Digital Health and Informatics Innovations for Sustainable Health Care Systems
J. Mantas et al. (Eds.)

doi:10.3233/SHTI240694

Modelling Events in Biomedical Decision Support Systems Using Ontologies

Marko MILETIC[a] and Murat SARIYAR[a,1]
[a]*Bern University of Applied Sciences, Switzerland*
ORCiD ID: Murat Sariyar https://orcid.org/0000-0003-3432-2860

Abstract. Biomedical decision support systems play a crucial role in modern healthcare by assisting clinicians in making informed decisions. Events, such as physiological changes or drug reactions, are integral components of these systems, influencing patient outcomes and treatment strategies. However, effectively modeling events within these systems presents significant challenges due to the complexity and dynamic nature of medical data. Especially the differentiation between events and processes as well as the nature of events is often unclear. This paper explores approaches to modeling events in biomedical decision support systems, considering factors such as ontology-based representation. By addressing these challenges, we strive to provide the means for enhancing the functionality and interpretability of biomedical decision support systems concerning events.

Keywords. CDSS, machine learning, BFO, UFO, IAO, events

1. Introduction

In medical practice, there are numerous decisions that must be made based on events. For instance, when managing adverse drug reactions, swift and informed reactions are paramount. Take, for example, a patient undergoing treatment for hypertension. If they experience adverse cardiac events in response to a newly prescribed medication, such as *Mavacamten* [1], understanding the timing, severity, and recurrence of these events becomes crucial for subsequent decision-making. Hence, accurate characterization of events facilitates the identification of potential risk factors and informs future clinical decision-making. A clinical decision support system (CDSS) is designed to facilitate informed decision-making in healthcare. However, there is no universally accepted standard available for modeling and considering events in a CDSS. While HL7 FHIR offers the *Event* pattern that can be integrated into various resources, it falls short of addressing central issues surrounding event modeling [2]. For instance, the definition of events as "the performance of some activity" is overly restrictive, and the associated attributes fail to provide clarity on distinguishing between occurrences as processes or events. Such deficiencies are commonly encountered in event modeling.

In the context of our work on modeling events in intraoperative neuromonitoring (IOM), we have encountered analogous challenges [3]. Events were defined as "real-world property instances" in that work. In IOM, meticulous documentation of events

[1] Corresponding Author: Murat Sariyar, Bern University of Applied Sciences, Quellgasse 21, CH2502 Biel/Bienne, Switzerland; E-mail: murat.sariyar@bfh.ch.

such as alterations in neural responses or surgical interventions is paramount for ensuring patient safety and optimizing surgical outcomes. As existing data standards often fail to address these complexities, we turned to ontologies and realized that significant problems exist there as well. Notably, the Ontology of Adverse Events presents conceptual issues (OAE [4]). For instance, defining adverse events as “a pathological bodily process that occurs after a medical intervention” blurs the distinction between events and processes, posing difficulties in defining their temporal attributes. This differentiation is crucial for event documentation, which relies heavily on precise references to events and their properties. For CDSS, the fundamental question regarding events is whether inference about real sequences is necessary (e.g., determining if a specific signal change triggers brain dysfunction) or if only additional attributes for machine learning algorithms are required. In the former case, a deeper understanding of event nature is imperative, while in the latter case, it may be unnecessary.

In the following, we will first concentrate on modeling events within the upper-level ontologies Basic Foundational Ontology (BFO [5]) and the Unified Foundational Ontology (UFO [6]), and then derive implications for modeling events in CDSS applications. Our objective is to establish best practices aimed at error prevention and fostering clear modeling using insights from ontological modeling. The scientific nature of this investigation stems from its systematic analysis and synthesis of ontological principles and frameworks.

2. Methods

BFO serves as the foundational framework for numerous biomedical ontologies, having inspired over 250 ontology-driven projects. Its generality and adaptability make it an ideal starting point for designing subject-specific ontologies. BFO introduces a fundamental distinction between continuants, such as material entities, which endure over time, and occurrents, like disease stages, which unfold over time and exist only through their constituent parts. In this context, only occurrents are relevant. In BFO 2.0, an occurrent is an entity that unfolds in time or is the instantaneous boundary of such an entity (e.g., a beginning or an ending), or it represents a temporal or spatiotemporal region. In simpler terms, it encompasses processes, events, or the time aspect related to an occurrent. A process is defined as an occurrent with proper temporal parts that inheres in some material entity. Conversely, an event marks the temporal boundary of a process, often determined by decisions, as instantaneity cannot be precisely delineated. Processes remain unchanged because they are the changes themselves, while events remain unchanged as there is no unfolding process to effectuate a change. Despite BFO's clear distinction between processes and events, many specialized ontologies conflate the two. For instance, OAE defines adverse events as “a pathological bodily process that occurs after a medical intervention”, a definition more suited to UFO.
An alternative to BFO is UFO, which forms the basis of OntoUML [7]. While it may not enjoy the same prominence in the biomedical domain as BFO, UFO stands out for its highly philosophically grounded concepts. Comprising three modules, UFO offers a comprehensive framework: UFO-A focuses on endurants (equivalent to continuants in BFO) for structural conceptual modeling, UFO-B addresses perdurants (equivalent to occurrents in BFO), and UFO-C encompasses social and intentional aspects. UFO-B distinguishes between atomic and complex events, abstaining from using the term “process” for the latter. Atomic events denote manifestations of unique object

