

CO-NNECT: A Framework for Revealing Commonsense Knowledge Paths as Explications of Implicit Knowledge in Texts

Maria Becker, Katharina Korfhage, Debjit Paul, Anette Frank

Department of Computational Linguistics, Heidelberg University

mbecker|korfhage|paul|frank@cl.uni-heidelberg.de

Abstract

In this work we leverage commonsense knowledge in the form of knowledge paths to establish connections between sentences, as a form of *explicitation of implicit knowledge*. Such connections can be direct (singlehop paths) or require intermediate concepts (multihop paths). To construct such paths we combine two model types in a joint framework we call CO-NNECT: a *relation classifier* that predicts direct connections between concepts; and a *target prediction model* that generates target or intermediate concepts given a source concept and a relation, which we use to construct multihop paths. Unlike prior work that relies exclusively on *static* knowledge sources, we leverage language models finetuned on knowledge stored in ConceptNet, to *dynamically* generate knowledge paths, as explanations of implicit knowledge that connects sentences in texts. As a central contribution we design manual and automatic evaluation settings for assessing the *quality* of the generated paths. We conduct evaluations on two argumentative datasets and show that a combination of the two model types generates meaningful, high-quality knowledge paths between sentences that reveal implicit knowledge conveyed in text.

1 Introduction

Commonsense knowledge covers simple facts about the world, people and everyday life, e.g., *Birds can fly* or *Cars are used for driving*. It is increasingly used for many NLP tasks, e.g. for question answering (Mihaylov et al., 2018), textual entailment (Weissenborn et al., 2018), or classifying argumentative functions (Paul et al., 2020). In this work, we leverage commonsense knowledge in the form of single- and multihop knowledge paths for establishing connections between concepts from different sentences in texts, and show that these

paths can *explicate* implicit information conveyed by the text. Connections can either be direct, e.g. given the sentences *The car was too old* and *The engine broke down*, the concepts *car* and *engine* can be linked with a direct relation (singlehop path) $car \rightarrow \text{HASA} \rightarrow engine$; or indirect – here intermediate concepts are required to establish the link, as between *Berliners produce too much waste* and *Environmental protection should play a more important role*, where the link between *waste* and *environmental protection* requires a multihop reasoning path: $waste \rightarrow \text{RECEIVESACTION} \rightarrow recycle \rightarrow \text{PARTOF} \rightarrow environmental\ protection$.

We show that two complementary model types can be combined to solve the two subtasks: (i) for predicting singlehop paths between concepts, we propose a *relation classification* model that is very precise, but restricted to direct connections between concepts; (ii) for constructing longer paths we rely on a *target prediction* model that can generate intermediate concepts and is thus able to generate multihop paths. However, the intermediate concepts can be irrelevant or misleading. To our knowledge, prior work has applied *either* relation classification *or* target prediction models. We propose CO-NNECT, a framework that establishes Commonsense knowledge paths for CONNECTING sentences by *combining* relation classification and target prediction models, leveraging their strengths and minimizing their weaknesses. With CO-NNECT, we obtain *high-quality knowledge paths* that explicate implicit knowledge conveyed by the text.

We focus on commonsense knowledge in ConceptNet (Speer et al., 2017), a knowledge graph (KG) that represents concepts (words or phrases) as nodes, and relations between them as edges, e.g., $\langle oven, \text{USED FOR}, baking \rangle$. As instances of the model types we use COREC (Becker et al., 2019), a multi-label relation classifier that predicts *relation*

types and that we enhance with a pretrained language model; and COMET (Bosselut et al., 2019), a pretrained transformer model that learns to generate *target concepts* given a source concept and a relation. In contrast to models that retrieve knowledge from static KGs (Mihaylov et al., 2018; Lin et al., 2019), both models are fine-tuned on ConceptNet and applied *on the fly*, to dynamically generate knowledge paths that generalize beyond the static knowledge, allowing us to predict unseen knowledge paths. We compare our models to a baseline model that solely relies on static KGs.

We evaluate our framework on two English argumentative datasets, IKAT (Becker et al., 2020) and ARC (Habernal et al., 2018), which offer annotations that explain implicit connections between sentences. While knowledge paths have been widely used in NLP downstream tasks, a careful evaluation of these paths has not received much attention. As a central contribution of our work, we address this shortcoming by designing manual and automatic settings for path evaluation: we evaluate the relevance and quality of the paths and their ability to represent implicit knowledge in an annotation experiment; and we compare the paths to the annotations of implicit knowledge in IKAT and ARC, using automatic similarity metrics.

