

Challenges of Knowledge Graph Evolution from an NLP Perspective

Tabea Tietz^{1,2}, Mehwish Alam^{1,2}, Harald Sack^{1,2}, and Marieke van Erp³

¹ FIZ Karlsruhe – Leibniz Institute for Information Infrastructure, Karlsruhe, Germany

`firstname.lastname@fiz-karlsruhe.de`

² Karlsruhe Institute for Technology, Institute AIFB, Germany

³ KNAW Humanities Cluster, DHLab, Amsterdam, the Netherlands
`marieke.van.erp@dh.huc.knaw.nl`

Abstract. Knowledge graphs often express static facts, but concepts and entities change over time. In this position paper, we propose challenges that arise from the perspective of combining NLP and KG evolution in the digital humanities domain based on preliminary experiments.

Keywords: Knowledge Graph Evolution · NLP · Cultural Heritage

1 Introduction

Knowledge graphs (KGs) intend to represent what we consider true about (part of) the world. KGs are created at a certain point in time and can be considered static snapshots of the real world [8]. However, “*Knowledge lives. It is not static, nor does it stand alone*” [2]. Thus, concepts continuously change over time and can vary between social contexts and locations, i.e. we live in a world with infinite variation and variability.

These concept changes may be a result of technological developments, changing social constructs, political decisions, globalization etc. For example, our current understanding of *family* as a concept has changed drastically over the years, e.g. with same-sex marriages being allowed in more and more states. Likewise, the concept of a *country* can change in terms of tangible properties such as borders, official language(s) and rulers, but also in latent ways such as its citizen’s identities within the state and the perception of a country’s culture by foreigners. These concepts are manifested not only in our cultures norms and values, but also documented through photographs, newspapers, books, music, film, and ads. Digital humanities research often involves the understanding of cultural heritage data. Recently, novel methods involving Natural Language Processing (NLP) supported by KGs have entered the humanities research community [5]. Hence, the evolution of real-world concepts within a KG in combination with NLP is especially relevant. Therefore, *evolution* can be understood in two ways:

1. Natural language text can form the content of a KG through NLP. In this way, *evolution* refers to the text itself. Historical text as the mirror to soci-

eties in varying realities and contexts thereby defines what is being modeled in a KG. Here, NLP is thus an essential part of the process of KG evolution.

2. It can also be assumed that a KG is created or evolves independently of automated NLP processes. In this case, *evolution* means that classes, instances and values are created or altered by a source outside of the reality a text was authored or analyzed in. In this case, NLP is not part of the initial process of evolution, but applies whatever reality is defined in the KG to its source text.

There are a number of challenges in representing the fluidity of a concept within KGs, especially respecting their cultural, temporal and geographical contexts. The goal of this position paper is to describe the challenges that arise from the perspective of combining NLP and KG evolution in the digital humanities domain based on preliminary experiments on the concept of *apple pie*. Furthermore, we present some strategies on how these challenges may be addressed.

The remainder of the paper is structured as follows. Section 2 presents related work on knowledge graph evolution. The use case of apple pie recipes is briefly described in section 3. In section 4, the problems of KG evolution in combination with NLP is discussed based on the use case of apple pie recipes and strategies on how to tackle these issues are presented. Section 5 concludes this paper.

2 Related Work

Concept drift over time is studied in [9]. This analysis is based on theories of concept identity and concept morphing. The authors define the meaning of a concept in terms of intension, extension and label. The intension changes when properties are added or disregarded, the extension refers to the change of instances in the ontology, and a label changes when the name of a concept changes.

Once concept drift has been detected, maintaining KGs with respect to changing entities is the next challenge. [6] mostly focuses on the verification of changes to ensure high data quality. They found that evolutionary patterns in KGs are similar to social networks. Their results contribute to an improved KG editing process towards better efficiency and reliability. This work takes KG evolution from a different angle than the presented paper. The authors describe that errors in KGs occur due to vandalism and carelessness. However, the issue that a definition about a concept may be true at one point in time, but not in another is not addressed.

Knowledge graphs are dynamic and the facts related to an entity are added or removed over time. Therefore, multiple versions of the knowledge graph represent a snapshot of the graph at some point in time. Entities undergo evolution when new facts are added or removed. The approaches to solve the problem of automatically generating a summary out of different versions of a knowledge graph are limited. The authors in [8] propose an approach to create a summary graph capturing temporal evolution of entities across different versions of a knowledge graph in order to use the entity summary graphs for documentation generation, profiling or visualization purposes.

