

Regulatory Compliance as Procedural Knowledge

Erwin Filtz^{1,*}, Josiane Xavier Parreira^{1,†}

¹Siemens AG Österreich

Abstract

Regulatory compliance requires practitioners to interpret legal provisions and determine whether specific conditions are satisfied in a given situation. Depending on the outcome, different actions, obligations, or restrictions apply, and must be assessed to ensure that the relevant rules are followed. In this paper, we propose a new way of thinking about regulatory compliance, namely as procedural knowledge. Rather than treating compliance solely as the interpretation of legal text, we make explicit the underlying decision process by representing it as a structured procedure, composed of steps, transitions, conditions, and outcomes, that can be inspected and reused. We argue that approaches from procedural knowledge representation provide a natural foundation for capturing the operational structure underlying regulatory decision-making. Moreover, by explicitly capturing this process, we can guide both experts as well as less experienced users in executing all the compliance requirement checks. Moreover, it also provides a structured representation of the tasks carried out, thus enhancing auditability. We outline how such procedures can be modeled using the Procedural Knowledge Ontology (PKO), enabling compliance logic to be expressed in a structured and machine-interpretable format. To illustrate this idea, we present a small case example based on a subsection of the Austrian electricity law concerning the onboarding of new energy grid devices, showing how a regulatory decision process can be represented as a procedure composed of interconnected steps and decision points, how such a representation can support a guided system for carrying out compliance checks, and how the resulting assessment executions can be recorded for traceability and audit.

Keywords

Regulatory compliance checking, Procedural Knowledge Ontology, Knowledge Representation, Guided system, Process auditability

1. Introduction

Compliance checking is a routine but high-stakes activity for organisations operating in regulated environments. As the body of national and international regulation continues to grow, ensuring that organisational decisions and operations conform to applicable legal norms can become increasingly complex and time consuming. Non-compliance can have serious legal, financial, and operational consequences, which makes reliable compliance checking a central organisational concern. This challenge is particularly visible in domains such as energy systems, where regulatory decisions often have to be made in operational contexts. For example, when onboarding a new grid-connected device, operators may need to determine whether the connection request satisfies applicable regulatory and technical requirements, whether additional obligations or restrictions apply, and under which conditions the device can be admitted to the grid.

Compliance checking is particularly challenging because legal provisions rarely translate directly into a single operational answer. Once the relevant provisions have been identified, they must be interpreted in relation to the concrete case and translated into an assessment process: relevant conditions must be examined, their satisfaction must be established in the given context, and the resulting obligations, restrictions, or outcomes must be determined. Domain experts routinely perform this operationalization, drawing on both the regulatory text and their own experience. What usually remains implicit, however,

SAGE'26: The International Workshop on Semantic Architectures for Governance and Explainability co-located with the 22nd International Conference on Semantic Systems (SEMANTICS), Sept 15–17, 2026, Ghent, Belgium

*Corresponding author.

†These authors contributed equally.

✉ erwin.filtz@siemens.com (E. Filtz); josiane.parreira@siemens.com (J. X. Parreira)

id 0000-0003-3445-0504 (E. Filtz); 0000-0002-3050-159X (J. X. Parreira)



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

is the structured sequence of considerations that connects the legal provision to the final compliance outcome.

In this paper, we argue that regulatory texts can be interpreted as sources of procedural knowledge, i.e., they provide information that can be translated into an explicit assessment structure. Together with expert knowledge, they define how a compliance check should be carried out, which conditions need to be examined, and which outcomes may follow. Our central idea is to make this structure explicit by representing it as a procedure composed of steps, transitions, conditions, and outcomes. Doing so enables compliance assessments to be guided in a structured and consistent way, while also making the assessment process easier to trace and audit. We show how this perspective can be modeled using the Procedural Knowledge Ontology (PKO), and ground the approach in an example from an Austrian electricity regulation. The contributions of this paper are as follows:

- A conceptual reframing of regulatory compliance as procedural knowledge, bridging the gap between regulation as an implicit operational decision process and its explicit procedural representation;
- A procedural representation of compliance using the Procedural Knowledge Ontology (PKO), together with the idea of recording executions in a structured way to support traceability and auditability; and
- An illustrative example based on Austrian electricity regulation, showing how provisions concerning the onboarding of grid-connected energy devices can be represented procedurally, how such a representation can support a guided system for carrying out compliance checks, and how the resulting assessment executions can be recorded for traceability and audit.

