Handling Inconsistency in Business Rule Bases (Extended Abstract)

Carl Corea[1]

[1]University of Koblenz-Landau, Germany

Abstract

This thesis covers a wide range of aspects related to hand-ling logical inconsistency in business rule bases. In particular, this thesis presents methods and techniques for the *detection, analysis* and *resolution* of inconsistencies. The developed means were evaluated formally, integrated into open-source libraries such as bpmn.io and Camunda, and were evaluated in experiments with human participants. Here, it could be shown that the developed results allow human modellers to understand inconsistencies with a higher efficiency and less needed mental effort. The results of this thesis have been published in various conferences and journals, e.g., Artificial Intelligence (AIJ), and were awarded with the best paper award at the WI conference.

Keywords

Inconsistency Measurement, Busines Rules Management

1. Introduction

In theory, business rules are a backbone for governing compliant business processes. However, what happens if there are mistakes in the set of business rules, e.g., due to human modelling errors?

Business rules are usually created and maintained collaboratively, and over time. In this setting, modelling errors can occur frequently [1, 2, 3]. A central problem here is that of logical inconsistency, i.e., business rules that cannot hold at the same time [4]. For example, consider the following business rule base $\mathcal{B}_1$ from the financial sector, with the intuitive meaning that we have two contradictory business rules, stating that 1.) customers with a mental condition are eligible for a credit, and 2.) customers with a mental condition are not credit eligible (we refer the reader to the full dissertation for syntax and semantics).

$$\mathcal{B}_1 = \{mentalCondition \rightarrow creditEligible, mentalCondition \rightarrow \neg creditEligible\}.$$

Given a process instance where we encounter a customer with a mental condition, $\mathcal{B}_1$ is inconsistent in the classic-logical sense, as it would entail contradictory conclusions. In the worst case, such errors can lead to compliance breaches due to erroneous decision-making or even make it impossible to execute the processes. Such problems can however easily arise in practice, for instance, resulting from different modellers with different views on the same

Proceedings of the Best Dissertation Award, Doctoral Consortium, and Demonstration & Resources Track at BPM 2021 co-located with the 19th International Conference on Business Process Management (BPM 2021)
ccorea@uni-koblenz.de (C. Corea)

domain of interest, company mergers or acquisitions, or a lack of oversight [3, 5]. Importantly, mind that contradictions in rule bases of real-life complexity are not as easy to spot as in the above example, since they often arise from combinations of several constraints. In this context, companies need to be supported with means for the detection and analysis of inconsistencies in the scope of *business rules management* (BRM). In case of inconsistencies, methods are needed that analyze the inconsistency and present the users with a careful analysis as a strategic basis for re-modelling.

Consequently, the aim of this thesis is to support companies in handling inconsistent business rules by developing technical means for *detecting, analyzing* and *resolving* inconsistencies in business rule bases. In particular, measures to *quantify* the severity of inconsistency, means for *pin-pointing* the causes of inconsistency and *repair* algorithms are developed and evaluated. The following sections outline the addressed research gaps and contributions of this thesis.

2. Problem Statement

A central goal of this thesis is to help companies *understand* inconsistencies as a basis for determining suitable and plausible resolution strategies. Here, analyzing and resolving inconsistencies in business rules can be placed within the context of *business rule organizing*, which, in essence, is a BRM lifecycle phase geared towards ensuring error-freeness within a set of business rules [4].

Especially in the scope of understanding inconsistencies, a "simple" detection of inconsistencies is often not sufficient. To anticipate our empirical results, for some real-life rule bases we found situations where there were almost 28.000 contradictory subsets (as in $\mathcal{B}_1$) in a single rule base, many of them of great size and containing long transitive rule chains. In such cases, it is not feasible to "simply" present the user a list of detected problems, and means are needed to *prioritize* the detected problems. From our literature analysis in [4], it can however be observed that despite strong advocation in the literature, there are currently no sufficient means for quantitatively assessing the severity of inconsistencies in business rule bases. In result, the primary research problem addressed in this thesis is to develop means for the (quantitative) analysis of inconsistencies.

Furthermore, in the scope of inconsistency resolution, means are needed to *pin-point* highly problematic rules that should be attended to, e.g., removed. Identifying elements that need to be removed to "repair" a rule base has already been addressed in works such as [6], however, these works are geared towards an automated resolution. Here, as deleting business rules might be highly sensitive, automated approaches might yield implausible results. For example, if a system computes that inconsistency can be resolved by deleting only one rule, this result is of no use if that rule is business-critical and must be retained. Consequently, a second problem addressed in this thesis is to leverage quantitative root-cause analysis (via element-based inconsistency measures) to compute recommendations for inconsistency resolution. This extends current results on inconsistency resolution with a human-in-the-loop perspective and allows experts to consider different resolution strategies.

