On the Combination of Event Calculus and Empirical Semantic Drifts

Franz Krause[1], Xi Yan[2], Baptiste Darnala[3,4] and Michel Dumontier[5]

[1]Data and Web Science Group, University of Mannheim, Germany
[2]University of Hamburg, Germany
[3]Elzeard, Bègles, France
[4]LIRMM, University of Montpellier, France
[5]Institute of Data Science, Maastricht University, Netherlands

Abstract

The encoding and exploitation of semantics has been gaining popularity, as exemplified by the uptake of digital ontologies and knowledge graphs. However, the semantics of domain objects usually do not reflect how they evolved over time, i.e., which events their dynamic transitions are based on. While a number of methods have been proposed to trace events and their impacts on a domain, there is a paucity of approaches to effectively join them. Thus, we combine event calculus as an analytical approach for modeling causal relationships between events and effects with semantic drifts as an empirical approach for quantifying the impact of domain updates. We demonstrate how their respective weaknesses can be addressed and how their interaction can improve the representation of semantic transitions.

Keywords

Event Calculus, Semantic Drift, Dynamic Knowledge Graph

1. Introduction

Semantics, often referred to as the study of meaning and truth, represent the foundation of human cognitive abilities and thus most research fields [1]. After all, without semantics it is impossible to put things into context and draw conclusions about them. Semantics are thus examined in various real-world domains, e.g., biomedicine [2], manufacturing [3], or finance [4].

Within a domain, objects are assigned meaning based on their interactions with each other. Thus, knowledge graphs (KGs) structure knowledge based on ontological conceptualizations so that the semantics of an object can be inferred from its graph neighborhood [5]. Accordingly, given a respective KG $\mathcal{G}$, the corresponding domain axioms can be applied for logical inferences. However, regarding an ordered time set $\mathcal{T}$ and some dynamic KG $(\mathcal{G}_t)_{t\in\mathcal{T}}$, they lack the ability to incorporate individual semantic events as enablers of domain updates. For $t_i, t_k \in \mathcal{T}$ with $t_i < t_k$, domain axioms can only be applied to the already updated KG $\mathcal{G}_{t_k}$. The causes of these updates are thus neglected, i.e., semantic events are not included in the logical inferences.

SEMMES'23: Semantic Methods for Events and Stories co-located with ESWC, May 29 2023, Hersonissos, Greece
franz.krause@uni-mannheim.de (F. Krause); xi.yan@uni-hamburg.de (X. Yan); baptiste.darnala@elzeard.co (B. Darnala); michel.dumontier@maastrichtuniversity.nl (M. Dumontier)
0000-0003-2869-6815 (F. Krause); 0009-0004-7379-4080 (X. Yan); 0000-0001-7390-4850 (B. Darnala); 0000-0003-4727-9435 (M. Dumontier)

2. Related Work

To incorporate semantic events within logical inferences, some works already exist, which can be divided into analytical and empirical approaches. Analytical approaches, such as event calculus [6], extend first-order logic (FOL) rules for temporal reasoning [7, 8]. Domain updates between two timestamps t_i and t_k are initiated by events e_t with $t \in [t_i, t_k]$, and relationships between domain objects are regarded as fluent states that can be either true or false. Accordingly, effects are defined that take into account the semantics at time t_i and the events that have taken place inbetween t_i and t_k to conclude the fluent states within the updated domain image. As prior domain knowledge about events is directly incorporated, temporal reasoning is a proactive approach that can also be applied to temporal knowledge graph extensions [9, 10].

Contrarily, empirical approaches assume two self-contained semantic representations of the domain knowledge. Besides knowledge graphs, external representation types, such as word or text annotations, are explicitly allowed as well. Based on the representations for both timestamps t_i and t_k, an attempt is made to identify so-called semantic drifts [11], i.e., to identify domain objects whose semantics have drifted from t_i to t_k, and to quantify these drifts [12, 13]. In conclusion, semantic drifts represent a reactive approach to the subsequent identification of semantic transitions and thus also events and their effects as enablers of these transitions.