Our main contributions are: i) we propose CO-NNECT, a framework that combines two complementary types of knowledge path prediction models that have previously only been applied separately;¹ ii) we show that commonsense knowledge paths generated with CO-NNECT effectively represent implied knowledge between sentences; iii) we propose an evaluation scheme that measures the quality of the knowledge paths, going beyond many approaches that use knowledge paths for downstream applications without analyzing their quality.

2 Related Work

In this work we combine relation classification and target prediction for generating commonsense knowledge representations over text. **Relation classification** covers a range of methods and learning paradigms for representing relations. A variety of neural architectures such as RNNs (Zhang et al., 2018), CNNs (Guo et al., 2019), sequence-to-sequence models (Trisedya et al., 2019) or language models (Wu and He, 2019) achieved state-

of-the-art results. To our knowledge, Becker et al. (2019) is the only work that proposed a relation classification model specifically for ConceptNet relations, which we adapt for our work. Besides COMET (Bosselut et al., 2019), the model used in our approach, Saito et al. (2018) perform **target prediction** on ConceptNet using an attentional encoder-decoder model. They improve the KB completion model of Li et al. (2016) by jointly scoring triples and predicting target concepts.

Utilizing commonsense knowledge paths.

When using commonsense knowledge for question answering (Mihaylov et al., 2018), commonsense reasoning (Lin et al., 2019) or NLI (Kapanipathi et al., 2020), most approaches rely on paths retrieved from *static* knowledge resources. In contrast, we propose a framework that in addition makes use of *dynamic* knowledge provided by language models. Few other models have used knowledge paths **dynamically**, e.g. Paul et al. (2020), who enrich ConceptNet on the fly when classifying argumentative functions.

Wang et al. (2020) make use of **language models** for question answering. They generate multi-hop paths by sampling random walks from ConceptNet and finetune a language model on these paths to connect question and answers, improving accuracy on two question answering benchmarks. Bosselut et al. (2021) generate knowledge paths using a language model for zero-shot question answering, which they use to select the answer to a question, surpassing performance of pretrained language models on SocialIQA (a multiple-choice question answering dataset for probing machine’s emotional and social intelligence in a variety of everyday situations). Similarly, Chang et al. (2020) incorporate knowledge from ConceptNet in pretrained language models for SocialIQA. They extract keywords from question and answers, query ConceptNet for relevant triples, and incorporate them in their language models via attention. Their evaluation shows that their knowledge-enhanced model outperforms knowledge-agnostic baselines. Finally, Paul and Frank (2020) propose an attention model that encodes commonsense inference rules and incorporates them in a transformer based reasoning cell, taking advantage of pretrained language models and structured knowledge. Their evaluation on two reasoning tasks shows that their model improves performance over models that lack external knowledge. Hence, none of these sys-

¹The code for our framework can be found here: <https://github.com/Heidelberg-NLP/CO-NNECT>.

tems *directly* evaluates the **quality** of the generated paths, but measure the effectiveness of commonsense knowledge *indirectly* by evaluation on downstream tasks. We will address this shortcoming in our work by carefully evaluating the quality of the generated paths.

3 Enriching Texts with Commonsense Knowledge Paths

This section describes CO-NNECT, the framework we propose for enriching texts with commonsense knowledge, by establishing relations or paths between concepts from different sentences. Towards this aim, we apply relation classification and target prediction models in combination. We first characterize differences between the two model types and their instantiations, COREC-LM and COMET, describe how we adapt them to our task and evaluate them on ConceptNet to assess their performance (§3.1). We then show how we utilize the models to establish connections between concepts in texts (§3.2) and present a baseline model that uses ConceptNet as a static KG to establish commonsense knowledge paths (§3.3).

3.1 Comparing and Evaluating Model Types

Relation classification and target prediction both aim at representing relational commonsense knowledge, but the respective task settings are fundamentally different. We choose two models that have been developed for representing commonsense knowledge in CN: COREC, a relation classification and COMET, a target prediction model.

Relation Classification with COREC-LM. A relation classifier is ideally suited to predict *direct* relations between concepts, hence we can apply COREC (Becker et al., 2019), an open-world multi-label relation classification system, for this task. Given a pair of concepts c_s, c_t from sentences, it predicts one or several relations r_i from a set of relation types R_{CN} that hold between c_s and c_t . We enhance the original neural model with the pretrained language model DistilBERT (Sanh et al., 2019) to construct a classifier we call COREC-LM. We finetune this model on ConceptNet by masking the relations and use sigmoid as output layer to model the probability of each relation independently, accounting for ambiguous relations in CN.

