

# Unlocking Sustainability Compliance: Characterizing the EU Taxonomy for Business Process Management

Finn Klessascheck<sup>1,3</sup>[0000-0001-6961-1828],  
Stephan A. Fahrenkrog-Petersen<sup>2,3</sup>[0000-0002-1863-8390],  
Jan Mendling<sup>2,3</sup>[0000-0002-7260-524X], and Luise Pufahl<sup>1,3</sup>[0000-0002-5182-2587]

<sup>1</sup> Technical University of Munich, School of CIT, Heilbronn, Germany  
`firstname.lastname@tum.de`

<sup>2</sup> Humboldt-Universität zu Berlin, Berlin, Germany  
`firstname.lastname@hu-berlin.de`

<sup>3</sup> Weizenbaum Institute, Berlin, Germany

**Abstract.** To promote sustainable business practices, and to achieve climate neutrality by 2050, the EU has developed the *taxonomy of sustainable activities*, which describes when exactly business practices can be considered sustainable. While the taxonomy has only been recently established, progressively more companies will have to report how much of their revenue was created via sustainably executed business processes. To help companies prepare to assess whether their business processes comply with the constraints outlined in the taxonomy, we investigate in how far these criteria can be used for *conformance checking*, that is, assessing in a data-driven manner, whether business process executions adhere to regulatory constraints. For this, we develop a few-shot learning pipeline to characterize the constraints of the taxonomy with the help of an LLM as to the *process dimensions* they relate to. We find that many constraints of the taxonomy are useable for conformance checking, particularly in the sectors of energy, manufacturing, and transport. This will aid companies in preparing to monitor regulatory compliance with the taxonomy automatically, by characterizing what kind of information they need to extract, and by providing a better understanding of sectors where such an assessment is feasible and where it is not.

**Keywords:** Sustainability · Conformance Checking · EU Taxonomy · Business Processes.

## 1 Introduction

In light of the issue of climate change and unsustainable human activity [43], it is important to promote *sustainable* business practices, i.e., conducting business in a way that can meet the needs of the present generations without endangering those of the future [6]. This need has also been identified by governing bodies, such as the *European Union* (EU). As a consequence, the EU has defined the *taxonomy for sustainable activities* [10, 13], subsequently referred to by us as *taxonomy*. The taxonomy aims to create clear indicators of when business activities are contributing towards sustainability and when not, and to create financial incentives for proven sustainable

business practices and investments into them [29, 41]. For various *business practices*,<sup>4</sup> the taxonomy defines criteria along which a *substantial contribution* towards *sustainability goals* can be verified, and criteria which must not be violated. Increasingly, companies *will* face having to assess their business processes for compliance with the taxonomy [29]. However, assessing whether a business practice does or does not meet relevant criteria is, so far, a manual process: Some companies offer manual or semi-automatic questionnaire-based assessments; a *taxonomy calculator* provided by the EU relies exclusively on manual input in the form of an Excel sheet.<sup>5</sup>

To overcome the challenges of manually assessing whether a business activity meets the criteria of the taxonomy, it appears feasible to *check in a data-driven manner* whether the execution of a business practice complies with relevant taxonomy criteria or not. Since the definition of business practice [11] is closely related to that of business processes [44], we interpret business practices as “categories” of business processes, which allows us to investigate the taxonomy and its criteria from a *business process management* (BPM) standpoint.

*Conformance checking*, which is a technique of the *process mining* and BPM fields, aims to compare recorded business process executions in the form of an event log with a formal representation of the to-be process behavior, so that either, the process execution can be improved to more closely resemble the formal representation, or vice versa [7]. Automatic compliance monitoring can use conformance checking techniques with the goal of assessing whether a business process complies with regulatory constraints — such as those described in the taxonomy — during its execution, based on recorded event data [16, 21]. An overview of this is provided in Fig. 1. For applying conformance checking with the aim of monitoring compliance to the regulatory constraints described in the taxonomy, this taxonomy first needs to be *operationalized*, that is, translated into the form of a *prescriptive model*. Further, the prescriptive model and the event data used for conformance checking need to align w.r.t. what information they contain. For this, companies need to be aware of what data they need to capture during the execution of their business processes. Therefore, we see a need to better understand what data-capturing requirements the taxonomy imposes on business processes, and in how far the constraints contained therein are even applicable for conformance checking.



**Fig. 1.** Conceptual overview of compliance monitoring with conformance checking [7, 21]

<sup>4</sup> Referred to as *economic activities* in the taxonomy; to avoid confusion with the notion of business process activities we use the term *business practice*.

<sup>5</sup> See <https://viridad.eu>, <https://www.briink.com/solutions/esg-questionnaire-assistant> and <https://ec.europa.eu/sustainable-finance-taxonomy/wizard> [Accessed: 23/05/2024]

Related work has focussed on extracting constraints from text directly into process models (e.g., [1, 30, 47]) and on characterizing textual constraints based on i.a. their relation to other constraints, process models, or their relevance to a given business process (e.g., [39, 45, 46]). Our investigation, however, aims at an earlier stage of conformance checking that does not require operationalization into a concrete prescriptive model or a concrete business process against which conformance is to be checked. Rather, we aim to understand what kind of requirements the taxonomy imposes on event log data so that it can be captured and prepared for regulatory compliance checking with conformance checking w.r.t. the taxonomy. For this, we first need to understand what kinds of constraints the taxonomy embodies (i.e., which characteristics a prescriptive model would make prescriptions towards, such as certain activities that need to be executed in a specific order, certain thresholds that activities must not exceed, etc.), and whether all of them can be related to a process view, since so far, it is unclear to what extent the taxonomy can even be operationalized for conformance checking. Therefore, this work aims to address the following *research questions* (RQ):

**RQ1:** How can the EU taxonomy be operationalized with regard to business processes?

**RQ2:** Which constraints of the EU taxonomy can be used for automatically assessing whether a business process fulfills its respective sustainability criteria with conformance checking techniques?

Since the taxonomy contains constraints for around 80 types of business practices [28, 41] – which makes a manual characterization of the entire taxonomy infeasible – this work uses a *few-shot machine learning approach* to gain insights into what constraints the taxonomy consists of, and how they might be operationalized in practice. In doing so, we also explore the potential of novel approaches based on *large language models* (LLMs) for operationalizing regulations in the area of conformance checking.

The remainder of the article is organized as follows: Section 2 provides background on the taxonomy, regulatory compliance monitoring with conformance checking, and few-shot learning approaches. Section 3 provides related work on approaches for extracting rules from text for compliance monitoring. In Sec. 4, we present the research approach of this paper. Section 5 provides the results thereof, characterizing the constraints of the taxonomy and its potential for conformance checking uses. We further discuss their implications for practice in Sec. 6. Finally, Sec. 7 provides future work and concludes the article.

## 2 Background

Next, we will outline the general idea of the EU taxonomy, and we will give an overview of automatic compliance monitoring approaches within the business process literature.

