

THE RESEARCH UNIVERSITY PRESS

International
Computational Law Journal

VOLUME 2 / ISSUE 2 / CURATED SOURCE EDITION

# Governing Digital Legal Systems: Insights on Artificial Intelligence and Rules as Code

Hamish Fraser · Tom Barraclough

MIT Computational Law Report · Originally published 2024-10-14

Independent reader edition under CC BY 4.0. This attribution cover was added; original source PDF pages follow unchanged. Original authorship and source identity are retained. MIT, Stanford University, Stanford Law School and CodeX have not sponsored or endorsed this republication.

Source, license and version evidence accompany the online article.

**MIT Computational Law Report**

# Governing Digital Legal Systems: Insights on Artificial Intelligence and Rules as Code

**Hamish Fraser Tom Barraclough**

**Published on:** Oct 14, 2024

**URL:** https://law.mit.edu/pub/governingdigitallegalsystems

**License:** Creative Commons Attribution 4.0 International License (CC-BY 4.0)

# Executive summary

## Artificial intelligence and rules as code

### AI systems

The world is swamped with activity on artificial intelligence. In particular, a number of governmental and commercial operators are examining how artificial intelligence can be deployed in legal and regulatory contexts.

This more recent activity comes at a time when guidance, principles, best practice and other forms of soft regulation have already been established. These are now crystallising into binding legislation as well as other forms of harder regulatory restriction.

### Rules as code

More modestly, another approach to automated systems is catching the attention of commercial, academic and governmental audiences, known as “rules as code”. Rules as code is a label given to a deliberate approach to converting law and other forms of regulation into computer code and other digital formats. This enables law to be read by and operationalised in digital systems. Three key use cases for rules as code include:

1. **Deployment in automated decision-making systems.** Systems can receive data inputs representing key facts, apply algorithmic rules to that data in ways that approximate the law, and produce outputs that indicate how the law applies to that given set of facts. Deployment of rules as code in this way permits clear explanations for how the system operates to be given, as well as making it auditable and transparent.
2. **Modelling policy settings.** Digital models can be created that reflect current policy settings such as legal and financial thresholds for determining access to benefits or the imposition of taxes. These let policymakers and others assess the likely impact of changes to these thresholds and how they might apply in given scenarios.
3. **Clarifying legal and policy implementations.** In recent times, legal drafters and policy professionals have explored methods that create digital implementations of proposed legal rules in ways that permit inconsistencies, ambiguities, or gaps to be identified during the policy development and drafting process. This particular use has not been as widely adopted given the shift it requires in working practices for people whose workload is high, skill set is specialised, and tolerance for change and experimentation is low.

### Collision

We anticipate that AI systems and rules as code will increasingly collide with each other. They can currently be distinguished sharply by the relative absence of machine learning-based approaches to rules as code, compared with the dominance of statistical machine learning approaches in contemporary artificial intelligence products

and services. However, both rules as code and modern AI systems have common antecedents, and both of them can be applied for identical purposes (for example, in automated decision-making systems or policy modelling systems). There are also calls for increasing use of large language models to analyse legal materials and convert them into other formats, including computer code. All in all, it is safe to say that both areas are tightly intertwined.

# Governing digital legal systems

## Digital governance and AI regulation

In this paper, we argue that developments in how artificial intelligence systems should be governed must also be applied to the digital and operational components of rules as code systems. We call for greater attention to these governance processes as well as briefly illustrating how particular elements of AI governance frameworks could apply to rules as code systems. These include:

1. **Transparency** around the deployment and use of rules as code systems, including notification when they are being used.
2. **Explainability** of rules as code systems and their outputs, particularly where they purport to implement the law or have legal consequences.
3. **Risk assessment** and mitigation obligations, as are becoming common in other technology governance frameworks.
4. **Post-deployment monitoring and review** obligations to ensure systems are operating as intended, or in response to emerging risks, or legal amendments.
5. **Continuity** across implementation phases from design, to development, to deployment, and then on to post-deployment monitoring and review, including requirements around record-keeping and documentation.
6. **Multistakeholder input** as appropriate at each phase of development to enhance risk assessment and mitigation, and to identify errors and negative impacts.
7. **Accountability mechanisms** for people designing, developing and deploying systems, as well as review and appeal mechanisms for users and other stakeholders affected by the system.

## Looking broadly to digital legal systems

Next, we argue that the scope for "rules as code" should be considered much more expansively than its historical and contemporary framing. Rather than a relatively restricted practice of taking legal norms and documents, interpreting how they work, and modelling that interpretation in rule-based systems, instead we suggest that interpretations of multiple legal instruments can be implemented in various digital formats, such as data structures or datasets.

The effect of this is that rules as code requires neither rules nor code, and the domain is better thought of as relating to systems that implement interpretations of the law in digital formats. We refer to these as "digital

legal systems" because of their dual character. This class of digital legal systems is significant, and we describe one example from our own work that converted secondary legislation and an array of other primary legal sources into data structures, a "rule-builder", and an interactive user-facing display for the rules which is dynamically generated from basic comma-separated value files and easily accessible to non-technical staff.

## Governance of legal components

Finally we argue that the legal component of digital legal systems also requires dedicated governance practices. People designing and implementing digital systems are interpreting multiple sources of law and making decisions about how to incorporate their interpretation of the law into a digital output or system. This decision-making process should be conducted in ways that build confidence in the decisions that have been made, and allow those decisions to be revised over time.

To illustrate what this process of legal assurance might require, we use the example of a standard form agreement for the sale and purchase of real estate developed by a society of legal professionals in New Zealand. We draw on our own prior research and interviews about how this agreement was developed as well as how it has been governed over time. We identify key features of the agreement to understand how legal components of a digital legal system might operate, and what might be required to provide assurance that the interpretation of the law embedded in a system is reliable. Insights include the following:

1. The identity, qualifications, and status of people drafting and reviewing the legal component is one of the most important features for building confidence.
2. The legal component of a system will incorporate multiple legal sources and how they apply to a specific purpose or use case. Legal source material may not be explicitly referred to, but it will be in some cases.
3. Confidence in the agreement is supported by its accessibility to users, legal professionals and independent reviewers, particularly the judiciary. Judicial review of a legal instrument can add significantly to its reliability.
4. Multistakeholder input to a legal instrument can build confidence in it among key users, drive uptake, and ensure that issues are identified quickly through real world practice and experience.
5. Permissioning and version control can be critical. Integrity can be protected through intellectual property mechanisms, professional disciplinary mechanisms, and software that controls how amendments are made, and clearly signals any amendments to users.
6. Legal components should be responsive to legal, socio-economic, and technological change. If successful, they could be in operation for decades.

Features of this and other similar legal instruments present a useful source of insight for considering how the legal components of digital legal systems can be developed and managed over time and we encourage further investigation in this area.

## Recommendations and way forward

We conclude by inviting further collaboration and partnership in developing this work further as Syncopate develops associated products and services. We offer some insights from what we have learned so far, and suggest further areas for investigation by researchers, practitioners, regulators, and others.

---

# Artificial intelligence, rules as code, and digital legal systems

This report explains the concept of "rules as code" in order to suggest some features for how the digital and legal components of a rules as code system might be governed. When it comes to suggesting governance requirements, we look to two areas.

- On governing the system's digital component, we look to contemporary approaches to the governance of artificial intelligence.
- On governing the system's legal component, we look to governance of a widely used standard form sale and purchase agreement for real estate.