Target Prediction with COMET. To generate multihop paths that include (possibly novel) intermediate concepts, we apply COMET (Bosselut

et al., 2019), a transformer encoder-decoder based on GPT-2 (Radford et al., 2019). Input to the model is a source concept c_s and a relation r_i . Then the pretrained language model is finetuned using ConceptNet as labelled train set for the task of generating new concepts. Depending on the beam size, COMET can propose multiple targets per input instance.

Datasets. To compare model performances, we evaluate COREC-LM and COMET on the **CN-100k** benchmark dataset of Li et al. (2016), which is based on the OMCS subpart of ConceptNet. The dataset comprises 37 relation types such as ISA, PARTOF or CAUSES and contains 100k relation triples in the train set and 1200 in the development and the test set, respectively. CN-100k contains a lot of infrequent relations which are hard to learn and often overspecific (e.g. HASFIRST-SUBEVENT), and hence not useful for establishing high quality relations and paths between concepts. We therefore extract a subset that contains all triples of the 13 most frequent relations (**CN-13**).² CN-13 covers 90,600 triples for training, 1080 triples for development, and 1080 triples for testing.

Since our application task requires that the relation classifier also learns to detect that a given concept pair is *not* related, we extend the data for training and testing COREC-LM with a RANDOM class that contains unrelated concept pairs, which we add to CN-100k and CN-13.³

PoS Sequence Filtering. We apply a type-based PoS sequence filtering for COREC-LM and COMET, where the type is dependent on the predicted relation. The relation ISA, for example, is supposed to connect two noun phrases; in contrast, HASPREREQUISITE typically relates two verb phrases. We determine frequent PoS sequence patterns for specific argument types from the ConceptNet resource and use them to filter relation and path predictions.

Metrics. We evaluate COREC-LM in terms of weighted F1-scores, precision and recall, which is its genuine evaluation setting. For COMET we report precision scores for the first prediction with highest confidence score (hits@1); we further report hits@10 which gives information if the correct

²These are: ATLOCATION, CAUSES, CAPABLEOF, ISA, HASPREREQUISITE, HASPROPERTY, HASUBEVENT, USEDFOR, CAUSESDESIRE, DESIRES, HASA, MOTIVATED-BYGOAL and RECEIVESACTION.

³For details about the construction of the RANDOM class, cf. Appendix.

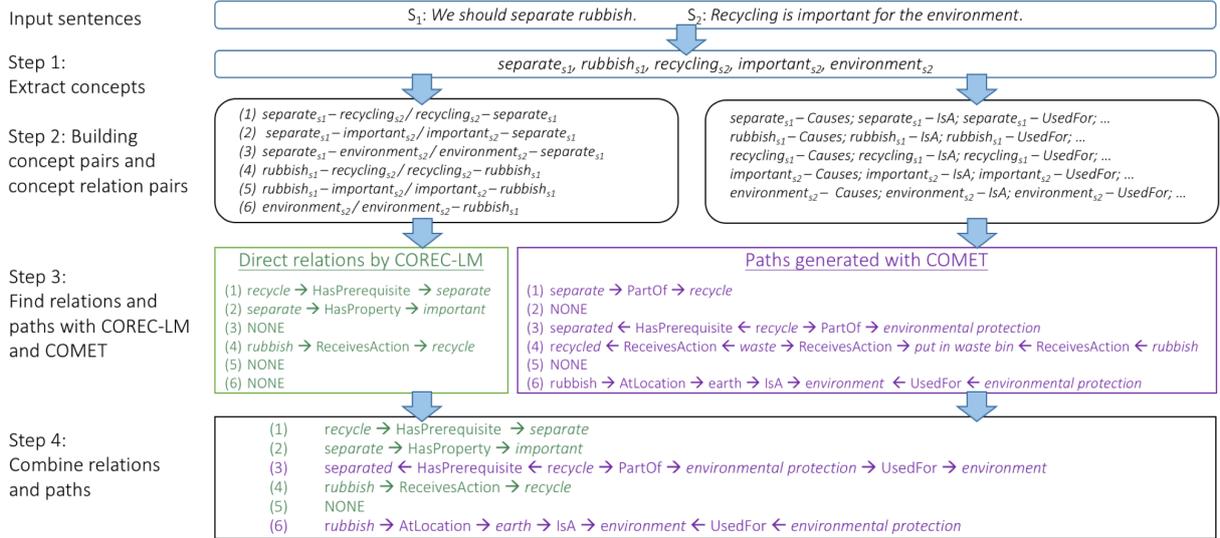


Figure 1: Our framework CO-NNECT: It finds single- and multihop paths between concepts, as explicitations of implicit knowledge that connects sentences.

triple is included in the first ten predictions (which will be important since we later use a beamsize of 10 for generating paths). In addition, we report accuracy using the Bilinear AVG model of Li et al. (2016) (COMET’s genuine evaluation setting), which is trained on CN-100k and produces a probability for a generated relation triple to be correct. Following Bosselut et al. (2019), we apply a beamsize of 1 and a threshold at 0.5 for judging a triple as correct.

