A Regulatory Agent Model using Case-Based and Language-Based Reasoning in Sequential Games

Andreas Korger[1], Joachim Baumeister[1]

[1]University of Würzburg, Am Hubland, D-97074 Würzburg

Abstract

Regulation often involves interdisciplinary considerations in disaggregated environments where knowledge is only partially accessible. Case abstraction can be exploited to address communication and collaboration tasks in a multilateral regulatory setting. In this work, we investigate how semantic clusters defined by the abstracted experience knowledge coded in regulatory datasets can be transformed into an agent-based model representing regulatory scenarios as sequential games. In this regard, we contribute to the field of regulatory case-based reasoning and suggest a sequential regulatory case-based agent model. We show how the combination of case-based and language-based reasoning fills gaps in context-aware regulatory agent modeling. The approach is put into practice in the domain of nuclear safety. Monte Carlo experiments suggest how the emergent information collected from configurable games can be used to support humans working in a regulatory environment, such as predicting safety budgets for disaggregated safety teams.

Keywords

Knowledge Management, Agent-Based Modeling, Multi-Agent Model, Semantics, Case-Based Reasoning, Large Language Models, Regulation, Sequential Game, Monte Carlo Experiment

1. Introduction

In a multilateral (involving several regulatory actors) and disaggregated regulatory scenario, there exist parties that aim for regulation and parties who are affected by this ambition for regulation. For instance, fire safety institutions, medical safety institutions, and governmental institutions stand for such interest groups with coherent and concurrent ambitions. The regulatory action of the one (regulatory measure) can become the impact (incident) of the other [1]. Disaggregated means that a common goal exists in a group, but centralized data management is unavailable. Practically, no common vocabulary is maintained, and information may be (partially) private, so data science methods cannot be easily applied to all the data involved.

The regulatory actions of the other actors have to be considered by each actor in their own regulatory decisions. Involving several parties can lead to a non-linear and complex process as every party aspires to reach different (unknown) regulatory goals for diverse (unknown) contextual configurations of the environment. In this work, we present a way to create sequential [2] regulatory Monte Carlo [3] experiments that simulate such regulatory scenarios as sequential games to conclude, e.g., for resource planning.

LWDA'24: Lernen Wissen Daten Analysen, 2024, Würzburg
a.korger@informatik.uni-wuerzburg.de (A. Korger); joba@uni-wuerzburg.de (J. Baumeister)

CEUR Workshop Proceedings (CEUR-WS.org)

The abstraction into semantic regulatory concepts mitigates these shortcomings of dissaggregation [4]. Efforts of regulatory semantification often do not incorporate all contextual aspects of these effects due to the costs involved in creating and maintaining the connected regulatory framework. We observed that case-based reasoning (CBR) methods open the door to support the semantification process of regulatory data in multilateral and disaggregated environments. CBR assists in climbing the knowledge pyramid from semantics over pragmatics to reasoning [5]. This enrichment allows actions to be added to knowledge to implement an active sequential game.

This aim for a regulatory case-based reasoning (RCBR) approach stands between basic CBR and process-oriented CBR [5]. The application of processes requires strict compliance of actors in the process and a certain degree of predictability. This is not given in many regulatory scenarios, and a fully process-oriented approach is not maintainable. In a regulatory case-based reasoning setting, it is not always clear which distinct goals are to be reached. Nevertheless, the relations and pseudo-causal order of regulatory concepts lead to pseudo-process characteristics of the system.

In this work, we exploit a well-maintained and aggregated regulatory corpus from the domain of nuclear safety to simulate a disaggregated environment. We aim to gain semantic and pragmatic information by observing the effects emerging in these experiments. For instance, budget allocation is challenging to address with normal semantification of texts, e.g., by extracting entities and relations [4]. Imagine a setting where a fixed budget (financial, workforce, time) has to be distributed over ten different safety teams with individual needs unknown by the other teams. We show that a solely data-driven agent-based model (ABM) can simulate individual demands and thus give pragmatic guidance, e.g., for demand planning.

1.1. Solution Approach

We set up on a regulatory case structure that can generalize to diverse regulatory domains [4]. Data clustering can be exploited to create (one possible) agent clustering. Agent policies can be derived from each data cluster by creating an individual similarity assessment as the agent's utility function. We employ large language models (LLMs) to mitigate contextual problems arising in disaggregated settings. To code new knowledge, we search for patterns in the simulated sequence of applied cases and suggest excerpts for similar environments.

1.2. Contribution and Research Questions

With this work, we contribute to integrating CBR and large language models for the semantic augmentation of existing regulatory data. We address the field of regulatory CBR and contribute to its characterization. We investigate how data clustering can be used to create a distinct agent-based simulation. We suggest how the emergent information gained from the simulation can be added to the existing experience knowledge.