In this part of the report, we explain the overlaps and parallels between rules as code and artificial intelligence, as a matter of substance, history, and contemporary overlap. We also explain why the integrity of rules as code systems rests heavily on careful governance of their legal components in order to be reliable.

## Contemporary regulation of artificial intelligence

Readers in 2024 will need no background on the wave of enthusiasm for new AI applications based on large language models following the release of OpenAI's ChatGPT product in November 2022. By contrast, some readers may have forgotten or overlooked the preceding wave of interest in regulating AI systems around 2014-2020. This flowed from the increasing adoption of deep learning neural networks for various applications, as well as statistical models for predictive analytics.

One consequence of the 2014-2020 period is that a strong normative foundation has developed for describing the principled and ethical expectations on people developing, designing and deploying AI systems, particularly where these are used in a government context to make decisions about the legal rights and responsibilities of citizens. This was particularly apparent in the inclusion of provisions around automated decision-making in the European Union's General Data Protection Regulation.

A related area of focus during this period was prompted by the potential risks associated with using large statistical models at a population level, even if these would not directly impact the rights and interests of individuals. Big data models were used to predict changes in behaviour over time, identify trends that may not have otherwise been apparent, or to model the impact of these big picture trends. When such big data models are used as the foundation for policy development or system design, there is a risk that flaws in the models or

the datasets they use produce inaccuracies or inequities that become reflected in overall policy design. This may mean such systems are ineffective, discriminatory, or simply harmful to particular groups, or to the population as whole.

The key point to note is that contemporary discussion about regulating artificial intelligence (including the ethics of AI, data ethics, AI safety, and responsible AI) had already quite thoroughly explored the fundamental expectations that should be placed on designers, developers and deployers of AI systems, even before the current large language model moment. These expectations focused heavily on automated decision-making systems, but they equally focused on the use of AI systems for large-scale modelling to the extent this informs decision-making about public policy or system design. These expectations are beginning to be embedded in national legal frameworks, as well as bilateral and multilateral agreements between states.

## What is "rules as code"?

Rules as code is a contemporary label that has been applied to what used to be described as "AI and Law". There is a historical association and overlap between rules as code type approaches, and artificial intelligence.

The core premise of rules as code is that legal rules can be converted into computer code (in the form of programming languages), with the effect that these rules can be analysed and implemented by computer systems. There is a dense body of historic academic scholarship, as well as a notable proliferation of commercial products, which applied these historic approaches to enterprise scale systems. At the time, statistical approaches to artificial intelligence founded on machine learning were treated as being less promising than rule-based or symbolic approaches to AI, which relied in part on deliberate encoding of knowledge and rules into computer systems, rather than leaving these to be uncovered through training of machine learning systems based on large datasets and significant computational power. Several examples during this period focused on "coding" specific statutes, as well as the development of domain specific languages and dedicated applications, with a clear focus on converting legal rules in statutes (and sometimes other materials) into specific lines of code.

## Contemporary interest in rules as code

Contemporary focus on rules as code has flowed from a period from around 2018 to 2020 when a government unit in New Zealand was investigating how to better deliver government services using digital systems. Serendipitously, the team was joined by Matti Schneider, who brought extensive experience as a software developer working in the French Government, which had developed "OpenFisca", an open source python application for converting government benefits and rules of taxation into code. The New Zealand team was also led for a time by Pia Andrews, a public service innovator from Australia with a history in financial regulation, an area which has explored digital implementation of rules through "RegTech", or regulatory technologies. From New Zealand, this work caught international interest at the 2018 Digital Nations Summit. This was supplemented by Pia Andrews' subsequent positions working with government agencies in Canada,

and then again in Australia. Since then, the Federal Government of Australia has taken steps to adopt rules as code practices into part of its core government systems, and the Government of Canada has also conducted work implementing rules as code, strongly influenced by the work of Jason Morris.

Other links between New Zealand and the longstanding work in the rest of the world flowed through members and associates of the core New Zealand team establishing links with existing networks and academics in the European Union (notably Monica Palmirani and the Lex Summer School, leading work on standardisation through LegalDocML and LegalRuleML) as well as the interest of the OECD, who completed a paper called "Cracking the Code", which examined the concept and explained it for a global audience. Subsequent work contextualised the contemporary rules as code movement against its historical antecedents in AI and Law and legal informatics, with notable reports and comments from Guvernatori and Huggins et al, as well as Greenleaf, Chung an Mowbray of AustLII in Australia, and Hildebrandt and Diver in the European Union COHUBICOL project. Other notable work continued in Stanford's CodeX programme, work by Matthew Waddington at the Jersey Legislative Drafting Office and the Singapore Management University's Center for Computational Law.

## The role of legal interpretation

The contemporary renaissance of rules as code approaches led to a new wave of entrants who brought genuine optimism about the possibilities and potential benefits of implementing legislation (or law more generally) as code. Many advocates also worked actively in operational settings in government, and came from backgrounds such as policy development, government service delivery, or software development. While some lawyers and legal drafters were included, few participants had the time or resources to investigate the historical antecedents of the law-as-code movement. This meant that the role of legal interpretation as a specific method of deriving meaning from written law was sometimes overlooked.

Some advocates presented utopian visions of direct implementation of code passed by Parliaments, overlooking the dystopian potential of excluding judicial oversight from the way law is interpreted, applied, and challenged in legal proceedings. Along with others, our 2021 report emphasised again the importance of legal interpretation as a component of the rule of law, and the need to factor this into the development of rules as code systems.

## "Highly reliable interpretations" implemented as code

To some, the fact that code cannot be law, or that converting law to code requires a process of legal interpretation, presents a fatal barrier to rules as code's potential as a movement. By contrast, Jason Morris forcefully argued that all operational implementations of the law have the status of "an interpretation", including advice by lawyers. Law has no single true meaning until it is interpreted.

Building on this conclusion, we suggested instead that advocates for development and adoption of rules as code models should instead focus their attention on what can be done to demonstrate that the legal interpretation embedded in their code is reliable. We referred to these as "highly reliable interpretations", and looked for examples of other types of instruments that implement law at scale in a manner that permits a high degree of confidence that the instrument is legally correct. The example we identified was a standard form sale and purchase agreement for residential real estate, which is jointly developed, managed and owned by a professional association of lawyers, and a professional association of real estate agents. We identified key features that made the agreement legally reliable and suggested how these might be applied in rules as code implementations.

## Potential uses of rules as code systems

There are two main use cases usually suggested for the implementation of rules as code, with a less prominent third.

1. **Expert systems and automated or semi-automated decision-making systems.**

   Rules as code models let people use computers to find out whether something is legal or not, or how the law applies to a given pattern of facts (as represented in a dataset). This can most clearly be seen in the tendency by rules as code practitioners and advocates to work toward social welfare contexts, including administration of benefit systems. Key examples of this emphasis on social welfare entitlements include the Digital Benefits Network in the United States, the work of the Digital Aotearoa Collective, the BenefitMe site, and the original work in the Service Innovation Lab.
2. **Policy modelling.**

   The settings of a particular policy initiative can be modelled in machine executable rules, and this allows people to tweak those policy settings digitally, or to present different scenarios in that model through different datasets. For example, benefits can be tweaked monetarily, or eligibility criteria can be tightened or loosened, or foreseeable scenarios like increases in unemployment rates can be assessed through different data points. This approach can be most clearly seen in work by PolicyEngine, although we suspect it is much more widespread within government agencies around the world.
3. **Testing implementation of ambiguous legal rules or policy settings.**

   While less widely adopted, modelling rules in code can also foster disambiguation of law and policy settings. For example, people responsible for implementing rules can ask people designing those rules, "When you said X, precisely what did you mean by that? How would that apply in this situation?" This has been repeatedly described as one of the core benefits of a "better rules approach" to policy development, but we have not yet heard an example of it being meaningfully deployed in policy development, although we

note that the computer readable legislation project in Jersey has been using it in legal drafting. In our experience, policy development has well entrenched patterns and work is often conducted under time pressure. Generally, people responsible for designing policy systems may be unable or unwilling to explain policy decisions in greater detail. There are also sometimes good reasons for separating design from implementation.

