

THE RESEARCH UNIVERSITY PRESS

International
Computational Law Journal

VOLUME 2 / ISSUE 1 / CURATED SOURCE EDITION

Identifying a Set of Autonomous Levels for AI-Based Computational Legal Reasoning

Lance Eliot

MIT Computational Law Report · Originally published 2021-12-07

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 •

Identifying a Set of Autonomous Levels for AI-Based Computational Legal Reasoning

Lance Eliot

Published on: Dec 07, 2021

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

ABSTRACT

Computational law is currently bereft of a defined set of autonomous levels for differentiating AI-based capacities in legal reasoning. This paper provides a strawman proposal for establishing a standard for such autonomous levels and seeks to activate a vital and vigorous discussion to aid in progress toward putting in place a suitable and industry-wide Levels of Autonomy (LoA) framework for AI-based Legal Reasoning (AILR) systems.

Introduction

Standards are crucial to the maturation and advancement of many endeavors. The lack of an established standard can undercut demonstrative progress in ways that might not at first glance seem apparent.

For example, until the standardization of freight containers, global trade was stymied by a plethora of different sizes and shapes of containers for shipping goods (see Lewis [4]). Incompatibilities existed across a vast network of interdependent shipping logistics since each port would need to be prepared for a wide variability or have to make a Hobson's choice of what type of containers would be received and processed. The same confounding issues confronted trucking companies, warehousing and handling facilities, and generally cascaded throughout all salient aspects of shipping.

By the seemingly simple act of identifying and putting in place acceptable standards for freight containers, the entirety of large-scale shipping was tremendously aided, including the garnering of new efficiencies, reduced costs, streamlined logistics, faster processing times, and so on. Though the act of creating and instituting a standard for shipping containers might appear to be a straightforward matter, there was a great deal of discourse and angst involved. Thus, notably, the path toward a globally acceptable standard is rarely an easy one.

In the field of computational law, there is not yet a standard for a burgeoning area of both theoretical research and practitioner application, namely the realm of AI-based legal reasoning systems. In brief, there is an existing lack of established levels-of-autonomy in the field of computational law. This omission or gap is not readily apparent, rarely overtly recognized, and insidiously permeates all manner of efforts underway.

Consider some illuminating examples.

Suppose an academic research lab professes that they have made a breakthrough in Machine Learning that can boost legal reasoning systems. This is likely admirable and advantageous, but to what degree has their innovative effort made progress toward achieving autonomous capacities? Without some form of measuring stick, there is no ready means to assess whether this advancement has been substantive and game-changing or might be more so incremental.

Envision a marketplace example in which a commercial vendor of LegalTech products boasts that they have attained a remarkable infusion of Natural Language Processing (NLP) into their latest legal-oriented wares. Again, the question arises as to how much of a demonstrative change this makes toward a legal reasoning system that can operate either semi-autonomously or fully autonomously.

All told, though it might not be obvious to the naked eye, the missing ingredient, or crux of this conundrum, is that without an identified and promulgated standard for legal reasoning system autonomy, the ability to gauge new advances and discern the magnitude or degree of progress is woefully undercut and altogether problematic. This is especially beguiling since there is at times a tendency to overstate asserted advances, and thus discerning the wheat from the chaff can be overly arduous and not immediately satisfiable (see the points articulated by Linna [5]).

The mainstay of the question that arises here involves the degree of intelligent-like behaviour that has been achieved and whether the latest applied system or theoretical research discovery provides an improvement over prior instances. plus If so, what is the magnitude of the advancement so incurred. Merely stating that there is AI involved is vacuous and insufficiently able to demark the work that has been done and the accomplishment resulting from that work.

A key reason for this difficulty is due to ambiguities of the AI moniker per se, proffering a rather broad and vague umbrella term that is ostensibly amorphous and lacking in any substantive delineation of what the advanced automation constitutes (see the background in Markou [6]). What is needed to rectify this ambiguity is a type of numeric Richter scale that denotes the level of AI that has been infused into a legal system.

