2026 European Conference on Computing in Construction
Corfu, Greece
July 12-15, 2026

A DECOUPLED ARCHITECTURE FOR DIGITALLY ACTIONABLE AND TRACEABLE REGULATORY WORKFLOWS IN THE AEC SECTOR

Hans Kristjan Aljas, Ergo Pikas, and Martin Thalfeldt
Tallinn University of Technology, Tallinn, Estonia

Abstract

The Architecture, Engineering, and Construction (AEC) sector faces complex regulatory frameworks like building energy performance certification under the European Energy Performance of Buildings Directive (EPBD). Traditional approaches suffer from fragmented documentation, procedural ambiguities, and poor traceability and auditability, causing outcome inconsistencies and automation challenges. This paper proposes a decoupled conceptual architecture for digitally actionable and traceable regulatory workflows, leveraging compliance modelling, semantic web technologies, and linked data. By separating semantic representation, reasoning, execution, and interaction layers, it improves compliance, interoperability, and adaptability. Demonstrated on Estonian Energy Performance Certificates (EPCs), it addresses key gaps and supports broader application beyond energy efficiency.

Introduction

Regulatory workflows in the AEC sector are characterised by high complexity. They simultaneously involve physical building elements and procedures, formal norms (directives, standards, laws), institutional and professional practices (national methodologies, auditing bodies, firms), and heuristic expert judgment (soft law, socio-cognitive interpretation).

A representative example is the Energy Performance of Buildings Directive (EPBD) and its national transpositions (European Parliament, 2024). Despite a common European framework, these produce poorly comparable Energy Performance Certificates (EPCs) (Jenkins et al., 2024; Sayfikar and Jenkins, 2025) riddled with methodological non-compliance and inconsistencies (Aljas et al., 2025; Carvalho et al., 2024; Schaffer et al., 2026), and create maintenance burdens for public authorities responsible for custom software pipelines (buildingSMART International, 2023). These phenomena reveal a fundamental tension between:

(1) the hybrid socio-technical nature of regulatory norms and
(2) the current paradigms used to digitise and automate them.

Importantly, this tension cannot be resolved through comprehensive standardisation or exhaustive semantic formalisation alone. Regulatory methodologies are shaped by institutional, geographical, and methodological heterogeneity, alongside practical and organisational constraints that limit the feasibility of total harmonisation (Jenkins et al., 2024; Sayfikar and Jenkins, 2025). As a result, regulatory workflows must remain operable under conditions of partial standardisation, evolving methodologies, and uneven digital maturity.

Existing approaches treat regulations either as static rules to be checked post-hoc aka “Compliance by Detection” or as logic to be programmed into custom software aka “Compliance by Design” (Casanovas, 2017). Both paradigms inevitably produce tightly coupled, brittle systems that struggle to evolve harmoniously with the norm, the context, or the technology. Furthermore, (Feng and Wang, 2025) note that overlapping stakeholder pressures in complex construction workflows expose the limitations of post-hoc compliance mechanisms.

This work advances Compliance through Design (CtD) — the governance-oriented, interpretive paradigm introduced by (Casanovas, 2017)—by providing a systematic operational framework for digitally actionable and traceable regulatory workflows.

The core artefact is a decoupled conceptual architecture with four layers (Semantic Workflow Model, Reasoning and State Engine, Execution and Orchestration Layer, Interaction Layer). The artefact functions as a knowledge-modelling framework that treats regulatory workflows as evolvable, linked-democratic ecosystems rather than monolithic code bases or conventional data pipelines. In practice, such an approach aims to reduce data and logic fragmentation, allowing stakeholders to focus on their core responsibilities or interests by minimising routine administrative tasks and extraneous cognitive load, and thereby reducing the potential for errors.

Taxonomy of compliance

In regulatory contexts, compliance encompasses adherence to laws and standards while extending to ethical and societal norms (Governatori, 2018; Hashmi et al., 2018). Compliance paradigms form a progression of increasing proactivity, addressing the limitations of reactive methods in complex environments.

Compliance by Detection operates reactively, identifying non-conformity during or after execution, which often necessitates costly redesign. Conversely, Compliance by Design shifts to prevention, incorporating rules into the design stage to ensure conformity in advance (Casanovas, 2017; Governatori, 2018). Legal Compliance by Design further refines this by addressing value-based requirements and the legality of the entire process rather than simple formal rules (Hashmi et al., 2018).

Compliance through Design (CtD) constitutes the most holistic evolution; it moves beyond rule-embedding to encompass the systemic interoperability required to model affordances—the action possibilities enabled by a system—alongside broader socio-legal conditions (Poblet et al., 2019). The paradigm acknowledges that compliance is not merely a technical constraint but a social process involving legal interpretation, institutional coordination, and the relationships between regulated entities, citizens, and the law (Casanovas, 2017; Hashmi et al., 2018). From this viewpoint, regulatory requirements operate within an ecosystem of practices and responsibilities, rather than being fully reducible to technical checks.