Taking the first two use cases – automated decision-making systems, and large scale policy modelling – it should be clear that contemporary discussions about safe and responsible use of artificial intelligence and machine learning systems are directly engaged. Specifically, there is now a well-established body of norms, principles, and regulation imposing limits on how automated decision-making systems can be designed, developed, deployed, monitored, maintained and reviewed.

Additionally, there are strong ethical considerations associated with the use of big data to inform policy development, or model policy options. Therefore AI and big data ethics frameworks can also be relevant to policy modelling approaches, even where these are not necessarily based on machine learning.

## Key conclusions

This brief summary is important for contextualising the suggestions we make in the next section. Key points from this summary are as follows:

1. A contemporary movement and practice is directed at converting law into code to be implemented in computer systems. This is referred to as “rules as code”.
2. The “rules as code” movement has historical roots in the development of artificial intelligence. Contemporary “rules as code” systems are distinct from contemporary AI models, because the latter rely predominantly on machine learning from large data sets using statistics, rather than manual programming. Despite this, rules as code is a modern manifestation of the “AI and Law” subject, and can be easily conceptualised within law as an AI system.
3. Contemporary developments in AI have crystallised into established regulatory frameworks. These frameworks apply to use of AI systems in decision-making, but they also apply to high level applications of models that do not have a direct impact on the rights and interests of individuals.
4. Rules as code systems have both digital and legal components. Both components must be subject to effective governance, and governance of the legal component is equally as important as governance of the digital component.

In the next section, we suggest key governance features for these legal and digital components, and we also suggest that the scope of rules as code is much broader than traditionally thought.

# Governing digital legal systems

Rules as code systems have a digital component and a legal component, both of which require dedicated governance. In this section, we draw on two examples to make more specific suggestions about how each could be governed.

## Governing digital components

In many respects, digital systems that implement the law are indistinguishable from the kinds of AI systems that are subject to so much public scrutiny and regulatory attention today, differing mainly in the way that an absence of statistical components may change the risks that arise and the governance arrangements that follow. While our discussion emphasises automated decision-making systems, we recall that these principles may equally apply in high level modelling approaches without direct impact on individuals.

### EU legislation on automated decision-making systems

Any rules as code system receiving data on individual cases in an operational context will probably require assessment against legislation, and two illustrative examples from the European Union include the General Data Protection Regulation, and the incoming Artificial Intelligence Act. In any given case, these specific acts may not apply, but they present a useful example because they apply in a significant market, have extra-territorial effect, and likely have some global impact through "the Brussels effect".

Under the GDPR, data controllers are obliged to tell data subjects whether automated decision-making systems will be used in the processing of their data, as well as the logic of those systems and anticipated consequences of the use of those systems (see art 13(2)(f), art 14(2)(g), and art 15(1)(h)). Article 22 of the GDPR gives a right not to be subject to decisions made solely by an automated system without explicit consent, or in other constrained circumstances (art 22(2)). In terms of their applicability to rules as code systems, the consequences of an automated decision are explicitly framed in terms of their "legal effects" (art 22(1)) and the importance of protecting the subject's "rights and freedoms and legitimate interests", including the right to seek human intervention, express a counter point of view, and to "contest the decision" (art 22(3)).

Under the EU AI Act (the Corrigendum 808 version as at 1 June 2024), an AI system is defined in a way that we believe easily includes many rules as code systems:

> *"a machine-based system designed to operate with varying levels of autonomy, that may exhibit adaptiveness after deployment and that, for explicit or implicit objectives, infers, from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments".*

High-risk AI systems are classified in Annex III by their implementation context, which we believe also includes likely applications of rules as code systems.

- Social welfare: “AI systems intended to be used by public authorities or on behalf of public authorities to evaluate the eligibility of natural persons for essential public assistance benefits and services … as well as to grant, reduce, revoke, or reclaim such benefits and services”.
- Migration: “AI systems intended to be used by or on behalf of competent public authorities … to assist … the examination of applications for asylum, visa or residence permits and for associated complaints with regard to the eligibility of the natural persons applying for a status …”.
- Justice systems: “AI systems intended to be used by a judicial authority or on their behalf to assist a judicial authority in researching and interpreting facts and the law and in applying the law to a concrete set of facts, or to be used in a similar way in alternative dispute resolution”. The concept of a judicial authority is not limited only to judges and tribunals.

High risk systems are subject to elevated obligations of record-keeping, risk assessment, monitoring and review, human oversight, and disclosure. A specific risk of automation bias is identified in the legislation, which would also anticipate unjustified deference by human decision-makers to the output of a rules as code system.

## Illustrative requirements

Contemporary artificial intelligence systems have been examined from a regulatory and governance perspective for years now. While various jurisdictions, organisations and associations frame these requirements in different ways, there are some clear high level themes that are broadly accepted (see for example the OECD AI principles, the 2021 UNESCO Recommendation on the Ethics of Artificial Intelligence, and the requirements of the forthcoming EU Artificial Intelligence Act). We suggest the following regulatory themes are a useful starting point for governing the digital components of rules as code systems and we expand on this list in Appendix 2.

- **Transparency**: People should be made aware when they are being subjected to automated decisions or other forms of adjudication or filtering based on a system using rules as code.
- **Explainability**: A rules as code system should be explainable by reference to the legal criteria it purports to implement and how these criteria have been inferred from legal sources. People should be able to request a meaningful explanation of a decision as it applies to their circumstances.
- **Risk assessment**: Like other technological governance frameworks, a risk assessment framework is likely to be the expected approach to digital legal systems, requiring identification, management, mitigation, and remediation of risks. Increasingly, key risks are framed by reference to fundamental human rights, or to health and safety. Legal risks will also be particularly important for digital legal systems. The intensity and depth of risk assessment can be proportional to the nature, likelihood, impact, and remediability of the risks involved for specific systems.
- **Multistakeholder input**: Consistent with other models of Internet governance and digital governance, multistakeholder input can have benefits for enhancing the operation of a system and mitigating potential

risks, especially to marginalised groups.

- **Continuity across implementation phases**: Digital legal systems should be assessed at different phases of design, development, deployment, and post-deployment monitoring and review. Considerations and risks may be different at distinct phases, or require the input of different groups. Importantly, decisions made downstream from system design could have significant impacts on the ultimate legality and reliability of the system actually deployed, so continuity across phases through administration, record-keeping, and documentation is important.
- **Monitoring and review**: Once a system has been deployed, there are ongoing obligations to monitor for potential risks and errors, and obligations to periodically review the system to ensure that it remains accurate and reliable. If designers and developers are not the same entity as the deployer of the system, there is some obligation to make deployers aware of indicators to be monitored and the frequency of review required.
- **Accountability, review and appeal**: While a system may be automated to varying degrees, the person or organisation responsible for deploying the system is accountable for its outputs and any impacts it has. There may also be accountability for designers and developers, depending on the circumstances. Systems must be subject to a process of review and appeal by people subject to the system's influence. This is especially important where the system purports to apply legal criteria.

