
Commonsense Knowledge Base Construction in the Age of Big Data

Simon Razniewski¹

Abstract: Compiling commonsense knowledge is traditionally an AI topic approached by manual labor. Recent advances in web data processing have enabled automated approaches. In this demonstration we will showcase three systems for automated commonsense knowledge base construction, highlighting each time one aspect of specific interest to the data management community. (i) We use *Quasimodo* to illustrate knowledge extraction systems engineering, (ii) *Dice* to illustrate the role that schema constraints play in cleaning fuzzy commonsense knowledge, and (iii) *Ascent* to illustrate the relevance of conceptual modelling. The demos are available online at <https://quasimodo.r2.enst.fr>, <https://dice.mpi-inf.mpg.de> and ascent.mpi-inf.mpg.de.

Keywords: Commonsense knowledge; Web information extraction; Data cleaning

1 Motivation

Knowledge and reasoning about general-world concepts are major challenges in AI. In recent years, these tasks are supported by a growing number of knowledge repositories, so-called *commonsense knowledge bases* (CSKBs), that store statements like LIONS LIVE IN GROUPS, or PAINTERS USE PENCILS. CSKBs share similarities with longer-established encyclopedic KBs (e.g., DBpedia, Yago, Wikidata) - wide scope, pragmatic construction, triple-centric data models, incompleteness and reliance on the open-world assumption, yet also differ notably in terms of semantic and construction complexity [LFL+19; WDRS21].

In terms of *semantic complexity*, commonsense assertions generalize across sets of instances, which makes it more complex to qualify their truth (e.g., LIONS ATTACK HUMANS - does that actually apply to most lions, in most places, most of the time?).

In contrast, instance-centric encyclopedic knowledge in most cases goes well with a binary truth notion (e.g., Angela Merkel is either born in Hamburg, or not).

In terms of *construction complexity*, CSKBs typically have to deal with a myriad of corroborating and conflicting candidate assertions, and cannot rely on a fixed set of relations that alleviates consolidation. For example, texts may state that cats are solitary, and that cats are loners (corroborating assertions), but also that lions live in packs (conflicting, given that lions are big cats).

CSKB construction has recently been revived at a global level [ISZ21; SLA+18; THLB19; ZKSR20], and receives growing interest also in the German data management community [HSW+20; JBCC19; LFD21; OZHH20].

In this demonstration session we showcase three systems for CSKB construction, Quasimodo [RR20; RRP+19], Dice [CRW20a; CRW20b] and Ascent [NRW21], which are part of a long-term research agenda on commonsense at the MPII.² Each system highlights one particular aspect of CSKB construction with high relevance to the data management community:

1. **Quasimodo: Knowledge extraction systems engineering:** Quasimodo is a multi-source CSKB construction system, exemplarily highlighting the important stages of CSKB construction, from source identification to extraction, normalization, validation signal gathering and supervised consolidation.
2. **Dice: Soft schema constraints for CSK cleaning:** In databases, schema constraints are typically clear-cut rules that are strictly enforced. Although strict rules are not applicable to CSK, in Dice we show how soft schema constraints can be utilized to clean noisy CSKB assertions within an approximate and joint reasoning framework.
3. **Ascent: Conceptual modelling for CSK:** Most CSKBs adopt a simple triple-based data model, thus gaining maximum flexibility, at the price of limits on expressivity. With Ascent we show how conceptual modelling can help to capture important relations among commonsense concepts, as well as qualifications for CSK triples.

2 Systems Description

Quasimodo is a CSKB construction framework designed to extract salient assertions, that is, assertions that would spontaneously come to the mind of humans. To achieve this, Quasimodo taps into search engine query logs via auto-completion suggestions as a non-standard input source. However, frequent queries – which are the ones that are visible through auto-completion – are often about sensational and untypical issues. Therefore, Quasimodo combine a recall-oriented candidate gathering phase with two subsequent phases for cleaning, refining, and ranking assertions. Pipeline runs can be inspected in the demonstration system, where one can see in detail how source text give rise to extraction candidates, how these are validated against corroboration sources, and how assertions are ranked and utilized in use cases. Uniquely, it is the first automated CSKB construction system focusing on salient assertions.

² <https://www.mpi-inf.mpg.de/departments/databases-and-information-systems/research/yago-naga/commonsense>

Dice: Database constraints like foreign keys, or value type restrictions, or regex checks are instrumental for ensuring data quality, yet not suited for the open nature of CSK. The Dice framework therefore uses relaxed soft constraints, such as that assertion of taxonomic parents should typically also apply to children, or that textually similar assertions more likely have a similar truth value, to consolidate noisy candidate sets of CSK assertions. Specifically, Dice encodes 17 kinds of these constraints, encodes statement scoring into a weighted MaxSAT problem, and uses approximate solutions found via linear programming in order to efficiently consolidate very large candidate sets. The demo system allows inspecting the underlying taxonomies, the grounded constraint system, and the resulting scores of assertions. Uniquely, it is the first CSKB consolidation system that jointly consolidates candidate assertions.

Ascent introduces a data model where CSK assertions are refined by qualitative facets, thus allowing a finer-grained qualification of assertion truth. It relies on a combination of judicious automated web querying to collect source texts, and dependency-parsing-based extraction that obtains not only S-P-O-triples, but also qualifications such as TIME, LOCATION or CAUSE. To deal with noise and redundancy, in a consolidation phase, assertions are iteratively grouped and semantically organized by an efficient combination of fast word2vec similarity, and classification based on a fine-tuned RoBERTa language model. It is one of the first systems to go beyond the pure triple-based data model for CSK.

