **Speed:** Data sharing is slowed or completely stopped because data owners have no digital transfer "handshake" recognized by their organization of how their data is protected and shared due to the issues above.
- **Scale:** When the risk matrix evolves, organizations have to make sure every risk profile across their organization enforces that change accordingly.
- **Technical soundness:** Legal and compliance professionals, who likely lack technical insight, have no way to understand the effectiveness of the technical controls that are mentioned in the risk profile in order to keep the risk level at a minimum.
- **Legal validity:** Legal and compliance guidance for risk assessment are being interpreted by a variety of stakeholders, including data users and data owners.

The most effective way to mitigate this anti-pattern is to reduce the margin of manoeuvrer for interpretation/translation de facto held by non-legal roles and make interpretations/decisions produced by lawyers executable through the activation of appropriate controls. These legal interpretations/decisions should also be traceable and contestable, to make sure they can be understood and adapted when context evolves. This can be achieved through compliance automation.

How automated policies can bolster compliance

Compliance automation is usually thought in terms of workflow capabilities and controls and actions planning, and the focus is usually set on data subject or consumer rights management and tracing of data and data usage. This type of automation can be highly sophisticated and in some cases have been said to involve artificial intelligence (AI), although the term is usually broadly defined in this context and refers to narrow AI rather than general AI.

Creating workflows and planning for controls and actions for rights management and data flows observation is however not enough to eliminate compliance anti-patterns in analytics and data science environments. This is because this does not ensure early integration of legal and technical expertise with a view to effectively frame or even tame practices within data science environments. In other words, this does not guarantee a true 'by design' approach to privacy and data protection within data science environments.

Fusing the three stages of the PSR process into one stage through the creation of fully-integrated automated policies is an effective way to address and mitigate the four dimensions of compliance anti-patterns, i.e., speed, scale, technical soundness and legal validity.

The building of automated policies requires a layered approach based upon three principles that are inspired from the Object Oriented Programming paradigm:[16]

- **Abstraction.** This approach starts with identifying requirements that are common to all types of projects or families of projects[17] and segregating them from project-specific requirements. By project, we understand a self-contained collaborative initiative that is planned to achieve a particular aim. In other words, this layered-approach is based on an abstraction process which makes it possible to progressively reduce the scope of the legal/compliance and technical assessment and ultimately focus upon what is truly necessary for each data science project.

- **Encapsulation.** This layered approach thus leads to a process of encapsulation, i.e., the succinct depiction of the essential legal/compliance and technical requirements of each data science project.

- **Inheritance.** Importantly, this layered approach makes it possible to create a process of inheritance between data science projects. On average, it's unlikely that new projects will be radically different from existing projects. Identifying common denominators will thus make it possible to simplify assessments, improve quality of decision-making and speed up execution through early activation of appropriate controls.

Let's illustrate this layered approach with an example and go back to our second compliance anti-pattern. Let's assume a common infrastructure has been set up within an organization in an attempt to standardize access to data.

16 "Object-oriented programming (OOP) is a computer programming model that organizes software design around data, or objects, rather than functions and logic. An object can be defined as a data field that has unique attributes and behavior." Margaret Rouse, Object-Oriented Programming (OOP) (2020), available at https://searchapparchitecture.techtarget.com/definition/object-oriented-programming-OOP.

17 Notably, Article 35 of the General Data Protection Regulation expressly states that "[a] single [risk] assessment may address a set of similar processing operations that present similar high risks."

Through the abstraction process, it's possible to separate failure modes[18] that relate to the infrastructure supporting data access, failure modes that relate to the object (data items or data sources) and failure modes that relate to the purpose of the processing activities.

The first category of failure modes should not require a new assessment each time a data access request is submitted.

The second category of failure modes can be treated through the combination of data and context controls aimed at de-identifying the data, which will make it possible to group projects into families (which could then be divided into subfamilies), e.g. projects requiring access to identifying data or projects having access to de-identified data only. At this stage, it should be possible to create a rule that will automatically discriminate between the different families or subfamilies of projects and apply a set of controls defined by a team comprising legal/compliance personnel and privacy technologists, whose skills are not necessarily the same as those of traditional data scientists.

The third category of failure modes are likely to be entirely project specific and will require ad hoc documentation. This is the case for example for the failure modes that relate to the requirement of fairness in data protection legislations (i.e. the processing shall be fair for the subjects of the data). As we argued in our previous whitepaper[19], no one control will do for this particular failure mode and the assessment will have to be ongoing taking into account model assumptions, model limitations as they evolve over time as well as the impact of the entire decision-making pipeline in which the model will be embedded. A data scientist working with her business sponsor is particularly well placed to compile this documentation, which can then be regularly reviewed by legal/compliance personnel as the project progresses. If the project leads to unacceptable results, the legal/compliance personnel should be able to terminate access to data and take appropriate mitigation measures as regards the model. At this stage, it should be possible to create a rule that will automatically terminate the project if a selected set of failure modes have not been documented.

Each time a rule or a combination of rules is potentially applicable to a variety of projects, is immediatly executable and has been produced through integration of relevant interdisciplinary expertise with a view to meet legal/business requirements, an automated policy is born. Rule creation should be done in one go without unnecessary fragmentation of the rule-making process, which is what usually happens when following PSR processes. Legal engineers are particularly well-placed to perform this integrating task as they combine legal expertise and system configuration skills.[20] In addition, automated policies should always have at least one owner and be contestable, which ultimately means that they should be motivated and adaptable, e.g. through the carving out of exceptions.

With automated policies, compliance scales and its net widens. Legal and technical assessments are more focused, integrated and therefore of better quality. Consequently, the speed at which data is accessed should increase.

18 For a definition of failure modes and an explanation as to how they can be leveraged for facilitating compliance in data analytics environments, see S. Stalla-Bourdillon et al., Data Protection by Process: How to Operationalize Data Protection by Design for Machine Learning, Immuta and FPF whitepaper (2019), available at https://www.immuta.com/data-protection-by-process-fpf-whitepaper/.
19 S. Stalla-Bourdillon et al., Data Protection by Process: How to Operationalize Data Protection by Design for Machine Learning, Immuta and FPF whitepaper (2019), available at https://www.immuta.com/data-protection-by-process-fpf-whitepaper/.
20 Ciara Byrne, Don't Call me a lawyer, I am a Legal Engineer, FastCompany (2019), available at https://www.fastcompany.com/90372705/dont-call-me-a-lawyer-i-am-a-legal-engineer.

Conclusion

Compliance automation, as we see it, is not an attempt to algorithmatize legal prescriptions, but an attempt to inject a high degree of accountability into business processes and thereby facilitate the demonstration of compliance in analytics and data science environments.

While most, if not all, attempts to code and replace the law, in particular through predictive analytics methods, have proved unsuccessful up until now, a by-design approach to compliance through the building of automated policies is a promising path towards safer data-driven practices.

This whitepaper aimed to outline a framework to operationalise compliance automation, and we welcome suggestions or comments to improve this framework. Please reach out to governance@immuta.com with feedback.

IMMUTA™

115 Broad Street, 6th Floor, Boston, MA 02110 | immuta.com | (800) 655-0982