For greater adoption of rules as code and digital legal systems, each of these themes could be subject to much greater investigation in a rules as code context. Work by the COHUBICOL project in the European Union has examined this in the greatest detail to our knowledge and work by Hildebrandt, Diver and colleagues is an essential starting point.

## Governing legal components

In addition to the digital components of rules as code systems and the way they are implemented, it is equally important to consider the governance processes that determine how the legal components of the system are created, implemented, and reviewed. In our 2021 research, we turned our minds to scenarios where "highly reliable interpretations" of the law are developed and managed in ways that could be emulated in rules as code (see here and here).

The example we identified was a specific kind of standard form sale and purchase agreement which is the default document for dealing in residential real estate in New Zealand. Notably, it governs high value legal transactions in an economy where residential real estate plays a significant role, with traditionally significant volumes of transactions.

The document is published by the Auckland District Law Society (recently renamed "the Law Association"), and is generally referred to as "the ADLS agreement". The agreement is effectively an interpretation of the law converted into another format, which is reusable, widely shared, and deployed with a high degree of confidence that the understanding of the law embedded within it is beyond any reasonable dispute.

We asked ourselves what processes are in place that give New Zealand, its citizens, and its legal community such confidence that the way the ADLS agreement deals in property is reliable? These same features could be applied to the legal components of rules as code systems to generate similar reliability.

## ADLS Sale and Purchase Agreement

From research and interviews around the ADLS agreement, we identified a number of features which generate confidence in its interpretation and implementation of a range of legal instruments that apply to a specific kind of activity. Our analysis in the 2021 report was extensive, and that analysis has grown as we translate key features for the rules as code context. We have included more detail in Appendix 3, but key features are as follows.

- **Identity and status of drafters and reviewers**: The agreement sits within the custody of a specific committee of a professional association of legal professionals. The committee members' experience, qualifications and identities are a relevant factor for the agreement's reliability: for example, members of the committee write academic texts on land law, and are engaged in legal practice themselves.
- **Multiple legal inputs**: The agreement deals with the sale of land at a contractual level, but it also includes provisions related to taxation, and has a real effect for the transfer of title to land. The agreement therefore incorporates a variety of primary legal sources, including multiple legislative instruments, and case law. It does not necessarily do so implicitly, but it occasionally includes explicit statutory references. The agreement performs a specific task based on an implied comprehension of those legal primary sources.
- **Judicial review**: The agreement is a contract that can be subject to dispute resolution in the courts. The agreement has been subject to litigation in the past, including specific comment by the New Zealand Supreme Court, which made editorial suggestions for the agreement's revision in the case at hand. This substantially increases confidence in the agreement's reliability and accuracy.
- **Open sourcing**: Anyone with access to the agreement can see and assess its text and legal provisions. This accessibility is supported by the fact that legal professionals are still involved in real estate transactions, supporting meaningful comprehension among users. Because the agreement is widely available, a wide variety of parties (including banks, for example) know what the agreement contains. There is also dedicated professional training on the agreement as a legal instrument for real estate agents, lawyers, and others.
- **Multistakeholder input**: The copyright in the agreement is jointly held by the relevant law society, as well as the Real Estate Institute of New Zealand. It therefore reflects multi-stakeholder input and testing based on representative membership bodies. This ensures that issues with implementing the agreement are identified rapidly, and that any changes have the benefit of multi-stakeholder review.
- **Permissioning and version control**: The agreement is subject to copyright, which is enforced. Users who are not members of professional associations must purchase specific copies of the form. The administering legal body creates and sells specific software that allows for alteration of the agreement, but any alterations are performed by "red-lining", deliberately signalling which provisions have been amended, and how they

have been amended. At the time of our research, a disciplinary case was being pursued against a real estate agent who had amended the agreement without authorisation in a way that was not obvious to the other party. The committee manages amendments to specific versions, and versions are numbered for clarity and subsequent reference.

- **Clear and concise use case**: The agreement performs a very specific function - a contract for sale and purchase of residential real estate between two parties. Any additional functions it has are consequential on performing this task. For example, it includes an extra-judicial dispute resolution framework as well as an interim resolution for situations where disputes arise. This reflects the way that settlement of property transactions can exist in a long chain of events, which can be disrupted by a dispute or other barrier to completion. Interestingly, the repeatability and specificity of the agreement's function, as well as its reliability, means that non-lawyers can perform tasks that might otherwise be reserved for lawyers, with consequential growth in industries like conveyancing, as well as use of the agreement on a conditional basis by non-lawyers.
- **Longevity**: The agreement was first established in the 1960s or 70s and is now in its 11th edition. The agreement has responded to technological changes over time, as well as the increasing complexity of the legal requirements associated with sale and purchase of residential property. It was originally limited to around six pages and printed on specific paper from the United Kingdom, a quirky example illustrating the inevitability of technological change.

We are confident that there are countless examples of legal instruments like this that can be examined for further insights. In the process of designing digital systems that implement the law, or work within the law's restrictions, a necessary part will be the creation of a natural language specification that either re-states or interprets the way various legal sources apply to the given task at hand. Digital components will be developed from this "interpretation", and the ability to build confidence in the interpretation's reliability, and its traceability to its digital implementation, is a critical area of further development. In Appendix 3, we provide a longer list of features and suggest how these may translate into assurance mechanisms for the legal interpretations that inform the development of digital systems.

## An expansive view: digital legal systems

While we have focused primarily on automated decision-making systems in this discussion, we think the law as code area is actually much broader. This presents significant academic, legal, regulatory and commercial opportunities.

- If embedding law in digital formats doesn't require a whole statute to be coded, or every conceivable legal rule to be able to be represented in a specific format, then the potential scope of the kinds of law that can be digitally implemented is massively expanded.
- In addition, technical methods do not have to be able to represent every conceivable legal scenario. Instead, different technical methods can be used for different purposes depending on the legal framework and

operational context, expanding the technical approaches available for development.
- Law also structures behaviour in ways that are more complex and nuanced than decision-making by a decision-maker. This means the digital outputs that can be useful are much broader than decision-making systems, or systems that determine whether actions are lawful or not. For example, law can require the collation of data in particular formats and at particular intervals, or the sharing of that data in specific ways, and digital legal systems can be used to perform this function.

We observe that there are already signs of this approach being taken in digital regulation in the European Union. For example, the EU Digital Services Act regulates digital platforms and search engines and one of its primary regulatory mechanisms is to require transparency and information disclosure. The Commission has begun creating and implementing various documents, templates, data structures and databases to support DSA implementation using digital systems. To our knowledge, there is no public indication that the kinds of digital legal governance processes we have outlined here have been implemented. For example, assurance and documentation processes could be used to demonstrate how the categories in the statement of reasons database had been created from the DSA itself as a legal instrument. This presents an interesting opportunity for further investigation.

## Digital legal systems

Rules as code doesn't have to deal with all law, and it doesn't have to deploy perfect or comprehensive technical methods. Instead, all it has to do is transparently implement how the designer thinks the law works and apply that in a specific operational context. The tools used for doing so depend on the situation at hand and what the ultimate system is designed to do. In many situations, implementation may be through a dataset, or a data structure, not code or programming languages.

Because of this, we think of this area as being about "digital legal systems" rather than rules as code, because neither rules nor code are necessary. The discipline relates to how to design systems that are both digital and legal in nature: they are digital because they are implemented in digital systems, and they are legal because they are produced to do things governed by the law in ways that accurately reflect what the law requires, and may have legal consequences for the people using them.