### 2.1 EU Taxonomy for Sustainable Activities

In order to create incentives for investments in sustainable technologies and to provide a transparent classification system for when business practices are sustainable, the

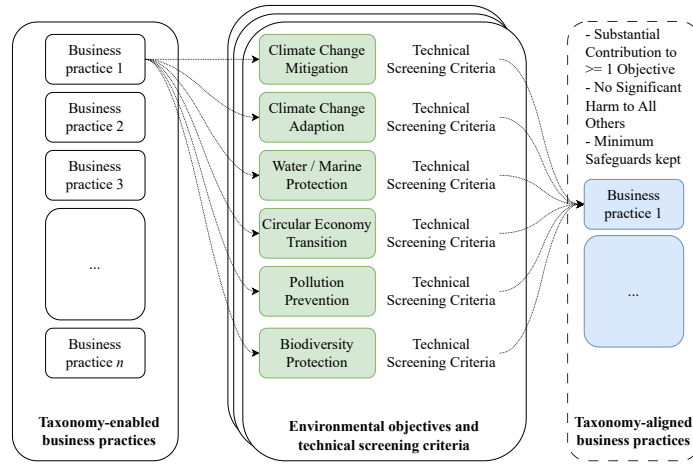
EU has established the *taxonomy for sustainable activities* [2, 10, 41]. The ultimate objective behind the taxonomy is to support the EU in transiting to climate neutrality by 2050 [29, 41].

In essence, the taxonomy defines six *environmental objectives* with which sustainable business practices are identified: 1) mitigating climate change; 2) adapting to climate change; 3) sustainably using and protecting water and marine resources; 4) transitioning to a circular economy; 5) preventing and controlling pollution; and 6) protecting and restoring biodiversity and the ecosystem [8, 28]. Not all possible business practices of all industries are covered by the taxonomy, but only those which are deemed to be able to make a *substantial contribution* towards climate neutrality, or are needed for other sectors to make a substantial contribution [41]. If a business practice is part of the taxonomy, it is called *taxonomy-enabled*, and can *potentially* make a contribution to one of the six environmental objectives. If a taxonomy-enabled business practice: 1) indeed contributes to one of the environmental objectives; 2) causes no *significant harm* (DNSH) to any of the six objectives; 3) meets *minimum safeguards* (such as the UN Guiding Principles on Business and Human Rights); 4) adheres to *technical screening criteria*, it is, in fact, *sustainable* according to the taxonomy [2, 8, 28]. Business practices that meet all of these criteria are also called *taxonomy-aligned*. In short, to be taxonomy-aligned, a taxonomy-enabled business practice *must* contribute to at least one of the six environmental objectives, *must not* cause significant harm to any of the others, and *must* meet minimum safeguards [2, 8, 28]. Notably, the minimum safeguards are not directly defined *in* the taxonomy, but rather, are references to taxonomy-external guidelines and regulations. Concretely, organizations need to ensure that they follow the *OECD Guidelines for Multinational Enterprises*, *UN Guiding Principles on Business and Human Rights*, the *Declaration of the International Labour Organisation on Fundamental Principles and Rights at Work* as well as the *International Bill of Human Rights* [13, Article 18]. Since the focus of our work is on compliance of business processes and not of organizations or supply chains, and minimum safeguards are often not assessed on the level of business processes [29], we forgo including these external constraints in our investigation. Figure 2 provides a schematic overview of the taxonomy and its concepts.

For assessing the taxonomy alignment of a taxonomy-enabled business practice, the taxonomy describes *technical screening criteria*, which describe under which circumstances an activity either makes a contribution or causes significant harm to an environmental objective [28, 41]. It should be noted that a business practice is not necessarily able to make substantial contributions to more than one economic objective, and hence may have only one set of technical screening criteria for one substantial contribution.

For an illustration of how the taxonomy documents technical screening criteria for environmental objectives, we refer to the EU’s *taxonomy compass* and the corresponding Excel file.<sup>6</sup> Generally, for each business practice, a set of possible environmental objectives to which a substantial contribution can be made is provided — for each objective, criteria for a substantial contribution, as well as DNSH, are documented.

<sup>6</sup> See <https://ec.europa.eu/sustainable-finance-taxonomy/taxonomy-compass/the-compass> and <https://ec.europa.eu/sustainable-finance-taxonomy/assets/documents/taxonomy.xlsx> [Accessed: 18/06/2024]



**Fig. 2.** Schematic overview of the EU taxonomy for sustainable activities and its concepts, derived from [10, 12]

By determining how many of their business processes align with the taxonomy, companies can report how much of their business output is generated via sustainable business practices [12, 29]. For this, companies need to know exactly which of their business practices are not just taxonomy-enabled, but also taxonomy-aligned — and in the future, these taxonomy disclosures will be subject to mandatory audits [29]. Notably, from 2025 onwards, companies that meet certain characteristics (i.e., more than 250 employees, more than EUR 25M balance sheet total, or more than EUR 50M net turnover) will be obligated to do so, with the threshold for having to report decreasing in subsequent years [29]. Small and medium-size enterprises will also be subject to a disclosure obligation [29]. Therefore, it appears prudent to investigate how automatic business process compliance monitoring can help companies in assessing the taxonomy-alignment of their business processes.

## 2.2 Automatic Business Process Compliance Monitoring

The automatic monitoring of business process compliance [27] allows organizations to ensure their business practices comply with regulations [18]. Through an analysis of process execution data, it is possible to check at runtime if a business process complies with specified rules [16, 21]. Figure 1 provides a conceptual overview of automatic compliance monitoring with conformance checking, with the EU taxonomy being one example of a regulation that can be operationalized. In order for an organization to adopt these techniques, the following steps need to be taken [7, 21]:

1. The relevant piece of regulation needs to be operationalized into a *prescriptive model*. This model embodies constraints towards one or more *process dimensions* [37], which are *control-flow*, *time*, *resources*, and *data* [7, 21], of the process under investigation.

2. An *event log* needs to be extracted from recorded process executions, describing the actual process behavior [36].
3. The prescriptive model and event log are used by a conformance checking *algorithm*, which checks whether or where the process deviated from the prescriptive model. This can serve as a starting point for further diagnosing or explaining deviations, and subsequently, remedying unwanted deviations [7, 21, 35].

It should be noted, however, that the prescriptive model and event log need to align in their respective process dimensions: If, for example, the prescriptive model imposes constraint on the data dimension of the investigated process, but the event log does not contain this information, the conformance check *cannot* yield the desired insights [7]. Hence, organizations need to know, potentially in advance, which process dimensions are, in fact, *relevant* for the conformance check. Based on this, they can appropriately capture the relevant execution data and extract it into the subsequent event log.

### 3 Related Work

In this paper, we investigate the properties of the EU taxonomy for sustainable business practices and its potential for being used for compliance monitoring with conformance checking, by extracting insights from the taxonomy’s regulatory texts and the compliance constraints described therein. In that respect, our work relates to other contributions that also deal with rule extraction for conformance checking/process mining purposes, and contributions that use machine learning techniques to do so.