Model Performances. COREC-LM achieves high F1-scores on CN-100k (76.5) and CN-13 (86.0).⁴ Scores are significantly lower when adding the RANDOM class (-7pp on CN-100k&CN-13), indicating that detecting unrelated concept pairs is not trivial. The results show that a strength of COREC-LM is its precision (90.1/CN-100k; 88.2/CN-13) – which we will leverage when combining models. COMET achieves high accuracy scores (92.3/CN-100k; 96.3/CN-13) according to the bi-score. For the much stricter metric hits@1 which judges a triple as correct only if it matches the respective triple in the testset, much lower scores are achieved (25/CN-100k; 23.5/CN-13), which is evident given the wide range of possible target generations. Higher scores for hits@10 (65.3/CN-100k; 65.9/CN-13) show that the chance for correct predictions significantly rises with increasing beam size.

In sum, COREC-LM and COMET both aim

⁴The original version of COREC (Becker et al., 2019) achieves F1 of 53.31/CN-100k; 72.33/CN-13.

at learning commonsense knowledge representations, but tackle different tasks and have different strengths and weaknesses. COREC-LM is very precise in its predictions, but is restricted to predicting *direct* relations between two given concepts. COMET is more powerful since it can genuinely generate *novel* target concepts and thus can generate multihop paths. However, it tends to be more imprecise, and bears the risk of generating *irrelevant or noisy* concepts. Hence, a combination of models seems beneficial, to predict high-quality single- and multihop paths between concepts.

3.2 Establishing Connections Using Relation Classification and Target Prediction

In the following we describe how we combine and apply COREC-LM and COMET in a joint framework, CO-NNECT, to establish high-quality knowledge paths between sentences. An overview is given in Fig. 1. In the first step we **extract relevant concepts** from the text. For this we integrate the concept extraction tool COCO-EX (Becker et al., 2021a), which extracts meaningful concepts from texts and maps them to concept nodes in CN, considering all surface forms.

Linking Concepts with Direct Relations. We construct all possible pairs of concepts extracted from S_1 and S_2 by taking the cross product $c_s \times c_t$, where c_s is a concept from S_1 , and c_t a concept from S_2 (Fig. 1, Step 2, left). We then apply COREC-LM trained on CN-13+RANDOM with a tuned threshold of 0.9 for predicting which rela-

tion $r_i \in R_{CN13}$ holds between the concept pairs, or whether no relation holds (RANDOM) (cf. Fig. 1, Step 3 (left) for examples).

Linking Concepts with Multihop Paths. COMET requires as input a source concept and a relation. For each concept pair c_s, c_t we build such inputs by combining c_s with each relation $r_i \in R_{CN13}$, yielding 13 pairs c_s, r_i which we input for target prediction (Step 2, right). To discover relation chains starting from S_2 , we apply the same process to c_t , using c_s as target concept. We also include *inverse relations*, which gives us greater flexibility for connecting entities, i.e., paths are allowed to contain inverted triplets (e.g. *baking* ← USED FOR ← *oven* → AT LOCATION → *kitchen*). To this end, we switch the order of concept pairs within a given relation r_i , relabel the relation as r_i^{-1} , and add the inverted relation pair to COMET’s training set.

Forward Chaining. For all pairs in the cross-product $c_s \times c_t$, for each input c_s, r_i and c_s, r_i^{-1} we generate the 10 most confident concepts c_{t_i} with COMET (beamsize 10) trained on CN-13 including inverse triples. We continue with all paths where the generated concept c_{t_i} has minimum cosine distance of 0.7 to the respective target concept c_t . We generate the next hop by using each c_{t_i} as a new source concept, combine it with each of the 13 original and inverse relations, generate novel target predictions, and again compare to the target concept. This similarity comparison guides the forward chaining process towards the chosen target concepts and helps detecting contextually relevant paths. We use ConceptNet numberbatch embeddings for the encoding of concepts; for multiword concepts we average the embeddings of all non-stopwords.

Terminating Paths. We terminate path generation as soon as the similarity between c_{t_i} and c_t is higher than 0.95 – here we expect the two concepts to express the same meaning. We restrict the path length to 3 hops and consider only completed paths for evaluation (Step 3, right in Fig. 1).