## An example of a digital legal system

From 2022-2024, we worked with New Zealand's drinking water regulator, Taumata Arowai. It had developed a legal instrument - secondary legislation - designed to require drinking water suppliers to test drinking water for particular substances and measurements, to record that information, to assess whether it falls within particular limits, and to share that information with the regulator. Suppliers were also required to perform assurance checks to ensure their systems for performing these measurements were functioning properly.

The rules that governed this process were published in a PDF, which included a range of legal rules, as well as a proliferation of different tables (also published in PDF). The tables indicated what should be measured, how frequently it should be measured, how frequently compliance should be assessed, when that information should be transmitted to the regulator, and sometimes the maximum allowable values for any particular measurement. The rules were a relatively self-contained document, but they also relied on other external legal instruments for setting the maximum values for some substances, as well as the way those substances should be measured. These were recorded in a separate set of regulations. The overall authority to request the information, the obligation to provide it, and the ability to specify the format of the information provided was set out in a separate piece of legislation.

None of these rules, tables, formats, or measurements were meaningfully accessible to a computer system. In particular, the referencing and numbering system used was not granular enough to be used in a dataset, and there were no digital assets that could be used to embed all of these parameters and associated data structures directly into a computer system. This is despite the fact that the largest, most high risk water suppliers would universally make use of automated digital systems for measuring drinking water, would create records using digital systems, and would then have to transmit those to the regulator in other digital formats.

From this web of legal documentation, we created a single data structure that could be used in digital systems to organise and receive all relevant measurements and data from any supplier. We created documentation and a referencing system that ensured all fields could be traced back to the source material and regularly verified whether we accurately understood the reporting process in situations where the law wasn't clear. From this data structure, it was possible to create an API specification, which allowed suppliers to report all data via API.

Because we had digitised the rules document, it was possible to create a "rules builder" that let people understand what their obligations were based on particular thresholds they met. It also allowed them to explore situations where they may opt for a higher level of compliance for whatever reason. All the relevant substances and measurements that were required under the rules and the regulations, as well as the units of measurement required, the frequency at which those things should be measured, and when they should be reported were also published as a dataset. This enabled people to view only the rules that were relevant to them, rather than having to process the full rules document.

We created a digital system for implementing the law as the regulator understood it, without requiring the use of particular programming languages, and without needing to translate an entire piece of legislation into a machine executable format. The ultimate outputs we created embedded an understanding of the law and "the rules" following a cautious and responsible multidisciplinary process, which could have been audited, explained and substantiated at any stage. The potential uses of the system also extend well beyond basic automated decision-making into a kind of traceable digital infrastructure. This system can also be published and transparently amended, or defended in legal proceedings. The digital legal system can be used as the

fundamental infrastructure for a whole array of digital products - developed by the regulator, the public, representative organisations, or commercial entities.

## Parallel drafting for digital legal systems

A forceful conclusion from anyone who has attempted to convert legislation into computer code is that it is a difficult task. It is usually made more difficult by ambiguities in the overall legislative scheme. Edge cases and common sense interpretations that are tolerable in natural language are generally intolerable to computer systems. Schemes that will inevitably be implemented digitally have not always been designed with this fact in mind. The process of coding the law quickly surfaces all of these issues. This led a cross-government group of people in New Zealand to conclude that a superior approach would be to perform parallel drafting, where natural language rules (in the form of legislation) were formulated as part of the policy development process, which would also be simultaneously modelled in computer code (the "better rules approach").

Parallel drafting is a novel process because typically legislative drafting comes once a policy has already been fully developed, and there is only limited scope for drafters to clarify ambiguities, or recommend corrections. Multidisciplinary parallel drafting enables service design and implementation considerations to be brought into the policy development process because of the way that legislative drafting and computer programming require more detailed attention to logical sequences and process flows. Essentially, policy practitioners working in multidisciplinary contexts with people responsible for implementation are required to articulate how their designs would actually work in practice, rather than simply leaving that to people responsible for implementing it. This leads to a more coherent system that was designed for implementation in digital systems, and minimises the risk that policy implementation undermines the original policy outcomes being pursued.

Parallel drafting could have significant benefits for digital legal governance because it will very likely make law easier to implement in digital systems, and because traceability between law and digital systems can be designed into the assets produced from the outset. Importantly however, any digital assets produced through parallel drafting will not have the same legal authority as the law itself. The code produced during the drafting process is simply an indication of what policy makers and legal drafters intended to achieve. There remains a risk that the words they used in the legislation take on different meanings over time, or in response to new circumstances, or because they may conflict with other laws and principles.

This capacity for natural language rules to evolve over time is desirable, and must be preserved from a rule of law perspective. However, for this reason, judicial interpretation of natural language legal instruments remains the authoritative source of law, and even parallel drafted digital instruments can diverge from that interpretation.

### Conclusion on digital legal systems

By adopting this framing of "digital legal systems", we can more accurately talk about a wider range of potential applications and work toward realising the significant opportunities such systems present. At the same time, our attention must always remain on the legal character of such systems and their potentially significant impacts. Effective governance processes, as well as innovative approaches to design and review over time are a critical part of realising this potential.

# Conclusion and next steps

To conclude, we offer some suggestions for further areas of investigation and research, as well as sharing our experience of the most important steps in developing digital legal systems.

## Developing body of practice and research

For researchers, students and academics, we think there is a productive area of inquiry looking at the following areas.

### Governance of legal documents and legal instruments

We are confident there will be longstanding best practice frameworks for governing the legal components of digital legal systems, such as documents or legal instruments that interpret how law applies in particular operational contexts. Documents or instruments of this kind are a desirable and necessary component for illustrating how law has been interpreted to form a specification that can be implemented in a computer program. Authoritative standard form legal agreements such as the sale and purchase agreement for residential property are one example. Sources of insight could come from declaratory judgment procedures, the management of multilingual legal documents, and practices around developing standard form documents and legal templates.

### Governance of digital systems

The software industry has deep experience in governing digital systems to ensure they are reliable and secure. Further governance insights could be drawn from frameworks for safe deployment of artificial intelligence and automated decision-making systems and how higher level principles might be applied to digital legal systems. Useful sources of insight could come from AI risk and assurance processes, management of business rules in rules engines, and the use of either decentralised or centralised software repositories such as GitHub, along with open source governance practices.

### Applied research

There would be merit in performing applied research demonstrating how digital systems for implementing the law in particular operational contexts can be traced back to legal documentation, including the law itself, and how this documentation can be made available for audit and assurance purposes. Potential sources of insight

may come from software documentation practices (including existing tooling used for this purpose), the development and assurance processes around the OpenFisca implementations published in relation to French law, as well as the practices used in business rule development and audit and assurance frameworks.

### Parallel drafting processes

There is more work to be done in the theory and practice of parallel drafting approaches as a governance mechanism where policy frameworks, legal rules, and software implementations are designed at the same time with a view to increasing traceability between them, and enhancing governance processes. Potential sources of insight may be drawn from drafting processes for multilingual legal documents. Interesting work is being done by technologists with experience at digital platform companies to use parallel drafting processes to enhance the reliability of LLMs for content moderation.

### Backwards compatibility

There will be numerous systems already in operation that perform a similar function, or have similar features, to rules as code digital legal systems. They will not have been designed in the way we describe here, and are unlikely to have been designed with the governance features we describe. A productive area of work would be to identify those systems, assess them against these criteria, and consider both whether these criteria are adequate, or whether the systems require reform.