2. Related Work

Substantial work has been done on case-based agent modeling and distributed case-based reasoning. This helps to integrate with state-of-the-art statistical language models. Asking large language models for safety advice implies risks, such as relying on hallucinated facts, as stated by Oviedo-Trespalacios et al. [6]. To automatically learn safety knowledge from past accidents is an ideal scenario, which is still facing several shortcomings as addressed by Niu et al. [7].

The concept of a conversational multi-agent environment is yet taken up for practical application as presented by Wu et al. [8]. A team of communicative agents is created based on LLMs to achieve the task of collaborative software engineering by Qian et al. [9]. The critical evaluation showed the pitfalls of the LLMs. Explainability of agent decisions is a crucial criterion of success in the interaction of humans with agents.

Silva et al. [10] present an evaluation of factors that influence the teaming of humans with agents. This work was tremendously helpful in the first design considerations of this work to avoid common pitfalls from the beginning. Our goal is to set up an agent with knowledge based on the data found in a corpus. A similar task is to transfer knowledge to an agent to reduce training costs, which was addressed by Castagna and Dusparic [11]. The trust of humans in AI-based moderation was investigated by Molina and Sundar [12]. A review of data-driven accident prevention systems was presented by Assadzadeh et al. [13].

Tran and Schönwalder [14] present a work to integrate different knowledge sources for fault management in communication networks into one case-based system and, therefore, introduce a new step of "retention" into the classical CBR cycle. Pla et al. [15] present a case-based tool to integrate different experts into a decision process. They developed a schema for the cooperative communication of involved actors. How to personalize intelligent tutoring systems using case-based reasoning and multi-agent systems was investigated by Gonzalez et al. [16]. Schneider et al. [17] use generative artificial intelligence for athletic training education by developing case-based scenarios with ChatGPT.

We analyzed current publications on collaborative crisis management to get insight into aspects important to a successful approach. In their work to create an agent-based platform for crisis management, Bhattacharya [18] faces a very similar scenario but presents a domain-specific approach in opposition to our generalization ambitions.

A survey of scientific publications of the last decade towards agent-based models of human response to natural hazards was collected by Mis et al. [19]. Their work helped assess important features of successful agent modeling in the domain. Analogously, critical factors for successful crisis management were aggregated by Bynander and Nohrstedt [20]. We assume these quality factors are also important to set up an agent-based model in our context. Parker et al. [21] present core assumptions. Yet, not up to current state-of-the-art still contributing worthy ideas to the subject is the work of Balaraman et al. [22].

3. Methods - Regulatory Multi-Agent-Model

Regulatory texts can represent a collection of applied cases and can be re-clustered into such [4]. We assume that the given semantic clustering of the corpus into topic-related units and cases can be considered analogously as a real-world clustering into coherently acting interest groups [4]. These and further assumptions are the foundation of the simplified regulatory agent model.

Assumptions: The experiment takes place in the "abstracted regulatory world". The real world is abstracted into cases consisting of incidents (problem), measures (solution), their context, and relations between them. We assume a sequential game [2] and that spatial information is neglected. There is no quality in the time steps, like, e.g., a 24-hour cycle.

Each agent is defined by a semantic cluster representing a confined team and regulatory topic. Measures and incidents last for the time step in which they are applied. Specific durations can override this. For instance, the incident *hurricane* will last longer than an ordinary storm. Each initial agent case base is generated from one semantic cluster.

We assume perfect knowledge of the agents regarding the regulatory actions of other agents; the others' actions are observable, but what stays hidden from the agents is the textual case definition of other agents. We assume incomplete knowledge; the other agents' individual similarity (utility) function is unknown to each agent.

The agent's actions are defined by the available measures. Each agent has a limited budget for executing measures at a specific time step. The agent's operations are to ADD, DELETE, and SUBSTITUTE measures. The scope of incidents and measures is the entire world. An agent is activated if a similar case to the current environment state is found in his case base. To save computational resources, agents do not directly communicate; they only observe the environment and then act. The agent learns by adding new cases to be incorporated into the next semantic cluster version (retain step).

Definition 1 (Regulatory agent). *Let $\mathcal{A}_{CB} = \{\mathcal{CB}_1, ..., \mathcal{CB}_n\}$ be n disjunct subsets of a case base $\mathcal{CB}$ with the previously described regulatory case structure. Let $A_i = \langle \mathcal{CB}_i, K_{Ai}, P_{Ai}, H_{Ai} \rangle, i \in [1, n]$ be an agent number i derived from the case cluster $\mathcal{CB}_i$ with knowledge K_{Ai}, the actions P_{Ai} created out of $\mathcal{CB}_i$ and K_{Ai}, and a set of configuration hyper parameters H_{Ai}.*

The environment has two purposes. It encapsulates what happens in the world apart from agents' behavior. Further, it tracks what agents contributed to the environmental history. All agents can observe it. This facilitates a scenario where agents observe each other and communicate directly. The environment is analogously defined as an agent derived from semantic clusters or pre-processed external data sources in the following way.

