

Commonsense Knowledge Mining from Term Definitions

Zhicheng Liang, Deborah L. McGuinness

Department of Computer Science, Rensselaer Polytechnic Institute
Troy, NY, USA
liangz4@rpi.edu, dlm@cs.rpi.edu

Abstract

Commonsense knowledge has proven to be beneficial to a variety of application areas, including question answering and natural language understanding. Previous work explored collecting commonsense knowledge triples automatically from text to increase the coverage of current commonsense knowledge graphs. We investigate a few machine learning approaches to mining commonsense knowledge triples using dictionary term definitions as inputs and provide some initial evaluation of the results. We start from extracting candidate triples using part-of-speech tag patterns from text, and then compare the performance of three existing models for triple scoring. Our experiments show that term definitions contain some valid and novel commonsense knowledge triples for some semantic relations, and also indicate some challenges with using existing triple scoring models.¹

Introduction

A variety of natural language related tasks, e.g. question answering (Lin et al. 2019) and dialog systems (Young et al. 2018), are able to achieve better performance by introducing commonsense knowledge. However, the knowledge collection process is difficult because commonsense knowledge is assumed to be widely known, thus rarely stated explicitly in natural language text. A Commonsense Knowledge Graph (CSKG) is usually represented as a directed graph, where nodes represent concepts and edges denote some pre-defined relations between concepts. Most existing, large-scale CSKGs are built by expert annotation, e.g. Cyc (Lenat 1995), or by crowdsourcing, e.g. ConceptNet (Speer, Chin, and Havasi 2017) and ATOMIC (Sap et al. 2019). With respect to the broadness and diversity of commonsense knowledge, these CSKGs typically suffer from low coverage.

To increase coverage, there have been two lines of research to infer commonsense knowledge automatically. One is CSKG completion, which aims to learn a scoring model that distinguishes between triples expressing commonsense knowledge and those that do not (Li et al. 2016; Saito et al. 2018; Malaviya et al. 2020). A learned scoring model can

be used to estimate the plausibility of candidate concept-relation-concept triples that are either constructed between existing concepts or extracted from the raw text of new sources. Triples scored above a certain threshold are considered to be valid and then used to augment a CSKG.

The other line is CSKG generation, which aims to generate a new node (or, concept), t_2 , in arbitrary phrases, and connect it with an existing node, t_1 , using a pre-defined relation, R , in order to construct a new triple (t_1, R, t_2) (Saito et al. 2018; Bosselut et al. 2019). One weakness of these generative models is the low novelty rate of generated concepts. As reported in the ConceptNet experiments by Bosselut et al. (2019), only 3.75% of generated triples contain novel object nodes given the subject and the relation of a triple, which limits the ability to augment a CSKG on a large scale. In contrast, the CSKG completion approach is more promising because it allows the introduction of more novel/unseen concepts to the existing graph, by mining diverse candidate triples from external text resources.

Previous work on commonsense knowledge mining exists. Blanco, Cankaya, and Moldovan (2011) extracted commonsense knowledge using concept properties and metarules. Li et al. (2016) extracted knowledge triples from Wikipedia and used their trained model to score the triples. Jastrzebski et al. (2018) further evaluated the novelty of these extracted triples and introduced an automated novelty metric that correlates with human judgement. Davison, Feldman, and Rush (2019) leveraged a pre-trained language model to mine commonsense knowledge. Zhang et al. (2020) mined commonsense knowledge from linguistic patterns of raw text. However, as a potential source of commonsense knowledge, dictionary term definitions have yet to be explored.

Dictionary term definitions are compiled to provide precise descriptions of the properties of terms, concepts, or entities in our daily life. Based on the assumption that *concepts have properties which imply commonsense* (Blanco, Cankaya, and Moldovan 2011), we further assume that some commonsense knowledge could be extracted or inferred from these definitions. For example, the term definition of “bartender” is *One who tends a bar or pub; a person preparing and serving drinks at a bar*, from which one could infer commonsense triples such as $(bartender, IsA, person)$, $(bartender, AtLocation, bar)$, $(bartender, AtLocation, pub)$, $(bar-$

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹Our code and data are available at: <https://github.com/gychant/CSKMTermDefn>

tender; CapableOf, preparing and serving drinks). Among these triples, the second triple is already included in ConceptNet, and the others have semantically similar counterparts, e.g. (bartender; CapableOf, fill glass with drink), (bartender; CapableOf, mix drink), and (bartender; CapableOf, shake drink).