## Key governance questions in our experience

From several years experience, as well as observing rules as code projects and discussing them with participants, we advise careful consideration of the following questions at the outset and expand on these in Appendix 1.

1. **Purpose and use case**: what does the system you're producing actually do? Scoping the work has important follow on effects for governance.
2. **Resourcing**: what level of human and financial resourcing is available to support the project? What skills and expertise are available?
3. **Governance**: who is empowered to make legal and regulatory decisions, what review and oversight procedures are required, and does the accountable decision-maker fully understand the relevance and function of the digital and legal components?
4. **Documentation and assurance**: people developing the system will be empowered to work faster and with greater certainty if they know how to record decisions they are making and revisit them if necessary. How will key decisions be captured as they relate to legal, interpretive, and digital components?
5. **Copyright and ownership**: who will own the legal and technical documentation produced, or the code outputs? Copyright can be an important mechanism for maintaining integrity of the system, even if it is ultimately free and open source.

Asking and answering these questions at the outset will considerably improve the chances of a project's success and longevity.

### Focus on use case and purpose

In our experience, an important step is identifying a specific use case for rules as code implementation, or a digital legal system.

- Identifying a use case is a critical component of narrowing down the scope of the legal and technical work required - for example, it allows teams to rule out whole parts of a piece of legislation, or even whole statutes.
- The end use-case for a system will also shape the interpretation decisions being made because the impact and effects of interpretive decisions and design choices can be more reliably predicted when anchored in reality, rather than the abstract.
- Establishing a specific use case is also a critical component in taking a rules as code project from being an interesting, experimental, or academic exercise to one that has real world consequences. This means it can be costed and justified when it comes to resource and time investment.
- A specific intended use case for any eventual system also permits firmer answers to be given to most governance questions that may arise, including who will be accountable for the system and how it will be reviewed over time.

In combination, questions around end use, available resourcing, and governance and accountability also heavily influence how the project is approached and who is accountable for delivering it. Failure to anticipate these questions can be fatal to a project's longevity.

## Invitation to partner and collaborate

Syncopate is working to develop methods, products and services that maximise the potential of digital legal systems and build confidence in the way they're developed. We want to work with clients, partners and others and you can find out more by following us on social media, or getting in touch on the Syncopate Lab website.

---

# Appendix 1: Starting points for developing, implementing and governing digital legal systems

From our experience developing these systems and reflecting on how they should be governed, we advise close consideration of the following questions for existing and prospective practitioners, or for people who may commission the development of such systems.

| | | |
|---|---|---|
| Purpose and use case | What does your system actually do? What will be the operating context, and who is using it? What are its inputs and what are its outputs? | Determines level of risk and assurance required as well as level of investment that is justified. Also influences scope of legal material involved and technical design choices which will have budgetary ramifications. |
| Resourcing | What level of human and financial resourcing is available to support the project? What skills and expertise are available? | Formulating a legal interpretation that is capable of digital implementation requires legal and technical skills which can be scarce and expensive. High risk systems must also be designed and developed properly, with governance considerations in mind. Few shortcuts are possible. |
| Governance | Who is empowered to make specific decisions? How will your governance process be incorporated from the start? What review process is required for changes to legal specification, technical specification, or alignment assurance? | In our experience, spontaneous collaborative development of rules as code models can be difficult, and seldom include governance models. |
| Documentation and assurance | How will all key design decisions be captured? How can those decisions be traced to implementation decisions? If those decisions are revised, how can the implications of those revisions be traced efficiently? | Documentation and assurance processes should be implemented at the outset and subject to quality assurance and oversight. Transparency is a fundamental consideration especially for code in automated decision-making or high risk contexts. |
| Copyright and ownership | Who will own the legal and technical documentation produced, or the code outputs? | Copyright is an important mechanism for exercising control over the integrity of the system, even if the system is intended to be made freely available and open source. Copyright is also important where an individual or organisation may be legally accountable for the system’s impacts. |

# Appendix 2: AI regulation and digital governance

In many respects, digital systems that implement the law are indistinguishable from the kinds of AI systems that are subject to so much public scrutiny and regulatory attention today, differing mainly in the way that an absence of statistical components may change the risks that arise and the governance arrangements that follow. In the left hand column we have laid out key components of many AI regulatory frameworks and statements of ethical principles. In the right hand column we have suggested how this might flow through to governance of the technological components of a rules as code, or digital legal system.

| **AI governance theme** | **Example implementation in rules as code or digital legal systems** |
|---|---|
| Transparency | People should be made aware when they are being subjected to automated decisions or other forms of adjudication or filtering based on a system using rules as code. |
| Explainability | A rules as code system should be explainable by reference to the legal criteria it purports to implement. People should be able to request a meaningful explanation of a decision in their circumstances. They should also be able to test it and ask for a review of any decision made. |
| Purpose | Systems must be designed for a specific purpose, and that purpose should be clear to the user or deployer so that risks of harm can be avoided from deploying the system in ways it was not intended to be deployed. |
| Risk assessment | Like other technological governance frameworks, a risk assessment framework is likely to be the expected approach to digital legal systems. |
| Multistakeholder input | Consistent with other models of Internet governance and digital governance, multistakeholder governance can have benefits for enhancing the operation of a system and mitigating potential risks. |
| Design | Obligations to consider risk and safety exist at the design phase, not just at the point the system is implemented. Design with multistakeholder input is important for generating insight into the way a system may have inequitable impacts, or disproportionate impacts on a particular group. |

| | |
|---|---|
| Development | The way a system is developed can introduce errors or risks that were not apparent during the design phase. Equally, developing a system without adequate knowledge of the potential risks involved in deploying the system, as anticipated in the design phase, can lead to risk mitigation measures being undermined. |
| Monitoring and review | Obligations on people designing, developing and deploying digital legal systems do not necessarily cease at the point the system is implemented. There are ongoing obligations to monitor for potential risks and errors, and obligations to periodically review the system to ensure that it remains accurate and reliable. If designers and developers are not the same entity as the deployer of the system, there is some obligation to make deployers aware of indicators to be monitored and the frequency of review required. |
| Accountability | While a system may be automated to varying degrees, the person or organisation responsible for deploying the system is accountable for its outputs and any impacts it has. There may also be accountability for designers and developers, depending on the circumstances. |
| Bias and non-discrimination | Discrimination may be a lesser risk without the use of statistical models, but systems can still be designed and implemented in ways that discriminate. For example, interacting with a system may require a certain kind of identification or access to a particular kind of computer system that excludes certain groups. Equally, systems may require the use of personal information in ways that are discriminatory, even if they otherwise implement the law reliably. |
| Proportionality based on risk | The stringency and comprehensiveness of risk assessment and mitigation measures can scale up or down depending on the level, likelihood, and kind of risk a system creates. Higher impact systems should be subject to higher levels of oversight and governance. |

| | |
|---|---|
| Diversity and accounting for marginalised communities | Certain risks will be invisible or underweighted to some people depending on their experiences and expertise, and those same risks may have a disproportionate impact on marginalised groups. These groups may be under-represented in teams and organisations working on digital legal systems. As a result, incorporating marginalised groups in the design, development, deployment, monitoring, and review processes can be an important component of effective governance. |
| Product disclosure and product liability | Software products are no different from any other kind of product insofar as liability is concerned. Disclaimers of legal liability may be required and certain kinds of risks ought to be disclosed to purchasers or end users to mitigate risk of harm. |
| Legality | Systems should operate within the law, and even if they reliably reflect the law they are intended to implement, there is still a risk they may infringe other laws, such as copyright, privacy, or data protection law, as well as other administrative law related to decision-making by public authorities. |
| Review and appeal | Systems must be subject to a process of review and appeal by people subject to the system's influence. This is especially important where the system purports to apply legal criteria. |
| Security | Digital legal systems must be designed with security in mind, including in ways that prevent unauthorised amendment or use. |
| Continuity | Digital legal systems should be assessed at different phases of design, development, deployment, and post-deployment monitoring and review. Considerations and risks may be different at distinct phases, or require the input of different groups. Importantly, perhaps even more than other digital systems, decisions made downstream from system design could have significant impacts on the ultimate legality and reliability of the system actually deployed, so there is a need for continuity across each phase. |

| | |
|---|---|
| Standardisation | There will be some benefit to standardisation of system design and governance, including potential use of industry technical standards, or governance standards across jurisdictions. Standards can also play an important role in adding detail to otherwise broad legal requirements. |