Definition 2 (Environment). *Let $\mathcal{E}_{CB} = \{\mathcal{CB}_{n+1}, ..., \mathcal{CB}_m\} : \forall \mathcal{CB}_j \in \mathcal{E}_{CB}, \mathcal{CB}_j \notin \mathcal{A}_{CB}$ be disjunct subsets of a case base $\mathcal{CB}$ with the previously described case structure. Let $\mathcal{E}_j = \langle \mathcal{CB}_j, H_{Ej} \rangle, j \in [n+1, m]$ be an environment number j derived from $\mathcal{CB}_j$, a set of configuration hyper parameters H_j. Let $e_t \in \mathcal{E}_j$ be a case presented by the environment at discrete time t.*

We define the following similarity measures for these case, agent, and environment definitions. The similarity assessment done by each agent when a new state $e_t \in E_j$ of the environment is presented is called primary similarity. The following notations are used in the definitions and simulation algorithm:

1. A_{SELF} = currently active agent
2. A_{OTHERS} = currently non active agents
3. $M(A_{it})$ = all measures an agent i placed at time step t
4. $I(A_{it})$ = all incidents an agent i recognized at time step t
5. $M(E_t)$ = all actions observable by all agents at time t

Definition 3 (Prior Similarity Assessment). *Let $sim_{PRIOR}(A_{it}, e_t)$ be the similarity of the Agent i to the environment state e_t at time t plus the context CTX_{A_it} transformed into a vector embedding by V(x). The agent is only activated if $max(sim_{PRIOR})$ of all agent's cases c_{a_i} exceeds a threshold δ. With $\delta \geq 0 \in H_{Ai}$ and dot(x) the dot-product of embedding vectors, similarities are defined as:*

$$sim_{PRIOR}(A_{it}, e_t) = max(\forall\, c_{a_i} dot(V(c_{a_i} + CTX_{Ai}), V(e_t))) \tag{1}$$

$$Activation(A_i) = true\, for\, sim_{PRIOR}(A_i, e_t) \geq \delta \tag{2}$$

The cluster-context-specific similarity is derived by embedding the case description together with a cluster-specific context CTX_{Ai}. This context can be, for instance, the document title, a document summary, and semantic knowledge about the cluster. Consequently, every agent not only has its own case base but also its own (slightly different) similarity measure to calculate similarities. Further, the context can encompass simulation-specific information like past actions, past environmental states, and predictions for the future that possibly influence the current similarity assessment of the agent.

For the second step of similarity assessment, we borrow the concept of surprise from information theory [23]. If an agent is surprised by the actions of other agents, he is more likely to consider changes in his own actions. The logarithmic calculation of the surprise value leverages the similarity gap between the prior similarity and the posterior similarity assessment.

Definition 4 (Posterior Similarity Assessment). *Let $sim_{POSTER}(A_{it}, e_{t-1})$ be the similarity incorporating the actions of all other Agents A_{OTHERS} done at t-1. Given $sim_{POSTER} > sim_{PRIOR}$ the agent will consider the secondary retrieval if the surprise about the best case max($sim_{POSTER}(A_{it}, e_{t-1})$) exceeds a threshold $\sigma \geq 0 \in H_{Ai}$ and add the new measures retrieved for t-1 to the solution retrieved in t.*

$$sim_{POSTER}(A_{it}, e_{t-1}) = max(\forall\, c_{a_i} dot(V(c_{a_i} + CTX_{A_i+M}), V(e_{t-1}))) \tag{3}$$

$$Surprise(A_t, e_{t-1}) = -\ln(1 - sim_{POSTER}(A_t, e_{t-1})) \tag{4}$$

With the previous definitions and assumptions, a first instantiated algorithm version can be formalized by running the multi-agent model. It is initialized with the number of semantic clusters used to create agents and the clusters used to define the environment. The embedding process is not part of this algorithmic description as it is encapsulated in the retrieval process and further discussed in Section 4.

At the beginning of each new step, the agents compare the actions taken one step before to those of the other agents. A new retrieval incorporating measures of the other agents into the posterior similarity assessment is executed. If the surprise calculated from the similarity gap between prior and posterior similarity exceeds a threshold, then the secondary case solution is preferred over the primary case solution, and additional measures are applied.

We facilitate this algorithm to a duration of measures of one time step. In a world where measures last only one step, measures can only be added and need not be deleted or substituted. The other agents then react to the changed actions of the agent in the next step. This leads to a continuous indirect interaction of all the agents.