Combining Approaches. With our framework CO-NNECT we leverage the potential of the complementarity of the two model types by combining COREC-LM and COMET in a straightforward way. Our hypothesis is that a system that admits both single and multihop connections for establishing links between concepts offers the greatest flexibility. We further hypothesize that direct relations should be

preferred over indirect multihop paths, since the latter could include irrelevant or misleading intermediate nodes. Thus, we discard all multihop paths for each concept pair for which COREC-LM predicted a direct connection (Fig. 1, Step 4, pair 4). If COMET *and* COREC-LM produce a singlehop path, we also prefer COREC-LM’s prediction, relying on the model’s high precision (pair 1 in Step 4). We keep the paths generated by COMET for concept pairs for which no direct relation could be established (i.e., COREC-LM predicted RANDOM or no prediction above its threshold, pair 3&6), assuming that in such cases intermediate concepts are required to establish a link. If only one of the models establishes a link, we keep this connection (pair 2), and if none of the models finds a link, we assume that the concepts are not (closely) connected (pair 5).

3.3 Static Baseline Model

We compare COREC-LM and COMET against the model of Paul and Frank (2019) that uses ConceptNet as a static KG. The system extracts paths between pairs of concepts from sentence pairs, hence conforms well to our setting. Following Paul and Frank (2019), starting from concepts in a sentence pair (§3.2), we construct a subgraph $G' = (V', E')$ of the ConceptNet graph, where V' comprises all concepts c_i in $\langle S_1, S_2 \rangle$. The system then finds all shortest paths p' from ConceptNet that connect any concept pairs in V' , and includes them in G' . It then includes, for any concepts in G' , all directly connected concepts from ConceptNet together with their edges. This yields a small sub-graph from ConceptNet that contains concepts and relations relevant for capturing conceptual links across the sentence pair. To select *relevant* paths, G' is filtered by computing scores for vertices and paths using PageRank and Closeness Centrality score, and we constrain path lengths to 3 hops.

4 Revealing Implicit Knowledge through Knowledge Paths: Experiments and Evaluation

In this section we evaluate the paths generated by our proposed models. We first present our datasets and statistics on established connections (§4.1), and then evaluate the quality of the paths manually (§4.2) and automatically (§4.3).

Sentence 1	Sentence 2	COREC-LM	COMET	Ranked Subgraphs	Gold Implicit Information
The morning-after pill ought to be available over-the-counter in pharmacies.	Even though with contraceptives there can be mistakes and problems.	<i>pill</i> IsA <i>contraceptive</i>	<i>pill</i> IsA <i>contraceptive</i>	<i>pill</i> IsA <i>contraceptive</i>	The morning-after pill is a contraceptive.
I as an employee find it very practical to be able to shop at least on weekends.	Plus, the state wants me to spend my money.	NONE	<i>spend</i> HasSubevent <i>go to mall</i> HasSubevent <i>shop</i>	<i>spend</i> Derived-From <i>spend money</i> RelatedTo <i>shop</i>	When someone is shopping, he is spending money.
Today I will delete my Facebook account.	I'm tired of my data boosting the market value of that corporation.	<i>Facebook</i> HasA <i>data</i>	NONE	<i>Facebook</i> RelatedTo <i>book</i> RelatedTo <i>data</i>	With a facebook account you make your data available.
Felons should be allowed to vote.	A person who stole a car at 17 should not be barred from being a full citizen for life.	NONE	<i>felon</i> CapableOf <i>steal</i>	NONE	Grand theft auto is a felony.
Comment sections have not failed.	With good moderation, useful and interesting discussions can be had.	<i>comment</i> Causes <i>discussion</i>	<i>comment</i> At-Location <i>debate</i> IsA <i>discussion</i>	NONE	Moderation is usually good.

Figure 2: Example generations from our three model types (first three instances from IKAT, last two from ARC).

4.1 Datasets and Statistics

The **IKAT** dataset (Becker et al., 2020) is based on the English Microtexts Corpus of short argumentative texts (Peldszus and Stede, 2016). For all sentence pairs that are adjacent or argumentatively related, annotators added the implicit knowledge that connects them, using short sentences. **IKAT** contains 719 such sentence pairs, from which we extracted 60,294 concept pairs. The **ARC** dataset (Habernal et al., 2018) contains arguments taken from online discussions in English, consisting of a claim and a premise, and an annotated implicit warrant that explains why the claim follows from the premise. We evaluate our models on the **ARC** test set that comprises 444 argument pairs, from which we extracted 21,898 concept pairs; and the corresponding warrants.

Example generations for both datasets from our three model types – COREC-LM, COREC, and ranked CN-graphs – appear in Fig. 2, where the first three sentence pairs come from **IKAT**, and the last two from **ARC**.