As such, having a definitive and standardized set of Levels of Autonomy (LoA) for AI-based Legal Reasoning (AILR) systems would usefully and succinctly provide a rigorous means to denote its capabilities. In short, the everyday use of a universal scale would demonstrably aid in unravelling and rationalizing the claims made by all;

and in doing so can achieve a definitive indication of the latest results forthrightly achieved.

Computational law needs a variant of the now-venerated freight container standardization.

That being said, and to clarify, establishing levels-of-autonomy is a far more intangible, arguable, and controversial matter than the sizes and shapes of shipping crates. One should anticipate that reaching concurrence on LoA for AILR will be a heated and quite contested machination. Despite a presumed scenario of protracted discourse, or perhaps because of it, the time seems ripe to proceed down the path toward levels-of-autonomy while the field is still in its nascent state and for which such standardization can imbue early-on benefits akin to those analogous to the freight containers standardization.

Background On Levels Of Autonomy In The Law

It is customary in the legal reasoning context to divide the conjoining of AI and computer-based systems into two focuses. The first being the application of automation to the act of legal reasoning, which primarily then serves as an adjunct or augmentation to human legal reasoning efforts. The second being the goal of achieving autonomous legal reasoning that consists of computer-based systems able to perform legal reasoning unaided by human legal reasoners and that can operate autonomously with respect to the practice of law. As per Galdon et al [2]: “Automation is defined as a system with a limited set of pre-programmed supervised tasks on behalf of the user. Autonomy, on the other hand, is defined as a technology designed to carry out a user’s goals without supervision with the capability of learning and changing over time.”

Law practices and legal professionals routinely today make use of automation in the performance of their needed tasks involving legal activities. A modern-day law office might use e-Discovery software as part of their case discovery pursuits, along with crafting new contracts via the use of an online cloud-based contract management system.. Generally, sensible adoption of law-related computer-based systems has significantly aided lawyers and legal staff in undertaking their efforts (such systems are collectively referred to as LegalTech) and has been cited as boosting efficiency and effectiveness accordingly. The automation being used for these purposes is not considered autonomous, though advancements in these systems are being fostered by infusing AI capabilities to someday achieve autonomous operation (see Surden [10]).

A significant body of research exists on attempts to generically clarify what constitutes autonomy or autonomous operations. There is much debate regarding the particulars of autonomy and different viewpoints ascribe differing qualities to the matter. For example, Sifakis [9] defines that “autonomy is the capacity of an agent to achieve a set of coordinated goals by its own means (without human intervention) adapting to environment variations.” A fully autonomous computer-based system in this computational law context would be one that can perform legal reasoning on its own, doing so without the aid of a human, and essentially perform legal reasoning that is on par with that of a human that is properly-versed in legal reasoning.

In devising a levels-of-autonomy framework for AI-based legal reasoning, we can reuse prior efforts of devising an LoA and reapply judiciously the results into the particulars of computational law. One of the most widely accepted and well-known levels-of-autonomy efforts consists of the Society for Automotive Engineers (SAE) J3016 Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles [8], which is globally recognized and known worldwide as a definitive standard LoA for Autonomous Vehicles (AVs) that are self-driving cars.

The SAE LoA posits these six levels of autonomy:

> Level 0: No Driving Automation
> Level 1: Driver Assistance
> Level 2: Partial Driving Automation
> Level 3: Conditional Driving Automation
> Level 4: High Driving Automation
> Level 5: Full Driving Automation

Some immediately notable aspects include a numbering scheme that ranges from a low-to-high indication. Starting with zero, there are six designated levels, the naming of the levels is intended to approximately reflect succinctly the nature of the levels, and each level per the details of the standard is considered separate and distinct from the other levels. This same overall format and convention will be reused and adapted for the LoA AILR.