Algorithm 1 Regulatory Multi Agent Simulation with only ADDING Agents

Require: $\delta, \sigma \geq 0$ for each agent
  initialize population n $agents[A_1, .., A_n] \leftarrow A_{CB}$
  initialize $environmentE \leftarrow E_{CB}$ (one environment)
  initialize $Sim_{PRIOR} = 0$, $Sim_{POSTER} = 0, t = 0$
  initialize $e_{-1} = []$
  while $t < maxsteps$ **do**
    for $A_i \in [A_1, .., A_n]$ **do**
      $M(A_{OTHERS}(t-1)) \leftarrow e_{t-1}$ ▷ check what other agents have done
      bestCase = retrieve($A_i, e_{t-1}, CTX(A_{OTHERS}(t-1))$
      $M(A_{SELF})$ = M(bestCase)
      Sim_{POSTER} = Sim(bestCase)
      if surprise(Sim_{POSTER} - Sim_{PRIOR}) $> \delta$ **then**
        POSTERIOR ACTIVATION = true
        $e_t \leftarrow$ ADD M(A_{SELF}) ▷ lately add posterior measures
      end if
      $e_t \leftarrow E$ ▷ sample new environment state
      bestCase = retrieve($A_i, e_t, CTX(t)$)
      $M(A_{SELF})$ = M(bestCase)
      Sim_{PRIOR} = Sim(bestCase)
      if (Sim_{PRIOR}) $> \sigma$ **then**
        PRIOR ACTIVATION = true
        $e_t \leftarrow$ ADD M(A_{SELF}) ▷ add prior measures
      end if
    end for
    $t = t + 1$
  end while

4. Results

We evaluated the presented agent model on a nuclear safety dataset created from a regulatory corpus. This corpus was previously annotated and transformed into an according dataset summing up to about 200,000 cases.

The corpus consists of 143 publicly available documents containing, in total, about 14,000 pages of English text published by the IAEA (International Atomic Energy Association), which is a sub-organization of the United Nations [24]. The IAEA aims to regulate the domain of nuclear safety internationally and gives advice and support to national authorities. All documents can be accessed via their website [24]. Further information like the used regulatory knowledge graph can be accessed via our website and github [25, 26].

Additionally, we used other datasets containing messages of people communicating during a disaster like an earthquake, hurricane, and flooding [27]. We pre-processed these datasets in the same manner as the nuclear safety documents, annotated regulatory concepts, and transferred them into a queryable case base. These datasets served as external environments.

The size of the case base clusters created from the nuclear safety documents ranges from approximately some hundred to some thousand cases per document cluster. Using only one nuclear safety document as an environment allows a limited maximum step size without repeating cases. The disaster datasets contain more than 10,000 ordered entries each, which allows for a larger maximum simulation step size and chronological sampling.

4.1. Semantic Pre-Processing

From the textual corpus, a terminology was extracted and manually classified into incidents and measures with several hierarchical layers. It contains concepts like `fire`, `manualFireFighting`, and `fireProtection`.

Adding to the contextual influence in the agent's case base created by the cluster-specific LLM embeddings, each agent has a specific knowledge K_i defined. This knowledge is derived from the regulatory concepts covered by the respective semantic cluster and contextualized as an LLM embedding.

Example 1 shows a detailed representation of one case and the highlighted features contained in the textual description of the case.

(1) **Case Example:** Text = "For reactors equipped with vessel closure plugs to retain the fuel in position, special design features should be provided to ensure that the probability of *ejection of the closure plug* is low. In the absence of such special features, the consequences of the *failure or the ejection of a single closure plug* should be evaluated as for a *missile*."
Case ID = S1.1001,
Context="Rules and Regulations for the safe transport of radioactive materials.",
Incidents = `ejection`, `failure`, `missile`,
Measures= `evaluate`, `ensure`, Cluster = S1

4.2. Observable Experiment Variables

As the agents are created from a given human-made regulatory framework, the goal is not to train the agents to improve in the sense of a reinforcement agent. One experiment goal is to add the next layer of semantification to gather (time series-related) information that is not collectible by NLP methods. The other goal is to suggest cases to human users to improve further regulatory texts to be created. The following variables are therefore observed in the experiments to get insight and draw conclusions.

1. $RF_{PRIOR}(A_i) = \sum_{t=1}^{maxsteps} Activation_{PRIOR}(A_i)/maxsteps$
2. $RF_{POSTERIOR}(A_i) = \sum_{t=1}^{maxsteps} Activation_{POSTERIOR}(A_i)/maxsteps$

The variables RF_{PRIOR} and $RF_{POSTERIOR}$ calculate the (relative) activation frequency and how often the agent needs to apply a regulatory action, respectively. We interpret this as related to an essential intervention by the safety team represented by the agent. Thus, the activation rates are related to potential costs accumulating over the simulation phase.

3. $F_{CONCEPT}(t) = \sum_{i=1}^{n} \mid M(A_i) \mid + \mid I(A_i) \mid$
4. $RF(M_i) = \sum_{t=1}^{maxsteps} \mid M_{it}(A_j) \mid /maxsteps$
5. $RF(I_i) = \sum_{t=1}^{maxsteps} \mid I_{it}(A_j) \mid /maxsteps$

The variable $F_{CONCEPT}$ gives information about the frequency of the occurrence of regulatory concepts, which is not derivable directly from texts. We defined the agents with a specific budget for processing incidents and measures at a particular time step. If $F_{CONCEPT}$ exceeds this budget, a critical situation can be assumed for that agent. The variables $RF(M_i)$ and $RF(I_i)$ calculate the relative frequency of a specific measure or incident in the whole experiment and give insight into how likely such concepts could occur in similar regulatory settings. Depending on these characteristics, regulatory safety teams can focus and prepare for likely measures and incidents.