We aim to examine the performance of existing machine learning approaches for mining the described commonsense knowledge triples from such term definitions, i.e. examining their capability of distinguishing valid and invalid triples extracted from the raw text, and to understand the potential and feasibility for mining commonsense automatically from this particular kind of resource.

Approach

In this section, we introduce how we extract candidate triples of commonsense knowledge from term definitions, as well as the models we use to evaluate their plausibility scores, which measure the level of validity of a triple.

Candidate Extraction

We use term definitions from the English version of Wiktionary, a freely-available multilingual dictionary.² Yet, our approach is agnostic to a particular definition resource. As in ConceptNet, the subject and object of a commonsense knowledge triple can be arbitrary phrases. Instead of generating candidates from simple N-grams, Li et al. (2016) extracted candidates using frequent part-of-speech (POS) tag patterns of concept pairs for each pre-defined relation. This method encourages candidates to respect the target commonsense knowledge graph being extended. We employ a similar approach to extracting candidates from term definitions. First, we parse the nodes in ConceptNet using spaCy to obtain their POS tags.³ Next, we choose the top k most frequent POS tag patterns for each relation and apply them to match text spans from term definition text. For instance, the phrase “preparing and serving drinks” can be extracted from the term definition of “bartender” using the POS tag pattern “VERB,CCONJ,VERB,NOUN” for the CapableOf relation. Finally, we construct knowledge triple candidates for each relation using a term as the subject and an extracted phrase from the definition of this term as the object.

Triple Scoring

We adopt three state-of-the-art triple scoring models for computing the plausibility of candidate triples. Each of them assigns a real-valued score to a given triple:

- Bilinear AVG (i.e. average) (Li et al. 2016) defines the plausibility score of a triple as (t_1, R, t_2) as $u_1^\top M_R u_2$ where $M_R \in \mathcal{R}^{r \times r}$ is the parameter matrix for relation R , u_i is a nonlinear transformation of the term vector v_i , and v_i is obtained by averaging the word embeddings of the original term, t_i . Li et al. (2016) used this model to score triples extracted from Wikipedia, since it performs better when scoring novel triples.

- KG-BERT (Yao, Mao, and Luo 2019) treats triples in knowledge graphs as textual sequences by taking the corresponding entity and relation descriptions, and learns the triple scoring function by fine-tuning a pre-trained language model.
- PMI model (Davison, Feldman, and Rush 2019) represents triple plausibility using pointwise mutual information (PMI) between head entity, h , and tail entity, t . Specifically, it translates a relational triple into a masked sentence and estimates the PMI using a pre-trained language model, computed by $\text{PMI}(t, h|r) = \log p(t|h, r) - \log p(t|r)$. The final plausibility score is obtained by averaging $\text{PMI}(t, h|r)$ and $\text{PMI}(h, t|r)$. This approach is unsupervised in the sense that the model does not need to be trained on a particular commonsense knowledge base to learn model weights. Davison, Feldman, and Rush (2019) find that it outperforms supervised methods when mining commonsense knowledge from new sources.

Experiment

We provide details of candidate triple mining and our execution of the three models on the candidate triples, and then include some analysis of our results.

Candidate Triple Mining

ConceptNet 5 (Speer, Chin, and Havasi 2017), the latest version of ConceptNet, is built from multiple sources, the Open Mind Common Sense (OMCS) project, Wiktionary, DBpedia, etc., thus introducing some domain-specific knowledge. To focus on commonsense knowledge only, we collect Wiktionary definitions for the terms that appeared in the English core of ConceptNet (total 162,363 terms).⁴ We further filter out definitions related to morphology, i.e. the ones containing “plural of”, “alternative form of”, “alternative spelling of”, “misspelling of”, resulting in 13,850 term definitions. We choose 12 representative ConceptNet relations as extraction targets. By applying the top 15 frequent POS tag patterns for each relation, we extracted around 1.4 million candidate triples in total for these relations (see Table 3 of the Appendix for detailed statistics).