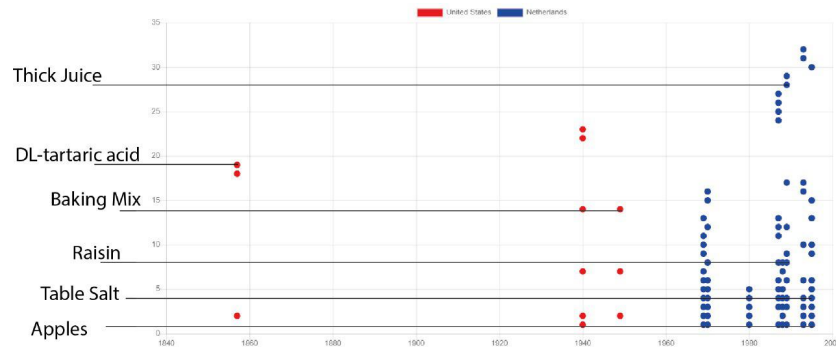


Fig. 1. Timeline visualization of apple pie ingredients in American and Dutch newspapers from 1857 to 1995 (<https://pimpmypie.github.io/>)

3 Use Case

The goal of this position paper is to investigate challenges of KG evolution from an NLP perspective and to provide future visions with respect to digital humanities research. These challenges are based on the preliminary analysis of the concept of *apple pie* recipes extracted from historical Dutch and American newspapers.

In order to study the evolution of apple pie recipes over time, data from Dutch and American newspapers was collected. As [3] remarked, recipes from newspapers reflect tastes and viewpoints in a certain time period and can offer understanding of food cultures. This makes newspapers an invaluable data source to study evolution compared with e.g., cookbooks, which provide a static collection of recipes. For this contribution, the ingredients of apple pie recipes and their corresponding quantities in different contexts (i.e., time, location) were analyzed. Since recipes from historical newspapers are not easily accessible, a small selection of recipes from digitized newspapers was made to provide a proof of concept and illustration of ideas. This selection includes recipes published in one of the four Dutch newspapers *Trouw*, *Het Parool*, *Volkskrant* and *NRC Handelsblad*, or one of the three American newspapers *Evening Star*, *Wilmington Morning Star* and *Pacific Commercial Advertiser* in the period from 1857 until 1995, resulting in 347 apple pie recipes. The recipes were transformed to a structured format, including the available context information (e.g. date, location and language of the publication). Finally, 12 recipes with publication dates spread over the time period 1857-1995 were investigated in a preliminary analysis. The recipes with extracted ingredients are visualized in figure 1 and available on Github.⁴

The concept of apple pie is seemingly simple: it should always contain apple, a kind of flour, a sweetener and a fat – so what is there to evolve? The challenges arising from our preliminary analysis are presented in the following section.

⁴ <https://pimpmypie.github.io/>

4 Challenges and (Possible) Strategies

During the preliminary analysis of the recipe data, we identified the following challenges:

Spatio-temporal context When extracting knowledge from historical sources, several spatio-temporal contexts often have to be taken into account. For example, an article published in an American newspaper in 1995 that describes a typical Hungarian apple pie recipe from the 1950s entails multiple contexts. Here we can distinguish the spatio-temporal metadata of the concept itself (in this example the recipe) and the metadata of its source (i.e. newspaper article). These provenance information will enable to trace the evolution of the concept over time and geographic regions.

Cultural context What is considered as true in one cultural setting may not be in another. For instance, the traditional apple strudel could be considered as a type of apple pie in some cultures, however, in the area formerly belonging to the Austro-Hungarian empire, a clear distinction is made between both desserts, even though the ingredient list is rather similar. Contextual information such as cookbook indexes (or more generally, taxonomies) can help resolve this issue.

Units Extracting and understanding ingredient units presented in the texts was found to be a major challenge in this use case. In modern sources, this involves modern units and their conversion between e.g. the imperial and the metric system of units (kilogram, pound, litre, cups) and in historical sources, this also includes units (usually) not customary on this day, e.g. ell, zentner. Furthermore, less tangible units are sometimes used, e.g. “a load of butter” or “two deep plates of apples”, which provides a greater challenge for the automated detection and interpretation of values and quantities. There are resources for historical measures which can be employed, but imprecise quantities will require human interpretation.

Language When attempting to generate a KG from historical text sources that captures the fluidity of concepts, a number of challenges for NLP arise. In our use case, we found figures of speech, such as metaphors (likely to appear in newspapers), that complicated the process of detecting apple pie recipes correctly. For instance, the following recipe was found: “*Take 1000 kilos of bombs, a few hundred hand grenades, as many boxes of cartridges, go to Vienna with them, make a coup there and wait until you get arrested. Then the apple strudel will be ready [...]*”⁵. Without correctly detecting the metaphor, bombs, grenades and cartridges would be added to the KG as ingredients. Previous research has started to investigate the combination of deep learning and KGs to detect metaphors [1].

Furthermore, the meaning of concept terms may change over time which can be captured with the help of latent representation of the words and

⁵ <http://anno.onb.ac.at/cgi-content/anno?aid=kik&datum=19190907&query=%22Apfelstrudel+Rezept%22~10&ref=anno-search&seite=7>

represented in the KG. One of the initial approaches construct time series of word usage using word embeddings (where one embedding space is generated for each point in time) [4]. In [7], the authors propose an approach based on three components; the first component takes time as an input and generates a time vector, the second component generates a word vector (independent of the time) and the third component combines the time and the word vector.