The shift toward greater proactivity is technically enabled by declarative paradigms. Unlike imperative approaches that specify rigid, step-by-step procedures, declarative modelling focuses on what must occur by defining relationships, preconditions, and effects (Governatori, 2018). This allows regulatory requirements to be expressed without prematurely committing to a fixed resolution strategy or technical details.

Linked Data ecosystems complement this approach by supporting norm sharing and knowledge aggregation across institutional boundaries, enabling compliance mechanisms to remain interpretable, adaptable, and aligned with broader governance structures (Poblet et al., 2019).

Structural mismatches in regulatory automation

Despite increasing digitalisation, regulatory workflows in the AEC sector remain difficult to automate in a traceable and sustainable manner. This is not primarily a data availability problem but a structural mismatch between regulatory intent, information representation, and technical execution.

Regulatory intent is expressed through natural-language instruments designed to establish legal authority and support professional interpretation. In technical domains, this intent is commonly conveyed through condensed, presentation-oriented formats—such as tables, annexes, and figures—that facilitate communication while remaining aligned with legal drafting practices. Empirical evidence highlights the limits of treating these representations as directly formalisable: (Hettiarachchi et al., 2025) show that, from an initial set of 20 674 sentences in English and Finnish building regulations, only 6% could be annotated as self-contained rules suitable for machine-readable formalisation. This observation is consistent with legal informatics research, which distinguishes between sources of authority and documentary resources, noting that regulatory meaning is not fully embedded in such artefacts and must be reconstructed through interpretive knowledge-acquisition processes (Casanovas, 2017).

Empirical analysis of design-stage EPCs in Estonia shows how this complex interpretation process propagates into methodological inconsistencies and mismatches between declared design solutions, assumed input parameters, and practical outcomes (Aljas et al., 2025).

EPBD provides an example of how fixed technical thresholds are combined with heuristic judgement, spanning multiple knowledge domains:

"Where technically and economically feasible, non-residential buildings with an effective rated output for heating systems, air-conditioning systems, systems for combined space heating and ventilation, or systems for combined air conditioning and ventilation of over 290 kW, must be equipped with building automation and control systems." - (European Parliament, 2024).

This hybrid formulation leaves essential aspects of compliance underspecified, including actor responsibilities, data sources, and the sequencing of decisions and practical actions.

Current digital tools primarily focus on operationalising these static methodologies by automating calculations, enforcing parameter constraints, and guiding data entry and reporting. While this improves usability and administrative efficiency, it also embeds a single, implicit interpretation of the regulation within individual software systems, resulting in tightly coupled implementations with limited transparency, adaptability, and traceability (buildingSMART International, 2023; Sayfikar and Jenkins, 2025).

Semantic approaches, including Linked Building Data, Building Information Graphs, regulatory-focused initiatives such as Regulatory Information Requirements, and building permitting-oriented ontologies such as OBPA or AEC3PO, contribute valuable partial formalisations of regulatory-relevant data, structures, and artefacts (buildingSMART International, 2023; Dridi et al., 2024; Hagedorn et al., 2025; Lygerakis et al., 2022; Wang and Sacks, 2025). However, these approaches do not explicitly model the procedural aspect that translates regulatory intent into executable workflows. As a result, compliance logic is either applied post-hoc through validators (Compliance by Detection) or embedded imperatively into software (rigid Legal Compliance by Design), which are commonly implemented across fragmented or siloed toolchains, with associated maintenance and coordination costs reported in the literature (buildingSMART International, 2023; Governatori, 2018).

At the workflow level, these fragmented implementations produce overlapping responsibilities and disconnected steps, imposing manual coordination and repeated data entry on users, while obliging software developers to replicate procedural logic and manage fragile integrations across tools. Such challenges persist even where extensive standardisation or semantic modelling is applied, reflecting well-documented institutional, societal, and

technological heterogeneity in regulatory practice (Carvalho et al., 2024; Jenkins et al., 2024).

Taken together, these observations reveal a procedural blind spot: the absence of an explicit, declarative representation of how regulatory intent is translated into executable workflows. Addressing this gap requires an architectural separation that preserves interpretive logic, supports traceability across abstraction levels, and enables adaptable automation.

The decoupled, goal-driven architecture

To address the procedural blind spot identified in regulatory automation, this work proposes a decoupled conceptual architecture that restructures regulatory workflows around explicit procedural representation. The architecture separates regulatory intent, logical reasoning, technical execution, and agentic interaction into distinct but linked layers, as shown in Figure 1.