The remainder of the paper is structured as follows. Section 2 summarizes related work. Section 3 presents the modeling of regulatory compliance as procedures exemplified with an Austrian electricity regulation. Section 4 concludes the paper by discussing the implications and benefits of the approach and outlining directions for future work.

2. Related Work

Research on compliance checking has long investigated how legal and regulatory knowledge can be represented in machine-interpretable form. LegalRuleML provides a standard for representing legal norms, including legal sources, defeasibility, temporal aspects, and deontic concepts such as obligations and permissions [1, 2]. More specific related work on data protection (GDPR) compliance combines legal document representation, legal ontologies, LegalRuleML, BPMN, and rule engines to support compliance checking [3]. These approaches are well suited for representing normative content and enabling legal reasoning. However, their main focus is typically on the formal representation and evaluation of legal rules. In contrast, our work focuses on the operational structure by which a compliance assessment is carried out in practice: the ordered sequence of assessment steps, conditions, decisions, and outcomes that connect a legal provision to a concrete compliance decision.

Another line of work concerns the semantic representation of processes and workflows. PROV-O provides a widely used ontology for representing provenance information across systems and contexts [4]. It is useful for documenting how results were produced, but it does not by itself provide a domain-specific model for reusable procedural definitions and their guided execution. Other approaches model business processes and workflows more explicitly. Rospocher et al. [5] formalise BPMN in OWL, enabling consistency checking and query answering over BPMN process descriptions. Käfer and Harth [6] propose an ontology for specifying, monitoring, and executing workflows in Linked Data environments, including operational semantics via rules. More recently, Krause et al. [7] discuss the representation of dynamic BPMN process executions in knowledge graphs, addressing the gap between static process models and observed process instances. Harth et al. [8] further survey process descriptions

on the Web and argue for graph-based representations that combine process descriptions, observations, and temporal reasoning. These works provide important foundations for process-aware knowledge graphs. Our work differs in that it does not aim to introduce a general workflow language or executable process calculus. Instead, we use procedural knowledge representation to capture the assessment logic that domain experts derive from regulatory text, with particular emphasis on guiding compliance assessments and recording their execution for traceability and auditability.

The Procedural Knowledge Ontology (PKO) is especially relevant because it explicitly models both procedures and their executions [9]. PKO was designed for industrial procedural knowledge management, where procedures are often tacit, distributed across documents, or embedded in expert knowledge. Furthermore, PKO explicitly distinguishes between procedure definitions and executions, which enables both the representation of expert-derived logic and individual compliance decisions for auditability. This makes it a natural fit for our perspective: regulatory compliance assessments can similarly be understood as expert-derived procedures whose structure is often implicit. Previous PERKS work has applied PKO to industrial microgrid operation, where procedures and execution traces were used to improve transparency around operational behaviour [10]. In the present paper, we shift the focus from transparency in microgrid operation to regulatory compliance: the procedure in this new concept does not describe an operational device behaviour, but the structured decision process required to assess whether a regulatory provision applies and what consequences follow.

3. Regulatory Compliance as Procedural Knowledge

As already mentioned, regulatory compliance requires practitioners to interpret legal provisions and determine whether specific conditions are satisfied in a given situation. Based on that assessment, different obligations, permissions or restrictions might apply. Although regulations define conditions and consequences, they rarely describe an explicit assessment procedure on how a compliance decision should be reached in practice.

In reality, the sequence of activities to be carried out due to a regulation is implicitly derived by domain experts. They identify the relevant assessment points contained in the legal text, determine their logical dependencies and establish an ordering that guides the compliance process. As a consequence, compliance assessments are often carried out as structured decision processes that remain largely undocumented or only individually documented and embedded in expert knowledge.