*Constraint Extraction from Text.* For extracting compliance constraints from regulatory documents, Dragoni et al. [9] propose a pipeline that combines multiple *natural language processing* (NLP) approaches. With their proposed pipeline, rules (in this case in the form of *obligations*, *permissions*, and *prohibitions*, see Hashmi et al. [19]) can be extracted into formal representations for a given regulatory text. Using semantic annotations, a process model can then be checked for whether it complies with the formal representations – this check is further described by Governatori et al. [15]. Moreover, Winter and Rinderle-Ma [47] describe an approach to generate process model fragments from regulatory documents by extracting constraints and their relation. Similarly, van der Aa et al. [1] describe an automatic approach for extracting declarative process models from natural language text that describes a process. Barrientos et al. [3] design an approach for extracting temporal constraints from natural language texts and determining violations thereof in an event log. Focussing on resource compliance, Mustroph et al. [30] extract compliance requirements from natural language with GPT-4, which are then matched to and verified against an event log. Further, Mustroph et al. [31] describe how generative AI, or more specifically, GPT-4, can be used to pre-process resource-related regulations into social network graphs. Based on these, they are able to detect compliance violations of process executions.

*Characterization of Constraints.* In terms of characterizing constraints present in regulatory documents, Winter et al. [48] provide a technique for characterizing

regulations based on text mining and clustering algorithms that can derive significant sentences for a regulatory document, which can be manually translated into e.g. process models. Similarly, Winter and Rinderle-Ma [46] design an approach to group constraints from textual documents and detect relations between them, e.g. based on similarity. Moreover, Winter et al. [45] provides a method for matching parts of regulatory documents with process models, which then allows a compliance assessment of the matched regulatory constraints and the process model. Aiming to facilitate a better understanding of regulatory documents, Sai et al. [38] propose an approach to parse legal definitions and relations between terms from regulatory documents into a knowledge graph, so that regulatory documents can be better understood and analyzed. Further, in order to compare regulatory documents with their translation into company-internal requirements, Sai et al. [40] provide an NLP-based which can detect deviations between the two documents, and helps to detect root causes of textual deviations. Finally, Sai et al. [39] investigate an automated approach for identifying passages of regulatory texts that are relevant for a business process based on the processes textual description. They find that expert judgement cannot be replaced by generative AI, but see the potential of AI uses for taking into account vaster amounts of context, which the authors deem to be advantageous in more complex settings.

In contrast to these contributions, which primarily focus on extracting constraints directly or characterizing them w.r.t a process model, event log or other texts, we exclusively focus on characterizing the constraints for the process dimensions they constrain. This would facilitate data extraction of recorded process execution and translation of the taxonomy into concrete constraints for subsequent analyses. However, with our categorization as a starting point, relevant extraction and matching approaches can be chosen, and the relevant data can be stored during process execution.

Notably, there is currently no automated approach that helps companies to understand the requirements posed on business process executions by regulatory texts w.r.t. the process dimensions so that the log and the subsequent prescriptive model pertain to the relevant perspectives and contain relevant information for compliance monitoring with conformance checking. In particular, the taxonomy has not yet been considered in this light — however, concepts from existing approaches can inform the design of new mechanisms, such as ours.

## 4 Mapping Sustainability Regulation to Conformance Constraints

In order to operationalize the taxonomy for business processes, we need to identify process constraints within the taxonomy. This allows us to transform the abstract rules for business practices into specific problems that can be addressed using automatic conformance monitoring.

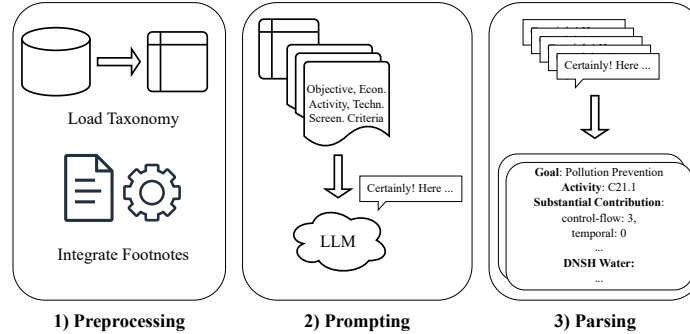
We assume that it would be best to use an industry expert with conformance checking knowledge to classify each rule of the taxonomy. However, since such experts are rare and the taxonomy covers many varying industries, an alternative solution becomes necessary. Because of this, we decide to use LLMs, since they are trained on

a wide variety of texts from different domains. Furthermore, our task can be framed as a text classification problem, and it was shown that LLMs can be used as *Few-Shot-Classifiers* [5]. Meaning that an LLM, which is trained for one task such as next text token prediction, can be with sufficient instructions utilized to perform another machine learning task, in our case text classification of the taxonomy. In particular, previous work has shown that LLMs can be successfully be used for process analysis tasks [17].

In order to characterize the taxonomy, which consists of several regulatory texts for each business practice and environmental objective it considers, we need to apply an LLM to each regulatory text and extract relevant information about the constraints. For this, we developed a pipeline, which we outline in the subsequent Section 4.1, drawing on the capability of LLMs to be applied in this manner.

#### 4.1 Overview

The overall approach for characterizing the constraints, or technical screening criteria, of the taxonomy for the process dimensions they pertain to, is illustrated in Fig. 3. The pipeline we describe consists of three stages: First, the taxonomy is *preprocessed* (see Section 4.2). Second, for each set of technical screening criteria of each business practice and environmental objective, an LLM is *prompted* to identify the types of constraints contained therein (see Section 4.3). Third, the LLM’s output is *parsed*, and the numbers and types of constraint per business practice and environmental objective is extracted (see Section 4.4). Finally, the resulting data is collected and can be *analyzed*.



**Fig. 3.** Schematic overview of the taxonomy constraint characterization pipeline

#### 4.2 Preprocessing

For the preprocessing step, we begin to read the entire taxonomy, which is available in the form of an Excel file.<sup>7</sup> For each environmental objective, the taxonomy file

<sup>7</sup> <https://ec.europa.eu/sustainable-finance-taxonomy/assets/documents/taxonomy.xlsx> [Accessed: 18/06/2024]



contains a sheet, in which each environmental activity as well as 1) the corresponding technical screening criteria for a substantial contribution to the objective, and 2) the technical screening criteria for DNSH to all other objectives are listed. Further, footnotes pertaining to each activity are contained in a column in each sheet. In order to make all relevant information available to the subsequent prompting step, the footnotes are appended to each technical screening criteria block of an activity if they are referenced therein.