Regarding the actual development of the envisaged measures, the scientific field of *inconsistency measurement* [7] studies concrete measures that allow to quantify the *degree* of inconsistency with a numerical value and therefore represents a good candidate for an application in this use-case. However, while the application of these results seems promising, due to several conceptual mismatches between the fields of inconsistency measurement and BRM, a straightforward application is not feasible. Put simple, all existing measures always return a value of 0 for problems as in $\mathscr{B}_1$, indicating satisfiability (which is true, e.g. in all cases where the customer does not have a mental condition, but not what is needed here). In result, current measures cannot be used to provide meaningful insights for business rule bases, e.g., at design-time. Therefore, a major problem addressed in this work is to adapt and extend results from inconsistency measurement to allow for a plausible application in BRM. In such, this thesis bridges the gap between the fields of knowledge representation and BRM and provides methods and techniques for practitioners and scholars alike.

3. Contributions

In order to address the research problems raised above, this thesis makes the following contributions.

Contribution 1: Novel means to detect potential inconsistencies in business rules at design-time. To detect "potential" inconsistencies as in $\mathscr{B}_1$ at design-time, we introduce the notion of *quasi-inconsistency*[8], which describes rules that will always be activated together, but have contradictory outcomes. We formalize the notion of quasi-inconsistency [8, 9], propose algorithms to compute all quasi-inconsistent subsets of a rule base [9, 10] and investigate the complexity of central aspects regarding the detection of quasi-inconsistency [8].

Contribution 2: Measuring Inconsistency in Business Rule Bases. We propose new rationality postulates, i.e., desirable properties, that should be met by inconsistency measures applied in a BRM context [8, 11]. We then show by counterproof that all central families of existing inconsistency measures cannot be plausibly applied in this context, as they uniformly violate the raised rationality postulates. Consequently, we show how existing measures can be adapted for application in design-time and run-time rule management. Here, as a main contribution, we show (by means of technical proofs) that the proposed measures satisfy the new postulates while maintaining other desirable properties of their "original" inconsistency measure counterpart [8, 11]. A strong focus is also set on the development of novel element-based inconsistency measures, that can assess the causes of inconsistency (also referred to as culpability measures). We develop novel culpability measures [10, 11], and show how these measures can be used to prioritize problematic rules in the scope of recommender systems [5].

Contribution 3: Resolving Inconsistency based on culpability measurement. Based on the developed element-based inconsistency measures, we propose an approach to resolve (quasi-)inconsistencies in rule bases [10]. Our approach uses the computed rankings to determine rules that should be deleted in order to resolve (quasi-)inconsistency. To enable a close human-in-the-loop integration, the proposed approach implements a semi-automated resolution in order to allow for a stepwise inconsistency resolution, as opposed to a fully automated resolution (which might perform implausible deletion operations). Our proposed approach for inconsistency resolution based on culpability measurement was awarded with the best paper award at the WI conference 2019 [10].

Contribution 4: Algorithms and Library. To allow for a seamless adaptation of our results, the results from Contribution 1-3 are implemented as an open-source library. Here, a strong focus is set on leveraging "theoretical" results from the field of knowledge representation to actual rule standards that may be encountered in practice. We therefore support the rule standards of FCL, Declare and DMN. For instance, the tools in [9, 12, 13] can be used to resolve quasi-inconsistency in Declare, resp. DMN. Likewise, the tool in [14] allows to analyze and resolve run-time inconsistencies in Camunda. The latter tool for monitoring consistent decision-making in Camunda was awarded with the Debeka Innovation Award 2019 by Debeka, a large German insurance company.

The results from Contributions 1-4 are evaluated by means of analytical evaluation (e.g. proofs for the compliance with rationality postulates, computational complexity analysis, proofs for completeness and soundness of algorithms) and run-time experiments (using synthentic data-sets and real-life data-sets such as declarative rules mined from BPI challenge logs). Furthermore, we evaluate the plausibility of applying our results "in practice", explained in Contribution 5.

Contribution 5: Effects of quantitative measures in BRM. We investigate the effects of the developed results in experiments with human participants. To this aim, we test different groups in understanding inconsistencies in business rules, where some groups have access to the element-based inconsistency measures developed in this thesis, and others do not. By means of measures such as performance-time and eye-tracking, we show that our developed results help human modellers to better understand inconsistencies, in a faster time, with better understanding accuracy and with less mental effort needed [15]. Also, we investigate different visualization techniques for actually presenting the developed metrics to users and show that a ranking-based visualization as developed in this thesis helps users to understand inconsistencies in a faster time and with less mental effort needed as opposed to other visualization techniques [16] (cf. also the extended version in [17]).