To close this gap, we argue that this information can be understood and represented as procedural knowledge. Rather than focusing exclusively on the formal representation of legal information, our focus is on making the sequence of assessment steps, conditions, decisions and outcomes explicit by putting them into a structure, making them executable, repeatable and auditable. To illustrate this idea, we use a compact example from an Austrian electricity regulation concerning grid access.

3.1. Grid access under ElWG

The example we are using is based on the Austrian *Elektrizitätswirtschaftsgesetz (ElWG)*. The ElWG regulates important aspects of electricity generation, distribution, storage and consumption, including the rights and obligations of actors seeking access to the electricity grid. In particular, we focus on articles §100 to §105 of the ElWG, which concern grid access requests. Depending on the circumstances of the case, access to the grid may be granted (with or without limitations), refused (e.g., due to insufficient grid capacity) or become subject to dispute-resolution procedures. Although the different outcomes are defined in the legal text, there is no specific procedure given on how or in which sequence the individual provisions need to be checked. Therefore, domain experts must translate the regulatory provisions into an ordered sequence of considerations, actions or checks, that enables a consistent compliance decision. We argue that a procedural representation of this knowledge is fitting, which we illustrate below, and it has many benefits, which will be highlighted next.

3.2. From Legal Text to Procedural Representation

To argue for our proposed approach, the selected ElWG articles were translated into a procedural representation by identifying the legally relevant assessment points contained in the text and organizing them into an explicit decision structure.

The mapping begins by identifying the relevant decisions, conditions and outcomes described in the regulatory provisions. Domain experts determine which assessments must be performed, which information is required for each assessment and how the outcome of one assessment influences subsequent steps. In our case study, the modelling process followed four steps: (i) identification of legally relevant decision points and outcomes; (ii) extraction of assessment conditions from the regulatory text; (iii) ordering of dependencies between assessments; and (iv) validation of the resulting procedure by domain experts. The objective was not to replace legal interpretation but to externalise the assessment logic routinely applied by practitioners. The resulting structure can then be represented as a procedure consisting of ordered steps, branching conditions and alternative outcomes. While the resulting procedural representation is validated by domain experts, LLMs can support the identification of candidate steps, conditions and outcomes from regulatory text. However, we want to emphasize that the LLM only has a supporting role and still requires expert review. For example, consider the following two points from the ElWG ¹:

§ 101 ElWG (1): "Where grid access cannot be granted without exceeding network capacity limits, the network operator may apply peak curtailment to wind or photovoltaic installations commissioned from 2027 onwards, provided no exception applies."

§ 101 ElWG (2): "An exception to peak curtailment applies where the installation bears the costs of the direct line itself, or where the installed capacity does not exceed 7 kW. In such cases, grid access shall be granted without power limitation."

Listing 1 illustrates how the regulatory provisions are mapped to PKO procedure steps and requirements. The procedure first evaluates whether sufficient grid capacity is available; if the requirement is satisfied, execution continues with the assessment of the technical connection requirements, otherwise the process transitions to a refusal outcome. This explicit representation captures both the sequential flow of compliance checks and the alternative paths arising from negative assessment results.

Representing compliance logic in this manner makes the assessment structure explicit and inspectable. What previously remained implicit in expert reasoning becomes available as a structured artifact that can be reviewed, reused and adapted when regulations evolve. For example, we have used the PKO model of the ElWG selection to create a guided execution and recording application, described next.

3.3. Guided Execution and Recording

Once represented in a procedural manner, the compliance logic can support the execution of compliance assessments in concrete cases. Rather than requiring practitioners to repeatedly derive the assessment process from legal text, the procedural representation provides an explicit structure that can be followed during execution.

We used the PKO representation of the ElWG to implement a guiding system, shown in Figure 1. Such a system can assist users by presenting the relevant assessment steps, required information, applicable conditions, and possible outcomes in the appropriate sequence. Such support is especially valuable for users who possess the necessary domain knowledge to perform the assessment but may lack extensive experience in interpreting the underlying regulation.