### 4.3 Prompting

Subsequently, the approach iterates through the taxonomy in the following manner: For each of the six climate objectives, and for each business practice that can make a substantial contribution to that objective, the taxonomy contains: 1) a text that describes the technical screening criteria with which a substantial contribution can be determined, and 2) up to five texts that describes the technical screening criteria with which significant harm to the other objectives can be determined. For each of these texts, the approach prepares a prompt to an LLM, with which the types of constraints and numbers thereof present in the respective technical screening criteria can be determined. The prompt includes a small task description and a brief explanation of the process constraint types, as well as the text passage that is currently being characterized.

We differentiate the process constraints along two aspects: (i) We classify what **process dimension** [7, 37] is targeted by a constraint. We categorize constraints into one of the following dimensions: *control-flow*, *temporal*, *resource*, and *data*. Alternatively, a constraint can be *irrelevant* from a process perspective; (ii) Furthermore, we distinguish the **granularity** of a constraint: meaning a constraint can either be targeted towards an aspect *within* an activity or *between* activities.

An example of a *temporal* constraint *between* activities could be the requirement to perform one specific activity within one week after another activity was performed, while an example of a *resource* constraint *within* an activity is an activity that needs to be performed by a person with a specific certification. Constraints classified as pertaining to the *control-flow* *within* an activity can be understood as *activity existence* constraints.

The prompt to retrieve this information is structured as follows:

- Briefly, the overall objective is described;
- The different types of process constraints are described, as well as the difference in granularity;
- Examples consisting of excerpts of the taxonomy and their potential classification are provided;
- The exact task of characterizing a section of the taxonomy is described;
- Requirements regarding the output are described;
- A placeholder is provided where, during pipeline execution, a description of the business practice can be inserted; and
- A placeholder for the text passage of the taxonomy for which the characterization needs to be done is provided, which will be filled during pipeline execution.

The entire prompt template is available online.<sup>8</sup> As a result, we receive the following information as a response from the LLM, giving us insights into the process constraints covered by a particular set of criteria:

- # of **activity existence** constraints of process activities
- # of **control-flow constraints** *between* process activities
- # of **temporal constraints** *within* process activities
- # of **temporal constraints** *between* process activities
- # of **resource constraints** *within* process activities
- # of **resource constraints** *between* process activities
- # of **data constraints** *within* process activities
- # of **data constraints** *between* process activities
- # of **process-irrelevant** constraints

#### 4.4 Parsing

Each response of the LLM to a prompt of one text (that describes a set of technical screening criteria) is parsed individually. The prompt instructs the LLM to return the results in a particular *JSON*-like notation. This is depicted in Listing 1.1.

**Listing 1.1.** Excerpt of the prompt template, describing the required response structure

```
{'control-flow': {
  'within_activities': [no. of activity existence constraints],
  'between_activities
    ': [no. of control-flow constraints between activities]},
'temporal': {
  'within_activities
    ': [no. of temporal constraints within activities],
  'between_activities
    ': [no. of temporal constraints between activities]},
'resource': {
  'within_activities': [no. of resource constraints within activities],
  'between_activities
    ': [no. of resource constraints between activities]},
'data':{
  'within_activities': [no. of data constraints within activities],
  'between_activities': [no. of data constraints between activities]},
'irrelevant': [no. of process-irrelevant constraints]}
```

In addition to some further processing steps (such as replacing comment-like symbols), this structure and the number of the respective constraints are extracted from the response. Additionally, we store the entire response text, as it often contains the LLMs “explanation” for the constraints that were identified. As a result, we know the number of constraints and types of one set of technical screening criteria for one business practice and one environmental goal. The parsing step is repeated for all further sets of technical screening criteria and prompt responses of each activity. Ultimately, we end up with six datasets, one for each environmental objective. Each dataset contains information on the number of constraints imposed on each business practice by the substantial contribution and DNSH criteria, which, once collated, subsequently allows further analyses.

<sup>8</sup> <https://github.com/fyndalf/unlocking-sustainability-compliance-replication-package>

## 5 Experimental Validation

In this section, we use our approach to validate the extraction of conformance constraints from the EU taxonomy. First, we outline our experimental setup in [Section 5.1](#). Next, we apply our approach to the EU taxonomy and report the results in [Section 5.2](#). Finally, we show the validity of our approach in [Section 5.3](#).

### 5.1 Experimental Setup

We implemented the pipeline using Python 3 and Pandas in a Jupyter Notebook. The entire implementation, as well as all data we used and generated, is made available online for reproducibility purposes.<sup>9</sup> Additionally, the LLM we used in this study was Meta’s Llama3,<sup>10</sup> which is openly available. More precisely, we used *llama-3-8b-instruct*, which is a version of Llama3 with 8 billion parameters that is fine-tuned for following instructions. Our prompting followed the official Llama3 documentation. For guaranteeing a quasi-deterministic output, we followed guidance on model-specific settings (such as *temperature*, *seed*). Instead of self-hosting Llama3, we opted for using a web-based service (<https://groq.com>), which at the time of writing offered free API access to a hosted instance of Llama3. However, our implementation can easily be adapted to access other hosts or LLMs.

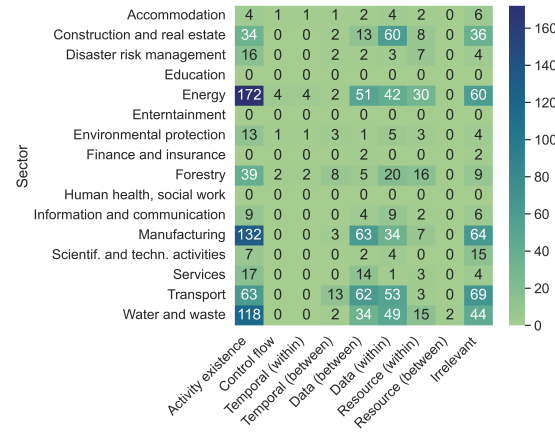
### 5.2 Results

*Constraint Types.* When analyzing the process constraints extracted with our pipeline, we observe the following distribution of constraint types: Out of a total of 1636 constraints we identified, we see a large focus on control-flow constraints (activity existence: 624; between: 8). For example, in order to make a substantial contribution to climate mitigation, market research and development business practices in the area of emission reduction technologies are required to have obtained a permit for operating a demonstration site, which can be interpreted as an activity existence constraint. Further, to contribute to the climate mitigation objective, business practices concerned with the construction, extension and operation of waste water collection and treatment must under some circumstances conduct an assessment of greenhouse gas emissions and subsequently disclose the result to investors (which can be understood as a control-flow constraint). This is followed by process-irrelevant constraints (323). For example, the substantial contribution criteria to climate mitigation of operating and providing personal mobility devices logistics requires that the vehicles are allowed to operate on the same infrastructure as bicycles and pedestrians, which is related rather to the environment in which the business practice is conducted, and not the business practice itself. The third-most present constraint type is the one of data constraints, both within process activities (284) and between them (255). For example, the substantial contribution criteria for climate mitigation of the manufacturing of

<sup>9</sup> <https://github.com/fyndalf/unlocking-sustainability-compliance-replication-package>

<sup>10</sup> <https://llama.meta.com/llama3/> [Accessed: 18/06/2024]

iron and steel describes greenhouse gas emission threshold for individual steps of the manufacturing process. As another example, the DNSH to pollution criteria for the biodiversity objective of conservation and environmental protection requires the use of fertilizers across the entire business practice to be minimized, which can be understood as a data constraint between activities. Less common are resource constraints (within: 96; between: 2) such as the substantial contribution criteria for the circular economy objective of the business practice of preparing end-of-life products and components for re-use, which makes requirements towards the tools and equipment being used. Finally, temporal constraints (within: 8; between: 36) are the least common: For example, the business practice of providing solutions for flood and drought risk prevention and protection is required to review the implemented solution periodically, in order to make a substantial contribution to the water protection objective.