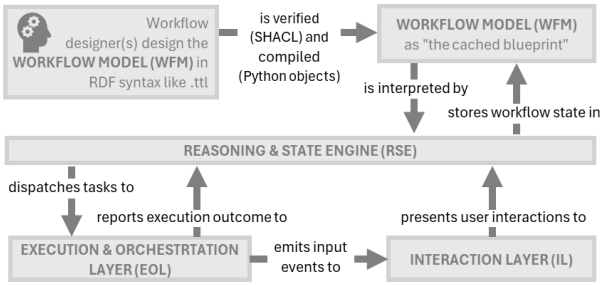

Figure 1: The 4-layer decoupled architecture

The workflow model: a design-time semantic blueprint

At the core of the architecture is a semantic Workflow Model (WFM) that serves as a design-time blueprint defining how regulatory completion can be achieved. The model is goal-driven: it defines the space of possible pathways through which explicit completion goals can be satisfied. Rather than prescribing a fixed execution sequence, the model specifies the set of steps, logic, and information dependencies that lead to these goals.

These elements are modelled as Resource Descriptive Framework (RDF) objects in a knowledge graph, making execution-relevant structure and meaning explicit through the use of two fundamental abstractions – Activities and Entities (*core layer* in Figure 2). Entities represent planned data items, which Activities can declare as inputs or outputs, forming a directed dependency structure that guides the execution, as exemplified in Figure 3. We modeled this core using a domain-independent semantic layer based primarily on the P-PLAN ontology.

The core model is complemented by our custom lightweight control vocabulary (*control layer* in Figure 2), which allows completion criteria and other regulatory logic to be explicitly represented within the Workflow Model as first-class, i.e., regulatory constraints and decision points are subclasses of as Activities and Entities.

Where beneficial, the core model can be extended with domain-specific knowledge graphs, such as Linked Building Data (LBD) vocabularies or local regulatory reference data. The semantic depth and operational specificity (as noted in Figure 2) of the WFM is thus configurable: different deployments can provide only the level of detail that execution handlers and users can interpret. Increased semantic depth enables tighter alignment between domain data and regulatory logic, allowing constraints and dependencies to be expressed explicitly where needed, but without requiring uniform detail across all deployments.

Scope of activities and hybrid control

The notion of an Activity is intentionally broad, reflecting the hybrid socio-technical nature of regulatory workflows. Operationally, Activities may be fully digital—such as automated data acquisition from APIs, parameterised calculations, consistency checks (e.g. BIM/LBD-aligned post-hoc conformity checks), logic resolution—or physical–digital, where a physical activity or professional judgment is acknowledged and documented through the system. Examples include confirmation of an on-site inspection, upload of supporting documents, or an explicit declaration that a regulatory requirement has been satisfied or acknowledged. Activities also distinguish between steps that can be executed automatically and those that require agentic input, while modeled uniformly within the same dependency structure.

Decoupled architectural layers

The architecture operationalises the WFM through four coordinated layers, following annotations in Figure 1.

The workflow model (WFM), as discussed, provides the design-time blueprint by encoding Activities, Entities, including goals and control points as a declarative semantic knowledge graph. At runtime startup, these semantic definitions are loaded, enriched with lightweight reasoning, and validated against design-time constraints (via Shapes Constraint Language aka SHACL) to ensure model conformity; they are then compiled into cached blueprint structures in two steps (the following terminology is aligned with Figure 2): first, the foundational workflow meta-model is superimposed on the project-layer blueprint instances to generate an in-memory workflow graph—i.e., cached Activity/Entity records with a compiled dependency map. Second, the *design & specialisation vocabulary* is retained as lightweight annotations (semantic types and dispatch configuration) that enable generic routing to execution and interaction handlers. However, *domain vocabulary* and *executable bindings* from the *project layer* are not required for the generic reasoning and orchestration mechanisms, but can be leveraged to implement custom functionalities (e.g., domain-specific functions, queries, or service adapters). In this way, the WFM functions as the authoritative specification of the workflow's structure and meaning, while remaining decoupled from the concrete mechanisms that execute steps, solicit inputs, and perform external tasks.

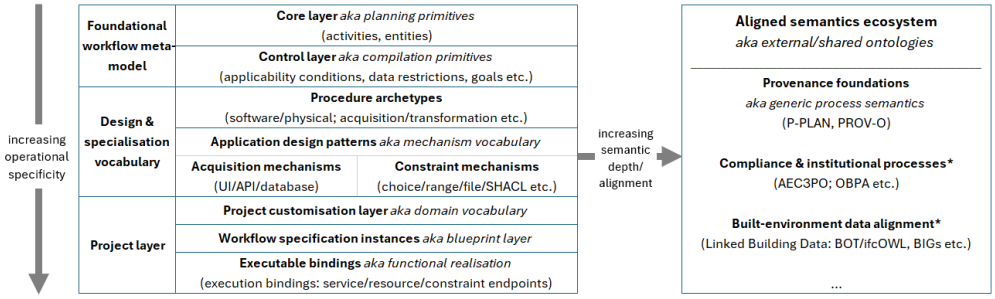

Figure 2: The Workflow model's semantic decomposition. (*) marks frameworks not implemented in the current instantiation.