In the following, both approaches are discussed in more detail and summarized in a compact manner. Thereby, we focus on their advantages and in particular their disadvantages with respect to their real-world applications. Based on these findings, we elaborate to what extent the combination of event calculus and semantic drifts can counteract their respective drawbacks and thus faciliate the incorporation of semantic events and transitions in dynamic domains.

3. Event Calculus

As an extension of FOL, event calculus is based on propositions that can be either true or false. A proposition $p(\omega_1, .., \omega_n)$ is composed of some predicate p that asserts a logical relationship among non-logical objects $\omega_1, .., \omega_n \in \Omega$, such as entities or concepts of a domain. Here, Ω represents the set of all available domain objects and $n \in \mathbb{N}$ denotes the arity of a property, i.e., the number of objects within its logical expression. Accordingly, KGs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with vertices $\mathcal{V} = \Omega$ and edges $\mathcal{E}$ can be interpreted as sets of binary propositions, i.e., $p(\omega_1, \omega_2) = True$ implies the directed edge $(\omega_1, p, \omega_2) \in \mathcal{E}$. For example, $president(\omega_1, \omega_2)$ is binary (i.e., $n = 2$) and true if a person ω_1 is the president of a country ω_2. Moreover, domain rules and composite propositions can be constructed via the connectives $\wedge, \vee, =, \neg, \Rightarrow, \Leftarrow$ and the quantifiers $\forall, \exists$.

Accordingly, event calculus is based on the consideration of propositions as **fluents**, namely conditions that can change over time. These fluents are reified, i.e., they are formalized as non-logical objects so that they can serve as inputs for functions with range $\mathcal{T}$. For example,

$$president(\omega_{DT}, \omega_{US}) = [t_{17}, t_{21}] \subseteq \mathcal{T}$$

describes the presidency term of ω_{DT} = Donald Trump in the ω_{US} = United States inbetween the timestamps t_{17} := January 20th, 2017 and t_{21} := January 20th, 2021. Even though different interpretations of event calculus can be found in the literature [14], most of them introduce the additional functional predicates *holdsAt*, *happens*, *initiates*, and *terminates*. The predicate *holdsAt* is used to determine whether a proposition holds at a timestamp $t \in \mathcal{T}$, e.g.,

$$holdsAt(president(\omega_{DT}, \omega_{US}), t) = True$$

holds for all $t \in [t_{17}, t_{21}]$. While *happens* indicates whether a semantic **event** e takes place at time $t \in \mathcal{T}$, the predicates *inititates* and *terminates* define its **effects**, i.e., how fluent states are affected by this event [15]. Events are encoded as non-logical objects and thus build the foundation for controlling dynamics in event calculus. In our example, the change in office (cio) in the US at time t_{21} represents the event e_{cio} that effectively terminated Donald Trump's presidency and initiated the tenure of ω_{JB} = Joe Biden, i.e., the fluent states are updated due to

$$happens(e_{cio}, t_{21}) \ \wedge \ terminates(e_{cio}, president(\omega_{DT}, \omega_{US})) \ \wedge \ initiates\big(e_{cio}, president\big(\omega_{JB}, \omega_{US}\big)\big) .$$

These events can be implicit as above, but also explicit, i.e., an event can explicitly characterize changes in fluent states, e.g., by considering the inauguration of Joe Biden on January 20th, 2021 as a single event. Such explicit events are commonly referred to as **actions** as they actively affect the structure of the domain knowledge. Accordingly, event calculus analyzes which actions need to be performed as effects of semantic events. Each effect is thus to be interpreted as an action. An overview of the dynamic transitions in event calculus can be found in Figure 1.

Figure 1: The main components of event calculus: Each effect is an action and each action is an event.

3.1. The Problem of Event Calculus

Since event calculus requires prior knowledge about a domain, it represents an analytical approach for modeling semantic events and their effects on fluent states. During the definition phase of the respective calculus, future events and effects must already be considered, which constitutes its major drawback. For example, the impacts of events may vary over time or they might even be unknown at the time of definition. Thus, for a dynamic domain, we want to trace the consistency and completeness of a given calculus. Empirical approaches are required for measuring the impact of events so that actions can be verified and missing effects can be indicated. For this purpose, we adopt approaches for determining so-called semantic drifts.