**Fig. 4.** Constraint types per industry sector across all environmental objectives and screening criteria

*Industry Sectors.* Looking at the individual industry sectors, we can further observe several patterns in the constraint types. Figure 4 depicts the sectors and constraints identified across all environmental objectives and screening criteria. First, we observe that energy, manufacturing, water, and transport are the industries most often constrained as to their contribution towards one or more climate goals. Noticeably, we see that activity existence criteria (such as permits that need to be obtained, assessments that need to be conducted) and data constraints (such as greenhouse gas limits, energy usage limits, etc.) play a vital role in assessing the taxonomy alignment in these sectors. Second, we see that resource constraints between activities, as well as control flow and temporal constraints play less of a role across all sectors. In the finance and insurance sector, we only identified two process-relevant constraints and two irrelevant ones. Finally, we see that three sectors, namely education, entertainment and human health/social work, have no identified constraints at all. A manual investigation into the taxonomy reveals that for both sectors, the taxonomy only

provides substantial contribution criteria to the climate adaption goal, and no DNSH criteria. The substantial contribution criteria are provided in a very abstract manner, and we were unable to manually identify fine-grained process constraints.

*Environmental Objectives.* Next, we analyze the constraint types belonging to the screening criteria of different environmental goals (meaning, the respective substantial contribution criteria of one goal and the associated DNSH criteria for all other goals). Figure 5 illustrates the number of constraints of each type per environmental objective. Here, we see that climate adaption and climate mitigation — which are the two environmental objectives with the highest number of business practices for which they govern taxonomy alignment — also contain the highest number of constraints across all types. Interestingly, we see that constraints related to the goals of pollution prevention and water protection are not at all characterized w.r.t. control-flow or temporal aspects within activities, while some constraints for the biodiversity, climate adaption, and climate mitigation use these process dimensions. Across all environmental objectives, however, we see a general focus on activity existence and data constraints, with resource constraints within activities also being represented.

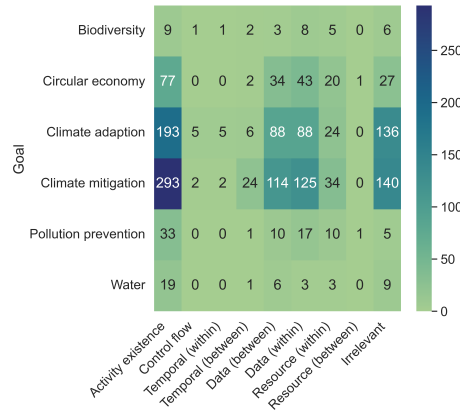


Fig. 5. Number of constraints per type and environmental objective

### 5.3 Validation

For validating our approach and the resulting insights, a research assistant with a background in environmental and resource management and one of the authors with a background in computer science and BPM manually assessed a sample of the taxonomy and constraint characterization (i.e., all 357 characterizations for the environmental objectives of water protection, circular economy, pollution prevention, biodiversity protection). We compared each characterization and the underlying response of the LLM with the taxonomy’s original text and assessed whether the result was *entirely plausible* (we deem that all constraints have been found and classified

accurately), *largely plausible* (we deem that all constraints have been found, with a slight deviation in the constraint types; such as when a constraint that can be read as an activity existence constraint requiring an activity to be executed has instead been read as a resource constraint requiring the activity to be executed by a specific role or resource), *somewhat plausible* (we deem that most relevant constraints have been found, but their classification is debatable), or *implausible* (we deem that central constraints have not been found, or constraints have been clearly mischaracterized). Conflicts regarding the plausibility assessment were resolved in discussions. This mode of validation allows us to assess the results of our approach without creating a “gold standard” result (as other studies do, see e.g., [39]), since we lack the regulatory expertise which would be necessary for creating such a standard.

Table 1 shows the share of constraints we assessed regarding their plausibility. We see that 340 of 357 characterizations are *at least* assessed as somewhat plausible. Distributing the plausibility on a four-step scale (three being entirely plausible, zero being implausible), we see an **average plausibility of 2.74**, i.e., more than largely plausible. This means that, in general, we expect a characterization to be at least largely plausible. While we observed implausible characterizations, they seem to largely stem from references to taxonomy-external regulatory texts and standards which have not been considered and ambiguous terminologies (such as “times” as a frequency instead of referring to a temporal aspect). Overall, we infer that the classification approach can serve as a starting point for creating prescriptive models, and that it provides largely plausible constraint characterizations, which may need to be supplemented with a manual investigation. This is particularly the case when external regulations are involved.

**Table 1.** Plausibility assessment of 357 constraint characterizations

Assessment	Entirely Plausible	Largely Plausible	Somewhat Plausible	Implausible
Characterizations	308	24	8	17

## 6 Discussion

After presenting and validating our results, we now discuss them further. As we have seen, business practices in the industry sectors of energy, manufacturing, transport, and water and waste constitute a large part of the process-relevant constraints we identified and thus appear well-suited to be investigated for their taxonomy alignment with conformance checking. In some sectors, particularly finance, education, entertainment, human health and social work, we were unable to identify a high number of constraints that would have been operationalizable for conformance checking. Therefore, we conclude, that these sectors appear less promising for applications of compliance monitoring with conformance checking.

In general, we can apply conformance checking to around 80% (i.e., 1313 of 1636) of the constraints which we identified in the taxonomy. For all other constraints, and in particular, in sectors where we had difficulties automatically identifying constraints, conducting manual compliance monitoring and taking further expert knowledge into account appears indispensable. Hence, we have answered RQ2.

Further, as we have shown, the approach we have designed can aid in the creation of prescriptive models for compliance monitoring with conformance checking by helping end users to better understand: 1) what types of constraints are likely to be present in a single relevant piece of regulation (since it can be difficult to determine the concrete constraint type, or whether constraint is actually operationalizable), and 2) what types of constraints comprise a larger set of regulations (as it helps in choosing and implementing correct techniques). As a subsequent step, end users then need to operationalize the constraints into a prescriptive model for automatic compliance monitoring with conformance checking, drawing on existing approaches for this, in relation to actual business processes. This allows the taxonomy to be operationalized for business processes. Therefore, we have addressed RQ1.

