

Challenges in Legal Process Discovery*

Hugo A. López¹

Department of Computer Science
University of Copenhagen**

Abstract. One of the main promises of process conformance is the opportunity to align normative processes (i.e. how the process *should* behave) and event logs (i.e. how does the process *actually* behaves). Results of conformance checking are valid as long as normative processes correspond to actual norms. Recent developments advocate the use of Natural Language Processing (NLP) to process model discovery from texts. We present a series of challenges in textual process discovery that limit its applicability to real norms. The challenges emerges from experiences with legal practitioners in the digitalization of administrative processes in Danish and Italian municipalities, and they need to be solved in order to provide accurate normative processes that reflect the intent of laws.

Key words: Natural Language Processing, Normative Processes, Process Discovery, Process Conformance

1 On (legal) process compliance and normative models

More than thirty years separate us from Sergot's seminal work on Compliance by Design [1]. In principle, the compliance problem remains the same: given a law, ensure that all executions of a process behave in accordance with the rights and obligations there established, and that it never allows the violations ruled in the law. Specifically, this requires us 1) to provide a formal representation of laws, 2) to establish a formal link between legal policies and events in a process, and 3) to verify that the execution of events does not violate the specification. Process compliance is close to the heart of BPM, where processes operate in environments heavily regulated such as manufacturing, banking or healthcare. However, while decades of research have looked at the semantic representation of norms and compliance verification, fewer efforts have been placed in the elicitation of normative policies for compliance. Laws are *declarative artefacts* that contain events, decisions, rights, obligations, violations and relations between them. Assuming a representation framework for the specification of laws (either in terms of logics [1, 2, 3, 4], or models [5, 2, 6, 7, 8]), how are formal policies related to the legal articles in a law?

Rather than focusing on expressiveness considerations in compliance languages (that have been covered in different surveys, e.g. [9]), our focus is on

* Copyright © 2021 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

** Work supported by the Innovation Fund Denmark project *EcoKnow* (7050-00034A)

mechanising the mapping between a (textual) law and its formal representation. Legal specialists are not trained in compliance languages and they need support to generate formal representations of the law. This requires a triple effort: first, they need to parse the law and identify which fragments can be formalized, second, for each candidate rule, encode the fragments in terms of a formal specification that is semantically equivalent (modulo theory) to the original text. Finally, such interpretations need to be validated so there is a congruence between the possible worlds encoded in the specification, and the legal interpretation. Thus our problem, in short, is:

How can we discover normative models that preserve the intended semantics of laws in a time-efficient way?

The challenges reported in this paper result from 4 years of interactions between the author and municipal sectors in Denmark and Italy in the formalization of laws using DCR graphs in the EcoKnow project [10]. The challenges presented here are not particular to a modelling notation and thus they correspond to any modelling technique used to generate normative models.

2 Elicitation of normative models via NLP

NLP is a key enabling technology for solving our question. Instead of discovery from event logs, NLP allows the identification of processes from texts. This is necessary for the discovery of normative models as ambiguity, length and complexity of laws hinder discovery results. While discovery techniques present encouraging results for (short, imperative) process descriptions, there is no evidence of its application to laws. Moreover, initial application of discovery methods in industrial settings reveals that there is still work to mature the technologies [11]. The challenges emerged from interactions with case workers, lawyers and consultants with domain expertise in danish and italian laws. They resulted in the elicitation of an annotation guideline as an initial step to create corpora containing manually assigned process-law pairs [12]. Our initial experiments in the Danish laws included annotations of 55 articles from the danish administrative acts for family and social services [13]. In the italian case, we focused on municipal laws governing the release of construction permits [14]. Both groups performed textual annotations of laws into process elements, with a follow-up including interviews and think-out-loud sessions. The most interesting challenges follow:

2.1 Challenges for general textual process discovery

Challenge 1: The process in the law Any law contains several pages of text combining technical and non-technical information. Technical information refers to rules and procedures required to generate a legal outcome. In a sense, this is related to the (deontic) logics, whose aim is *“the study of those sentences in which only logical words and normative expressions occur essentially. Normative*

expressions include the words ‘obligation’, ‘duty’, ‘permission’, ‘right’, and related expressions” [15]. A discovery algorithm needs to select only the technical information in the law, filtering non-process information to avoid false-positives.