4. Semantic Drifts

For two (not necessarily consecutive) timestamps $t_i, t_k \in \mathcal{T}$, semantic drifts are introduced to measure the impact of semantic updates inbetween both timestamps [16, 17, 18]. Analogous to Section 3, entities and concepts of a domain are considered as non-logical objects $\omega \in \Omega_t$ with semantic representations $\pi_t(\omega) \in \Pi_t$ for some $t \in \mathcal{T}$. The representation set Π_t may include sets of propositions (cf. Section 3), but also external representations, such as textual annotations.

Accordingly, a semantic drift measure $\phi_{t_i \succ t_k} : \Omega_{t_i} \cap \Omega_{t_k} \to \mathbb{R}_{\geq 0}$ is derived from some distance measure $\psi : \Pi_{t_i \succ t_k} \times \Pi_{t_i \succ t_k} \to \mathbb{R}_{\geq 0}$ by means of $\phi_{t_i \succ t_k}(\omega) := \psi(\pi_{t_i}(\omega), \pi_{t_k}(\omega))$. Here, $\Pi_{t_i \succ t_k}$ represents a shared representation space. Thus, for objects $\omega, \omega' \in \Omega_{t_i} \cap \Omega_{t_k}$, the inequality $\phi_{t_i \succ t_k}(\omega') > \phi_{t_i \succ t_k}(\omega)$ indicates a larger semantic drift in ω'. However, it must be assumed that $\Pi_{t_i \succ t_k}$ is a space with some well-defined distance measure ψ, which is generally not the case.

4.1. Semantic Drifts in Numerical Embedding Spaces

To solve this problem, embedding mappings $\gamma_t : \Pi_t \to \Pi_t^\star$ can be used to embed the given semantic representations within a numerical representation space $\Pi_t^\star$ equipped with well-defined distance measures. For example, KG embedding methods like TransE [19] or RDF2Vec [20] assign numerical representations to the nodes of a KG [21]. Similarly, Natural Language Processing (NLP) introduces language models like Word2Vec [22], BERT [23], or T5 [24] to convert text into numerical embeddings [25]. Typically, $\Pi_t^\star$ is chosen to be a real-valued embedding space, e.g., $\Pi_t^\star = \mathbb{R}^{d_t}$ with $d_t \in \mathbb{N}$. To merge the numerical representations of both timestamps $t_i, t_k \in \mathcal{T}$ within a joint embedding space $\Pi_{t_i \succ t_k}^\star$, embedding alignments can be performed via

$$\alpha_{t_i} : \mathbb{R}^{d_{t_i}} \to \mathbb{R}^{d_{t_i \succ t_k}} \text{ and } \alpha_{t_k} : \mathbb{R}^{d_{t_k}} \to \mathbb{R}^{d_{t_i \succ t_k}}$$

with $\Pi_{t_i \succ t_k} = \mathbb{R}^{d_{t_i \succ t_k}}$ and $d_{t_i \succ t_k} \in \mathbb{N}$, to approximate different representations of identical objects. Considering the alternative representations $\pi_t^\star(\omega) := \alpha_t(\gamma_t(\pi_t(\omega)))$, these are adjusted through

$$\pi_{t_i}(\omega) \approx \pi_{t_k}(\omega) \iff \psi\left(\pi_{t_i}^\star(\omega), \pi_{t_k}^\star(\omega)\right) \approx 0$$

for some predefined distance measure $\psi : \mathbb{R}^{d_{t_i \succ t_k}} \times \mathbb{R}^{d_{t_i \succ t_k}} \to \mathbb{R}_{\geq 0}$ like the cosine or the euclidean distance, so that outliers are defined as semantic drifts. For the sake of completeness, it should be mentioned that only one or even no alignment may be performed. For example, embeddings can be aligned in an existing embedding space $\mathbb{R}^{d_{t_i}}$ or $\mathbb{R}^{d_{t_k}}$ via $\alpha_{t_i} = id$ or $\alpha_{t_k} = id$. For some dynamic embedding methods, such as [26, 27], it is even possible to a priori assume $\gamma_{t_i}(\pi_{t_i}(\omega)) \approx \gamma_{t_k}(\pi_{t_k}(\omega))$ for $\pi_{t_i}(\omega) \approx \pi_{t_k}(\omega)$ and thus also $\alpha_{t_i} = \alpha_{t_k} = id$. In conclusion, distance measures can be applied to quantify the semantic drift of an object $\omega \in \Omega_{t_i} \cap \Omega_{t_k}$ via $\phi_{t_i \succ t_k}(\omega) := \psi(\pi_{t_i}^*(\omega), \pi_{t_k}^*(\omega))$.