The explicit procedure can additionally serve as the foundation for recording compliance assessments as procedure executions. Each execution captures the path taken through the procedure, the

¹The quotations are unofficial English translations by the authors. For the authoritative wording, readers should consult the original German version of the Austrian Electricity Industry and Organisation Act (ElWG), available at <https://ris.bka.gv.at/GeltendeFassung.wxe?Abfrage=Bundesnormen&Gesetzesnummer=20013066>.

```

@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix pko: <https://w3id.org/pko#> .
@prefix pplan: <http://purl.org/net/p-plan#> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix elwg: <https://example.org/elwg#> .

# Step: evaluate peak-curtailement conditions
elwg:S101_1 a pplan:Step ;
  rdfs:label "Pruefung Spitzenkappungsbedingungen"@de ,
  "Evaluate peak curtailment conditions"@en ;
  pko:stepNumber "101.1" ;
  pko:isStepOfMultiStep elwg:MS101 ;
  pko:nextAlternativeStep elwg:S101_1a ; # exception -> no curtailment
  pko:nextAlternativeStep elwg:S101_2 . # curtailment -> set limit

# Requirement driving the branch
elwg:Req101_Ausnahme a pko:Requirement ;
  rdfs:label "Ausnahmetatbestand: Kostentragung / Leistung <= 7 kW"@de ,
  "Exception condition: self-funded direct connection / capacity <= 7 kW"@en .

# Transition: exception applies
elwg:T101_Ausnahme a pko:Transition ;
  rdfs:label "Ausnahme greift -> keine Spitzenkappung"@de ,
  "Exception applies - no peak curtailment"@en ;
  pko:fromStep elwg:S101_1 ;
  pko:toStep elwg:S101_1a ;
  pko:hasRequirement elwg:Req101_Ausnahme ;
  pko:hasExplanation elwg:Exp_T101_Ausnahme .

elwg:Exp_T101_Ausnahme a pko:Explanation ;
  dcterms:description
    "Die Anlage traegt die Kosten der Direktleitung selbst oder hat
    eine Nennleistung von <= 7 kW -- Spitzenkappung entfaellt."@de ,
    "The installation bears the cost of the direct connection itself
    or has an installed capacity not exceeding 7 kW; in such cases,
    peak curtailment is not applied."@en .

# Outcome of the exception branch
elwg:S101_1a a pplan:Step ;
  rdfs:label "Keine Spitzenkappung (Ausnahmetatbestand)"@de ,
  "No peak curtailment (exception)"@en;
  pko:stepNumber "101.1a" ;
  pko:isStepOfMultiStep elwg:MS101 ;
  pko:nextStep elwg:GW101 .

```

Listing 1: PKO representation of the EIWG grid access procedure (excerpt: §101).

information provided during execution, the conditions evaluated, and the resulting outcome. The PKO RDF representation of such a procedure execution is shown in Listing 2.

Such execution records provide structured documentation of how a compliance decision was reached. This improves traceability and auditability by enabling organisations to reconstruct the assessment process and understand which factors contributed to a particular outcome.

4. Discussion and Outlook

Regulatory compliance requires practitioners to move from legal provisions to concrete decisions in specific situations. While legal texts define relevant conditions and consequences, the operational assessment logic that connects them often remains implicit in expert knowledge. In this paper, we proposed a conceptual reframing of regulatory compliance as procedural knowledge. Rather than

Phase 2 — § 101 Spitzenkappung

✓

§101.1 — Prüfung Spitzenkappungsbedingungen

✓

§101.1a — Keine Spitzenkappung (Ausnahmetatbestand)

💡

Die Anlage trägt die Kosten der Direktleitung selbst oder hat eine Nennleistung von ≤ 7 kW — der Ausnahmetatbestand greift, Spitzenkappung entfällt.