Number of links and hops. Table 1 gives statistics of the paths generated between concepts from sentence pairs from **IKAT** and **ARC** using our different models. We find that COREC-LM finds relations between around 22k from 66k concept pairs in **IKAT**, while COMET only generates paths between 3,660 pairs. This can be explained by the very high similarity threshold we imposed for guiding the forward chaining process towards the target concept, since our motivation was not to generate as *many* paths as possible, but paths that are *meaningful* and contextually *relevant*. When combining paths from COMET and COREC-LM, we find links for more than 24k concept pairs in **IKAT**. The highest number of links is established by ranking CN-subgraphs (50k linked concept pairs). For

ARC, which contains 22k concept pairs, COREC-LM finds links between around 10k and COMET around 2k concept pairs, while almost 15k pairs can be connected using ranked CN-graphs. In both datasets, the ranked CN-graphs contain on average 2.1 **hops** (relations) per path, while the paths generated by COMET are shorter (1.4 on **IKAT**/1.5 on **ARC**). In fact, COMET establishes many direct relations (69% of all paths are single hops), whereas the ranked CN-graphs are mostly two- (49%) or three-hop paths (37%).

Replacing Vague Relations in CN-Graphs. We find that in contrast to COREC-LM and COMET, the ranked CN-graphs are constructed using mostly the very general relation **RELATEDTO** (71%/IKAT; 72%/ARC), followed by the likewise vague relation **HASCONTEXT** (8% in both datasets).⁵ For determining the impact of vague relations on path quality, we replace all **RELATEDTO** and **HASCONTEXT** relations in the ranked CN-graphs with relations predicted by COREC-LM (trained on CN-13, threshold 0.9). For **IKAT**, we replace 43.4% of all **RELATEDTO** and 46.2% of all **HASCONTEXT** instances, in **ARC** we replace 70.7% of all **RELATEDTO** and 37% of all **HASCONTEXT** relations. We use this version when evaluating paths, in addition to the original ranked CN-graphs.

4.2 Manual Evaluation of Path Quality

Our statistics showed that most links between concepts can be revealed using knowledge paths retrieved from ConceptNet as a static KG, whereby these paths tend to contain multiple hops and a high amount of vague relations. Fewer links are established using the dynamic models COREC-LM and COMET, which produce shorter paths using

⁵For details on relation distributions cf. Appendix.

		COR	COM	CONN	CN
IKAT	linked pairs	21,934	3,660	24,063	50,003
	avg. hops	1	1.4	1.1	2.1
ARC	linked pairs	9,844	1,826	10,828	14,940
	avg. hops	1	1.5	1.1	2.1

Table 1: Statistics of paths generated by COREC-LM, COMET, their combination (CO-NNECT), and ranking CN-graphs (CN): number of concepts pairs between which a link was found, and average number of hops per path.

only specific relation types from CN-13. Since our aim is to construct high-quality, meaningful knowledge paths that help to explain implicit information (rather than establishing as many links as possible), we now examine the quality and relevance of the knowledge paths. We set up an annotation experiment, providing annotators with 100 sentence pairs from each dataset, with marked concepts (one from S_1 and one from S_2) and the path generated between these concepts by (i) COREC-LM, (ii) COMET, (iii) ranked paths from CN, and (iv) ranked paths with replaced vague relations (CN-r).

Annotation Setup. For each sentence pair, our annotators evaluated if 1) the path is a meaningful and relevant explanation for the connection between the two sentences (very relevant/relevant/neutral/not relevant/misleading); if 2) the path represents implicit information not explicitly expressed in the sentences (yes/no); and 3) which model generates the path that is most helpful and expressive for understanding the connection between the sentences. 4) To evaluate the *combination* of COREC-LM and COMET in CO-NNECT, we generate a subset for each dataset that includes all sentence/concept pairs for which COREC-LM predicted a singlehop path *and* COMET generated a multihop path (10 pairs per subset). For these instances we ask in addition whether the multihop paths include unrelated, unnecessary or uninformative intermediate nodes (yes/no), misleading intermediate nodes (yes/no); or intermediate nodes that are important for explaining the connection and missing in the direct relation predicted by COREC-LM (yes/no).⁶ Annotations were performed by two annotators with a linguistic background. We measure IAA using Cohen’s Kappa and achieve an agreement of 81%. Remaining conflicts were

⁶The annotation manual together with example annotations can be found here: <https://github.com/Heidelberg-NLP/CO-NNECT/blob/main/manual.pdf>

	IKAT				ARC			
	COR	COM	CN	CN-r	COR	COM	CN	CN-r
Predictions	74	64	88	88	78	60	76	76
Relev. +2	70	50	36	40	63	49	30	34
+1	19	27	22	24	25	28	28	32
0	8	18	27	21	8	9	29	22
-1	3	5	10	10	2	6	4	3
-2	0	0	5	5	2	8	9	9
Impl. yes	80	78	57	67	87	81	57	62
Knowl. no	20	22	43	33	13	19	43	38
Best Link	65	64	28	34	76	70	7	14

Table 2: Manual evaluation of paths from COREC-LM, COMET, ranked CN-graphs (CN), and CN-graphs with replaced vague relations (CN-r); all numbers in %.: How many concept pairs could be linked (line 1), are the links relevant and meaningful (2-6), do the links represent implicit knowledge (7-8), how often a link was chosen to be most helpful for understanding the connection (9).

resolved by an expert annotator.