6. $Sem_{I/M}(M_i, A_j) = \sum_{t=1}^{maxsteps} Diff(M_i(A_{OTHERS}), M_i(A_j))/F_{Mi}$

As we introduced before, regulatory concepts can be ambiguous. The measure of one can be the incident of another. The variable $Sem_{I/M}(M_i, A_j)$ gives a semantic estimate of how to classify the regulatory measure M_i into the classes of `incident` and `measure` from the perspective of the agent A_j. If an agent A_j also applies a particular measure at time step t applied by any other agent (A_{OTHERS}) then this measure is classified as a measure also for A_j because he acts in the same way. In contrast, the semantics of an incident are assumed if the others' measures are not applied by the agent. For instance, if a fire alarm is initiated as a measure reacting to a fire by another agent, then all agents who do not apply this measure might experience a high impact of the measure and have to respond to the others' fire alarm as if it were an incident. From this ratio, we learn if a regulatory concept tends to affect other agents more or less. (In time, step $t + 1$, measures can even have an incidental character for A_j himself.)

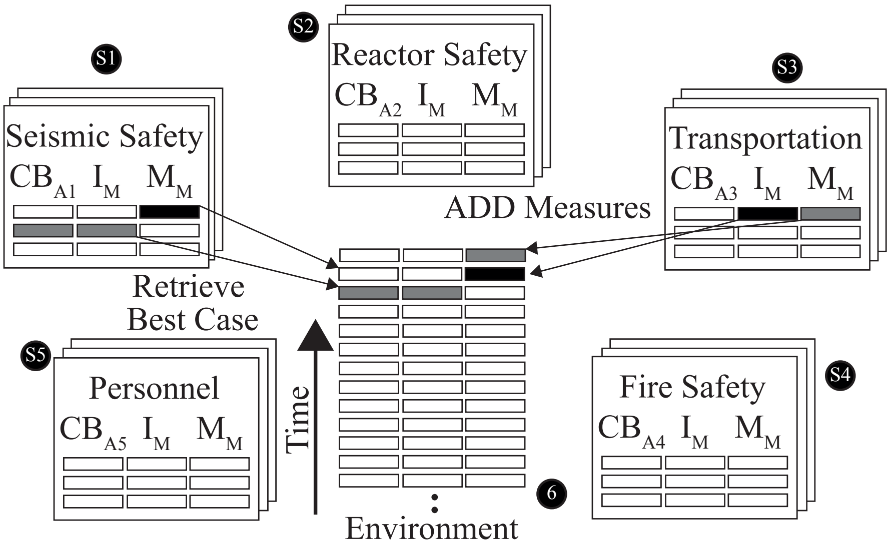

Figure 1: Representation of a regulatory scenario with five agents (S1-S5) and one environment (6) showing symbolic retrieval and placement of measures.

4.3. Experiments

In the following, we present a collection of experiments as depicted in Figure 1 that describe the calibration of hyperparameters, individual characteristics of each agent configuration, and comparable aspects of the simulation outcome. The available dataset was divided into an experimental data environment with training, validation, and test splits, leaving 10% of all sentences for each; validation and testing.

We initially tested the algorithms with a small selection of three documents that were known in terms of document content to investigate the behavior with human insight. The calculation of similarities and vector embeddings is computationally intense. Therefore, we made a selection of 26 documents. Then, we pre-calculated a vector embedding for each case, integrating the document title as context, and saved the embeddings for later use in the simulation. As statistical language models, we used a locally hosted smaller spacy [28] transformer model for development purposes and the openAI API to run the final simulations using an ADA2 embedding model [29].

The calibration of the ABM in new regulatory settings is a crucial initial question. Our experiments used the same threshold values for δ and σ for all agents. We then simulated some steps to observe the activation rates of the agents. We chose the thresholds so that most agents had primary activation rates between 50 and 80 percent and secondary activation rates around 10 percent. If the thresholds are chosen too low, the agents are activated in every step, and if they are too high, the agents are never activated. For a finer tuning of δ and σ or specific calibration for each agent, a deeper understanding of the regulatory setting and collaboration with domain experts seems reasonable.