4.2. Semantic Drifts based on Representation-based Distance Measures

In contrast to embedding-based methods for determining semantic drifts, other approaches exist which omit the prior embedding of the domain objects and instead consider representation-based distance measures $\psi : \Pi_{t_i \succ t_k} \times \Pi_{t_i \succ t_k} \to \mathbb{R}_{\geq 0}$, where $\Pi_{t_i} = \Pi_{t_k} = \Pi_{t_i \succ t_k}$ is always assumed. For example, $\Pi_{t_i \succ t_k}$ could represent all valid graph neighborhoods of some node representation within a KG, or it could represent the set of all english text fragments. Such formalisms can be found in [28, 29], among others, which consider compositions of semantic representations and corresponding distance measures. These distance measures are typically derived from similarity measures $\sigma : \Pi_{t_i \succ t_k} \times \Pi_{t_i \succ t_k} \to [0, 1]$, i.e., $x, y \in \Pi_{t_i \succ t_k}$ are semantically indistinguishable if $\sigma(x, y) = 1$ holds and unequal for $\sigma(x, y) = 0$. Accordingly, semantic drifts can be defined via

$$\phi_{t_i \succ t_k}(\omega) = \psi\left(\pi_{t_i}(\omega), \pi_{t_k}(\omega)\right) = 1 - \sigma\left(\pi_{t_i}(\omega), \pi_{t_k}(\omega)\right) \in [0, 1].$$

Since arbitrary $\mathbb{R}$-valued distances measures can be derived almost analogously, we restrict ourselves to such $[0, 1]$-valued distance measures in the following without loss of generality.

For example, textual object annotations $\pi_{t_i}^{label}(\omega)$ and $\pi_{t_k}^{label}(\omega)$ for some $\omega \in \Omega_{t_i} \cap \Omega_{t_k}$ can be compared via text comparison methods like the Monge-Elkan similarity [30]. Analogously, KG-based representations can be considered, e.g., the number of adjacent nodes or the sets of common edges can be determined for both timestamps t_i and t_k to define the semantic drift as their difference or by applying set similarity measures like the Jaccard index [31], respectively.

Representation-based distance measures are also applied in [32], where additional KG-based aspects like URIs, superclasses and subclasses, and equivalent classes are incorporated. Overall, such approaches are always based on heuristics that can be directly applied to some graph structure or external representation (e.g., text) to measure the semantic drifts of domain objects.

4.3. The Problem of Semantic Drifts

Compared to the analytical approach of analyzing semantic transitions in event calculus, semantic drifts represent an empirical method for quantifying the impact of semantic transitions on domain objects. Order statistics of the drift scores $\{\phi_{t_i \succ t_k}(\omega) : \omega \in \Omega_{t_i} \cap \Omega_{t_k}\}$ can be defined and compared, e.g., the $m < |\Omega_{t_i} \cap \Omega_{t_k}|$ domain objects can be determined that drifted the most. However, the selection of some well-defined threshold $\tau \in \mathbb{R}_{\geq 0}$ with

$$\omega \text{ drifted significantly inbetween } t_i, t_k \in \mathcal{T} :\iff \phi_{t_i \succ t_k}(\omega) \geq \tau \tag{1}$$

is not trivial at all and solutions need to be elaborated for this problem.

5. Combining Event Calculus and Semantic Drifts

In this chapter, we discuss how event calculus and semantic drifts can be combined for the modeling and analysis of semantic transitions in dynamic domains. According to Section 3, transitions between two timestamps $t_i, t_k \in \mathcal{T}$ are always based on events e_t with $t \in [t_i, t_k]$ that implicitly or explicitly update the domain knowledge. In the following, we first reveal to what extent prior knowledge from an existing event calculus can improve the quality of semantic drifts. Subsequently, we show to what extent semantic drift measures can be applied to counteract the drawbacks of formalisms based on event calculus.