Spitzenkappung ab 2027: Wind- oder PV-Anlage

Yes

No

Wind- oder PV-Anlage, Inbetriebnahme ab 2027, kein Ausnahmetatbestand

§101.2 Leistungsbegrenzung festlegen — *does not apply*

Figure 1: The compliance wizard guiding a user through the ElWG grid-access assessment procedure. Phase 2 shows the sub-steps of § 101 (Spitzenkappung) with completion tracking: § 101.1 and § 101.1a are checked off, and the PKO explanation associated with the taken transition is surfaced inline as a contextual hint. The yes/no decision point below determines whether the curtailment conditions apply, causing § 101.2 to be filtered out as not applicable.

```

# Procedure execution instance
exec:Run1 a pko:ProcedureExecution ;
  pko:hasExecutedProcedure elwg:ProcElWG_Netzzugang ;
  pko:includesStepExecution exec:SE_S101_1, exec:SE_S101_1a ;
  prov:startedAtTime "2025-07-08T09:00:00Z"^^xsd:dateTime .

# Step execution 1 -- evaluate curtailment conditions
exec:SE_S101_1 a pko:StepExecution ;
  pko:hasExecutedStep elwg:S101_1 ;
  pko:isIncludedInProcedureExecution exec:Run1 ;
  pko:nextStepExecution exec:SE_S101_1a ;
  prov:startedAtTime "2025-07-08T09:00:00Z"^^xsd:dateTime ;
  prov:endedAtTime "2025-07-08T09:02:15Z"^^xsd:dateTime .

# Step execution 2 -- exception confirmed, no curtailment
exec:SE_S101_1a a pko:StepExecution ;
  pko:hasExecutedStep elwg:S101_1a ;
  pko:isIncludedInProcedureExecution exec:Run1 ;
  prov:startedAtTime "2025-07-08T09:02:15Z"^^xsd:dateTime ;
  prov:endedAtTime "2025-07-08T09:04:50Z"^^xsd:dateTime .

```

Listing 2: Example of procedure execution record for §101.

viewing compliance solely as the interpretation and evaluation of legal rules, we argue that compliance assessments can be understood as structured procedures composed of assessment steps, conditions, decision points and outcomes.

Using the Austrian *Elektrizitätswirtschaftsgesetz (ElWG)* as an illustrative example, we demonstrated how regulatory text can be translated into an explicit procedural representation using the Procedural Knowledge Ontology (PKO). The example showed that even a relatively compact regulatory subsection contains operational assessment logic, including applicability determinations, condition evaluations, alternative paths, and resulting outcomes. By making this logic explicit, procedural models can support

the guided execution of compliance assessments and enable the structured recording of assessment executions. Such representations make assessment logic transparent and provide a foundation for improved traceability and auditability of compliance decisions.

The contribution of this paper should be understood as an initial proof of concept rather than a comprehensive evaluation. The analysis focuses on a limited regulatory example and does not attempt to demonstrate completeness across broader legal domains. However, the modeling approach can be extended to larger regulatory corpora with richer decision paths.

The distinction between a procedural definition and its execution is particularly valuable in the compliance domain. The procedural definition captures the generalized assessment logic derived from expert interpretation, while execution records document how this logic was applied in a specific case. Such records can support traceability and auditability by providing structured evidence of which assessment path was followed, which information was considered, and how a particular compliance outcome was reached. This capability becomes increasingly important in highly regulated environments where organisations must justify and document their decisions.

Our approach also creates new opportunities for human-centred support systems. Once compliance knowledge is represented procedurally, systems can guide users through the assessment process by presenting relevant tasks, conditions, and decision points at appropriate stages. Such support may be particularly valuable for practitioners who possess the required domain knowledge but have limited experience interpreting regulatory texts. Rather than replacing experts, procedural representations help disseminate expert knowledge and support more consistent execution of compliance activities.

Moreover, this work might have further implications for AI-assisted compliance support. Procedural representations can provide AI agents with explicit assessment structures that define which steps should be performed, which information is required, and how intermediate outcomes constrain subsequent actions. Rather than relying on autonomous generation of compliance reasoning, future systems could use such representations as transparent frameworks for guiding interactions, explanations, and decision support. In this way, AI agents could support the execution of compliance assessments while remaining grounded in procedures defined and validated by domain experts.

Another promising direction concerns the role of execution records. Beyond their immediate value for auditability, collections of execution traces may serve as an organisational memory of past compliance activities. Over time, they could provide contextual information about recurring decision patterns, frequently encountered situations, and common assessment paths. Such procedural histories may become a valuable resource for future compliance support systems and for analysing how compliance processes evolve in response to regulatory change.

Overall, viewing regulatory compliance as procedural knowledge provides a basis for transforming implicit expert reasoning into explicit, executable, and auditable compliance processes.

