Dice: A Joint Reasoning Framework for Multi-Faceted Commonsense Knowledge*

Yohan Chalier, Simon Razniewski, and Gerhard Weikum

Max Planck Institute for Informatics
{ychalier, srazniew, weikum}@mpi-inf.mpg.de

Abstract. We present a web prototype for the Dice framework for joint consolidation of noisy multi-faceted commonsense knowledge (CSK) [1]. In the demonstration session, participants will be familiarized with the multi-faceted knowledge representation formalism used in Dice, and get the opportunity to inspect grounded constraint systems and resulting inferences for two CSK collections, ConceptNet and Quasimodo. The web prototype is available at https://dice.mpi-inf.mpg.de.

1 Introduction

Commonsense knowledge (CSK) is a potentially important asset towards building versatile AI applications, such as visual understanding for describing images or conversational agents like chatbots. In delineation from encyclopedic knowledge on entities like Trump, Sydney, or FC Liverpool, CSK refers to properties, traits and relations of everyday concepts, such as elephants, coffee mugs or school buses. Prior works on acquiring CSK, such as ConceptNet [9], WebChild [10], TupleKB [5], Quasimodo [8] and others [4], have compiled statements that associate concepts with properties that hold for most or some of their instances.

Yet they compute each concept and statement in isolation from others, and the only quantitative measure (or ranking) is a confidence score capturing whether the statement is valid. There is no information about whether a property holds for all or for some of the instances of a concept, and there is no awareness of which properties are typical and which ones are salient from a human perspective.

Dice overcomes these limitations by introducing a multi-faceted model of CSK statements and methods for joint reasoning over sets of inter-related statements [1]. Dice captures four different dimensions of CSK statements: plausibility, typicality, remarkability and salience, with scoring and ranking along each dimension. For reasoning and ranking, it relies on soft constraints, to couple the inference over concepts that are related in a taxonomic hierarchy. The reasoning is cast into an integer linear programming (ILP), and leverages the theory of reduction costs of a relaxed LP to compute informative rankings. This demonstration will showcase the Dice reasoning system. Participants will engage with the multi-faceted knowledge representation, explore grounded constraint systems, and inspect resulting statement sets.

2 System Overview

Knowledge representation. While binary truth assignments are sensible for encyclopedic knowledge (e.g., Trump either was born in NY, or was not), they fall much short of CSK, which generalizes across individuals and time. To model refined semantics of CSK statements, DICE captures four facets of concept properties: (i) *Plausibility* indicates whether a statement makes sense at all (like the established but overloaded notion of confidence scores). (ii) *Typicality* indicates whether a property holds for most instances of a concept (e.g., not only for cubs). (iii) *Remarkability* expresses that a property stands out by distinguishing the concept from closely related concepts (like siblings in a taxonomy). (iv) *Saliency* reflects that a property is characteristic for the concept, in the sense that most humans would spontaneously list it in association with the concept. This way, statements are better contextualized, loosely similar as by use of qualifiers in traditional entity-centric KBs [6].

Joint reasoning framework. DICE identifies inter-related concepts by their neighborhoods in the WebIsALOD [3] concept hierarchy or via embeddings, and utilizes a set of weighted soft constraints that allows us to jointly reason over the four dimensions for sets of candidate statements. The soft constraint reasoning framework is then cast into an integer linear program (ILP), for which, by harnessing the theory of reduced cost for LP relaxations, rankings for each of the facets are computed. The constraint system is hereby bootstrapped with unidimensional scores from existing CSK collections, probabilistic interpretations of these scores [8,7], and scores based on textual entailment [2].

Our constraint system includes logical clauses such as

$$\text{Plausible}(s_1, p) \wedge \text{Related}(s_1, s_2) \wedge \neg\text{Plausible}(s_2, p) \wedge \ldots \Rightarrow \text{Remarkable}(s_1, p)$$

where ... refers to enumerating all siblings of s_1, or highly related concepts. The constraint itself is weighted by the degree of relatedness; so it is a soft constraint that does allow exceptions.

For example, our framework can draw the following probabilistic inferences:

- Macaques eating bananas makes it likely that also stump-tailed macaques eat bananas (parent-child propagation)
- Penguins not flying is remarkable when most taxonomical siblings do fly (sibling interdependencies)
- Being able to swim correlates with being able to dive (statement similarity)
- Lions attacking humans being salient implies the event being at least plausible (dimension interrelation)

Implementation and web interface. The reasoning was performed using the Gurobi math solver library. The solving of slices of the full constraint system took several hours on a 500 GB RAM machine, we therefore computed the results offline on two existing unidimensional CSK collections, ConceptNet and Quasimodo. Further details on DICE are in [1], and the code can be found at https://github.com/ychalier/dice

We now provide a web interface to explore the underlying constraint systems and resulting statement sets. It is available at https://dice.mpi-inf.mpg.de.