5.1. Improving Semantic Drifts via Event Calculus

As mentioned in Section 4.3, one major drawback of semantic drifts is that conclusions regarding the significance of a semantic drift $\phi_{t_i \succ t_k}(\omega)$ can only be drawn in relation to those of other objects in $\Omega_{t_i} \cap \Omega_{t_k} \setminus \{\omega\}$. While the determination of a globally valid threshold $\tau \in \mathbb{R}_{\geq 0}$ from Equation 1 seems impossible, an event calculus can be applied to derive a transition-specific threshold $\tau_{t_i \succ t_k}$ that only examines the updates between $t_i, t_k \in \mathcal{T}$. Considering the example from Section 3, the inauguration of a person could be defined as a significant event with respect to his or her semantics (if not otherwise impacted). Thus, by specifying $t_i = t_{17}$ and $t_k = t_{21}$,

$$\omega \text{ drifted significantly inbetween } t_i, t_k \in \mathcal{T} :\iff \phi_{t_i \succ t_k}(\omega) \geq \tau_{t_i \succ t_k} := \phi_{t_i \succ t_k}(\omega_{JB})$$

can be defined so that all domain objects that drifted at least as much as Joe Biden are determined as objects whose semantics changed significantly. For this, knowledge is required about the semantic events underlying the dynamic transitions, which is made possible by event calculus.

5.2. Improving Event Calculus via Semantic Drifts

In Section 3.1, two major drawbacks of formalisms based on event calculus are identified.

I Events are descriptive, i.e., their semantic impacts can be encoded in their effects, but their significance can not be measured in a quantitative manner.

II Events and their effects are defined prior to deployment. Thus, the calculus could be incomplete and might need to be updated in the future.

As the outcome of an event is encoded within its effects, i.e., their actions on a collection of domain properties (e.g., edges in a KG), we can identify objects whose semantics were updated as the cause of an event. Thus, the introduction of semantic drifts provides a possible solution to the first drawback since the drifts of affected objects can be measured and aggregated so that the impacts of different events can be compared, as exemplified in the following.

Example. We assume the example from Section 3 regarding the change in office in the United States on $t_k := t_{21} =$ January 20th 2021 that is defined as the semantic event e_{cio}. Further, we assume that $t_i :=$ January 19th 2021 and t_k are consecutive. To analyze the impact of this event in more detail, we split up e_{cio} into two subevents that represent Donald Trump's resignation and Joe Biden's inauguration, respectively. We achieve this by means of

$$e_{cio,\alpha} \mathrel{\widehat{=}} \text{Donald Trump's resignation} \quad \text{and} \quad e_{cio,\beta} \mathrel{\widehat{=}} \text{Joe Biden's inauguration.}$$

Finally, we define an additional timestamp t_j with $t_i < t_j < t_k$, so that $e_{cio,\alpha}$ happens at t_j and $e_{cio,\beta}$ happens at t_k. Since ω_{US} is directly affected by both subevents, we want to quantify their impacts on it. Semantic drift scores $\phi_{t_i \succ t_j}(\omega_{US})$ and $\phi_{t_j \succ t_k}(\omega_{US})$ can be utilized to answer this question. In this context, it is important to note that, in the case of embedding-based semantic drifts, alignments should be performed within the shared embedding space $\Pi^{\star}_{t_j}$, i.e., by considering $\alpha_{t_j}(\cdot) = id$, to ensure comparability of the drift scores. △

Similarly, semantic drifts can counteract the second drawback mentioned above. Since they are based on empirical observations of semantic representations within the domain, we can identify how much domain objects have drifted semantically, to determine whether the calculus is capable of describing the dynamic transitions and the corresponding domain updates.