3 Demonstration Experience

Extraction systems engineering (Demo: <https://quasimodo.r2.enst.fr>)

Demo attendees will be guided step-by-step through the system design of the Quasimodo CSKB construction pipeline. Each component will be exemplified with sample concepts, and we will critically review strengths and limitations of the implemented approach. Furthermore, attendees will be shown how the resulting knowledge can be utilized in two use cases, question answering, and the game of Taboo .

Soft constraints for data cleaning (Demo: <https://dice.mpi-inf.mpg.de>)

Attendees will be shown the web-extracted taxonomy used for building the constraint systems, as well as the constraint templates. They can then inspect resulting constraint systems for individual subjects, and how changes in input weights affect output scores. Furthermore, attendees will be introduced to the multidimensional scoring mechanism of Dice, and its advantages and challenges compared with existing unidimensional scoring mechanisms.

Conceptual modelling for CSK (Demo: <https://ascent.mpi-inf.mpg.de>)

Attendees will be first introduced to the shortcomings of existing CSKBs in terms

of their organisation of subjects and assertions, using the generic example of ELEPHANT TRUNK, which is out of dictionary for existing CSKBs, while TRUNK alone is confused with trunks of different animals, as well as even car trunks. Attendees will then be shown how Ascent retains TRUNK and similar concepts as dedicated aspects of ELEPHANT, as well as how it retains taxonomic as well as state- and activity-based subgroups like INDIAN ELEPHANT, BABY ELEPHANT, FORAGING ELEPHANT. Attendees will be shown how the attached assertions differ substantially, as well as how qualitative assertions of time, location or degree refine generic triples. Finally, attendees will be shown the role that state-of-the-art pretrained language models play in the statement consolidation.

Acknowledgment

I am grateful to the principal authors of the respective demo systems, Julien Romero, Yohan Chalier and Tuan-Phong Nguyen, as well as to all other collaborators.

References

- [CRW20a] Y. Chalier, S. Razniewski, G. Weikum. “Dice: A Joint Reasoning Framework for Multi-Faceted Commonsense Knowledge”. In: *ISWC system demonstrations*. 2020 (cit. on p. 2).
- [CRW20b] Y. Chalier, S. Razniewski, G. Weikum. “Joint Reasoning for Multi-Faceted Commonsense Knowledge”. In: *AKBC (2020)* (cit. on p. 2).
- [HSW+20] S. Heindorf, Y. Scholten, H. Wachsmuth, A.-C. Ngonga Ngomo, M. Potthast. “CauseNet: Towards a Causality Graph Extracted from the Web”. In: *CIKM*. 2020 (cit. on p. 2).
- [ISZ21] F. Iliovski, P. Szekely, B. Zhang. “CSKG: The CommonSense Knowledge Graph”. In: *ESWC (2021)* (cit. on p. 2).
- [JBCC19] S. Jebbara, V. Basile, E. Cabrio, P. Cimiano. “Extracting common sense knowledge via triple ranking using supervised and unsupervised distributional models”. In: *Semantic Web Journal* (2019) (cit. on p. 2).
- [LFD21] V. Losing, L. Fischer, J. Deigmoeller. “Extraction of Common-Sense Relations from Procedural Task Instructions using BERT”. In: *Global Wordnet Conference*. 2021 (cit. on p. 2).
- [LFL+19] Y. Li, A. Feng, J. Li, S. Mumick, A. Halevy, V. Li, W.-C. Tan. “Subjective databases”. In: *VLDB* (2019) (cit. on p. 1).
- [NRW21] T.-P. Nguyen, S. Razniewski, G. Weikum. “Advanced Semantics for Commonsense Knowledge Extraction”. In: *WWW* (2021) (cit. on p. 2).

- [OZHH20] J. Omeljanenko, A. Zehe, L. Hettinger, A. Hotho. “LM4KG: Improving Common Sense Knowledge Graphs with Language Models”. In: *ISWC*. 2020 (cit. on p. 2).
- [RR20] J. Romero, S. Razniewski. “Inside Quasimodo: Exploring Construction and Usage of Commonsense Knowledge”. In: *CIKM system demonstrations*. 2020 (cit. on p. 2).
- [RRP+19] J. Romero, S. Razniewski, K. Pal, J. Z. Pan, A. Sakhadeo, G. Weikum. “Commonsense Properties from Query Logs and Question Answering Forums”. In: *CIKM* (2019) (cit. on p. 2).
- [SLA+18] M. Sap, R. LeBras, E. Allaway, C. Bhagavatula, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, Y. Choi. “Atomic: An atlas of machine commonsense for if-then reasoning”. In: *AAAI* (2018) (cit. on p. 2).
- [THLB19] A. Talmor, J. Herzig, N. Lourie, J. Berant. “CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge”. In: *NAAACL* (2019) (cit. on p. 2).
- [WDRS21] G. Weikum, L. Dong, S. Razniewski, F. Suchanek. “Machine Knowledge: Creation and Curation of Comprehensive Knowledge Bases”. In: *Foundations and Trends in Databases* (2021) (cit. on p. 1).
- [ZKSR20] H. Zhang, D. Khashabi, Y. Song, D. Roth. “TransOMCS: From Linguistic Graphs to Commonsense Knowledge”. In: *IJCAI* (2020) (cit. on p. 2).