Taking a broader look at relevant BPM techniques, we believe that a particular focus on greenhouse gas emissions as data constraints, especially in the sectors of manufacturing, transport and energy, gives new importance to BPM approaches focussed on assessing emissions of business processes on process and activity levels (see, e.g., [22, 33, 34]). Broader still, the taxonomy itself has been the subject of various criticisms. On the one hand, the taxonomy’s underlying notion of sustainable development [2] has been criticized as ambiguous and an *oxymoron* [20], and counterproductive to *actual* sustainability [32]. On the other hand, the taxonomy has also been described as *too* restrictive and as a bureaucratic burden that would be unable to benefit the overall economy [25]. Hence, the role played by the taxonomy in promoting sustainability is still subject to scholarly debate.

Nonetheless, this is one of the first papers to bring an understanding of the taxonomy to the business process management and enterprise computing communities. We have striven to provide conceptual clarity and impulses for future research on the taxonomy and its potential for sustainable business practices. Further, we believe that the pipeline we developed can potentially be applied to other regulatory frameworks as well.

*Threats to Validity.* There are several threats to the validity of our study. First, we have not validated the constraint characterization in its entirety, and have rather focussed on plausibility instead of completeness. However, our experimental validation showed that the characterization is generally plausible, and can serve as a *starting point* for further manual investigation. Second, for compliance monitoring with conformance checking, the prescriptive process model into which regulations are operationalized needs to be shown to be regulatory compliant as well (see [16]). This is a concern explicitly not addressed in our approach, as we have investigated in how far a prescriptive process model can be created at all. Moreover, technical limitations of LLMs, such as “confabulations”, “hallucinations” and inherent biases (see, e.g., [4, 26, 42]), apply to our study as well. However, our application is concerned with a very technical lens and is less of a generative application scenario. By following existing knowledge on prompt engineering in the BPM discipline, we sought to curb the impact of these limitations. Further, while our approach may produce results with some inaccuracies, the classification can still serve as a starting point for investigating individual constraints in depth. Finally, other techniques in the compliance checking space exist that may allow further kinds of constraints to be extracted from regulations against which

processes can be checked. However, as our investigation has been concerned only with conformance checking applications — which are limited to the four process dimensions of control-flow, time, resources, and data — we did not consider them further.

## 7 Conclusion and Future Work

To conclude, in this paper we have investigated if the EU taxonomy for sustainable activities can be operationalized for automatic compliance monitoring. For this, we have developed a pipeline that uses few-shot learning with an LLM, to identify and characterize the types of constraints applicable for conformance checking. We saw that many constraints of various industries can, in fact, be operationalized for this, which will allow companies to automatically monitor compliance with regard to the taxonomy. We have demonstrated that operationalizing the EU taxonomy into constraints for conformance checking may be partially automated; in this paper, we provide a starting point for such an automation. Besides this technical contribution, we have also introduced the taxonomy to the business process management and enterprise computing communities. The characterization pipeline may provide beneficial for assessing the capability of other complex regulatory frameworks as well.

Future work includes translating our classification approach into real-world application scenarios — further investigating how stakeholders can be supported in creating prescriptive models in line with the EU taxonomy, and how event log capturing and extraction benefits from our constraint classification would be a valuable contribution. We also aim to conduct more in-depth empirical evaluations of our mapping approach by comparing automatically generated results with results generated by taxonomy experts. Moreover, since we focus exclusively on business process compliance and have abstracted away from constraints regarding their broader context, we have explicitly excluded the analysis of minimum safeguards. However, approaches that enable compliance monitoring of constraints across organizations exist (such as [14, 23, 24]), and fruitful future work might lie in operationalizing the regulations and guidelines relevant to the taxonomy’s minimum safeguard criterion for these approaches. Finally, we believe that providing concrete guidance to end users in the form of a “handbook” or constraint patterns based on our preliminary findings reported herein would be a relevant addition.

**Acknowledgments.** This study was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grants no. 465904964, 496119880 and 531115272, by the German Federal Ministry of Education and Research (BMBF) under grant no. 16DII133 (Weizenbaum Institute), and by the Einstein Foundation Berlin under grant no. EPP-2019-524. Further, the authors would like to thank Man Tuen Chan for his support in validating the approach presented in this paper.

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.



## References

1. van der Aa, H., Di Ciccio, C., Leopold, H., Reijers, H.A.: Extracting Declarative Process Models from Natural Language. In: Giorgini, P., Weber, B. (eds.) *Advanced Information Systems Engineering*, vol. 11483, pp. 365–382. Springer International Publishing, Cham (2019). [https://doi.org/10.1007/978-3-030-21290-2\\_23](https://doi.org/10.1007/978-3-030-21290-2_23)
2. Alessi, L., Battiston, S., Melo, A.S., Roncoroni, A.: The EU sustainability taxonomy: A financial impact assessment (KJ-NA-29970-EN-N (online)) (2019). <https://doi.org/10.2760/347810>
3. Barrientos, M., Winter, K., Mangler, J., Rinderle-Ma, S.: Verification of Quantitative Temporal Compliance Requirements in Process Descriptions Over Event Logs. In: Indulska, M., Reinhartz-Berger, I., Cetina, C., Pastor, O. (eds.) *Advanced Information Systems Engineering*, vol. 13901, pp. 417–433. Springer Nature Switzerland, Cham (2023). [https://doi.org/10.1007/978-3-031-34560-9\\_25](https://doi.org/10.1007/978-3-031-34560-9_25)
4. Bender, E.M., Gebru, T., McMillan-Major, A., Shmitchell, S.: On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. pp. 610–623. ACM, Virtual Event Canada (Mar 2021). <https://doi.org/10.1145/3442188.3445922>
5. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
6. Brundtland, G.H.: Our Common Future—Call for Action. *Environmental Conservation* **14**(4), 291–294 (1987). <https://doi.org/10.1017/S0376892900016805>
7. Carmona, J., Van Dongen, B., Solti, A., Weidlich, M.: *Conformance Checking: Relating Processes and Models*. Springer International Publishing, Cham (2018). <https://doi.org/10.1007/978-3-319-99414-7>
8. Conea, A.M.: EU taxonomy: Qualifying as green. *LESIJ-Lex ET Scientia International Journal* **29**(2), 26–39 (2022)
9. Dragoni, M., Villata, S., Rizzi, W., Governatori, G.: Combining Natural Language Processing Approaches for Rule Extraction from Legal Documents. In: Pagallo, U., Palmirani, M., Casanovas, P., Sartor, G., Villata, S. (eds.) *AI Approaches to the Complexity of Legal Systems*, vol. 10791, pp. 287–300. Springer International Publishing, Cham (2018). [https://doi.org/10.1007/978-3-030-00178-0\\_19](https://doi.org/10.1007/978-3-030-00178-0_19)
10. European Commission: EU taxonomy for sustainable activities - European Commission, [https://finance.ec.europa.eu/sustainable-finance/tools-and-standards/eu-taxonomy-sustainable-activities\\_en](https://finance.ec.europa.eu/sustainable-finance/tools-and-standards/eu-taxonomy-sustainable-activities_en)
11. European Commission: NACE Rev. 2. Office for Official Publications of the European Communities, Luxembourg, revision 2, english edition edn. (2008)
12. European Commission: A User Guide to Navigate The EU Taxonomy for Sustainable Activities (2023), <https://ec.europa.eu/sustainable-finance-taxonomy/assets/documents/Taxonomy%20User%20Guide.pdf>
13. European Commission: Regulation (EU) 2020/852 of the European Parliament and of the Council of 18 June 2020 on the establishment of a framework to facilitate sustainable investment, and amending Regulation (EU) 2019/2088 (2020), <https://eur-lex.europa.eu/eli/reg/2020/852/oj>
14. Fdhila, W., Rinderle-Ma, S., Knuplesch, D., Reichert, M.: Decomposition-based verification of global compliance in process choreographies. In: *2020 IEEE 24th International Enterprise Distributed Object Computing Conference (EDOC)*. pp. 77–86 (2020). <https://doi.org/10.1109/EDOC49727.2020.00019>