Concept Modeling The above described challenges also raise the question about how broad or narrow concepts should be modeled in a KG to be able to capture concept change. Is the ontology modeled too specific (e.g. the recipe has to include specific ingredients), recipes from economically weak years (in which these ingredients were not available) would not be considered even though they would yield to interesting results from a digital humanities viewpoint. For example, in figure 1, the US recipe from 1857 includes citric acid instead of apples, which could provide hints on economic shortages. On the other hand, underspecified models may introduce false positives, i.e. recipes falsely detected as apple pies as the example (metaphor) above emphasizes. Furthermore, it is a challenge to determine the properties that define concepts, also with respect to changes over time. If apple pie and apple strudel have a rather similar list of ingredients, also (possibly) the baking procedure and technique plays a vital role in defining the concept, which would have to be modeled in a KG.

Evaluation Finally, an evaluation of an ontology capturing concept change in a KG based on natural language text descriptions should provide measures on how well the model reacts to changes in the data. However, as described above, evaluating whether or not something is to be regarded as an apple pie (or any other concept) depends on many aspects, e.g. cultural background as well as the spatio-temporal setting. Hence, a concise ground truth is a challenge to create and possibly only tendencies can be given. One strategy on how to deal with this is to create a crowd-truth which states the viewpoints on a cultural heritage object by a greater amount of human evaluators from varying backgrounds as well as domain experts. In future work, a crowdsourcing campaign on apple pie (and further use cases) is envisioned.

5 Discussion and Conclusion

The real world is constantly changing and knowledge that was considered true at one point in time in a specific cultural and spatial setting may not be true in another context. That means contexts evolve. On the other hand, there are KGs, which are created and maintained to continuously compose knowledge. However, often KGs are static and only reflect one snippet of reality. This static representation of the real world is a problem when attempting to understand historical descriptions of concepts (e.g., in newspapers), because linking historical concepts to today’s understanding of the same concept may distort its meaning.

In this paper, the problem is addressed on the foundation of preliminary experiments on the concept of apple pie. The take home message of this contribution is that modeling KG evolution even for simple and contained concepts like

apple pie provide complex challenges for ongoing research in the Semantic Web community including the various contexts that need to be taken into account, ambiguities in language and used units as well as the granularity of the model and evaluation. Most of the challenges that were detected are generalizable to numerous concepts within the digital humanities domain.

In future work, more use cases (apart from apple pie) will be analyzed and methods on how to represent KG evolution will be evaluated on the foundation of the challenges and strategies as presented in section 4.

Acknowledgement. This work was made possible by the *International Semantic Web Research Summer School*⁶ in Bertinoro, July 2019. The authors would like to thank the Summer School directors, tutors, the organizing team and the fellow students, especially Mortaza Alinam, Wouter van den Berg, Lientje Maas, Fabio Mariani and Eleonora Marzi.

References

1. Alam, M.: Can knowledge graphs and deep learning approaches help in representing, detecting and interpreting metaphors? Workshop on Deep Learning for Knowledge Graphs (DL4KG) co-located with ESWC 2019 **Vol-2377** (2019)
2. Bonatti, P.A., Decker, S., Polleres, A., Presutti, V.: Knowledge Graphs: New Directions for Knowledge Representation on the Semantic Web (Dagstuhl Seminar 18371). Dagstuhl Reports **8**(9), 29–111 (2019). <https://doi.org/10.4230/DagRep.8.9.29>
3. van Erp, M., Wevers, M., Huurdeman, H.: Constructing a recipe web from historical newspapers. In: Int. Semantic Web Conference. pp. 217–232. Springer (2018)
4. Kulkarni, V., Al-Rfou, R., Perozzi, B., Skiena, S.: Statistically significant detection of linguistic change. In: Gangemi, A., Leonardi, S., Panconesi, A. (eds.) Proceedings of the 24th International Conference on World Wide Web, WWW 2015, Florence, Italy, May 18-22, 2015. pp. 625–635. ACM (2015)
5. Meroño-Peñuela, A., Ashkpour, A., Van Erp, M., Mandemakers, K., Breure, L., Scharnhorst, A., Schlobach, S., Van Harmelen, F.: Semantic technologies for historical research: A survey. Semantic Web **6**(6), 539–564 (2015)
6. Nishioka, C., Scherp, A.: Analysing the evolution of knowledge graphs for the purpose of change verification. In: 2018 IEEE 12th International Conference on Semantic Computing (ICSC). pp. 25–32 (Jan 2018). <https://doi.org/10.1109/ICSC.2018.00013>
7. Rosenfeld, A., Erk, K.: Deep neural models of semantic shift. In: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers). pp. 474–484 (2018)
8. Tasnim, M., Collarana, D., Graux, D., Orlandi, F., Vidal, M.E.: Summarizing entity temporal evolution in knowledge graphs. In: Companion Proceedings of The 2019 World Wide Web Conference. pp. 961–965. ACM (2019)
9. Wang, S., Schlobach, S., Klein, M.: Concept drift and how to identify it. Journal of Web Semantics **9**(3), 247 – 265 (2011), semantic Web Dynamics Semantic Web Challenge, 2010

⁶ <http://semanticwebschool.org/>