Regarding scalability, the size of the case base has a major impact on the runtime, which leaves room for optimization by exploiting the structural knowledge about the cases for faster retrieval. Additionally, the complexity of similarity measures was limited by (our) computational

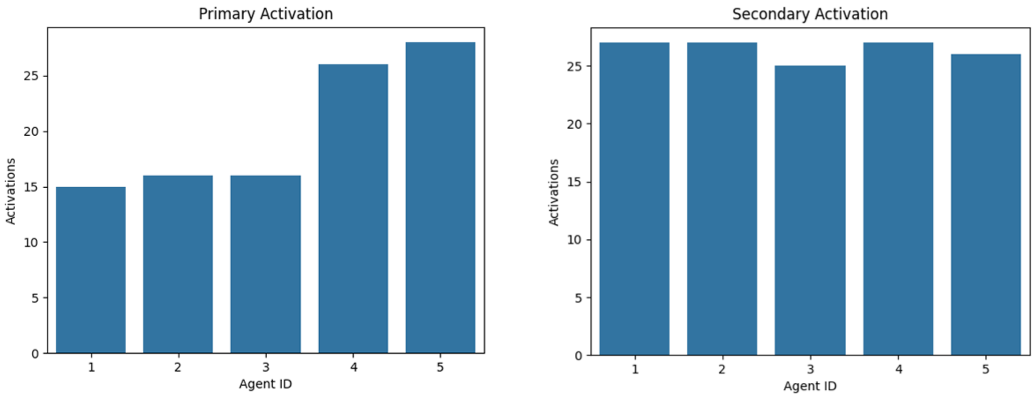

Figure 2: The variables F_{PRIOR} and $F_{POSTERIOR}$ for five agents in an experiment of 50 steps against an earthquake environment randomly sampled [27].

resources.

We ran the simulations on an environment with 840 short messages during a significant earthquake in Italy in 2016. An earthquake is a major threat to nuclear safety and a relevant event in that context. We sampled this environment randomly, presenting a random message at each time step.

4.3.1. Budget Planning

Figure 2 shows an experiment of 50 steps done with a setting of five agents created from: 1194(ID:1), 1115(ID:2), 1122(ID:3), 1158(ID:4), and 1448(ID:5) [24]. Documents 1158 and 1448

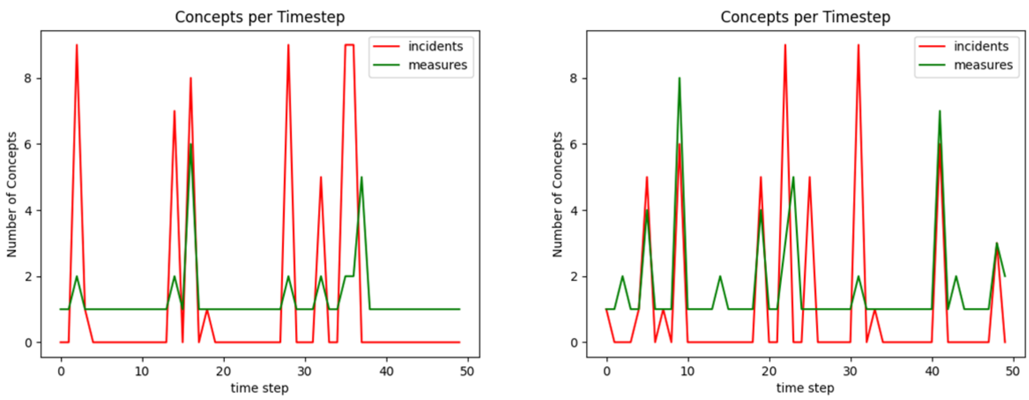

Figure 3: Variable $F_{CONCEPT}$ for two experiments against an earthquake environment with two different agent configurations of three and five agents.

are related to seismic aspects of nuclear safety. The diagrams show the cumulated primary activation and secondary activation. The primary activation rate indicates the relation of the agents with ID:4 and ID:5 to seismic hazards. An interesting result was the rate of secondary activation of the agents. The clear distinction of the agents in the prior similarities blurs in the posterior similarities because the reactions of all agents are respected, leading to more even activation rates. Regarding our previously explained calibration strategy, a secondary activation rate of about 50 percent seems too high, so the threshold σ should have probably been chosen smaller.

4.3.2. Regulatory Density

We assume the more regulatory actions occur within a particular time interval, the more focus should be on these phases. A dense regulatory activity signifies that more problems could arise and that the constellation leading to more activities is worth investigating more deeply. The development $F_{CONCEPT}$ of the regulatory concepts activated in all agents is displayed for two different environment settings in Figure 3. The diagrams show that there are time intervals with noticeably more action.

4.3.3. Incident-Measure Semantification

Table 1 shows the rate $RF(M_i)$ of the three selected regulatory concepts `design`, `requirements`, and `limit`. Table 2 shows the rate $Sem_{I/M}(M_i, A_j)$ how often a certain classification into incident and measure was derived for these regulatory concepts taken from the same earthquake experiment as described before.

Table 1
Relative frequency of selected measures for three agents.

