Towards a Scalable Architecture for Legal Ontologies Integrated into Digital Twins of Administrative Law

Florian Schnitzhofer[1], Christoph G. Schuetz[1]

[1]Institute of Business Informatics – Data & Knowledge Engineering, Johannes Kepler University Linz, Altenberger Str. 69, 4040 Linz, Austria

Abstract

Administrative-law provisions are still published almost exclusively in natural language, forcing every stakeholder to translate identical rules into bespoke code bases—a practice that invites inconsistency, hampers transparency, and inflates maintenance costs. Recent work on *Digital Twins for Administrative Law* (DTAL)[1] suggests that legislation be issued together with machine-readable ontologies and executable logic, yet guidance on how to architect such systems remains scarce. In this work we propose a layered reference architecture that separates (i) the natural-language statute, (ii) a core ontology expressed in OWL, (iii) a configuration layer for mutable policy parameters, and (iv) an executable-rule layer exposed through a RESTful and MCP façade. Grounded in design science research, we implemented a proof-of-concept twin of the Upper-Austrian tourism-levy statute and qualitatively evaluate the twin with legal, software, and public-administration experts. Early results[2] suggest that ontology-driven twins can reduce duplicate implementations, streamline updates, and enhance legal certainty, thereby strengthening the Rule of Law in automated decision-making.

Keywords

Legal Ontologies, Semantic Interoperability, Automated Decision-Making, Digital Twins

1. Introduction

Administrative law forms the legal basis for government activities and public services such as taxation, welfare, and licensing. Yet, administrative law is disseminated almost entirely in natural language. Each public or private actor must, therefore, interpret and encode identical provisions into separate systems, leading to redundancy, opacity, and divergent outcomes [3, 4]. Frequent amendments amplify the problem: every minor change in tax rates or thresholds forces cascades of software patches across heterogeneous platforms.

From a rule-of-law perspective the stakes are high. Inconsistent software implementations can jeopardize equality before the law, frustrate judicial review, and erode public trust [5, 6]. Traditional development pipelines, where jurists paraphrase legislation, engineers translate prose into code, and jurists then audit outputs, must be repeated for every deployment [4]. Digital-twin engineering, by contrast, promises a canonical, executable representation of the statute that synchronizes with real-world data in (near) real time [7, 8].

This paper refines the following research question:

> *Which software-architecture choices best support hierarchical decision-support systems grounded in ontology-based digital representations of administrative law?*

Focusing on the Upper-Austrian tourism levy, we argue that a layered Digital Twin for Administrative Law (DTAL) can reduce contradictory interpretations by different stakeholders, enable rapid updates, and serve as a building block for a future system of systems covering entire legislative corpora.

Developers Workshop: Semantics - 21st International Conference on Semantic System, September 03–05, 2025, Vienna, AT
schnitzhofer@dke.uni-linz.ac.at (F. Schnitzhofer); schuetz@dke.uni-linz.ac.at (C. G. Schuetz)
https://www.dke.uni-linz.ac.at/ (F. Schnitzhofer); https://www.dke.uni-linz.ac.at/ (C. G. Schuetz)
0009-0005-3338-0366 (F. Schnitzhofer); 0000-0002-0955-8647) (C. G. Schuetz)

2. State of the Art

The concept of digital twin (DT) has progressed from manufacturing origins, where they couple virtual models with physical artefacts for real-time analytics, to indispensable enablers of smart factories, aerospace maintenance, and city infrastructures [7, 8, 9, 10]. A mature DT stack offers bi-directional data flow, life cycle traceability, and hierarchical scaling from unit to system-of-systems levels [11]. In parallel, ontology engineering (OE) provides formal, machine-readable conceptualizations; upper-level ontologies such as DOLCE and UFO standardize generic categories, while domain-specific legal ontologies—including FOLaw, LRI-Core, and LKIF-Core—capture statutes, roles, and normative relations [12, 13, 14, 15, 16]. Systematic reviews report a sharp rise in legal-ontology output over the last three decades, propelled by needs for semantic retrieval, compliance checking, and system interoperability [17].