Model Settings

For Bilinear AVG, we simply adopt the trained model released by Li et al. (2016).⁵ For KG-BERT,⁶ we train the model using the full training set of 100 thousand triples that were used to train Bilinear AVG.⁷ Since this 100k training set contains only positive triples, KG-BERT generates negative triples by replacing the head entity or tail entity of a positive triple with a random entity in the training set. Note that if a corrupted triple with a replaced entity is already among

⁴We use the English triples with ConceptNet 4 as source.

⁵Available at https://ttic.uchicago.edu/~kgimpel/comsense_resources/ckbc-demo.tar.gz

⁶Available at <https://github.com/yao8839836/kg-bert>

⁷Available at https://ttic.uchicago.edu/~kgimpel/comsense_resources/train100k.txt.gz

²https://en.wiktionary.org/wiki/Wiktionary:Main_Page

³<https://spacy.io/>

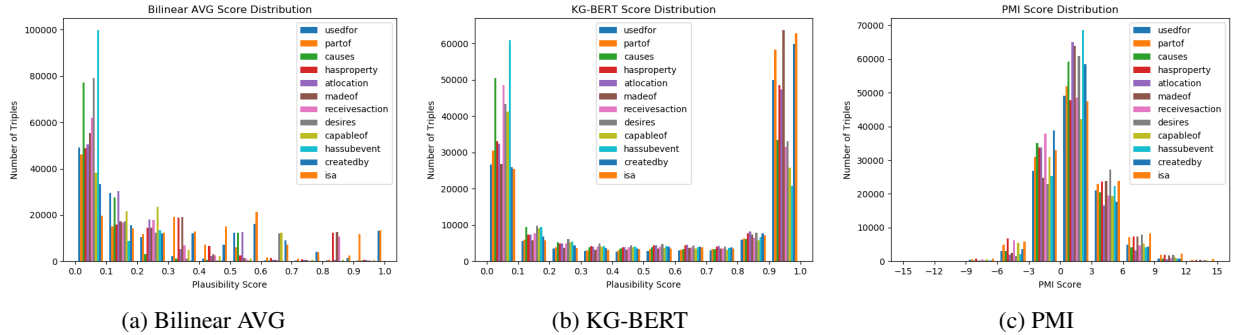


Figure 1: Score distribution of collected triples (best viewed in color).

the positive triples, it will not be treated as a negative triple. We train KG-BERT for 10 epochs and use the model that achieves the best accuracy on the development set released by Li et al. (2016).⁸ For PMI model, we directly run it on the candidate triples without training. Specifically, Bilinear AVG and KG-BERT assign each candidate triple a plausibility score in the range of $[0, 1]$. After running these models, we separately rank the candidate triples of each relation by their scores in descending order.

Analysis

We first analyze the distribution of scores assigned to the collected candidate triples by the three models we considered. The histogram plots with 10 bins are shown in Figure 1. We observe high variability in score distribution across these models, indicating that they do not have an evident consensus towards the plausibility of the candidate triples. It is also supported by the Kendall’s tau coefficients (Kendall 1938) computed between each pair of the models, whose absolute values are all less than 0.01.

To evaluate the inferred knowledge, we adopt two metrics in the literature: *Validity* and *Novelty*. Validity describes how many generated/extracted triples are plausible and is measured by human evaluation. Bosselut et al. (2019) conducted automatic evaluation using the pre-trained Bilinear AVG model developed by Li et al. (2016). Secondly, novelty is measured by the percentage of generated triples that are unseen in the training set (Li et al. 2016; Saito et al. 2018). We also evaluate it against the English subset of ConceptNet. Of the two, validity is more important, yet novelty is also a desirable property when augmenting a CSKG.

1. Novelty When we measure novelty of a triple, we treat concepts in the triple as strings. Preprocessing includes removing stop words in the strings, and treating them as a bag of words after lemmatization and stemming. This approach is stricter than exact match. Similarly, Bosselut et al. (2019) measured novelty using word token based minimum edit distance of generated phrase objects to the nearest one in the training set. We find that around 99% of the candidate triples of each relation are novel w.r.t either triples in

the training set, or triples in the English core of ConceptNet 5. Notably, ConceptNet 5 already contains some triples whose concept pairs are automatically extracted from Wiktionary definitions and have been assigned a vague relation *RelatedTo*, yet only for concept pairs both having a Wiktionary entry. Even with respect to the full English subset of ConceptNet 5, the novelty rates of candidate triples are above 99% for nearly all relations (except *IsA* with 94.7%). Notably, novelty may be affected by the presence of synonyms. To approximate novelty of triples based on semantic distance, Jastrzebski et al. (2018) used the sum of Euclidean distances between the averaged word embeddings of the heads and the tails, respectively, of two triples. We tested this approach using our data but it seems not a good proxy for novelty. We leave further novelty analysis as future work.