dispositions – properties that only manifest in specific situations due to triggering events and can also fail to manifest. Complex events consist of events as parts. Similar to BFO, changes are considered as perdurants, yet perdurants cannot themselves undergo change. Any perceived alteration in a perdurant is either a variation – different temporal parts of an event having incompatible properties – or a change occurring to some underlying endurant, which is the focal point of that event.

3. Results

The following three general recommendations can be inferred for modeling events in CDSS based on the observations made so far:

- Adopt standardized frameworks: Utilizing well-grounded concepts ensures clarity, consistency, and accuracy in representing events. Standardized frameworks also facilitate interoperability between different modeling systems.
- Distinguish between processes and events: Processes unfold over time and may involve multiple events, while events typically represent specific occurrences or transitions within processes. Equating processes with events may lead to misinterpretation, such as misconstruing emerging properties.
- Consider temporal aspects and relations of events: This involves assessing duration, timing, and relationships with other occurrences. Confusing processes with events can obscure temporal relationships between different stages of a process or with other events.

Specific recommendations for modeling practices in the CDSS context should be tailored to two distinct use cases. First, a CDSS may be employed to infer dependencies between real-world occurrences. Second, a CDSS may integrate events as additional predictors to enhance the prediction of a target variable. Even if the target variable itself is an event, the second use case presents differences. In one scenario, the focus lies on analyzing the dynamics of events and their interactions, while in the other, the objective is to predict a specific outcome. This dichotomy resembles the differentiation between a generative and a discriminative classifier, which respectively model $P(x, y)$ and $P(y|x)$. However, it's not merely another distribution being modeled; additional variables may be involved, and various types of causality may be relevant. Distinguishing whether an event x is causally related to an event y or if the underlying material entities have a causal relationship should be approached differently, despite the fact that, due to linguistic ambiguity, the emphasis sometimes appears to be solely on events [8]. When determining if an event x causes an event y, the focus is on analyzing the temporal sequence and potential mechanisms leading from one event to the other. This involves understanding how the occurrence of x influences the likelihood or occurrence of y. On the other hand, when examining if underlying material entities have a causal relationship, it involves exploring the structural, physiological, or other relevant characteristics of the entities to ascertain if they interact in a way that produces a causal effect.

Employing a CDSS to infer temporal dependencies of events in the first generic use case is relatively uncommon. The primary focus here is to capture the temporal dependencies of events and then determine whether processes really need to be treated separately. An illustrative example within the realm of decision-making concerning adverse drug events involves a CDSS focused on drug metabolism. In this context, a pharmacokinetic profile is established, akin to a digital twin, to assess the metabolization process of drugs. Events emerge as manifestations of dispositions reflecting the type of metabolization exhibited by a patient, such as rapid or slow processing rates. Relevant components are liberation, absorption, distribution, metabolism, and excretion. In such a scenario, it becomes evident how much BFO is informed by biomedical use cases and thus relevant for them. Actually, all these components are processes, as is a potential adverse drug reaction. If one focuses on the temporal sequences, then one wants to understand these processes and their interconnections. However, if only the phase transitions are important, according to BFO, they can be modeled as fiat boundaries of processes. Here, a certain degree of ambiguity must be converted into pseudo-precision through decision-making. This allows for the analysis of sequence order without focusing on temporal trends, enabling the inference of events from other events. Focusing on the sequence is already a step towards the second use case, but we are dealing still with the interconnection of real-world entities and an emphasize on theoretical understanding.

In the second generic use case, events are registered to be used as predictors in a prediction algorithm. A notable challenge arises by utilizing upper-level ontologies, as they tend to guide towards treating events as real-world entities without the capacity to assign them the role of "predictors". Of course, this can be done independently of any standard, but then one departs from the safe confines of a consistency-promising framework, to which we want to stick. Here, the interest lies not in real-time execution of processes but in the elapsed time, making the equation of events and processes less problematic conceptually (the focus is on the information itself). Hence, a promising approach is to change the perspective on events by not considering them as real-world entities, but by representing them as data items. Information entropy, for example, does exactly that: instead of examining the nature of events, it analyzes their "surprise effect" and calculates the number of bits needed to store the information. This shifts theoretical considerations about events outward and puts emphasis on the documentable attributes of events. The Information Artifact Ontology (IAO [9]) offers a BFO-based framework to model information about entities rather than the entities themselves. In this scenario, establishing references becomes crucial, achieved through structured descriptions. Systems like Ceusters et al.'s referent tracking [10] can initiate a chain of references tracing how information was captured, ensuring robust data modeling and integration.