Research on the intersection of artificial intelligence and law first framed legal reasoning as expert rules or logic programs [18], then as argumentation over precedent cases [19]. Contemporary strands aim to publish legislation in executable form: Rule-markup standards (LegalRuleML), domain-specific languages, and smart-contract templates embed deterministic logic alongside traditional prose [20, 21, 22]. Yet, smart-contract tooling, optimized for private bargains, struggles with public-law nuances, while data-driven language models pose rule-of-law concerns by producing probabilistic rather than authoritative outcomes [5, 3].

Schartum's model, where jurists verbalize rules, technologists translate them, and both iterate for validation, still forces every stakeholder to re-implement identical statutes, amplifying cost and inconsistency [4]. A real world example of such duplication can be exemplified on Austria's tourism-levy statute, with tens of thousands of enterprises encoding the same calculation logic independently.

3. Towards a Layered DTAL Architecture

Drawing on methods used in industry for digital twins [23] and on blockchain-based smart contracts [24, 25], our research posits a layered structure for a digital legal twin with the following four layers.

1. **Statutory Text**: the natural-language law as traditionally published.
2. **Ontology**: a formal structure encapsulating the law's semantics.
3. **Configuration Model**: adjustable parameters which would include the typical axioms of the ontology (e.g., tax rates, exemption thresholds).
4. **Executable Logic**: the computational rules for enforcement (e.g., formulae, conditionals) using parameters of the Configuration Model.

Such architecture is inspired by the separation-of-concerns principle central to smart-contract methodologies, which encourage modular design to allow incremental updates and transparent governance [22]. While certain amendments to legislation only change numeric thresholds or percentage rates, others require foundational revisions to interpretive definitions. By distinguishing between ontological structures and configuration parameters, this framework aims to reflect legislative updates systematically, facilitating improved legal compliance and reduced duplication of coding efforts.

3.1. Selecting a Structure of the Ontology

Existing legal ontologies such as FOLaw [14], LRI-Core [15], and LKIF-Core [16] are primarily designed to support the way jurists and their expert systems analyze, categorize, and interpret laws. Another standard that directly targets the expression of normative rules is LegalRuleML, an OASIS specification that represents the legal logical structure of statutes, regulations, and case law [20]. While these ontologies are well-suited for legal reasoning and knowledge representation, they are not optimized for ensuring rule of law in administrative law through automated decision-making.

A fundamental distinction of the DTAL approach for administrative law is that the ontology does not merely serve as a taxonomic classification of legal concepts but rather as a dynamic, multi-layered