# Appendix 3: ADLS sale and purchase agreement and legal governance

| **Observation about ADLS agreement** | **Relevance for rules as code governance** |
|---|---|
| The Agreement is drafted in natural language, but it reflects a workable operational interpretation of multiple legal instruments that increases the parties' compliance with and knowledge of the law. | • Wording of law can and must be reformulated at times into other words or formats, and this can be done in a satisfactory way for large numbers of high stakes transactions.<br>• When law is converted into other formats, for a particular use, it should be done in a way that is tightly aligned to a specific task, rather than in a generic or abstract sense.<br>• The users of an encoded interpretation should be able to comprehend to some extent, and make sense of it in their own minds, even if expert advice is required to meaningfully achieve that. |
| The Agreement draws on a wide range of primary legal sources. It is not only a reflection of land law, but also taxation law (in the way that GST is incorporated into land sales). | • For any legal task, there will be more than one source of law influencing behaviour.<br>• Creating encoded models of the law will, in all likelihood, require multiple statutes and sources of law to be incorporated into the interpretation/implementation, and focusing on just one statute is likely to be insufficient. |
| There is wide confidence in the reliability of the Agreement because of the way that it is produced and because of the qualifications of the people who produce it and monitor it. | • Encoded interpretations should be produced according to a transparent and justifiable process that can be used to persuade people they have been produced safely and reliably. The qualifications, expertise and experience of the people producing it will be relevant. |

| | |
|---|---|
| The Agreement is capable of being assessed by the judiciary and updated to reflect statutory amendment, judicial interpretation, and the impact of case law. | • Encoded interpretations must be able to be updated in order to remain reliable.<br>• Encoded interpretations will need to be subject to some form of ongoing monitoring and update schedule to ensure they remain current, with mechanisms in place for challenging or querying them in a way that is not unduly inhibitive.<br>• Non-developers will need some method of assuring themselves that the interpretation is an accurate reflection of the law (as well as that it is functioning as intended). |
| The Agreement includes its own dispute resolution mechanism. Parties who agree there is a dispute can refer it to experienced property lawyers for resolution. | • Encoded interpretations should anticipate, and perhaps specify, mechanisms for disputing how they operate and their accuracy and reliability. These mechanisms should remain current so long as the interpretation is to remain available for use. |
| Copyright in this case is an essential legal device for controlling how the agreement is used or modified. It is used to ensure that the utility of the standard form is not undermined. | • Both text and software can be subject to copyright and intellectual property protections. While open source code may be desirable for a number of reasons, when it comes to the accuracy and reliability of an encoded interpretation, there may be some merit in some person or organisation retaining intellectual property in order to preserve the integrity of the interpretation.<br>• Copyright in encoded interpretations could apply not only to the code produced, but also to any other documents that specify how the code should operate. |
| The agreement does not exhaustively state the law, nor is it held up as having greater authority than the other primary legal sources (or even secondary legal sources in the form of academic commentary) that inform its drafting. It is reproducible and scalable in the way that many copies of it can be produced and used rapidly. | • Encoded interpretations will be subordinate to primary legal sources.<br>• Encoded interpretations are not comprehensively complete statements of the law, they only embed particular legal rules and conclusions relevant to the particular task at hand. They may include implied legal matters, as well as explicit ones.<br>• Encoded interpretations can benefit not just from primary sources of law (statutes, case law, regulations) but also from secondary analysis from experts, such as from text books or written analysis.<br>• The widespread use of the ADLS agreement demonstrates that there is utility and demand for repeatable, reproducible, reliable, reusable instruments and artefacts that enable high volume transactions to be performed more efficiently. |

| | |
|---|---|
| The agreement is drafted and revised by a committee of the ADLS convened for that purpose. Membership of the committee comprises legal practitioners and academics with significant authority on the area of land law in New Zealand, including the author of the leading academic text. | • There is some benefit to having more than one person involved in production of an encoded interpretation.<br>• To facilitate coordination, some form of procedure or organising body is desirable, and this body itself is likely to be subject to procedural rules and accountability mechanisms.<br>• Membership in this body need not be restricted to lawyers, but to generate credibility, it is desirable to have lawyers involved. If lawyers with particular expertise and credibility in a specific subject are included, this is beneficial for the reliability of the encoded interpretation. |
| The New Zealand Supreme Court has made comment on the drafting of the Agreement. In fact, judicial comment has led to amendments to the Agreement in revised editions. | • When producing encoded interpretations that have legal effect, designers and developers must anticipate that the interpretation will be challenged in judicial proceedings.<br>• When encoded interpretations are challenged and subject to judicial assessment, their function and accuracy must be in a format that permits judicial assessment.<br>• One component of anticipating judicial assessment for legal accuracy is that the details of the encoded interpretation may not be able to be kept secret. |
| The Agreement has been in use since the 1960s or 70s. It is now in its 11th edition. One of the early versions of the Agreement was only six pages long. It was limited to six pages because it was printed on specific paper imported from the United Kingdom, because of the way it could be folded. Over time, the agreement has had to account for a wider number of legal instruments and greater prescriptiveness in dealing with the rights and obligations of the parties. | • Highly reliable encoded interpretations can and should be designed for long term use. If successful, they could be used - with amendment and review - for decades.<br>• Any process of amendment must include a procedure for sign-off and approval. It must also include some form of numbering system that allows different versions to be distinguished from each other.<br>• Encoded interpretations are likely to grow in complexity over time in response to changing socio-economic circumstances, technological development, and adoption in changing contexts.<br>• Ideally, the technology used for recording an encoded interpretation should be capable of conversion to another technological format to account for changes in technology over time. |

| | |
|---|---|
| The Agreement is produced by the Auckland District Law Society (now called the Law Association), not the New Zealand Law Society. The ADLS is the sole remaining regional law society after a period of consolidation, although it has members from across New Zealand and is not limited to the Auckland region. The NZLS has regulatory functions, whereas the ADLS does not. | • Producing a highly reliable encoded interpretation need not be a task done by a legally appointed authoritative national body.<br>• Highly reliable encoded interpretations can be produced solely for the benefit of the people who produce them: they need not be oriented solely to the public interest. |
| The ADLS agreement was originally drafted only by the ADLS. Subsequently, the Real Estate Institute of New Zealand became involved in the production of the agreement based on the experiences of its members and its members' interests in a useful document. | • Representative bodies of professionals or a particular industry can be useful participants and drivers of highly reliable encoded interpretations.<br>• Users of an encoded interpretation can bring important insights to how it works, whether it's working, and areas for potential improvement to reflect real world practice. |
| When interviewees were asked why anyone treats the agreement as being legally reliable, they pointed to the reputation and qualifications of the members of the committee. They also emphasised that the agreement has been reviewed judicially in disputes between parties to a transaction over the years. | • Speaking to existing users of highly reliable interpretations like the ADLS agreement can generate important insights into what generates confidence in their reliability.<br>• The qualifications and status of the people producing the interpretation are relevant, as are the existence and effectiveness of quality assurance and oversight processes. |
| The Agreement is subject to copyright jointly held by REINZ and ADLS. The copyright is enforced. It is not clear how the revenues generated by this copyright interest are used, but members of the Real Estate Institute of New Zealand gain access through their membership in the Institute. | • Intellectual property in encoded interpretations could be owned by industry bodies and membership organisations, including jointly between such organisations.<br>• Encoded interpretations can present a revenue source, presumably to facilitate maintenance, development and oversight. |