The differentiation between the two use cases can be succinctly summarized by two sets of concepts: (i) {explanation, theoretical understanding, property, occurrent}, which centers on comprehending the broader dynamics and theoretical underpinnings of events, and (ii) {documentation, identification, property-instance, continuant}, which emphasizes the identification of event instances primarily for the purpose of labeling and communicating events. For (i), we need means in the ontology to identify events as such and theoretically describe them in their real behavior. This is no easy task, as shown by recent work by Guarino et al. [11]. In the case of (ii), the corresponding ontology does not need to concern itself with the identification of events based on their essential properties. It is trusted that relevant specialists can handle this based on their experience. The ontology then only needs to focus on capturing the specific attributes and ensuring a consistent description. If such real-world verification of such a description is necessary, one must stick to (i). It is important in practice to clarify what is relevant, as there are very different requirements for (i) and (ii). In particular, the OAE seems not to have made this differentiation, making it irrelevant for many use cases. This is problematic, as many CDSS developers encounter OAE when modeling adverse events, and without a thorough examination of event modeling, they may face corresponding issues, at least when it matters to be consistent.

4. Discussion and Conclusions

Modeling events in the context of CDSS is not a novel endeavor, but incorporating upper-level ontologies presents challenges due to the need for a high degree of consistency and clarity. Adequately addressing these challenges makes ontology-based machine learning feasible, significantly aiding feature engineering. The semantic interoperability afforded by this approach also enhances the integration of additional data, particularly in multicenter applications. The extent of a CDSS expert's understanding of the nature of events varies by use case; however, domain experts should focus on the substantive modeling work, while CDSS experts translate this into appropriate mathematical models.

When events need to be understood in their real-world contexts, it often involves processes and requires more complex models, such as those used in pharmacokinetics, where understanding relationships is more critical than making time-sensitive decisions. In such cases, the CDSS might be part of a larger framework. Conversely, when the focus is on predictions, there is less emphasis on theory development and more on time-critical decisions. Here, distinguishing features of events are captured and incorporated into the predictive model, which is typically simpler than mathematical models for dynamic system behavior involving events and processes. In summary, we have provided guidance for modeling events in CDSS using upper-level and related ontologies, emphasizing the critical differentiation between models for dynamic system behavior and those for predictions. These distinctions necessitate distinct modeling strategies, which are not yet fully represented within the ontology and CDSS domains. The next step will involve crafting a systematic guideline for ontology-based modeling of events.

References

[1] Nag S, Gollapudi SK, del Rio CL, et al. Mavacamten, a precision medicine for hypertrophic cardiomyopathy: From a motor protein to patients. Sci Adv 2023; 9: eabo7622.

[2] Elkin PL, Brown SH. Fast Healthcare Interoperability Resources (FHIR). In: Elkin PL (ed) Terminology, Ontology and their Implementations. Cham: Springer International Publishing, 511-65.

[3] Neuenschwander S, Romao P, Holm J, et al. Developing an Ontology for Documenting Adverse Events While Avoiding Pitfalls. Inform Technol Clin Care Public Health 2022; 166–169.

[4] He Y, Sarntivijai S, et al. OAE: The Ontology of Adverse Events. J Biomed Semant 2014; 5: 29.

[5] Arp R, Smith B, Spear AD. Building Ontologies with Basic Formal Ontology. MIT Press, 2015.

[6] Guizzardi G, Benevides AB, Fonseca CM, et al. UFO: Unified Foundational Ontology. Appl Ontol 2022; 1: 167–210.

[7] Guizzardi G, Fonseca CM, Benevides AB, et al. Endurant Types in Ontology-Driven Conceptual Modeling: Towards OntoUML 2.0. In: Trujillo JC, Davis KC, Du X, et al. (eds) Conceptual Modeling. Cham: Springer International Publishing, 2018, pp. 136–150.

[8] Davidson D, Davidson D. Essays on Actions and Events: Philosophical Essays Volume 1. New Edition, New Edition. Oxford, New York: Oxford University Press, 2001.

[9] Ceusters W. An Information Artifact Ontology Perspective on Data Collections and Associated Representational Artifacts. Qual Life Qual Inf 2012; 68–72.

[10] Ceusters W, Smith B. Strategies for referent tracking in electronic health records. J Biomed Inform 2006; 39: 362–378.

[11] Guarino N, Baratella R, Guizzardi G. Events, Their Names, and Their Synchronic Structure. Appl Ontol 2022; 17: 249–283.