The reasoning & state engine (RSE) represents the stateful operational structure of a workflow run. It operates over a directed dependency model whose first-class elements are Activities and Entities of the WFM. During execution, the RSE may further extend this structure by materialising latent semantic conditions—such as unresolved alternative production paths or exceptional situations—into explicit control-point Activities (e.g. *:MultipleProducerChoice* on Figure 3) whose outcomes gate subsequent progression. Therefore, all state transitions are exposed as structured events and can be synchronised to provenance records, enabling traceable reconstruction of how outcomes were produced and which decisions governed the execution path.

The execution & orchestration layer (EOL) carries out the Activities that the RSE has determined to be ready for execution. It receives an Activity together with its corresponding semantic type(s) and configuration, routes it to an appropriate execution mechanism via pluggable handlers, and either executes the Activity and commits its outcomes back into the runtime state, or emits an explicit interaction request and transitions the Activity into a waiting state until the Interaction Layer fulfills it. By standardising execution behind handler contracts, the EOL provides a stable execution interface that accommodates both internal functions and external services, while preserving decoupling between workflow semantics and execution mechanisms.

The interaction layer (IL) functions as the unified presentation boundary for the system. It consumes execution-state and interaction events forwarded by the EOL, and renders them through interaction primitives—such as data entry requests, explicit pathway choices, and progress events—so the interface adapts to different workflow configurations. For human experts, the IL can procedurally generate appropriate UI elements like fields, dropdowns, choice dialogs and file uploads. Alternatively, for machine agents, these primitives can be exposed through structured interfaces for programmatic consumption (e.g., an API or an MCP-style server), enabling artificial agents to consume required input/choice requests and submit decisions against the same underlying workflow logic.

Instantiation of the architecture: Estonian EPC workflow

This section presents a proof-of-concept instantiation of the proposed architecture using the Estonian Energy Performance Certificate (EPC) methodology for new buildings, grounded in the EPBD framework and ISO 52000-series standards (Minister of Economic Affairs and Infrastructure of Estonia, 2023; Dijk and Hogeling, 2019). The methodology requires prediction of primary energy use based on typical building energy modelling components such as envelope characteristics, HVAC systems, ventilation rates, infiltration rate, component efficiencies, lighting, renewable energy contributions etc. It allows using defaults from national tables or documented values from design/measurements, with professional validation or additional proof required where substitutions occur.

The implementation focuses on modelling local regulatory *domain vocabulary*, procedural dependencies, and decision logic, rather than on full numerical coverage or exhaustive system integrations. In particular, the current scope emphasises conditioning of calculations, alternative compliance pathways, and explicit control points that govern how EPC calculations may be completed at building design stage.

While the implementation covers all four layers of the architecture, the WFM's *operational specificity* (vertical arrow on Figure 2) is currently limited to non-functional placeholders for integrable data sources (e.g. *:GetBimModelFromBuildingRegistry* on Figure 3) such as the Estonian National Building Registry (later potentially via the E-ehituse OpenAPI services), or for Semantic BIM pipelines (Land and Spatial Development Board of Estonia, 2026; Pauwels and Terkaj, 2016; Wang and Sacks, 2025), with parallel options to explicitly re-declare such data entities (e.g. *:UploadBimModel*). The result is a dynamically conditioned, form-like execution environment composed of dropdowns, numerical inputs, checkboxes, acknowledgements, and calculation steps.

At runtime, available fields, selectable options, and required acknowledgements are progressively introduced or constrained as upstream Entities become resolved. Early building-level declarations—such as building category, size class, and primary construction type—function as conditioning inputs that determine which regulatory requirements apply, which calculation methods are admissible, and which exceptions or simplifications may be used. For example, small buildings may qualify for simplified calculation procedures, certain building categories apply adjusted EPC thresholds (e.g. multiplicative factors for specific construction types), and building type selections govern lookup strategies for default parameters such as ventilation rates.

This conditioning is reflected directly in the user interface. Selections made at higher levels constrain which subsequent inputs are exposed and which options are selectable, ensuring that only methodologically admissible choices are presented. Data restrictions—such as expected data type, allowed value lists, or numeric ranges—are materialised in the front-end as input field configurations (e.g. dropdown enumerations, bounded numerical fields, or categorical selectors).

At points where the Estonian methodology permits alternative means of satisfying a requirement, the workflow introduces explicit procedural choices (*:MultipleProducerChoice*) that resolve how an Entity is grounded. A common pattern is the substitution between normative defaults and documented values. Selecting a default immediately resolves the Entity using standard tables (*:LookupInfiltrationCoeff*), while inputting a custom value (*:InputInfiltrationCoeff)* may introduce additional dependent steps, such as acknowledgements or the provision of supporting material (*:InfiltrationTest-Report*). The method by which the value was obtained is preserved in Entity state and provenance records, allowing downstream logic to access previous compliance-relevant decisions.