Regulatory Concept	Agent ID:4	Agent ID:2	Agent ID:1
`design`	0.66	0.58	0.54
`requirements`	0.22	0.24	0.24
`limit`	0.66	0.66	0.70

Table 2
Semantification ratios for selected measures and three agents.

Regulatory Concept	Agent ID:4	Agent ID:2	Agent ID:1
`design`	I: 0.27 / M: 0.73	I: 0.79 / M: 0.21	I: 0.78 / M: 0.22
`requirements`	I: 0.91 / M: 0.09	I: 0.50 / M: 0.50	I: 0.92 / M: 0.08
`limit`	I: 0.24 / M: 0.76	I: 0.70 / M: 0.30	I: 0.83 / M: 0.17

5. Conclusion and Future Work

The semantification of entities and their relations is well investigated. Yet, the evidence is clear that regulatory semantification requires more efforts to address further relevant characteristics of regulatory concepts.

This work showed how a further layer of semantics and pragmatics can be put on top in regulatory scenarios. We, therefore, exploited case-based methods as well as agent-based modeling. Initially, we presented an abstract regulatory case structure to realize domain-independent regulation, which allows the integration of different regulatory actors in a sequential regulatory scenario. These components were merged with an adjusted similarity assessment into an agent-based model. The model was used to simulate a disaggregated regulatory scenario where agents (actors) have limited knowledge about the motivation and decisions of other regulatory actors but have to react to their actions.

The theoretical considerations were carried out through a series of experiments in nuclear safety, which has different regulatory subdomains.

A primary outcome of this work was that the presented architecture was very flexible. From our perspective, new datasets can be easily integrated without significant pre-processing efforts, mainly due to the consistent case-based design and, of course, the backbone functionality of LLMs. A stacked design allows the exchange of components of similarity assessment, such as different statistical language models. This was important in terms of scalability. Good results could also be achieved with a less complex transformer model, allowing the processing of more data simultaneously. This is especially important in addressing real-time use cases where data must be processed live, for instance, during a disaster.

The promising results of the experiments lead to further research directions. As we work on a humanly maintained corpus, it is desirable to investigate how the emergent information of regulatory simulations can be integrated (retained) into the existing regulatory framework in a semantic manner. The semantification of simulations into a regulatory knowledge graph needs more exploration. At the same time, the refinement of regulatory case-based reasoning is important as the case structure allows it to be added to a knowledge graph representation. As agents will play an increasing role shortly, efforts to improve case-based theories in the direction of autonomous self-regulation and distributed case-based reasoning will be needed. Further, the focus is on developing the agent-based model. First, it will have a more sophisticated configurational component to adjust it to more complex real-world scenarios to gain deeper insight into patterns of regulatory multi-agent-based phenomena. Second, it will be capable of a few-shot evaluation by humans to set configurational parameters within a given human work budget and facilitate model calibration. Finally, the growing capabilities of generative models will allow for the addition of spatial components and visualization of the simulations. In this context, we see diverse use cases for the training and education of humans working in safety-related environments.

References

[1] A. Korger, J. Baumeister, Case-based generation of regulatory documents and their semantic relatedness, in: e. a. Kohei Arei (Ed.), Future of Information and Communication Conference San Francisco, volume 1130 of *Advances in Information and Communication*, Springer, Heidelberg, 2020, pp. 91–110.

[2] C. E. Shannon, Xxii. programming a computer for playing chess, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 41 (1950) 256–275. URL: https://doi.org/10.1080/14786445008521796. doi:10.1080/14786445008521796. arXiv:https://doi.org/10.1080/14786445008521796.

[3] N. Metropolis, S. Ulam, The monte carlo method, J. Am. Stat. Assoc. 44 (1949) 335.

[4] A. Korger, J. Baumeister, Case-based sample generation using multi-armed bandits, in: S. Massie, S. Chakraborti (Eds.), Case-Based Reasoning Research and Development: 31st International Conference, ICCBR 2023, Aberdeen, UK, Springer-Verlag, Berlin, Heidelberg, 2023, pp. 118–133.

[5] R. Bergmann, Experience Management, Springer, Berlin, Heidelberg, 2002.

[6] O. Oviedo-Trespalacios, A. E. Peden, T. Cole-Hunter, A. Costantini, M. Haghani, J. Rod, S. Kelly, H. Torkamaan, A. Tariq, J. David Albert Newton, T. Gallagher, S. Steinert, A. J. Filtness, G. Reniers, The risks of using chatgpt to obtain common safety-related information and advice, Safety Science 167 (2023) 106244. doi:10.1016/j.ssci.2023.106244.

[7] Y. Niu, Y. Fan, X. Ju, Critical review on data-driven approaches for learning from accidents: Comparative analysis and future research, Safety Science 171 (2024) 106381. doi:10.1016/j.ssci.2023.106381.

[8] Q. Wu, G. Bansal, J. Zhang, Y. Wu, S. Zhang, E. Zhu, B. Li, L. Jiang, X. Zhang, C. Wang, Autogen: Enabling next-gen llm applications via multi-agent conversation framework, ArXiv abs/2308.08155 (2023). URL: https://api.semanticscholar.org/CorpusID:260925901.