15. Governatori, G., Hashmi, M., Lam, H.P., Villata, S., Palmirani, M.: Semantic Business Process Regulatory Compliance Checking Using LegalRuleML. In: Blomqvist, E., Ciancarini, P., Poggi, F., Vitali, F. (eds.) *Knowledge Engineering and Knowledge Management*, vol. 10024, pp. 746–761. Springer International Publishing, Cham (2016). [https://doi.org/10.1007/978-3-319-49004-5\\_48](https://doi.org/10.1007/978-3-319-49004-5_48)
16. Groefsema, H., van Beest, N.R.T.P., Governatori, G.: On the Use of the Conformance and Compliance Keywords During Verification of Business Processes. In: Di Ciccio, C., Dijkman, R., Del Río Ortega, A., Rinderle-Ma, S. (eds.) *Business Process Management Forum*, vol. 458, pp. 21–37. Springer International Publishing, Cham (2022). [https://doi.org/10.1007/978-3-031-16171-1\\_2](https://doi.org/10.1007/978-3-031-16171-1_2)
17. Grohs, M., Abb, L., Elsayed, N., Rehse, J.R.: Large Language Models Can Accomplish Business Process Management Tasks. In: De Weerd, J., Pufahl, L. (eds.) *Business Process Management Workshops*, vol. 492, pp. 453–465. Springer Nature Switzerland, Cham (2024). [https://doi.org/10.1007/978-3-031-50974-2\\_34](https://doi.org/10.1007/978-3-031-50974-2_34)
18. Hashmi, M., Governatori, G., Lam, H.P., Wynn, M.T.: Are we done with business process compliance: State of the art and challenges ahead. *Knowl Inf Syst* **57**(1), 79–133 (2018). <https://doi.org/10.1007/s10115-017-1142-1>
19. Hashmi, M., Governatori, G., Wynn, M.T.: Normative requirements for regulatory compliance: An abstract formal framework. *Inf Syst Front* **18**(3), 429–455 (Jun 2016). <https://doi.org/10.1007/s10796-015-9558-1>
20. Johnston, P., Everard, M., Santillo, D., Robèrt, K.H.: Reclaiming the definition of sustainability. *Environ Sci Pollut Res Int* **14**(1), 60–66 (2007). <https://doi.org/10.1065/espr2007.01.375>
21. Klessascheck, F., Knoche, T., Pufahl, L.: Reviewing conformance checking uses for run-time regulatory compliance. In: van der Aa, H., Bork, D., Schmidt, R., Sturm, A. (eds.) *Enterprise, Business-Process and Information Systems Modeling*. pp. 100–113. Springer Nature Switzerland, Cham (2024)
22. Klessascheck, F., Weber, I., Pufahl, L.: SOPA: A Framework for Sustainability-Oriented Process Analysis and Re-design in Business Process Management (2024), <http://arxiv.org/abs/2405.01176>
23. Knuplesch, D., Fdhila, W., Reichert, M., Rinderle-Ma, S.: Detecting the effects of changes on the compliance of cross-organizational business processes. In: Johannesson, P., Lee, M.L., Liddle, S.W., Opdahl, A.L., Pastor López, Ó. (eds.) *Conceptual Modeling*. pp. 94–107. Springer International Publishing, Cham (2015)
24. Knuplesch, D., Reichert, M., Kumar, A.: A framework for visually monitoring business process compliance. *Information Systems* **64**, 381–409 (2017). <https://doi.org/https://doi.org/10.1016/j.is.2016.10.006>, <https://www.sciencedirect.com/science/article/pii/S0306437915301770>
25. Kooths, S.: EU Taxonomy: Mission Impossible. *The Economist's Voice* **19**(2), 243–249 (2023). <https://doi.org/10.1515/ev-2022-0028>
26. Ladkin, P.B.: Involving LLMs in legal processes is risky: An invited paper. *Digital Evidence and Electronic Signature Law Review* pp. 40–46 (Jun 2023). <https://doi.org/10.14296/deeslr.v20i.5610>
27. López, H.A., Debois, S., Slaats, T., Hildebrandt, T.T.: Business Process Compliance Using Reference Models of Law. In: Wehrheim, H., Cabot, J. (eds.) *Fundamental Approaches to Software Engineering*, vol. 12076, pp. 378–399. Springer International Publishing, Cham (2020). [https://doi.org/10.1007/978-3-030-45234-6\\_19](https://doi.org/10.1007/978-3-030-45234-6_19)
28. Malecki, C.: The EU taxonomy regulation: Giving a good name to sustainable investment. *Env Liability* **26**(4), 149–156 (2021), <https://ssrn.com/abstract=4235527>