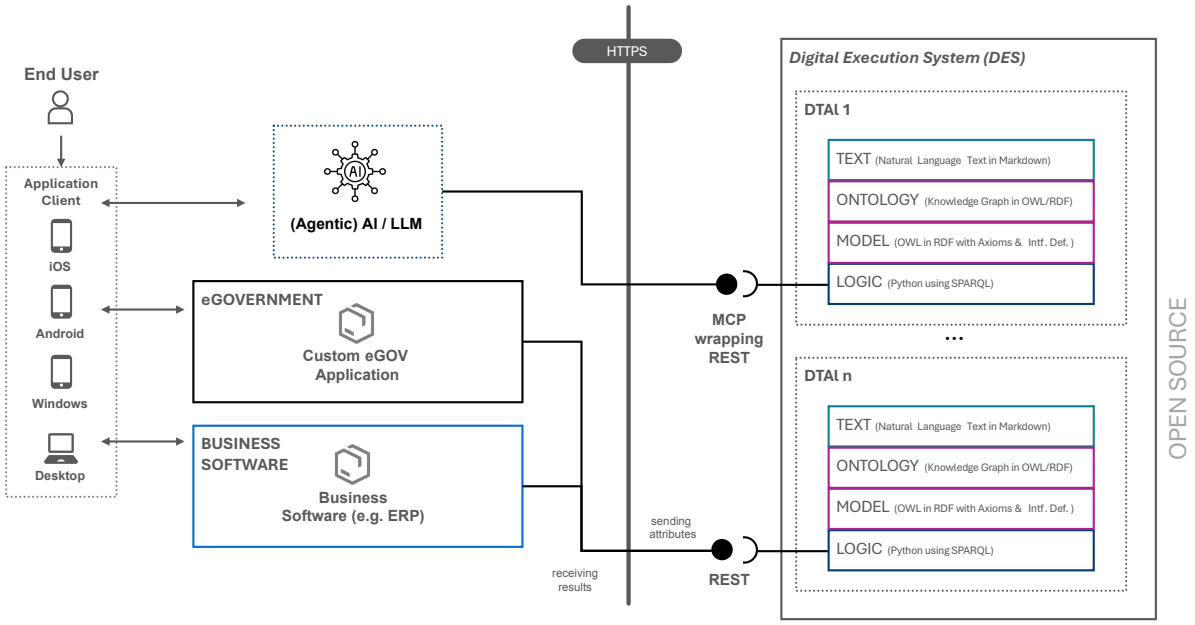

Figure 1: Software Architecture for Digital Twins of Administrative Law

semantic bridge connecting (1) legislative text, (2) real-world administrative data, and (3) the computational logic required for implementation. To guide further development, a literature review based on the methodology of *Legal Ontologies over Time: A Systematic Mapping Study* [17] for the years 2017 to 2025 is currently being conducted. This review aims to evaluate the structure, scope, and reusability of existing legal ontologies and to assess their potential integration into the DTAL framework.

3.2. Use Case: Upper-Austrian Tourism Levy

To illustrate the viability of this multi-layered framework, a tourism levy imposed in an Austrian province serves as a concrete case study. The *Oö. Tourismusgesetz 2018* levies an annual contribution on enterprises whose turnover derives from tourism. While assessment is centralised, roughly 29 000 businesses must embed the calculation in their enterprise systems or resort to manual spreadsheets.

The following mathematical formulation captures the legal and economic logic of the contribution scheme. Let U be taxable turnover (two-year lag), $p_{g,c}$ the rate for sector–municipality class pair, $M_{g,c}$ the minimum contribution, and $H = 4\,280\,000$ the maximum base. The payable amount T is then given by the following equation.

$$T = \max\Big(\min\big(p_{g,c} \cdot U,\ H\big),\ M_{g,c} \Big). \tag{1}$$

Tables 1 and 2 in the appendix list statutory parameters.

In our proof-of-concept implementation, the ontology is represented using OWL (Web Ontology Language), as per W3C standards [26]. We also integrated the European Legislation Identifier (ELI) [27] standard.

The logic layer is the core computational engine of the proof-of-concept law digital twin, implemented in Python (Stack Overflow Developer Survey 2024) [28] to leverage its flexibility and robust ecosystem. In an ongoing use-case scenario, the authors trialled a proof-of-concept DTAL implementation using Prolog, but it proved difficult to grasp for both computer scientists and legal practitioners (evaluation results will be published soon).

The logic layer of the Digital Twin offers one or more services through an interface, for example a RESTful API, making it accessible to external systems. Additionally, the interaces are published via the Model Context Protocol (MCP) to connect the DTAL with (agentic) AI applications. All API information and configuration details are published in the configuration model.

The source code of the use case is published open source on Github: https://github.com/FlorianSchnitzhofer/digital-twins-administrative-law-tourism-levy

4. Evaluation Plan

Three alternative prototypes in Python and Prolog as well as an XML-based prototpye will be implemented. A focus-group workshop (n = 5–10) with Austrian e-government officials, legal-tech vendors, and senior jurists will combine live walkthroughs, scenario testing, and a Delphi questionnaire. Data sources include video transcripts and observation checklists. Findings will iteratively refine architectural guidelines and produce a practitioner checklist for future DTAL projects.