In the presented thesis, we also show how all above results can be combined into a unified framework for handling inconsistencies in business rule bases. The above contributions have been published in a total of 11 publications [4, 5, 8, 9, 10, 11, 12, 13, 14, 15, 16]. All implemented tools are openly available and can be used "out-of-the-box".

To conclude, this thesis presents novel means for the detection, analysis and resolution of inconsistencies in business rule bases and evaluates the feasibility and plausibility of applying the developed results. The author and colleagues continue to pursue this topic, e.g., measuring inconsistency in multi-sets of rule bases [18], or understanding the behavioral consequences of deleting certain constraints in declarative process models during inconsistency resolution [19].

References

[1] S. Sadiq, G. Governatori, Managing regulatory compliance in business processes, in: Handbook on Business Process Management 2, Springer, 2015, pp. 159–175.

[2] K. Batoulis, A. Nesterenko, G. Repitsch, M. Weske, Decision management in the insurance

industry: Standards and tools, in: Proceedings of the BPM 2017 Industry Track, CEUR-WS.org, 2017, pp. 52–63.

[3] K. Smit, M. Zoet, M. Berkhout, Verification capabilities for business rules management in the dutch governmental context, in: Proceedings of the International Conference on Research and Innovation in Inf. Systems, IEEE, 2017, pp. 1–6.

[4] C. Corea, P. Delfmann, A taxonomy for busines rule organizing approaches in regard to business process compliance, Journal of Conceptual Modeling (2020).

[5] C. Corea, P. Delfmann, Supporting brm with inconsistency analysis, in: Proc. of the BPM 2018 Industry Track, CEUR-WS.org, 2018, pp. 141–147.

[6] C. Di Ciccio, F. M. Maggi, M. Montali, J. Mendling, Resolving inconsistencies and redundancies in declarative process models, Inf. Systems 64 (2017) 425–446.

[7] M. Thimm, Inconsistency measurement, in: Proc. of the 13th Int. Conference on Scalable Uncertainty Management, Springer, 2019, pp. 9–23.

[8] C. Corea, M. Thimm, On quasi-inconsistency and its complexity, Artificial Intelligence (2020).

[9] C. Corea, P. Delfmann, Quasi-inconsistency in declarative process models, in: BPM Forum 2019, Springer, 2019, pp. 20–35.

[10] C. Corea, M. Deisen, P. Delfmann, Resolving inconsistencies in declarative process models based on culpability measurement, in: 14. Internationale Tagung der Wirtschaftsinformatik (WI 2019), AISeL, 2019, pp. 139–153.

[11] C. Corea, M. Thimm, Towards inconsistency measurement in business rule bases, in: Proc. of the 24th European Conference on AI, IOS Press, 2020, pp. 704–711.

[12] C. Corea, J. Blatt, P. Delfmann, A tool for decision logic verification in DMN decision tables, in: BPM 2019 Demo Track, CEUR-WS.org, 2019, pp. 169–173.

[13] F. Hasic, C. Corea, J. Blatt, P. Delfmann, E. Serral, Decision model change patterns for dynamic system evolution., Knowledge and Information Systems (2020).

[14] C. Corea, P. Delfmann, A tool to monitor consistent decision-making in business process execution, in: BPM 2018 Demo Track, CEUR-WS.org, 2018, pp. 76–80.

[15] S. Nagel, C. Corea, P. Delfmann, Effects of quant. measures on understanding inconsistencies in business rules, in: Proceedings of the HICSS 2019, ScholarSpace, 2019, pp. 146–155.

[16] C. Corea, S. Nagel, P. Delfmann, Effects of visualization techniques on understanding inconsistencies in automated decision-making, in: Proceedings of the HICSS 2020, ScholarSpace, 2020, pp. 198–207.

[17] S. Nagel, C. Corea, P. Delfmann, Cognitive effects of visualization techniques for inconsistency metrics on monitoring data-intensive processes, Information Systems Management (2020) 1–16.

[18] C. Corea, M. Thimm, P. Delfmann, Measuring inconsistency over sequences of business rule cases, in: Proceedings of the 18th International Conference on Principles of Knowledge Representation and Reasoning (KR 2021)., 2021, pp. 1–5.

[19] C. Corea, S. Nagel, J. Mendling, P. Delfmann, Interactive and minimal repairof declarative process models, in: Proceedings of the BPM Forum 2021 co-located with the 19th International Conference on Business Process Management (BPM 2021)., CEUR-WS, 2021, pp. 1–16.