Results. Table 2 shows the results of our annotation experiment. On **IKAT**, 89% of the paths established by COREC-LM and 77% of the relations predicted by COMET were annotated as very relevant (+2) or relevant connections (+1), which only applies for 58% of the ranked CN-paths. 15% of the ranked CN-paths were annotated as not relevant (-1) or misleading (-2), which can be explained by noisy intermediate nodes; and 27% as vague (0), which can be explained by the large amount of un-specific relations. When replacing RELATEDTO and HASCONTEXT (CN-r), the amount of paths annotated as vague slightly decreases, and the amount of paths labelled as relevant and very relevant increases.

Moreover, paths generated by COREC-LM and COMET were found to yield better implicit knowledge representations than ranked CN-paths (line 8-9, Table 2), while we find that replacing vague relations in the CN-paths improves their ability of representing implicit knowledge. Finally, 65% of relations predicted by COREC-LM and 64% of paths generated by COMET were chosen as explaining the connections between sentences best, which is only the case for 28% of the CN-paths, and slightly better for the replaced version of the CN-paths (34%).

On **ARC**, the high amount of CN-paths annotated as vague (29%) again indicates uninformative connections and can be reduced when replacing vague with more specific relations. Relations predicted by COREC-LM were found to be less relevant for connecting sentences in ARC than in IKAT, but 87% of them were evaluated as appropriate expressions of implicit knowledge. 76% of the

relations predicted by COREC-LM were evaluated as best connections, which applies only for 7% of CN-paths and 14% of CN-paths with replaced relations. For COMET we find overall comparable results between IKAT and ARC.

Regarding the **combination** of COREC-LM and COMET addressed with question 4, according to our annotators 50% of the multihop paths in the IKAT subset include misleading nodes and *all* of them include irrelevant or uninformative nodes. Still, compared against the direct relations predicted by COREC-LM, annotators state for 30% of the multihop paths from COMET that they contain intermediate concepts that are important for explaining the connection. On the ARC subset, 40% of the multihop paths include misleading and 60% include irrelevant nodes, and only 20% contain important intermediate concepts that are missing in the direct relation. For each subset, annotators preferred the shorter path over the multihop path in 90% of the given sentence pairs. Comparing singlehop paths generated by COMET to direct relations predicted by COREC-LM for the same concept pairs, our annotators preferred the relation predicted by COREC-LM in 64% of the cases, in 29% the link was annotated as equally good, and only in 7% COMET’s generation was preferred.

To summarize, according to our manual evaluation, the dynamic models COMET and COREC-LM are better suited for generating meaningful knowledge paths that express implicit knowledge between sentences than ranked paths from the static CN knowledge graph, even though replacing vague by more specific relations slightly improves results. When comparing multihop paths to direct relations established between the same concept pairs, we find that longer paths tend to contain irrelevant or even misleading nodes, and that direct relations are preferred by human annotators. These findings support our proposed joint framework CO-NNECT, which gives preference to direct relations and utilizes multihop paths only if no direct connection between concepts can be revealed.

4.3 Automatic Evaluation Against Gold

Our goal is to generate meaningful paths that convey implicit knowledge between sentences. In our automatic evaluation we compare the set of model-generated paths between all concept pairs from two related sentences to the implicit knowledge annotated in IKAT and ARC for these sentences, using

similarity metrics.

Since the generated relation and path representations differ from the annotated natural sentences, we approximate a common representation as follows: We **encode the golden annotations of implicit knowledge** – usually short sentences – using three settings: (i) **Silver Paths**: we encode their relational knowledge, by extracting all concepts from each golden implicit knowledge sentence (*My dog has a bone* \rightarrow *dog, bone*) using the CN-extraction tool COCO-EX (Becker et al., 2021a), and predict the relations between them using COREC-LM, trained on CN-13 (*dog, HASA, bone*). If a sentence contains more than two relations, we concatenate the predicted relation triples. (ii) IKAT provides manual annotations of ConceptNet relations for the golden implicit knowledge sentences, which we use as **Gold Paths** (*The tree is in the garden* \rightarrow *tree ATLOCATION garden*). (iii) **Gold-NL**: Here we use the implicit knowledge (in natural language) as provided in the datasets: IKAT’s implicit knowledge sentences and ARC’s implicit warrants.