5. Conclusion

By coupling natural-language statutes with authoritative ontologies and executable logic, DTAL offers a pragmatic path toward consistent, transparent, and cost-efficient administrative workflows. The Upper-Austrian tourism-levy prototype demonstrates technical feasibility and sets the stage for broader system-of-systems integration. Future work will conduct empirical evaluations, extend the ontology, and explore formal certification pathways to embed DTAL artifacts into legislative promulgation processes.

Declaration on Generative AI

Generative AI tools were used solely to enhance wording and translate text.

References

[1] F. Schnitzhofer, The self-driving state: Automated decision-making in modern governance, in: Proceedings of the International Symposium for Law and IT 2025, Vienna, Austria, 2025. 19–22 February 2025.

[2] F. Schnitzhofer, Integrating legal ontologies and digital twin technology: Enhancing rule-of-law-compliant automated decision-making in administrative law, in: Proceedings of the 22nd Extended Semantic Web Conference (ESWC 2025), PhD Symposium, Portorož, Slovenia, 2025. https://2025.eswc-conferences.org/phd-symposium-program/ (Accessed: 2 June 2025).

[3] D. Citron, F. Pasquale, The scored society: Due process for automated predictions, Washington Law Review 89 (2014) 1–33.

[4] D. W. Schartum, From Legal Sources to Programming Code: Automatic Individual Decisions in Public Administration and Computers under the Rule of Law, Cambridge Law Handbooks, Cambridge University Press, 2020, p. 301–336.

[5] J. Raz, The rule of law and its virtue, in: The authority of law: Essays on law and morality, Oxford University Press, 1979. URL: https://doi.org/10.1093/acprof:oso/9780198253457.003.0011. doi:10.1093/acprof:oso/9780198253457.003.0011. arXiv:https://academic.oup.com/book/0/chapter/158984424/chapter-pdf/40055041/acprof-

[6] H. Kelsen, General Theory of Law and State, Harvard University Press, Cambridge, MA, 1945. URL: https://api.pageplace.de/preview/DT0400.9781351517997_A30394699/preview-9781351517997_A30394699.pdf.

[7] M. Grieves, Digital twin: Manufacturing excellence through virtual factory replication (2015).

[8] F. Tao, Q. Qi, L. Wang, A. Nee, Digital twins and cyber–physical systems toward smart manufacturing and industry 4.0: Correlation and comparison, Engineering 5 (2019) 653–661. URL: https://www.sciencedirect.com/science/article/pii/S209580991830612X. doi:https://doi.org/10.1016/j.eng.2019.01.014.