It is worth noting that core guiding principles underlying the formulation of the SAE standard were officially stated as: “1. Be descriptive and informative rather than normative. 2. Provide functional definitions. 3. Be consistent with current industry practice. 4. Be consistent with prior art to the extent practicable. 5. Be useful across

disciplines, including engineering, law, media, public discourse. 6. Be clear and cogent and, as such, it should avoid or define ambiguous terms."

Note that the aforementioned fifth guiding principle indicates specifically that the SAE standard was intended to be used across disciplines, including the legal domain.

An important feature of the SAE standard that might not be immediately apparent is the concept of an Operational Design Domain (ODD). An ODD is defined by the SAE standard as: "Operating conditions under which a given driving automation system or feature thereof is specifically designed to function, including, but not limited to, environmental, geographical, and time-of-day restrictions, and/or the requisite presence or absence of certain traffic or roadway characteristics."

The significance of this crucial concept is that it allows for a subdividing of a domain into those portions that might be amenable to autonomous capabilities or that may soon be amenable. Without such a proviso, it would tend to hamstring a set of levels in an LoA to require that either autonomy is entirely and completely the case at a given level or it is not at all at that level. This kind of take-it-or-leave-it conundrum was a stumbling block to the acceptability of some other LoAs and represented a subtle, but vital, form of progression in the formulation of the official SAE LoA.

The ODD concept will be instrumental in providing a similar benefit for the LoA of this proposed framework, as will be discussed in the next section.

One additional aspect to be covered briefly, particularly when discussing an LoA for the law, entails whether it might be feasible to reuse an existent accepted overarching ontology of the law. Thus, just as reusing an LoA offers merits, so too would reusing an overarching ontology of the law. For clarification, the meaning of ontology in this context is, as Neches et al. states, [7]: "An ontology defines the basic terms and relations comprising the vocabulary of a topic area as well as the rules for combining terms and relations to define extensions to the vocabulary." As legal scholars are aware, there is not a single unified ontology of the law, though many efforts have been undertaken to form such a taxonomy (for a review see Ghosh [3]). Though no widespread standardized legal ontology is yet established, the LoA AILR is devised to accommodate such an ontology if or when such an allied standard is formulated.

Proposed Levels-of-Autonomy for AI-Based Legal Reasoning Systems

When defining levels of autonomy, there are a multitude of factors that should be employed, doing so to systematically arrive at a parsimonious set that is logically

sound and inherently robust. Any notable facets that are omitted or skirted, whether inadvertently or by intent, could undermine the veracity of the definition; thereby weakening or entirely vacating the utility of the resulting taxonomy.

Utilized here is a bounded set of ten specific characteristics that are significant overall, and for which have been contributory in deriving the levels of autonomy for AI-based Legal Reasoning. Note that each such characteristic is valuable on its own merits and the listing of them in a numbered or sequenced fashion is not done to showcase priority or ranking. Instead, it is merely showcased for ease of reference.

Those ten key characteristics are:

- Scope
- Sufficiency of Reason
- Completeness
- Applicability
- Usefulness
- Understandability
- Foolproofness
- Observe Occam’s Razor
- Differentiable
- Logical Progression

For further details on how these apply to the proposed framework, see Eliot [1].

The proposed LoA for AILR is as follows:

> Level 0: No Automation for AI Legal Reasoning
> Level 1: Simple Assistance Automation for AI Legal Reasoning
> Level 2: Advanced Assistance Automation for AI Legal Reasoning
> Level 3: Semi-Autonomous Automation for AI Legal Reasoning
> Level 4: Domain Autonomous for AI Legal Reasoning
> Level 5: Fully Autonomous for AI Legal Reasoning

A brief summary of each LoA is given next and further elaborated in Eliot [1].

Level 0 is considered the no automation level. Legal reasoning is carried out via manual methods and principally occurs via paper-based methods.