2. Validity We sample 50 triples out of the high-scored triples of each relation and manually evaluate their validity as well as novelty using the aforementioned metric. The results are summarized in Table 1. Specifically, we sample triples from the ones scored at least 0.9 for Bilinear AVG and KG-BERT (see the *Qual.* columns for the numbers of “qualified” triples meeting this selection criterion), and sample from the top 1,000 high-scored triples for the PMI model since the scores are not in the range of $[0, 1]$. We report the proportion of valid triples in the samples (see the *V.* columns), and also the proportion of triples being valid and novel (see the *V.N.* columns). Table 2 of the Appendix lists some of the valid and novel triples from the samples we manually evaluated. The results show that the models do mine some valid and novel triples and the performance varies on different relations. Bilinear AVG achieves relatively high accuracy on relations e.g. *HasProperty*, *UsedFor*, *AtLocation*, while getting only 6% accuracy on *CausedBy*. KG-BERT performs better on some relations e.g. *UsedFor*, *CapableOf*, than other relations. PMI achieves better accuracy on the relation *IsA* compared with the other two models. From the distribution of accuracy, we infer that term definitions of Wiktionary contain more commonsense knowledge for the relations with relatively higher accuracy, including *AtLocation*, *CapableOf*, *HasProperty*, *IsA*, *MadeOf*, *UsedFor*. This matches our observation that dictionary term definitions describe more on such basic properties of a concept.

⁸KG-BERT achieves 77.5% accuracy on the dev set after 3 training epochs.

Relation	Bilinear AVG			KG-BERT			PMI	
	Qual.	V.	V.N.	Qual.	V.	V.N.	V.	V.N.
AtLocation	632	0.44	0.30	47,315	0.04	0.04	0.10	0.10
CapableOf	585	0.30	0.28	25,771	0.26	0.26	0.12	0.12
Causes	293	0.18	0.16	33,511	0.02	0.02	0	0
CreatedBy	13,156	0.06	0.06	59,836	0.02	0.02	0.02	0.02
Desires	67	0.22	0.18	33,126	0	0	0.04	0.04
HasProperty	635	0.44	0.38	48,603	0.06	0.06	0.08	0.08
HasSubevent	173	0.10	0.08	20,719	0.06	0.06	0.02	0.02
IsA	13,537	0.36	0.28	62,819	0.24	0.16	0.46	0.34
MadeOf	487	0.26	0.24	63,781	0.14	0.12	0.10	0.08
PartOf	11,909	0.14	0.12	58,410	0.16	0.16	0.04	0.04
ReceivesAction	238	0.12	0.12	31,491	0.20	0.20	0.06	0.06
UsedFor	491	0.50	0.32	49,913	0.36	0.34	0.18	0.18

Table 1: Manual evaluation results of high-scored triples.

Discussion

Our evaluations indicate the feasibility to mine commonsense triples using term definitions, for which we provide evidence using Wiktionary. We also find that the three models under evaluation have weaknesses in scoring triples. This is because the triples for manual evaluation are all sampled from the high-scored triples, yet for all relations, less than 50% of the triples are valid. To some extent, the actual validity of candidate triples could be ruled out since invalid triples should have been assigned low scores (thus not in our samples for manual evaluation) if a model is sufficiently discriminatory. We do not conclude on which model is the best due to the large discrepancy of their score distributions.

Our manual evaluation also uncovers some interesting model behaviors. Bilinear AVG tends to assign very high scores (> 0.99) to the *UsedFor* relation given concept pairs that actually have the *IsA* relation. KG-BERT tends to assign very high scores to the *CreatedBy* relation given concept pairs that actually have the *CapableOf* relation. Regarding the validity of the candidate triples we extracted from Wiktionary, we can get some sense from the results of manual evaluation. For KG-BERT that have a large amount of triples scored above 0.9, as shown in Figure 1b and Table 1, we could roughly estimate the number of valid triples for each relation using the accuracy numbers in Table 1, e.g. around $50,000 \times 0.34 = 17,000$ plausible *UsedFor* triples.