| | |
|---|---|
| We note that there was a case recently where a real estate agent used software to illegitimately modify the terms of the copyright agreement in a way that was not obvious to other parties to the transaction. The modification was only noticed by a lawyer shortly before the agreement was signed and led to disciplinary consequences for the agent. | • Unauthorised amendment to encoded legal interpretations can have significant consequences for people using them.<br>• Unauthorised amendment to encoded legal interpretations could be something that, in particular contexts, leads to professional disciplinary consequences.<br>• Encoded legal interpretations could, if amenable to unauthorised amendment, be used to mislead people about the state of the law, potentially with fraudulent consequences.<br>• There is a role for individual users' own lawyers or representatives in identifying errors, inaccuracies, and unauthorised amendments to encoded interpretations. This should be factored into any quality assurance and oversight process and, again, may require the terms of an encoded interpretation to be open to external scrutiny (ie, not secret). |
| The ADLS publishes software that facilitates legitimate amendment of digital versions of the Agreement. This software replicates the way amendment would occur with a paper copy. In other words, an amended term is struck through with a line and the revised text is included alongside the original text. This makes it obvious to the reader when the original document has been amended. | • Dedicated software can be produced by the organisation who exerts custody or stewardship over an encoded interpretation in order to permit, within conditions, amendment to that encoded interpretation.<br>• Where any amendment is made to an encoded interpretation, there should be deliberate signals included by default to make it clear to users how that interpretation has been altered - including by retaining the original code, and presenting it alongside the updated code. |
| The core benefit of the Agreement is that it makes it possible for users (including lawyers representing clients) to immediately know the contents of the agreement being contemplated by the parties. This saves time for the practitioner, and money for the client. It also allows bodies of practice and expertise to be developed around that specific agreement, including seminars and lectures, or commercial products. The predictability of the agreement also flows through into organisational workflows in legal practice. | • When reliable encoded interpretations are created, it can generate massive time savings, as well as significant increases in legal certainty, because the content of the interpretation is predictable and reliable without any further examination being required.<br>• Encoded interpretations could form the basis for whole workflows in businesses and professional organisations that use them. Where encoded interpretations become particularly well entrenched, dedicated expertise in that particular code model could itself be a valuable skill-set, and a matter of dedicated training and education. |

| | |
|---|---|
| Around 2008-2009, REINZ formed the view that the Agreement could be improved by re-drafting it using plain language drafting techniques. The intent was that non-lawyers would have a better understanding of the Agreement. There have been similar attempts over the years to produce a plain language version of the Agreement. Ultimately, even where they have been completed, the plain language re-drafts have never been widely adopted. Interviewees also noted that there had been some suggestion that the fundamental general clauses in the Agreement could be adopted or acknowledged in legislation. This has not been pursued. | • At times, we can expect that people will wish to convert encoded interpretations to other formats, or to refine them so they are more easily able to be understood, explained, and deployed, including by non-experts.<br>• When encoded interpretations are subject to a process of refinement or revision, even if they are intended to function identically, there may be reluctance to accept a revised version. Where encoded interpretations are particularly entrenched, it may be difficult to displace them, even if that would be desirable, and where any risks can be managed.<br>• Where encoded interpretations become particularly entrenched, there is a chance that they may be considered for incorporation by reference into primary or secondary legislation. |
| Parties who agree there is a dispute can refer it to experienced property lawyers for resolution. This mechanism is designed to allow property transactions to settle without parties relinquishing their rights to pursue a remedy for any breach, and to resolve disputes without litigation. It reflects the complex transactional environment in which the Agreement is used, where chains of transactions might settle all at once, and a disruption in one of these may affect the ability for settlement to occur on a wholly unrelated transaction. | • Encoded interpretations may play a complicated role in a broader sequence of legally significant events, which could have consequential impacts for people unconnected to the transaction or activity being managed by the encoded interpretation itself. There may be a need to design mechanisms or holding patterns into the encoded interpretation in order to ensure that errors or disputes do not have large scale consequential impacts.<br>• Encoded interpretations ought to include a specific dispute resolution mechanism, and there may be benefit to having a first-instance extrajudicial dispute resolution mechanism in order to promote efficiency, predictability, and certainty. |
| The Agreement is updated and amended according to the procedures of the Committee. The Committee periodically seeks input from members of the profession and the public. There have been instances where members of the public or their legal representatives have suggested issues caused by the Agreement, or improvements for subsequent editions. Academic members of the committee have also produced papers on proposed revisions to the agreement. REINZ also has input based on the experience of its members in using the agreement. | • People developing or administering encoded interpretations may benefit from exposing them to public consultation processes, including to the public, to professional organisations, and to specialist practitioners and users. There can be genuine sources of valuable insight that may lead to amendments.<br>• Encoded interpretations may be subject to academic comment or scrutiny, and there may be complex discussion about whether an amendment should be made or not. |

| | |
|---|---|
| Conveyancers have status as a separate class of professionals who, along with lawyers, can effect changes to the title of land in New Zealand's land registration system. | • When sufficiently reliable and user-friendly, encoded interpretations could permit identified classes of people to gain higher levels of authority to use them in order to facilitate legally significant activities. Non-specialists may become capable of performing tasks normally reserved for specialists, and this could receive legal recognition, or found new businesses and industries. |
| No individual or agency carries any legal liability for the accuracy of the document. The responsibility for providing legal advice and meeting client obligations still lies with legal practitioners. Any person who transacts using the Agreement without taking legal advice risks legal dispute. | • Even where custody and authorship of a highly reliable encoded interpretation is taken on by a person or body, it is likely that the interpretation will still be subject to some kind of disclaimer or exclusion of liability, due to the fact that the circumstances and details of a particular user cannot be fully known.<br>• A highly reliable interpretation used in one context by a specific user may be accurate and safe, but when used in another context by a different user, may be unreliable.<br>• It is likely that, even if highly reliable encoded interpretations can be developed, individual users would benefit from dedicated specialist advice to protect their interests. |
| It is possible for additional clauses to be added to the agreement, but these might undermine the integrity of the agreement as a whole: for example, an additional clause might contravene one of the standard clauses in a way that does not clearly indicate how that inconsistency should be resolved. | • While individual users, or their specialist advisers, may be authorised to add code to the interpretation without amending existing code, it is impossible to exclude the possibility that these additions could render the original code unreliable. On that basis, even additions (intended to supplement the original without alterations or amendment) should be treated with caution. |

***Acknowledgement***

*This research/project is supported by the National Research Foundation, Singapore under its Industry Alignment Fund – Pre-positioning (IAF-PP) Funding Initiative. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore.*