29. McClellan, A., Gehrke, N., Picard, N., Schellhas, C.: EU Taxonomy Reporting 2023. Tech. rep. (2023), [https://www.pwc.de/de/content/20e6bff9-\*\*ea5a-4d03-b375-a6a58f7b8b46/pwc-eu-taxonomy-reporting-2023.pdf\*\*](https://www.pwc.de/de/content/20e6bff9-<b>ea5a-4d03-b375-a6a58f7b8b46/pwc-eu-taxonomy-reporting-2023.pdf</b)
30. Mustroph, H., Barrientos, M., Winter, K., Rinderle-Ma, S.: Verifying Resource Compliance Requirements from Natural Language Text over Event Logs. In: Di Francescomarino, C., Burattin, A., Janiesch, C., Sadiq, S. (eds.) *Business Process Management*, vol. 14159, pp. 249–265. Springer Nature Switzerland, Cham (2023). [https://doi.org/10.1007/978-3-031-41620-0\\_15](https://doi.org/10.1007/978-3-031-41620-0_15)
31. Mustroph, H., Winter, K., Rinderle-Ma, S.: Social Network Mining from Natural Language Text and Event Logs for Compliance Deviation Detection. In: Sellami, M., Vidal, M.E., Van Dongen, B., Gaaloul, W., Panetto, H. (eds.) *Cooperative Information Systems*, vol. 14353, pp. 347–365. Springer Nature Switzerland, Cham (2024). [https://doi.org/10.1007/978-3-031-46846-9\\_19](https://doi.org/10.1007/978-3-031-46846-9_19)
32. Purvis, B., Mao, Y., Robinson, D.: Three pillars of sustainability: In search of conceptual origins. *Sustain Sci* **14**(3), 681–695 (2019). <https://doi.org/10.1007/s11625-018-0627-5>
33. Recker, J., Rosemann, M., Gohar, E.R.: Measuring the Carbon Footprint of Business Processes. In: *Business Process Management Workshops*, vol. 66, pp. 511–520. Springer, Berlin, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-20511-8\\_47](https://doi.org/10.1007/978-3-642-20511-8_47)
34. Recker, J., Rosemann, M., Hjalmarsson, A., Lind, M.: Modeling and Analyzing the Carbon Footprint of Business Processes. In: vom Brocke, J., Seidel, S., Recker, J. (eds.) *Green Business Process Management*, pp. 93–109. Springer, Berlin, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-27488-6\\_6](https://doi.org/10.1007/978-3-642-27488-6_6)
35. Rehse, J., Pufahl, L., Grohs, M., Klein, L.: Process mining meets visual analytics: The case of conformance checking. In: Bui, T. (ed.) *Proceedings of the 56th Annual Hawaii International Conference on System Sciences, HICSS 2023*. pp. 5452–5461. *Proceedings of the Annual Hawaii International Conference on System Sciences, IEEE Computer Society* (2023)
36. Remy, S., Pufahl, L., Sachs, J.P., Böttinger, E., Weske, M.: Event log generation in a health system: a case study. In: *Business Process Management: 18th International Conference, BPM 2020, Seville, Spain, September 13–18, 2020, Proceedings 18*. pp. 505–522. Springer (2020)
37. Russell, N., van der Aalst, W., Ter Hofstede, A.: *Workflow Patterns: The Definitive Guide*. MIT Press, Cambridge, MA (2015)
38. Sai, C., Damaratskaya, A., Winter, K., Rinderle-Ma, S.: Identification and Visualization of Legal Definitions and Legal Term Relations. In: Sales, T.P., Araújo, J., Borbinha, J., Guizzardi, G. (eds.) *Advances in Conceptual Modeling*, vol. 14319, pp. 151–161. Springer Nature Switzerland, Cham (2023). [https://doi.org/10.1007/978-3-031-47112-4\\_14](https://doi.org/10.1007/978-3-031-47112-4_14)
39. Sai, C., Sadiq, S., Han, L., Demartini, G., Rinderle-Ma, S.: Which Legal Requirements are Relevant to a Business Process? Comparing AI-Driven Methods as Expert Aid. In: Araújo, J., De La Vara, J.L., Santos, M.Y., Assar, S. (eds.) *Research Challenges in Information Science*, vol. 513, pp. 166–182. Springer Nature Switzerland, Cham (2024). [https://doi.org/10.1007/978-3-031-59465-6\\_11](https://doi.org/10.1007/978-3-031-59465-6_11)
40. Sai, C., Winter, K., Fernanda, E., Rinderle-Ma, S.: Detecting Deviations Between External and Internal Regulatory Requirements for Improved Process Compliance Assessment. In: Indulska, M., Reinhartz-Berger, I., Cetina, C., Pastor, O. (eds.) *Advanced Information Systems Engineering*, vol. 13901, pp. 401–416. Springer Nature Switzerland, Cham (2023). [https://doi.org/10.1007/978-3-031-34560-9\\_24](https://doi.org/10.1007/978-3-031-34560-9_24)
41. Schütze, F., Stede, J.: The EU sustainable finance taxonomy and its contribution to climate neutrality. *Journal of Sustainable Finance & Investment* **14**(1), 128–160 (2024). <https://doi.org/10.1080/20430795.2021.2006129>

42. Sui, P., Duede, E., Wu, S., So, R.J.: Confabulation: The Surprising Value of Large Language Model Hallucinations (2024). <https://doi.org/10.48550/ARXIV.2406.04175>
43. UN Environment (ed.): Global Environment Outlook – GEO-6: Summary for Policymakers. Cambridge University Press, Cambridge (2019). <https://doi.org/10.1017/9781108639217>
44. Weske, M.: Business Process Management - Concepts, Languages, Architectures. Springer, Berlin, Heidelberg, 2nd edition edn. (2012). <https://doi.org/10.1007/978-3-642-28616-2>
45. Winter, K., van der Aa, H., Rinderle-Ma, S., Weidlich, M.: Assessing the Compliance of Business Process Models with Regulatory Documents. In: Dobbie, G., Frank, U., Kappel, G., Liddle, S.W., Mayr, H.C. (eds.) Conceptual Modeling, vol. 12400, pp. 189–203. Springer International Publishing, Cham (2020). [https://doi.org/10.1007/978-3-030-62522-1\\_14](https://doi.org/10.1007/978-3-030-62522-1_14)
46. Winter, K., Rinderle-Ma, S.: Detecting Constraints and Their Relations from Regulatory Documents Using NLP Techniques. In: Panetto, H., Debruyne, C., Proper, H.A., Ardagna, C.A., Roman, D., Meersman, R. (eds.) On the Move to Meaningful Internet Systems. OTM 2018 Conferences, vol. 11229, pp. 261–278. Springer International Publishing, Cham (2018). [https://doi.org/10.1007/978-3-030-02610-3\\_15](https://doi.org/10.1007/978-3-030-02610-3_15)
47. Winter, K., Rinderle-Ma, S.: Deriving and Combining Mixed Graphs from Regulatory Documents Based on Constraint Relations. In: Giorgini, P., Weber, B. (eds.) Advanced Information Systems Engineering, vol. 11483, pp. 430–445. Springer International Publishing, Cham (2019). [https://doi.org/10.1007/978-3-030-21290-2\\_27](https://doi.org/10.1007/978-3-030-21290-2_27)
48. Winter, K., Rinderle-Ma, S., Grossmann, W., Feinerer, I., Ma, Z.: Characterizing Regulatory Documents and Guidelines Based on Text Mining. In: Panetto, H., Debruyne, C., Gaaloul, W., Papazoglou, M., Paschke, A., Ardagna, C.A., Meersman, R. (eds.) On the Move to Meaningful Internet Systems. OTM 2017 Conferences, vol. 10573, pp. 3–20. Springer International Publishing, Cham (2017). [https://doi.org/10.1007/978-3-319-69462-7\\_1](https://doi.org/10.1007/978-3-319-69462-7_1)