Since our analysis shows relatively wide variability in the studied scoring methods, reliance on these methods may need a deeper evaluation of their individual strengths and weaknesses. Further, validity and novelty are two useful metrics, however, additional research is needed to really consider when mined valid and novel content is worth adding. For instance, we obtained a valid triple (*camper*, *AtLocation*, *tent*) from the definition of *camper* - “A person who camps, especially in a tent etc.” - while ConceptNet already has a similar one (*camper*, *CapableOf*, *sleep in a tent*). To determine whether the former one is worth adding is challenging. One practical approach is to use the performance gain of downstream applications, like question answering, as the criterion for decision making.

Conclusion and Future Work

We presented a study on the feasibility of mining commonsense knowledge triples from Wiktionary term definitions. We examined three models on the performance of scoring newly extracted triples. We showed that they do mine some valid and novel triples and the performance varies on different semantic relations, and also observed some model weaknesses, e.g. low accuracy for high-scored triples and high variability between these models. Our findings suggest careful model pre-evaluation for use in practice. We plan to improve scoring models and candidate extraction techniques, and study the impact of new triples on downstream tasks.

Acknowledgments This work is funded through the DARPA MCS program award number N660011924033 to RPI under USC-ISI West.

References

- Blanco, E.; Cankaya, H.; and Moldovan, D. 2011. Commonsense knowledge extraction using concepts properties. In *Twenty-Fourth International FLAIRS Conference*. Citeseer.
- Bosselut, A.; Rashkin, H.; Sap, M.; Malaviya, C.; Çelikyilmaz, A.; and Choi, Y. 2019. COMET: Commonsense Transformers for Automatic Knowledge Graph Construction. In *ACL*.
- Davison, J.; Feldman, J.; and Rush, A. 2019. Commonsense Knowledge Mining from Pretrained Models. In *EMNLP-IJCNLP*.
- Jastrzebski, S.; Bahdanau, D.; Hosseini, S.; Nourkhovitch, M.; Bengio, Y.; and Cheung, J. C. K. 2018. Commonsense mining as knowledge base completion? A study on the impact of novelty. In *Proc. of the Workshop on Generalization in the Age of Deep Learning*, 8–16.
- Kendall, M. G. 1938. A new measure of rank correlation. *Biometrika* 30(1/2): 81–93.
- Lenat, D. B. 1995. CYC: A large-scale investment in knowledge infrastructure. *Communications of the ACM* 38(11): 33–38.
- Li, X.; Taheri, A.; Tu, L.; and Gimpel, K. 2016. Commonsense knowledge base completion. In *ACL*, 1445–1455.
- Lin, B. Y.; Chen, X.; Chen, J.; and Ren, X. 2019. KagNet: Knowledge-Aware Graph Networks for Commonsense Reasoning. In *EMNLP-IJCNLP*.
- Malaviya, C.; Bhagavatula, C.; Bosselut, A.; and Choi, Y. 2020. Commonsense Knowledge Base Completion with Structural and Semantic Context. *AAAI*.
- Saito, I.; Nishida, K.; Asano, H.; and Tomita, J. 2018. Commonsense knowledge base completion and generation. In *CoNLL*.
- Sap, M.; Le Bras, R.; Allaway, E.; Bhagavatula, C.; Lourie, N.; Rashkin, H.; Roof, B.; Smith, N. A.; and Choi, Y. 2019. ATOMIC: An Atlas of Machine Commonsense for If-Then Reasoning. In *AAAI*, volume 33, 3027–3035.
- Speer, R.; Chin, J.; and Havasi, C. 2017. ConceptNet 5.5: an open multilingual graph of general knowledge. In *AAAI*, 4444–4451.
- Yao, L.; Mao, C.; and Luo, Y. 2019. KG-BERT: BERT for knowledge graph completion. *arXiv preprint arXiv:1909.03193*.
- Young, T.; Cambria, E.; Chaturvedi, I.; Zhou, H.; Biswas, S.; and Huang, M. 2018. Augmenting End-to-End Dialogue Systems With Commonsense Knowledge. In *AAAI*.
- Zhang, H.; Khashabi, D.; Song, Y.; and Roth, D. 2020. TransOMCS: From Linguistic Graphs to Commonsense Knowledge. In *IJCAI*.