Level 1 consists of simple assistance automation for AI legal reasoning

Examples of this category would include the use of everyday computer-based word processing, the use of everyday computer-based spreadsheets, access to online legal documents that are stored and retrieved electronically, and so on. By-and-large, today's use of computers for legal activities is predominantly within Level 1. Human-based legal reasoning is aided by this level of automation on a rather simplistic basis. It is assumed and expected that over time, the pervasiveness of automation will continue to deepen and widen, and eventually lead to legal activities being supported and within Level 2, rather than Level 1.

Level 2 consists of advanced assistance automation for AI legal reasoning

Examples of this category would include the use of query-style rudimentary Natural Language Processing (NLP), simplistic elements of Machine Learning (ML), and statistical analysis tools for case predictions, and etc. Human-based legal reasoning is aided by this level of automation on a more advanced basis, including that the automation can partake in rudimentary "legal reasoning" related tasks, though without any notable semblance of autonomy. Gradually, it is expected that these primitive AI-based systems for legal activities will increasingly make use of more advanced automation. LegalTech that was once at a Level 2 will likely be refined, upgraded, or expanded to include advanced capabilities, and thus be reclassified into Level 3.

Level 3 consists of semi-autonomous automation for AI legal reasoning

Examples of this category would include the use of advanced Knowledge-Based Systems (KBS) for legal reasoning, the use of Machine Learning and Deep Learning (ML/DL) for legal reasoning, advanced NLP, and etc. . There is a modicum of semi-autonomous capacity in the system undertaking a narrowly delineated form of legal reasoning,doing so under the guidance and in conjunction with a human legal reasoner. Today, such automation tends to exist in research efforts or prototypes and pilot systems, along with some limited instances of commercial legal technology that are inclusive of these capabilities. All told, there is increasing effort to add such capabilities into LegalTech. It is anticipated that many of today's Level 3 will inevitably be refined or expanded to then be classifiable into Level 4.

Level 4 consists of domain autonomous computer-based systems for AI legal reasoning

This level reuses the conceptual notion of Operational Design Domains (ODDs), as utilized in the autonomous vehicles, but applied to the legal domain. Essentially, this

entails any AI legal reasoning capacities that can operate autonomously, entirely so, but in some limited or constrained legal domain. Legal domains might be classified by functional areas, such as family law, real estate law, bankruptcy law, environmental law, tax law, etc.

Level 5 consists of fully autonomous computer-based systems for AI legal reasoning

In a sense, Level 5 is the superset of Level 4 in terms of encompassing all possible domains as ultimately defined for Level 4. It is conceivable that someday there might be this envisioned and fully autonomous AI legal reasoning capability; one that encompasses all of the law in all foreseeable ways (though this is quite a tall order and remains quite aspirational without a clear-cut path of how this might one day be achieved). Nonetheless, it seems to be within the extended realm of possibilities.

Making Use Of The Framework

Consider two brief examples of how these levels-of-autonomy for AI-based legal reasoning can be advantageously utilized.

Suppose that a vendor comes out with an augmented e-Discovery tool that claims to have NLPand utilizes ML, which seems impressive on a cursory basis.

But what level does this attain?

Upon undertaking an appropriate rating or assessment, assume that the e-Discovery amplification is classified as being at a Level 2. Therefore, this is considered advanced assistive automation rather than existing as a semi-autonomous or fully autonomous capacity.

Meanwhile, in ready comparison, suppose that a competing vendor has an e-Discovery tool that is rated as a Level 3. All else being equal, one can readily construe that the Level 3 product has a superior level of an autonomous facility than the Level 2 offers. Thus, the use of this LoA enables a kind of above-board playing field and readily facilitates head-to-head comparison.

This same benefit can be realized in the legal research sphere too.

Suppose a legal scholar criticizes that existing AI-based “legal reasoning” algorithms are weak at identifying suitable criminal sentencing recommendations. That might be a valid concern, though this could be based on having examined say Level 1 such

systems, and therefore provides only a narrow perspective. Other researchers might inadvertently misconstrue the result and assume that all AI-based legal reasoning systems are equally deficient, when in fact, it could be that there are (or will be) Level 2 and Level 3 systems that are more robust and have overcome the identified weaknesses.