Challenge 2: Adequate process representations Most textual discovery methods assume that their inputs are imperative processes. For instance, parsing techniques in [16, 17] assume constructs such as start and end nodes. This information does not exist in laws: some rights are inherent, and they are valid for as long as their clauses remain valid.

Challenge 3: Sentence ambiguity Most works on automated process discovery rely on syntax-driven parsing (SDP) where building the abstract syntax tree is important to generate the formula representation of the process [18, 19, 20]. However, SDPs do not consider ambiguities in the semantic of the sentence. For example:

If the agent has completed his additional support and the clerk has issued the money order, the clerk closes the claim.

and

The claim is closed by the clerk if additional support is completed by the agent and the money order is issued by the clerk.

Contain the same rule pattern with antecedents *(the agent) completes additional support* and *(the clerk) issues money order*, and consequent *(the clerk) closes the claim*. This challenge is of particular importance in laws: they are typically written in passive voice, complicating the identification of atoms and rules. The introduction of passive and active voice sentences created divergence in the way legal annotators produced atoms and rules, leading to challenges in interannotator agreement in similar way than law annotations [21].

Challenge 5: Textual process discovery metrics While there is a standard set of measures to benchmark process discovery algorithms for logs [22], there is no consensus on what will be the measures for textual process discovery. The algorithms in [16, 18, 19, 23] all differ in the target modelling language. While each work provides an evaluation in terms of precision, measurements like fitness cannot be applied since there are no explicit traces. In our experience with lawyers, normative models are tested for validity based on their abilities to replicate *legal precedents*, that is, previous cases from the law [8].

2.2 Challenges specific to process discovery of legal texts

Challenge 6: What is a legal event? Most works in process discovery using NLP start considering linguistic patterns in [16] to identify events and activities. Here it is assumed that activities are written in a verb-object pattern (e.g.: *pay compensation for loss of earnings*). Discovering legal events requires us to extend this notion to rights and obligations. Sometimes, event detection will acknowledge the formal recognition for which an event has been performed, as in *“Compensation shall be subject to the condition that **the child is cared for at home as a necessary consequence of the impaired function**”* [24]. Such forms do not correspond to linguistic patterns in the state of the art.

Challenge 7: Policy-formula mismatch A common textual process discovery technique is the identification of stopwords [25]. Our experience is that stopwords are a necessary but not sufficient condition for semantic rule discovery. Assume an interpretation in Linear Temporal Logic³. A response pattern is modelled as the LTL formulae $\mathbf{G}(A \rightarrow \mathbf{F}B)$ and it is associated with an *obligation*: In all cases, if A is executed, there exists an eventual execution of B . For example:

The municipal council shall pay compensation for loss of earnings to persons maintaining a child under 18 in the home whose physical or mental function is substantially and permanently impaired

While the formula $\mathbf{G}(\text{maintainChildUnder18}(X, Y) \wedge \text{physicalMentalImpairment}(Y) \rightarrow \mathbf{F}\text{PayCompensation}(X))$ is expressible in LTL, the interpretation will be contested by a lawyer as this allows a payment on the negation of the antecedents. Modifying the interpretation (e.g. adding a condition relation between the antecedents and the consequence) does not generalize for all uses of “shall”. Interpretations are context-dependent.

Challenge 8: Compositionality Laws are not monolithic artefacts, and creating specifications from them need compositional operators. For example, in:

CASS 48.-(1) Before the municipal council makes a decision under sections 51–63, section 65(2) and (3) and sections 68–71 and 75, the child or young person must be consulted on these matters.

The model can only be expressed if we can compose §48 with §51, 52, etc.