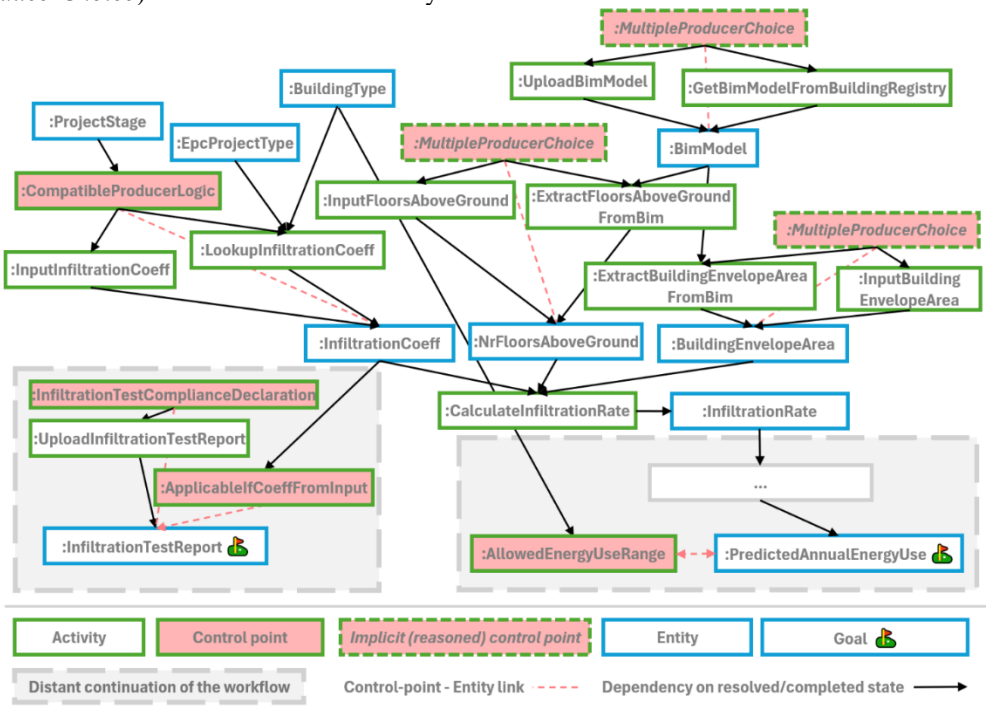

Figure 3: Simplified excerpt of the Estonian EPC Workflow Model showing main structural elements and their relations.

Not all regulatory requirements are digitally verifiable within the current scope. Where compliance depends on professional judgement, certification, or conformance to external standards that cannot be computationally assessed, explicit acknowledgement events are modelled in the Workflow Model and instantiated by the RSE as first-class Activities (*:InfiltrationTestComplianceDecla-ration)*. Although their outcomes are not derived numerically, they are committed to the runtime state and exposed through the Interaction Layer as explicit confirmations, preventing implicit assumptions and preserving traceability of regulatory decisions.

The implementation further addresses regulatory logic that applies across multiple instances of similar elements, such as technical systems or energy carriers. Their compatibility and co-dependence rules are expressed as aggregate conditions over sets of Entity instances in the Workflow Model. This constrains available options consistently across the workflow, preventing incompatible configurations at the point of selection. If upstream conditions change, previously valid choices are invalidated and surfaced through the Interaction Layer with clear visual feedback.

As the graphical interface is not the only possible data entry surface, the same restrictions are also enforced through implicit validation Activities attached directly to data Entities. Values are checked against their declared constraints regardless of how they are introduced, allowing inconsistencies to be detected even when upstream changes invalidate previously accepted inputs. This dual use of semantically declared data rules establishes a single source of truth, while supporting both immediate user-facing guidance and backend consistency validation at different stages of execution.

Where calculations are delegated to external procedures, such as *:CalculateInfiltrationRate*, these are represented as pluggable Activities bound to execution handlers in the EOL. Each such Activity consumes shared input Entities—such as building type (*:BuildingType*) or HVAC parameters—and produces a single declared output Entity, preserving the one-activity–one-output constraint of the Workflow Model. This keeps the dependency graph explicit for the RSE, while allowing external implementations to be introduced or replaced without affecting surrounding procedural logic or interaction behaviour. Figure 4 shows illustrative sections

from the Estonian EPC workflow expert system's UI in its present form (in development).