Appendix

In Table 2, we show some valid and novel knowledge triples from the samples we evaluated manually. We group them by the semantic relation and the model used to score them. In Table 3, we report the statistics of candidate triples.

Relation	Some Valid and Novel Examples from Human Evaluation		
	Bilinear AVG	KG-BERT	PMI
AtLocation	(camper, tent), (waiter, restaurant), (database, computer), (stove, room), (locker, store)	(paddler, canoe), (circle, figure)	(glioma, brain), (coalminer, coal), (collagen, extracellular), (thyroiditis, thyroid), (pneumothorax, chest)
CapableOf	(nose, smell), (owl, prey), (campfire, heat), (tablecloth, cover), (labor, work)	(salmonella, poisoning), (negotiation, achieving agreement), (shield, defense), (auto mechanic, repairing), (yeast, brew)	(pursuer, pursues), (massage therapist, massage therapy), (showroom, display), (passenger train, rail transport), (droplet, drop)
Causes	(damage, harm), (invest, development), (entertainment, enjoyment), (ponder, thought), (howl, sound)	(multiple sclerosis, depression)	N/A
CreatedBy	(corn earworm, helicoverpa zea), (kibbutz, economical sharing), (yale, heraldry)	(marine, military)	(noise pollution, excess noise)
Desires	(scientist, answer), (predator, prey), (graduate, degree), (workaholic, work), (judge, justice)	N/A	(sexist, practises sexism), (contestant, game show)
HasProperty	(fingerprint, unique), (reptile, cold-blooded), (chili, pungent), (beauty, attractive), (deck, flat)	(comet, celestial), (ratafia, bitter), (mashed potato, pulpy)	(forest, uncultivated), (copyright infringement, unauthorized use), (stockholder, owns stock), (beryllium, alkaline)
HasSubevent	(golf, hit), (bribe, exchange), (asthma, breath), (yawn, breath)	(meningitis, stiffness), (archery, shooting), (rebirth, birth)	(cheating, imposture)
IsA	(sailor suit, clothing style), (oil platform, large structure), (somnoplasty, medical treatment), (scanner, device), (immune response, integrated response)	(puma, lion), (mayor, leader), (royal, person), (bacteriostat, chemical)	(trombone, instrument), (fire truck, vehicle), (carnivore, animal), (cosmetic surgery, medical treatment), (katakana, Japanese syllabary)
MadeOf	(pickle, salt), (sushi, rice), (candle, wax), (cymbal, bronze), (vodka, grain)	(press release, statement), (daytime, time), (confetti, metal), (gas giant, methane), (casino, room)	(chocolate milk, milk), (coronary thrombosis, blood), (glassware, glass), (chopped liver, chicken liver)
PartOf	(expertise, knowledge), (barney, pejorative slang), (massage therapist, massage therapy)	(snooker table, snooker), (seabed, sea), (exotica, american music), (surcharge, price), (free market, market)	(chemistry, natural science), (barrier reef, adjacent coast)
ReceivesAction	(quantity, count), (speech, speak), (asset, value), (cleaner, clean), (plunger, remove)	(sausage, made), (experience, produced), (fishing rod, used), (ledger, record), (harpsichord, tuned)	(saddlery, saddler), (rowboat, rowing), (space station, habitation)
UsedFor	(message, communication), (mouthwash, clean), (ribbon, decoration), (tablecloth, protect), (article, report)	(hypothesis, observation), (immune system, response), (grader, maintenance), (nature reserve, conserve wildlife), (harpsichord, baroque music)	(meteorology, forecasting), (kidney, producing urine), (flush toilet, flush urine), (answer, question), (machine tool, machining)

Table 2: Examples of valid and novel triples.

Relation	AtLocation	CapableOf	Causes	CreatedBy	Desires	HasProperty
# of Candidates	121,184	105,488	122,879	124,064	122,748	122,382
Relation	HasSubevent	IsA	MadeOf	PartOf	ReceivesAction	UsedFor
# of Candidates	123,324	122,182	124,661	120,915	118,604	105,870

Table 3: Statistics of candidate triples.