3 Initial ideas towards solving the problem

We believe that there is still a great deal of work to do in order to make business processes mining from text work in industrial cases. Solving the first 5 challenges require a synergy between the NLP-BPM community to release resources (e.g.: corpora, models) for inspection and benchmarking. Moreover, we advocate for a co-created approach to textual process discovery. Rather than building the most accurate set of parsing rules, it is important to i) Embed expert knowledge on what is important in discovery, ii) Disambiguate texts, and iii) Implement techniques that can learn from users. From the language perspective, we see a growing interest in declarative process discovery techniques based on DECLARE [18], ATDP [26] and DCR graphs [23]. They can be adapted in the context of laws. Tools like the Process Highlighter [27] or the Model Judge [28] allow us to filter what should be captured as process information (challenge 1), capture knowledge such as what constitutes an event (challenge 6) and the semantics of legal rules (challenge 7). Tools might benefit from the integration with legal event detection [29] and norm-type classifiers [30, 31].

A second enabler to make discovery usable are language models trained in laws (challenge 3). While general-purpose models (e.g. BERT [32]) capture

³ The same pattern appears in other languages, e.g.: DCR graphs

context-dependency and short-ranging references, they might give wrong results due to the variety of texts they are trained on. Scalability is a factor to consider: the attention layers in transformer architectures scale quadratically. This limits the length of textual analysis and imposes limitations for discovery.

Third, dialogue systems can provide feedback to users about the semantics of each sentence. Surface patterns are indicators of intended meanings [31] that can be refined by presenting characteristic traces. This is fundamental to refine the mappings from surface text to formal semantics, thus reducing misinterpretations (challenge 7). Event-log generation from texts [26] might be able to contrast exemplary traces with user’s intended meaning.

Finally, a necessary step for discovery is the integration of parsers based on formal meaning representation (MR) frameworks [33]. MR aim to represent texts’ formal structure (e.g. graph-based meaning representations), reducing parsers imprecision in sentence variants with the same intended meaning (challenge 3).

References

1. Sergot, M.J., Sadri, F., Kowalski, R.A., Kriwaczek, F., Hammond, P., Cory, H.T.: The British Nationality Act As a Logic Program. *C.ACM* **29**(5) (1986) 370–386
2. Ly, L.T., Rinderle-Ma, S., Knuplesch, D., Dadam, P.: Monitoring Business Process Compliance Using Compliance Rule Graphs. In: *OTM. LNCS*, Springer (2011)
3. Ghose, A., Koliadis, G.: Auditing Business Process Compliance. In: *Service-Oriented Computing – ICSOC 2007. LNCS*, Springer (2007) 169–180
4. Governatori, G., Sadiq, S.: The journey to business process compliance. *IGI global* (2009)
5. Awad, A., Weidlich, M., Weske, M.: Visually specifying compliance rules and explaining their violations for business processes. *JVLC* **22**(1) (2011) 30–55
6. Ramezani, E., Fahland, D., van der Aalst, W.: Supporting domain experts to select and configure precise compliance rules. *LNBIP* **171** (2014) 498–512
7. Burattin, A., Maggi, F.M., Sperduti, A.: Conformance checking based on multi-perspective declarative process models. *Expert Syst. Appl.* **65** (2016) 194 – 211
8. López, H.A., Debois, S., Slaats, T., Hildebrandt, T.T.: Business Process Compliance Using Reference Models of Law. In: *FASE. LNCS*, Springer (2020) 378–399
9. Hashmi, M., Governatori, G.: Norms modeling constructs of business process compliance management frameworks: a conceptual evaluation. *AI & Law* (2017) 1–55
10. Hildebrandt, T.T., Andaloussi, A.A., Christensen, L.R., Debois, S., Healy, N.P., López, H.A., Marquard, M., Møller, N.L.H., Petersen, A.C.M., Slaats, T., Weber, B.: *Ecoknow: Engineering effective, co-created and compliant adaptive case management systems for knowledge workers*. In: *ICSSP, ACM* (2020) 155–164
11. Bellan, P., Dragoni, M., Ghidini, C.: A Qualitative Analysis of the State of the Art in Process Extraction from Text. **2776** 19–30
12. López, H.A.: *Extraction of Processes from Laws: Annotation Guidelines v1.1. shorturl.at/ahmnK* (2021)
13. The Danish Ministry of Social Affairs and the Interior: Consolidation Act on Social Services (September 2015) Executive Order no. 1053.
14. Regione Liguria: Legge regionale n.16 del 6 giugno 2008 e successive modifiche (2008) shorturl.at/kAP58.