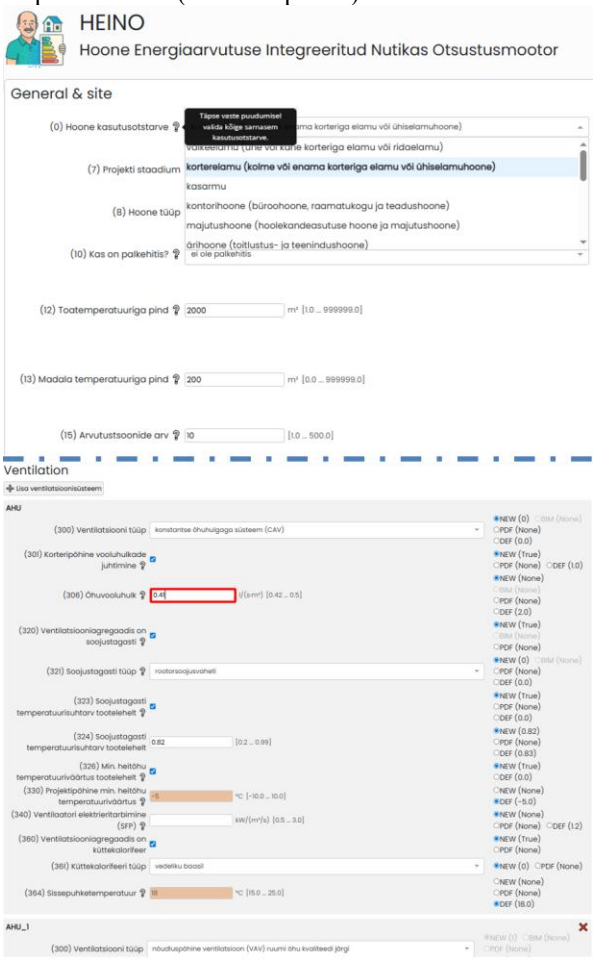

Figure 4: Screenshots from the current Estonian EPC workflow user interface (in development).

Discussion

This work frames regulatory automation as a problem of procedural articulation rather than document processing or isolated rule extraction. From this view, the central challenge is not digitising regulatory content, but coordinating intent, interpretation, and execution as workflows evolve, branch, and interact with external actors and systems. The proposed framework tackles this coordination structurally, though it cannot resolve every tension in regulatory and AEC digitalisation.

A recurring challenge is deciding how much formal structure is practical and desirable. Many regulatory steps involve physical actions, professional judgement, or context that resist full specification. The framework deliberately leaves such elements underspecified, treating them as obligations to acknowledge rather than compute, which prompts the question: how much structure supports traceability and coordination without over-constraining practice? Thus, workflow design remains an agent-guided task, with models reflecting institutional choices on where precision matters and where judgement should stay flexible.

Questions of maintenance and evolution follow naturally. Regulations change, interpretations shift, and data infrastructures mature unevenly. The Workflow Model does not erase these realities; it offers a stable reference point for managing them. Changes still demand agentic oversight and governance, but their effects become more localised and visible—underscoring that long-term viability depends as much on organisational processes and knowledge transfer as on technical design.

The framework delivers greatest value in workflows with strong interdependencies, sophisticated logic, and interaction with semantic or non-tabular data—the complex, evolving processes typical of energy performance certification or building permitting. In contrast, for workflows dominated by independent steps with minimal downstream impact—such as uploading parallel documents, running a fixed linear data pipeline, or handling predominantly physical sequences with little digital conditioning—the semantic modelling effort may yield limited incremental gains. Yet truly dependency-free cases are uncommon in regulated domains, where even seemingly separate or physical actions often carry traceability requirements, conditional applicability, or audit consequences that justify structured coordination.

The framework also highlights the limits of specialised, handcrafted solutions. Richard Sutton's "bitter lesson" from 70 years of AI research offers a useful lens: systems built around general, scalable structures or „meta-methods" ultimately outperform finely tuned domain heuristics, as the „contents of minds are tremendously, irredeemably complex" (Sutton, 2019). In regulatory automation, this suggests that durable progress more likely comes from simple, composable procedural primitives rather than from bespoke software. Early results with our framework—effective despite modest semantic depth—reinforce this view and motivate continued emphasis on foundational abstractions over narrowly optimised implementations.

The design remains intentionally open-ended, prescribing no single execution environment, interaction style, or data standard. This flexibility accommodates diverse implementations but raises challenges in interoperability, semantic alignment, and trust. Our semantic modelling (built on P-PLAN, PROV-O, a custom control vocabulary, and preliminary design & specialisation elements, as outlined in Figure 2) is still evolving; more sophisticated ontology alignment and refinement will appear in forthcoming work. Establishing and functionalising shared conventions through ontologies and data dictionaries remains an active area of research (buildingSMART International, 2023; Hjelseth et al., 2025; Utkucu and Sacks, 2025; Wang and Sacks, 2025). While the architecture supports incremental semantic enrichment, determining when and how to invest in deeper alignment is as much a policy decision as a technical one.

We see many directions demanding further exploration: real-world testing of long-running workflows to refine error handling and richer user-facing elements; semantic integration of permitting/compliance ontologies, including overlap resolution and cross-system

governance; and direct RDF graph interaction (designing/querying workflows) via natural-language and augmented retrieval tools (e.g., GraphRAG) to reduce technical overhead for non-experts. Future work is also expected to provide insights into the dilemma of deeper semantic standardisation vs reliance on self-contained tools or services, and explore semantic modelling strategies beyond the core layer that can most optimally cater for practical needs across the 4-layer architecture.