Acknowledgment

This work was partially funded by the European Union's Horizon research and innovation program project PERKS under grant agreement no. 101120323.

Declaration on Generative AI

During the preparation of this work, the author(s) used Claude Sonnet and M365 Copilot in order to: Drafting content, grammar and spelling check, paraphrase and reword. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

- [1] 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, 5th International Symposium, RuleML 2011- America, Ft. Lauderdale, FL, Florida, USA, November 3-5, 2011. Proceedings, volume 7018 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 298–312. doi:10.1007/978-3-642-24908-2_30.
- [2] Legalruleml core specification version 1.0, OASIS Standard, 2021. URL: <https://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/os/legalruleml-core-spec-v1.0-os.html>.
- [3] M. Palmirani, G. Governatori, Modelling legal knowledge for gdpr compliance checking, in: Legal Knowledge and Information Systems, volume 313 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2018, pp. 101–110. doi:10.3233/978-1-61499-935-5-101.
- [4] T. Lebo, S. Sahoo, D. McGuinness, K. Belhajjame, J. Cheney, D. Corsar, D. Garijo, S. Soiland-Reyes, S. Zednik, J. Zhao, PROV-O: The PROV Ontology, W3C Recommendation, World Wide Web Consortium, 2013. URL: <https://www.w3.org/TR/prov-o/>.
- [5] M. Rospocher, C. Ghidini, L. Serafini, An ontology for the business process modelling notation, in: P. Garbacz, O. Kutz (Eds.), Formal Ontology in Information Systems - Proceedings of the Eighth International Conference, FOIS 2014, September, 22-25, 2014, Rio de Janeiro, Brazil, volume 267 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, 2014, pp. 133–146. doi:10.3233/978-1-61499-438-1-133.
- [6] T. Käfer, A. Harth, Specifying, monitoring, and executing workflows in linked data environments, in: D. Vrandečić, K. Bontcheva, M. C. Suárez-Figueroa, V. Presutti, I. Celino, M. Sabou, L.-A. Kaffee, E. Simperl (Eds.), The Semantic Web – ISWC 2018, Springer International Publishing, Cham, 2018, pp. 424–440.
- [7] F. Krause, K. Kurniawan, E. Kiesling, H. Paulheim, A. Polleres, On the representation of dynamic bpmn process executions in knowledge graphs, in: F. Ortiz-Rodriguez, B. Villazón-Terrazas, S. Tiwari, C. Bobed (Eds.), Knowledge Graphs and Semantic Web, Springer Nature Switzerland, Cham, 2023, pp. 97–105.
- [8] A. Harth, T. Käfer, A. Rula, J. Calbimonte, E. Kamburjan, M. Giese, Towards representing processes and reasoning with process descriptions on the web, TGDk 2 (2024) 1:1–1:32. doi:10.4230/TGDk.2.1.1.
- [9] V. A. Carriero, M. Scrocca, I. Baroni, A. Azzini, I. Celino, Procedural knowledge ontology (PKO), in: E. Curry, M. Acosta, M. Poveda-Villalón, M. van Erp, A. K. Ojo, K. Hose, C. Shimizu, P. Lisena (Eds.), The Semantic Web - 22nd European Semantic Web Conference, ESWC 2025, Portoroz, Slovenia, June 1-5, 2025, Proceedings, Part II, volume 15719 of *Lecture Notes in Computer Science*, Springer, 2025, pp. 334–350. doi:10.1007/978-3-031-94578-6_19.
- [10] E. Filtz, J. X. Parreira, Increasing transparency and explainability in microgrid systems by leveraging procedural knowledge, in: D. Chaves-Fraga, I. Heibi, D. Garijo, D. Collarana, A. A. Salatino, S. Vahdati (Eds.), Joint Proceedings of Posters, Demos, Workshops, and Tutorials of the 21st International Conference on Semantic Systems co-located with 21st International Conference on Semantic Systems (SEMANTiCS 2025), Vienna, Austria, September 3-5, 2025, volume 4064 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2025. URL: <https://ceur-ws.org/Vol-4064/SENTIS-paper6.pdf>.