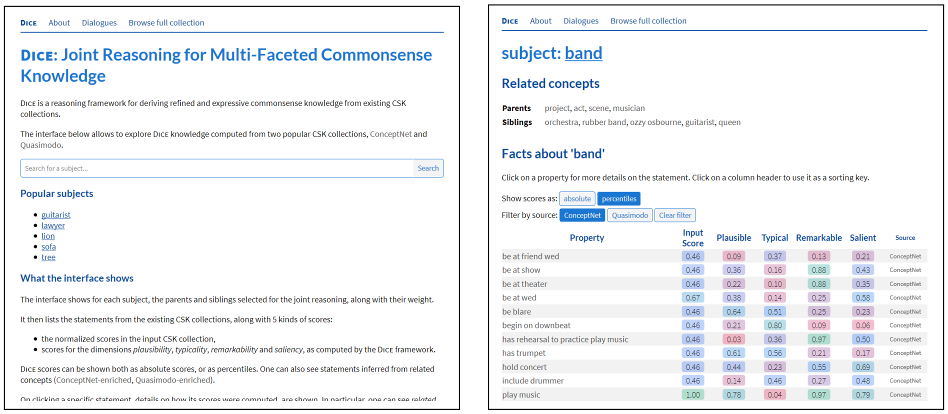

(a) Demo landing page.

(b) List of statements for subject *band*.

Fig. 1: Landing page and subject overview in the web-based demonstration.

3 Demonstration Experience

The demonstration session will consist of three parts.

1. Inspecting multifaceted KR. At the start, attendees will be familiarized with the multi-faceted KR used in DICE. From the landing page (Fig. 1 (a)), they will be shown statements for arbitrary CSK subjects of their choice (see Fig. 1 (b)), and get the chance to explore scores along the four facets, along with initial scores from ConceptNet and Quasimodo. For example, the statements for the subject *band* are available at https://dice.mpi-inf.mpg.de/subject/band. This way, attendees will learn about the difference of unidimensional and multi-faceted scoring.

2. Inspecting constraint systems. In the second part of the demonstration, attendees will explore constraint systems for individual statements. From the screen in (Fig. 1 (b)), attendees will already see taxonomic parents and siblings as used for the constraint system. Upon selecting a specific CSK statement in that screen, a new interface ((Fig. 2) opens, which on the left side (Fig. 2 (a)) shows related statements, numeric statement features, and resulting bootstrapping scores, and on the right side (Fig. 2 (b)) resulting constraint systems. For example, the constraint systems for *band: become famous* is available at https://dice.mpi-inf.mpg.de/fact/band/become-famous.

3. Chat templates In the third part, attendees will be shown a use case of multi-faceted CSK, dialogue. Provided with some chat templates, which allow to include CSK statements with high or low scores along the four dimensions, attendees will be able to write their own dialogues, and explore the impact of inserting statements of different types and scores into these (https://dice.mpi-inf.mpg.de/dialogues).

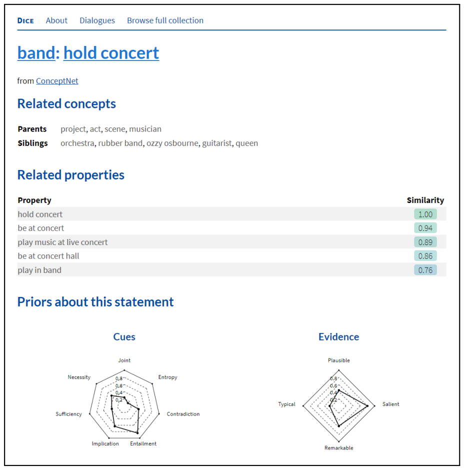

(a) Scores and neighbourhood for statement *band: hold concert*.

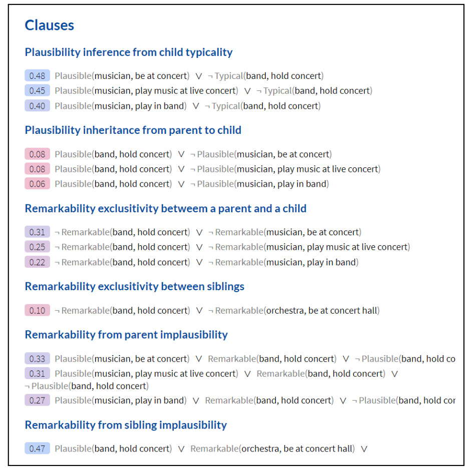

(b) Materialized clauses for statement *band: hold concert*.

Fig. 2: Statement details and constraint system.

References

1. Chalier, Y., Razniewski, S., Weikum, G.: Joint reasoning for multi-faceted commonsense knowledge. AKBC (2020)
2. Gardner, M., et al.: AllenNLP: A deep semantic natural language processing platform. Workshop for NLP Open Source Software (2018)
3. Hertling, S., Paulheim, H.: WebIsALOD: Providing hypernymy relations extracted from the web as linked open data. ISWC (2017)
4. Jebbara, S., Basile, V., Cabrio, E., Cimiano, P.: Extracting common sense knowledge via triple ranking using supervised and unsupervised distributional models. SWJ (2019)
5. Mishra, B.D., Tandon, N., Clark, P.: Domain-targeted, high precision knowledge extraction. TACL (2017)
6. Patel-Schneider, P.F.: Contextualization via qualifiers. Workshop on Contextualized Knowledge Graphs (2018)
7. Romero, J., Razniewski, S.: Inside Quasimodo: Exploring construction and usage of commonsense knowledge. CIKM (2020)
8. Romero, J., Razniewski, S., Pal, K., Pan, J.Z., Sakhadeo, A., Weikum, G.: Commonsense properties from query logs and question answering forums. CIKM (2019)
9. Speer, R., Havasi, C.: ConceptNet 5: A large semantic network for relational knowledge. Theory and Applications of Natural Language Processing (2012)
10. Tandon, N., de Melo, G., Suchanek, F.M., Weikum, G.: WebChild: harvesting and organizing commonsense knowledge from the web. WSDM (2014)