Ultimately, the framework is best viewed as scaffolding for ongoing exploration rather than a packaged product. By making procedural structure explicit, shareable, and inspectable while preserving institutional flexibility and technological diversity, it addresses—rather than ignores—the fragmented, dynamic, digital-physical realities of regulatory contexts. In such environments, this openness is likely essential for automation that proves meaningful and sustainable.

Conclusion

This paper addressed a persistent gap in the digitalisation of regulatory workflows: the absence of a stable structure for coordinating intent, interpretation, and execution across diverse actors, tools, and lifecycle stages. In fields like AEC, where compliance involves conditional procedures, partial automation, and professional judgement, existing approaches often fragment regulatory logic across documents, custom software, and informal practices—undermining traceability, adaptability, and sustainability.

To tackle this, the work introduced a declarative workflow model that reframes compliance as a procedural knowledge problem rather than a document- or rule-processing task. By encoding activities, entities, dependencies, and control points as semantic structures, the model defines what must be achieved and under what conditions—without dictating how steps are executed. This explicit, inspectable representation preserves regulatory meaning while keeping execution, interaction, and domain implementations flexible.

The four-layer architecture operationalises this model by decoupling semantic specification, state reasoning, execution orchestration, and interaction. The separation is practical: regulatory interpretation can evolve independently of software, execution mechanisms can change without rewriting semantics, and agentic judgement can be acknowledged as a first-class element rather than hidden or flattened into technical artefacts. The framework does not demand full formalisation—ambiguity, physical actions, and discretion are treated as legitimate procedural features, not automation failures.

Taken together, the contribution is not a finished system but a conceptual backbone for curating complex compliance processes. It enables incremental semantic enrichment, cross-domain coordination, and clearer accountability while remaining open to institutional practices and technological evolution. By prioritising simple, composable procedural abstractions over bespoke solutions, the work aligns with a broader shift toward scalable, general-purpose digital structures—offering a foundation for future research at the intersection of regulation, semantic technologies, and human–AI collaboration, where meaningful progress depends on making interpretation explicit, governable, and traceable rather than treating it as isolated engineering problems.

Acknowledgements

The authors acknowledge the use of large language models—including Grok 4.1 (Fast, Expert, Thinking modes), ChatGPT 5.1 & 5.2, and Google Gemini 3 (Pro & Flash)—used independently, and also within the Cursor software environment. These tools assisted in synthesising and communicating authors' original ideas and notes, referenced literature, codebase, and drafts.

All ideas, arguments, and conclusions remain the original work of the named authors. The text underwent rigorous manual review and iterative rewriting to ensure precise alignment with the authors' intent, cited sources, and scholarly standards, leaving no sloppy logic, over-interpretations or hallucinations. The models served only as aids for expression and refinement; full responsibility for content rests with the authors.

This work has been supported by the Estonian Centre of Excellence in Energy Efficiency, ENER (grant TK230) funded by the Estonian Ministry of Education and Research, European Commission through the European Union's Horizon Europe research and innovation program under grant agreement no. 101192930 (the ENTRANCE Project), and by the Estonian Research Council through the grant PSG963. This study was also co-funded by the European Union and Estonian Research Council via project TEM-TA78.

Abbreviations

AEC – Architecture, Engineering and Construction industry
AEC3PO – Architecture Engineering and Construction Compliance Checking and Permitting Ontology (Dridi et al., 2024)
BIGs – Building Information Graphs (Wang and Sacks, 2025)
BIM – Building Information Modelling
CtD – Compliance through Design (Casanovas, 2017)
EOL – Execution and Orchestration Layer (in proposed architecture)
EPBD – Energy Performance of Buildings Directive in European Union
EPC – Energy Performance Certificate
HVAC – Heating, Ventilation and Air Conditioning
IL – Interaction Layer (in proposed architecture)
LBD – Linked Building Data
MCP – Model Context Protocol by Anthropic
OBPA – Ontology for Building Permit Authorities (Hagedorn et al., 2025)
RDF – Resource Description Framework
RSE – Reasoning and State Engine (in proposed architecture)
WFM – Workflow Model (in proposed architecture)

References

Aljas, H.K., Sempelson, K.-S., Pikas, E., Thalfeldt, M., 2025. Framing Needs and Opportunities for Automating Energy Performance Certificates Validation. Presented at the 15th REHVA HVAC World Congress (CLIMA 2025), Milan, Italy.

buildingSMART International, 2023. Regulatory Information Requirements: Project status and findings from the perspective of RAVA3pro project (No. RR-2023-0001-PR). buildingSMART International.

Carvalho, C., Cleto, J., Monteiro, C., Fernandes, J., Fragoso, R., 2024. Energy Performance Certificates: policy needs and best practices | BUILD UP. European Commission.