Example. Analogous to the previous example, we assume the event e_{cio} and the timestamps t_i, t_k. Further, we consider some semantic drift measure $\phi_{t_i \succ t_k} : \Omega_{t_i} \cap \Omega_{t_k} \to \mathbb{R}_{\geq 0}$. Since the above-mentioned event calculus is rather restricted, it may disregard some semantic effects. Thus, semantic drifts can be applied to determine the order statistic of the objects' drifts or some threshold $\tau \in \mathbb{R}_{\geq 0}$ as in Equation 1 to determine if an object $\omega \in \Omega_{t_i} \cap \Omega_{t_k}$ drifted significantly. In our example, embedding-based approaches could implicitly or explicitly state that the semantics of the democratic and/or republican party were affected by this event, even though their semantic representations were not actively updated. Thus, semantic drifts can serve as indicators of the quality of an event calculus, so that it can be subsequently adjusted if necessary. △

6. Conclusion

In this work, we reviewed and analyzed event calculus and semantic drifts as approaches for identifying and understanding semantic transitions in dynamic domains. While event calculus is a proactive approach to the modeling of causal relationships between events and their effects, semantic drifts assume two self-contained images of a dynamic domain to reactively identify domain objects whose semantics have changed as an effect of a transition.

In this context, drawbacks are pointed out that impair or prevent their real-world applications. To counteract these drawbacks, we propose to conduct both approaches in a complementary manner, to combine the benefits of analytical modeling and empirical observations. On the one hand, semantic drifts are enriched by prior knowledge about the events underlying the transitions, so that it is made possible to assess whether drifts need to be regarded as significant or not. On the other hand, quantitative statements about impacts of events and their effects on domain objects enable the qualitative verification of temporal reasoning formalisms like event calculus. The proposed approach to combine both methods is exemplified and serves as a basis for future works to facilitate the modeling of dynamic knowledge based on semantic events.

Acknowledgments

- This work is part of the TEAMING.AI project which receives funding in the European Commission's Horizon 2020 Research Programme under Grant Agreement Number 957402.
- We acknowledge the financial support from the Federal Ministry for Economic Affairs and Energy of Germany in the project CoyPu (Project Number 01MK21007[G]).
- We acknowledge support from the National Office for Biodiversity with MesclunDurab grant and the I-NOV grants. This work was also partially achieved with support of the *Data to Knowledge in Agronomy and Biodiversity* (D2KAB – www.d2kab.org) project that received funding from the French National Research Agency (ANR-18-CE23-0017).

References

[1] L. Talmy, Toward a Cognitive Semantics. Volume 2: Typology and Process in Concept Structuring, Language, Speech, and Communication, MIT Press, 2000.

[2] S. K. Mohamed, V. Nováček, A. Nounu, Discovering protein drug targets using knowledge graph embeddings, Bioinformatics 36 (2020) 603–610.

[3] G. Buchgeher, D. Gabauer, J. Martinez-Gil, L. Ehrlinger, Knowledge graphs in manufacturing and production: A systematic literature review, IEEE Access 9 (2021) 55537–55554.

[4] X. Fu, X. Ren, et al., Stochastic optimization for market return prediction using financial knowledge graph, in: IEEE International Conference on Big Knowledge, ICBK, IEEE Computer Society, 2018, pp. 25–32.

[5] F. Krause, T. Weller, H. Paulheim, On a generalized framework for time-aware knowledge graphs, in: Proceedings of the 18th International Conference on Semantic Systems (SEMANTiCS), volume 55 of *Studies on the Semantic Web*, IOS Press, 2022, pp. 69–74.