15. Føllesdal, D., Hilpinen, R.: Deontic logic: An introduction. In: Deontic logic: Introductory and systematic readings. Springer (1970) 1–35
16. Friedrich, F., Mendling, J., Puhmann, F.: Process Model Generation from Natural Language Text. In: CAiSE. LNCS, Springer (2011) 482–496
17. Soares Silva, T., Toralles Avila, D., Ampos Flesch, J., Marques Peres, S., Mendling, J., Thom, L.H.: A Service-Oriented Architecture for Generating Sound Process Descriptions. In: EDOC. (2019) 1–10 ISSN: 2325-6362.
18. van der Aa, H., di Ciccio, C., Leopold, H., Reijers, H.A.: Extracting Declarative Process Models From Natural Language. In: CAiSE, Springer (2019)
19. Quishpi, L., Carmona, J., Padró, L.: Extracting Annotations from Textual Descriptions of Processes. In: BPM. LNCS, Springer (2020) 184–201
20. López, H.A., Marquard, M., Muttenthaler, L., Strømsted, R.: Assisted Declarative Process Creation from Natural Language Descriptions. In: EDOCW. (2019) 96–99
21. Witt, A., Huggins, A., Governatori, G., Buckley, J.: Converting copyright legislation into machine-executable code: Interpretation, coding validation and legal alignment. In: Proceedings of the 18th International Conference on Artificial Intelligence and Law (ICAIL), Association for Computing Machinery (ACM) (2021)
22. Buijs, J.C.A.M., van Dongen, B.F., van der Aalst, W.M.P.: On the Role of Fitness, Precision, Generalization and Simplicity in Process Discovery. In: OTM. Volume 7565. Springer (2012) 305–322
23. López, H.A., Strømsted, R., Niyodusenga, J., Marquard, M.: Declarative process discovery: Linking process and textual views. In: CAiSE Forum. Volume 424 of Lecture Notes in Business Information Processing., Springer (2021) 109–117
24. Debois, S., López, H.A., Slaats, T., Andaloussi, A.A., Hildebrandt, T.T.: Chain of Events: Modular Process Models for the Law. In: IFM. LNCS, Springer (2020)
25. Winter, K., van der Aa, H., Rinderle-Ma, S., Weidlich, M.: Assessing the Compliance of Business Process Models with Regulatory Documents. In: ER. LNCS, Springer (2020) 189–203
26. Sánchez-Ferreres, J., Burattin, A., Carmona, J., Montali, M., Padró, L., Quishpi, L.: Unleashing textual descriptions of business processes. SoSyM (2021)
27. López, H.A., Debois, S., Hildebrandt, T.T., Marquard, M.: The process highlighter: From texts to declarative processes and back. **2196** (2018) 66–70
28. Delicado, L., Sanchez-Ferreres, J., Carmona, J., Padro, L.: The Model Judge – A Tool for Supporting Novices in Learning Process Modeling. **2196** 91–95
29. Ferraro, G., Lam, H.P., Tosatto, S.C., Olivieri, F., Islam, M.B., van Beest, N., Governatori, G.: Automatic extraction of legal norms: Evaluation of natural language processing tools. In: JSAI-isAI, Springer (2019) 64–81
30. de Maat, E., Winkels, R.: Automated Classification of Norms in Sources of Law. In: Semantic Processing of Legal Texts. Volume 6036. Springer (2010) 170–191
31. Watl, B., Bonczek, G., Scepankova, E., Matthes, F.: Semantic types of legal norms in German laws: classification and analysis using local linear explanations. *AT & Law* **27**(1) (2019) 43–71
32. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805 (2019)
33. Ackermann, L., Neuberger, J., Jablonski, S.: Data-driven annotation of textual process descriptions based on formal meaning representations. In: CAiSE, Springer (2021) 75–90