Researchers would be able to utilize the scale as part of their legal research efforts, including applying the LoA to other authored studies to reveal hidden assumptions in prior or concurrent research.

Conclusion and Future Research

All in all, a measuring scale of this nature has applicability to both the day-to-day practice of the law and equally has applicability to legal scholarship.

Furthermore, this scale for AI-based legal reasoning serves as a wake-up call or catalyst for engaging in a timely and vital dialogue about how to best seek to rate or assess the emerging plethora of AI-enabled legal applications and claimed AI-advances for legal reasoning systems. Businesses that are gradually and inevitably going to be adopting these systems will need a convenient and apt means to compare and contrast competing products. Academics also are in need of a robust method for assessing how far along the advances in AI-powered legal reasoning capacities have progressed.

This is merely a proposed approach and intended as a strawman. Additional research and ongoing discussion will be needed to pursue the levels-of-autonomy. Whether an overall concurrence can be reached is unclear, but the debate itself will likely reap benefits throughout the field of computational law as efforts stridently continue to adopt AI capabilities and infuse intelligent-like advances into legal reasoning systems.

About the Author

Dr. Lance Eliot is a Stanford Fellow at Stanford University in the Stanford CodeX: Center for Legal Informatics and the Chief AI Scientist at Techbruim Inc. He previously was a professor at the University of Southern California (USC) where he headed a multi-disciplinary and pioneering AI research lab. Dr. Eliot is globally recognized for his expertise in AI.

References

1. Eliot, Lance (2020). AI and Legal Reasoning Essentials. LBE Press Publishing.

2. Galdon, Fernando, and Ashley Hall, Stephen Jia Wang (2020). “Designing Trust in Highly Automated Virtual Assistants: A Taxonomy of Levels of Autonomy,” Artificial Intelligence in Industry 4.0: A Collection of Innovative Research Case-Studies. https://www.researchgate.net/publication/342380935_Designing_trust_in_highly_automated_virtual_assistants_A_taxonomy_of_levels_of_autonomy
3. Ghosh, Mirna (2019). “Automation of Legal Reasoning and Decision Based on Ontologies,” Normandie University. https://tel.archives-ouvertes.fr/tel-02062174/document
4. Lewis, Barnaby (2017). “Boxing Clever: How Standardization Built A Global Economy.” ISO.org. September 11, 2017. https://www.iso.org/news/ref2215.html
5. Linna Jr., Daniel (2019). “The Future of Law and Computational Technologies: Two Sides of the Same Coin.” MIT Computational Law Report. December 6, 2019. https://law.mit.edu/pub/thefutureoflawandcomputationaltechnologies/release/2
6. Markou, Christopher, and Simon Deakin (2020). “Is Law Computable? From Rule of Law to Legal Singularity,” May 4, 2020, SSRN, University of Cambridge Faculty of Law Research Paper. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3589184
7. Neches, Robert, and Richard Fikes, Tim Finn, Thomas Gruber, Ramesh Patil, Ted Senator, William Swartout (1991). “Enabling Technology for Knowledge Sharing,” Fall 1991, Volume 12, Number 3, AI Magazine. https://ojs.aaai.org//index.php/aimagazine/article/view/902
8. SAE (2018). Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles, J3016-201806, SAE International.https://www.sae.org/standards/content/j3016_201401/
9. Sifakis, Joseph (2018). “Autonomous Systems: An Architectural Characterization.” arXiv: 1811:10277. https://arxiv.org/abs/1811.10277
10. Surden, Harry (2019). “Artificial Intelligence and Law: An Overview,” Summer 2019, Georgia State University Law Review. https://readingroom.law.gsu.edu/gsulr/vol35/iss4/8/

Header image generated with Wombo