Casanovas, P., 2017. Legal Compliance by Design (LCbD) and through Design (LCtD): Preliminary Survey.

Dijk, D. van, Hogeling, J., 2019. The new EN ISO 52000 family of standards to assess the energy performance of buildings put in practice. E3S Web Conf. 111, 04047. https://doi.org/10.1051/e3sconf/201911104047

Dridi, A., Patlakas, P., Lefrancois, M., Beach, T., Petr, H., Vakaj, E., 2024. Aec3po: A Knowledge Model for Machine Executable Construction Regulations. https://doi.org/10.2139/ssrn.4784123

European Parliament, 2024. Directive (EU) 2024/1275 of the European Parliament and of the Council of 24 April 2024 on the energy performance of buildings (recast) (Text with EEA relevance).

Feng, L., Wang, Z., 2025. Developing Proactive Compliance Mechanisms for Chinese International Construction Contractors: A PLS-SEM Analysis. Buildings 15, 1478. https://doi.org/10.3390/buildings15091478

Governatori, G., 2018. Declarative Approaches for Compliance by Design. Springer EBooks. https://doi.org/10.1007/978-3-319-76587-7_6

Hagedorn, P., Fauth, J., Zentgraf, S., Seiß, S., König, M., Brilakis, I., 2025. OntoBPR: An ontology-based framework for performing building permit reviews using standardized information containers. Adv. Eng. Inform. 66, 103369. https://doi.org/10.1016/j.aei.2025.103369

Hashmi, M., Koker, L. de, Casanovas, P., 2018. Legal compliance through design: preliminary results of a literature survey. CEUR Workshop Proc. 59–72.

Hettiarachchi, H., Dridi, A., Gaber, M.M., Parsafard, P., Bocaneala, N., Breitenfelder, K., Costa, G., Hedblom, M., Juganaru-Mathieu, M., Mecharnia, T., Park, S., Tan, H., Tawil, A.-R.H., Vakaj, E., 2025. CODE-ACCORD: A Corpus of building regulatory data for rule generation towards automatic compliance checking. Sci. Data 12, 170. https://doi.org/10.1038/s41597-024-04320-x

Hjelseth, E., Labonnote, N., Luczkowski, M., Teclaw, W., 2025. Enabling a fully integrated BIM environment using Semantics and Artificial Intelligence.

Jenkins, D.P., Sayfikar, M., Gomez, A., Fueyo, N., 2024. A Comparative Study of Energy Performance Certificates across Europe. Buildings 14, 2906. https://doi.org/10.3390/buildings14092906

Land and Spatial Development Board of Estonia, 2026. E-ehituse API services in Openapi format [WWW Document]. URL https://swaggerui.ehr.ee/3_d_kaksiku_api (accessed 1.28.26).

Lygerakis, F., Kampelis, N., Kolokotsa, D., 2022. Knowledge Graphs' Ontologies and Applications for Energy Efficiency in Buildings: A Review. Energies 15, 7520. https://doi.org/10.3390/en15207520

Minister of Economic Affairs and Infrastructure of Estonia, 2023. Methodology for Calculating Building Energy Efficiency (Hoone energiatõhususe arvutamise metoodika).

Pauwels, P., Terkaj, W., 2016. EXPRESS to OWL for construction industry: Towards a recommendable and usable ifcOWL ontology. Autom. Constr. 63, 100–133. https://doi.org/10.1016/j.autcon.2015.12.003

Poblet, M., Casanovas, P., Rodríguez-Doncel, V., 2019. Linked Democracy: Foundations, Tools, and Applications, SpringerBriefs in Law. Springer International Publishing, Cham. https://doi.org/10.1007/978-3-030-13363-4

Sayfikar, M., Jenkins, D.P., 2025. Energy performance certificate calculation methodologies across Europe and accommodating new performance indicators. Build. Serv. Eng. Res. Technol. 46, 45–59. https://doi.org/10.1177/01436244241282076

Schaffer, M., Marszal-Pomianowska, A., Wittchen, K.B., Kragh, J., Pomianowski, M.Z., Provera, A., 2026. Assessing Energy Performance Certificate Data over Time: Evidence from Denmark. Proc. 15th REHVA HVAC World Congr. - CLIMA 2025, Lecture Notes in Civil Engineering 2. https://doi.org/10.1007/978-3-032-06810-1_116

Sutton, R., 2019. The Bitter Lesson. incompleteideas.net. URL http://www.incompleteideas.net/IncIdeas/BitterLesson.html (accessed 2.7.26).

Utkucu, D., Sacks, R., 2025. Ontology for holistic building performance modelling and analysis. Autom. Constr. 175, 106197. https://doi.org/10.1016/j.autcon.2025.106197

Wang, Z., Sacks, R., 2025. Building Information Graphs (BIGs): remodelling building information for learning and applications. Data-Centric Eng. 6. https://doi.org/10.1017/dce.2025.10024