[9] E. Glaessgen, D. Stargel, The digital twin paradigm for future nasa and u.s. air force vehicles, 2012. doi:10.2514/6.2012-1818.
[10] Y. Lu, X. Xu, L. Wang, Smart manufacturing process and system automation – a critical review of the standards and envisioned scenarios, Journal of Manufacturing Systems 56 (2020) 312–325. doi:10.1016/j.jmsy.2020.06.010.
[11] W. Kritzinger, M. Karner, G. Traar, J. Henjes, W. Sihn, Digital twin in manufacturing: A categorical literature review and classification, IFAC-PapersOnLine 51 (2018) 1016–1022. URL: https://www.sciencedirect.com/science/article/pii/S2405896318316021. doi:https://doi.org/10.1016/j.ifacol.2018.08.474, 16th IFAC Symposium on Information Control Problems in Manufacturing INCOM 2018.
[12] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, A. Oltramari, WonderWeb Deliverable D18 Ontology Library (final), Technical Report, IST Project 2001-33052 WonderWeb: Ontology Infrastructure for the Semantic Web, 2003. URL: http://www.loa.istc.cnr.it/old/DOLCE.html.
[13] G. Guizzardi, Ontological Foundations for Structural Conceptual Models, Ph.D. thesis, 2005.
[14] A. Valente, Legal knowledge engineering: A Modelling Approach, Frontiers in artificial intelligence and applications, Sage Publications, 1995. URL: https://cir.nii.ac.jp/crid/1130000797194456576.
[15] J. Breuker, A. Valente, R. Winkels, Legal ontologies: A functional view (1998).
[16] R. Hoekstra, J. Breuker, M. Di Bello, A. Boer, The lkif core ontology of basic legal concepts, 2007, pp. 43–63.
[17] C. M. de Oliveira Rodrigues, F. L. G. de Freitas, E. F. S. Barreiros, R. R. de Azevedo, A. de Almeida Filho, Legal ontologies over time: A systematic mapping study, Expert Systems with Applications 130 (2019) 12–30. URL: https://www.sciencedirect.com/science/article/pii/S0957417419302398. doi:https://doi.org/10.1016/j.eswa.2019.04.009.
[18] R. Kowalski, M. Sergot, A logic-based calculus of events, New Generation Computing 4 (1985) 67–95. doi:10.1007/BF03037383.
[19] T. Bench-Capon, G. Sartor, A model of legal reasoning with cases incorporating theories and values, Artificial Intelligence 150 (2003) 97–143. URL: https://www.sciencedirect.com/science/article/pii/S0004370203001085. doi:https://doi.org/10.1016/S0004-3702(03)00108-5, aI and Law.
[20] M. Palmirani, G. Governatori, A. Rotolo, S. Tabet, H. Boley, A. Paschke, Legalruleml: Xml-based rules and norms, in: F. Olken, M. Palmirani, D. Sottara (Eds.), Rule-Based Modeling and Computing on the Semantic Web, Springer Berlin Heidelberg, Berlin, Heidelberg, 2011, pp. 298–312.
[21] D. Merigoux, N. Chataing, J. Protzenko, Catala: a programming language for the law, Proc. ACM Program. Lang. 5 (2021). URL: https://doi.org/10.1145/3473582. doi:10.1145/3473582.
[22] J. Hazard, H. Haapio, Wise contracts: smart contracts that work for people and machines, in: Trends and communities of legal informatics. Proceedings of the 20th international legal informatics symposium IRIS, 2017, pp. 425–432.
[23] S. Boschert, R. Rosen, Digital Twin—The Simulation Aspect, 2016, pp. 59–74. doi:10.1007/978-3-319-32156-1_5.
[24] T. A. Project, Accord project concepts, https://docs.accordproject.org/docs/accordproject-concepts.html, 2024. Accessed on: 30. Oktober 2024.
[25] J. G. Allen, 23Wrapped and Stacked: 'Smart Contracts' and the Interaction of Natural and Formal Language, in: Smart Legal Contracts: Computable Law in Theory and Practice, Oxford University Press, 2022. URL: https://doi.org/10.1093/oso/9780192858467.003.0002. doi:10.1093/oso/9780192858467.003.0002. arXiv:https://academic.oup.com/book/0/chapter/362294132/chapter-pdf/44144454/oso-978
[26] W3C, OWL 2 Web Ontology Language: Structural Specification and Functional-Style Syntax, Recommendation TR/owl2-syntax, World Wide Web Consortium (W3C), 2009. URL: https://www.w3.org/TR/owl2-syntax/, accessed on: 30. Oktober 2024.
[27] T. Francart, J. Dann, R. Pappalardo, C. Malagon, M. Pellegrino, The european legislation identifier, in: Knowledge of the Law in the Big Data Age, IOS Press, 2019, pp. 137–148.
[28] S. O. D. Survey, Most popular technologies, https://survey.stackoverflow.co/2024/technology, 2024.

Appendix

Table 1
Contribution rates (%).

Group	Class A	Class B	Class C	Class St
1	0.50	0.45	0.40	0.40
2	0.35	0.30	0.20	0.20
3	0.20	0.15	0.10	0.10
4	0.15	0.10	0.05	0.05
5	0.10	0.05	0.025	0.025
6	0.05	0.00	0.00	0.00
7	0.00	0.00	0.00	0.00

Table 2
Minimum contributions (EUR).

Group	Class A	Class B	Class C	Class St
1	69.00	51.00	34.50	34.50
2	51.00	34.50	34.50	34.50
3–5	34.50	34.50	34.50	34.50
6	34.50	0.00	0.00	0.00
7	0.00	0.00	0.00	0.00