[9] C. Qian, X. Cong, W. Liu, C. Yang, W. Chen, Y. Su, Y. Dang, J. Li, J. Xu, D. Li, Z. Liu, M. Sun, Communicative agents for software development, arXiv (2023). arXiv:2307.07924.

[10] J. C. da Silva, J. A. Matelli, E. Bazzo, Development of a knowledge-based system for cogeneration plant design: verification, validation and lessons learned, Knowledge-Based Systems 67 (2014) 230–243.

[11] A. Castagna, I. Dusparic, Expert-Free Online Transfer Learning in Multi-Agent Reinforcement Learning, arXiv, 2023. doi:10.3233/FAIA230291.

[12] M. D. Molina, S. S. Sundar, When AI moderates online content: effects of human collaboration and interactive transparency on user trust, Journal of Computer-Mediated Communication 27 (2022) zmac010. doi:10.1093/jcmc/zmac010.

[13] A. Assadzadeh, A. P. M. Arashpour, A. Rashidi, A. Bab-Hadiashar, S. Fayezi, A review of data-driven accident prevention systems: Integrating real-time safety management in the civil infrastructure context, in: International Symposium on Automation and Robotics in Construction, Banff Alberta, Canada, 2019. doi:10.22260/ISARC2019/0039.

[14] H. M. Tran, J. Schonwalder, Discaria—distributed case-based reasoning system for fault management, IEEE Trans. on Netw. and Serv. Manag. 12 (2015) 540–553. doi:10.1109/TNSM.2015.2496224.

[15] A. Pla, B. López, P. Gay, C. Pous, exit*cbr.v2: Distributed case-based reasoning tool for

medical prognosis, Decision Support Systems 54 (2013) 1499–1510. doi:10.1016/j.dss.2012.12.033.
[16] C. Gonzalez, J. Burguillo, M. Llamas Nistal, R. Fidalgo, Designing intelligent tutoring systems: A personalization strategy using case-based reasoning and multi-agent systems, ADCAIJ : Advances in Distributed Computing and Artificial Intelligence Journal 2 (2013). doi:10.14201/ADCAIJ2013244154.
[17] K. Schneider, D. Tomchuk, B. Snyder, T. Bisch, G. Koch, Incorporating Artificial Intelligence Into Athletic Training Education: Developing Case-Based Scenarios Using ChatGPT, Athletic Training Education Journal 19 (2024) 42–50. doi:10.4085/1062-6050-028.23.
[18] P. Bhattacharya, Data-driven scalable pipeline using national agent-based models for real-time pandemic response and decision support, The International Journal of High Performance Computing Applications 37 (2023) 4–27. doi:10.1177/10943420221127034.
[19] K. Mls, M. Kořínek, K. Štekerová, P. Tučník, V. Bureš, P. Čech, M. Husáková, P. Mikulecký, T. Nacházel, D. Ponce, M. Zanker, F. Babič, I. Triantafyllou, Agent-based models of human response to natural hazards: systematic review of tsunami evacuation, Natural Hazards 115 (2023) 1887–1908. doi:10.1007/s11069-022-05643-x.
[20] F. Bynander, D. Nohrstedt, Collaborative Crisis Management: Inter-Organizational Approaches to Extreme Events, Routledge, New York, 2019.
[21] J. B. H. H. O. R. Charles F. Parker, Daniel Nohrstedt, E. Baekkeskov, Collaborative crisis management: a plausibility probe of core assumptions, Policy and Society 39 (2020) 510–529. doi:10.1080/14494035.2020.1767337.
[22] V. Balaraman, H. Hayatnagarkar, M. Singh, M. Duggirala, Towards better crisis management in support services organizations using fine grained agent based simulation, in: PRIMA 2016: Principles and Practice of Multi-Agent Systems, volume 9862, Springer International Publishing, 2016, pp. 366–375. doi:10.1007/978-3-319-44832-9_24.
[23] C. E. Shannon, A mathematical theory of communication, Bell Systems Technical Journal 27 (1948) 379–423.
[24] IAEA, International Atomic Energy Association, 2024. URL: https://www.iaea.org.
[25] piri, PIRI Ontology Reference, 2024. URL: https://www.piri-safety.com.
[26] piri, PIRI Resources and References GitHub, 2024. URL: https://github.com/regdoc/piri.
[27] M. I. F. O. Firoj Alam, Umair Qazi, Humaid: Human-annotated disaster incidents data from twitter, in: 15th International Conference on Web and Social Media (ICWSM), 2021.
[28] M. Honnibal, I. Montani, S. Van Landeghem, A. Boyd, spaCy: Industrial-strength Natural Language Processing in Python, Zenodo (2020).
[29] openai, Open AI Model Reference, 2024. URL: https://www.openai.com.