[6] R. Kowalski, M. Sergot, A logic-based calculus of events, New Generation Computing 4 (1986) 67–95.
[7] F. van Harmelen, V. Lifschitz, B. Porter, Chapter 17 event calculus, in: Handbook of Knowledge Representation, volume 3 of *Foundations of Artificial Intelligence*, Elsevier, 2008, pp. 671–708.
[8] K. Van Belleghem, M. Denecker, D. De Schreye, On the relation between situation calculus and event calculus, The Journal of Logic Programming 31 (1997) 3–37.
[9] S. Gottschalk, E. Demidova, Eventkg: A multilingual event-centric temporal knowledge graph, in: Proceedings of the 15th Extended Semantic Web Conference (ESWC), Lecture Notes in Computer Science, Springer, 2018, pp. 272–287.
[10] J. Wu, M. Cao, et al., TeMP: Temporal message passing for temporal knowledge graph completion, in: Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), 2020, pp. 5730–5746.
[11] E. Davenport, B. Cronin, Knowledge management: semantic drift or conceptual shift?, Journal of Education for library and information Science (2000) 294–306.
[12] T. Tietz, M. Alam, et al., Challenges of knowledge graph evolution from an nlp perspective, in: 3rd Workshop on Humanities in the Semantic Web, 2020, pp. 71–76.
[13] S. Wang, S. Schlobach, M. Klein, Concept drift and how to identify it, Journal of Web Semantics 9 (2011) 247–265.
[14] R. Miller, M. Shanahan, Some Alternative Formulations of the Event Calculus, Springer, 2002, pp. 452–490.
[15] M. Shanahan, The event calculus explained, in: Artificial Intelligence Today: Recent Trends and Developments, Springer, 1999, pp. 409–430.
[16] A. Blank, P. Koch, Historical Semantics and Cognition, De Gruyter, 1999, pp. 1–14.
[17] E. Davenport, B. Cronin, Knowledge management: Semantic drift or conceptual shift?, Journal of Education for Library and Information Science 41 (2000) 294–306.
[18] T. G. Stavropoulos, S. Andreadis, et al., A framework for measuring semantic drift in ontologies, in: Joint Proceedings of the Posters and Demos Track of the 12th International Conference on Semantic Systems (SEMANTiCS), 2016.
[19] A. Bordes, N. Usunier, et al., Translating embeddings for modeling multi-relational data, in: Proceedings of the 27th Annual Conference on Neural Information Processing Systems, 2013, pp. 2787–2795.
[20] P. Ristoski, H. Paulheim, Rdf2vec: Rdf graph embeddings for data mining, in: Proceedings of the 15th International Semantic Web Conference, Springer, 2016, p. 498–514.
[21] Q. Wang, Z. Mao, et al., Knowledge Graph Embedding: A Survey of Approaches and Applications, IEEE Transactions on Knowledge and Data Engineering 29 (2017) 2724–2743.
[22] T. Mikolov, I. Sutskever, et al., Distributed representations of words and phrases and their compositionality, in: Proceedings of the 26th International Conference on Neural Information Processing Systems (NeurIPS), Curran Associates Inc., 2013, p. 3111–3119.
[23] J. Devlin, M.-W. Chang, et al., BERT: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics, ACL, 2019, pp. 4171–4186.
[24] C. Raffel, N. Shazeer, et al., Exploring the limits of transfer learning with a unified text-to-text transformer, J. Mach. Learn. Res. 21 (2020).

[25] A. Kutuzov, L. Øvrelid, et al., Diachronic word embeddings and semantic shifts: a survey, in: Proceedings of the 27th International Conference on Computational Linguistics, ACL, 2018, pp. 1384–1397.
[26] F. Krause, Dynamic knowledge graph embeddings via local embedding reconstructions, in: The Semantic Web: ESWC 2022 Satellite Events, Springer International Publishing, 2022, pp. 215–223.
[27] V. Hofmann, J. Pierrehumbert, H. Schütze, Dynamic contextualized word embeddings, in: Proceedings of the 11th International Joint Conference on Natural Language Processing, ACL, 2021, pp. 6970–6984.
[28] S. Wang, S. Schlobach, M. Klein, Concept drift and how to identify it, Journal of Web Semantics 9 (2011) 247–265.
[29] T. Stavropoulos, S. Andreadis, et al., SemaDrift: A hybrid method and visual tools to measure semantic drift in ontologies, Journal of Web Semantics 54 (2019) 87–106.
[30] A. E. Monge, C. P. Elkan, The field matching problem: Algorithms and applications, in: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, AAAI Press, 1996, p. 267–270.
[31] P. Jaccard, Distribution de la flore alpine dans le bassin des Dranses et dans quelques régions voisines, La Société Vaudoise des Sciences Naturelles 37 (1901) 241 – 272.
[32] G. Capobianco, D. Cavaliere, S. Senatore, Ontodrift: a semantic drift gauge for ontology evolution monitoring, in: MEPDaW@ISWC, 2020, pp. 1–10.