For **encoding the generated paths** we apply two settings: (i) we concatenate all concepts and relations within the paths; (**Generated Paths**) and (ii) we translate the relation triples and paths to (pseudo) natural language using templates provided by ConceptNet (e.g. c_s CAUSES $c_t \rightarrow$ *The effect of c_s is c_t* ; **Generated Paths-NL**).

We apply two **automatic similarity metrics**, comparing (a) Generated vs. Silver Paths, (b) Generated Paths-NL vs. Gold-NL, and (c) Generated vs. Gold Paths (only IKAT). (i) We encode each representation as described above using ConceptNet numberbatch embeddings (Speer et al., 2017) (for multiword concepts we average the embeddings of all non-stopwords), and compute cosine similarity between them, and (ii) we use BERTScore F1 (Zhang et al., 2020) to compare representations, which computes string similarity using contextualized embeddings. Both metrics lie in $[-1, 1]$.

Results. Table 3 shows that the paths generated by combining COREC-LM and COMET in our framework CO-NNECT achieve the highest similarity scores according to Numberbatch-Cosim on **IKAT** in setting (a) and (b), while for (c) we get the highest Cosim scores for ranked CN-graphs with replaced vague relations. According to BERTScore, either COREC-LM (setting a) or COMET (setting b) applied separately, or both applied in combination (setting c) achieve highest

	COR	COM	CONN	CN	CN-r
(a) Generated Paths vs. Silver Paths					
IKAT	.61/.85	.54/.82	.62/.84	.57/.78	.58/.80
ARC	.41/.84	.39/.82	.42/.86	.40/.77	.40/.78
(b) Generated Paths-NL vs. Gold-NL					
IKAT	.69/.81	.65/.83	.70/.81	.65/.75	.69/.76
ARC	.72/.81	.66/.82	.72/.81	.71/.75	.77/.76
(c) Generated Paths vs. Gold Paths					
IKAT	.57/.78	.49/.78	.58/.79	.66/.73	.67/.74

Table 3: Comparing generated paths to implicit knowledge annotations on IKAT and ARC, measured by Cosim/BERTScore (F1).

results on IKAT. On **ARC**, CO-NNECT achieves both highest Cosim and BERTScores in setting (a), while in (b) we get the best scores for CN-r according to Cosim, and the best scores for COMET according to BERTScore.

Summarizing our insights from automatic evaluations, we find that **COREC-LM** achieves high scores when applied separately *or* in combination with COMET (CO-NNECT). **COMET** applied in isolation does not yield the highest scores, but helps to boost COREC-LM’s performance in the joint CO-NNECT framework. **Ranked CN-graphs** achieve highest Cosim in two settings/datasets (ARC–b; IKAT–c), but we do not find significant improvements when replacing vague relations in CN-graphs (expect for Cosim in setting b). This can be explained by the fact that even though many RELATEDTO and HASCONTEXT instances could be replaced, for both datasets a large amount of vague relations still remain (56.6% of RELATEDTO/53.2% of HASCONTEXT in IKAT; 29.3% RELATEDTO/63% HASCONTEXT in ARC). Therefore, the vague relation types in the CN-graphs still remain problematic when representing implicit knowledge.

When comparing our **manual** evaluation results to the **automatic** scores, we find that the generations that were manually evaluated as most relevant and meaningful explanations of implicit knowledge are not always highest-ranked by automatic metrics, which points to two limitations of our automatic evaluation: Besides well-known issues regarding the reliability, interpretability, and biases of automatic metrics (Callison-Burch et al., 2006), we evaluate the generated paths against an annotated *reference* – paths or sentences – which is often only one among several valid options for expressing the implicit knowledge (cf. Becker et al. 2017). This means that a generated path may still be a relevant

explicitation of implicit information, even if *not* similar to the reference. Hence, automatic scores are to be considered with caution.

5 Conclusion

Our work aims to leverage commonsense knowledge in the form of single and multihop paths, to establish knowledge connections between concepts from different sentences, as a form of explicitation of implicitly conveyed information. We combine existing relation classification and target prediction models in a dynamic knowledge prediction framework, CO-NNECT, utilizing language models finetuned on knowledge relations from ConceptNet. We compare against a path ranking system that employs static knowledge from ConceptNet as a baseline and evaluate the quality of the obtained paths (i) through manual evaluation and (ii) using automatic similarity metrics, by comparing generated paths to annotations of implicit knowledge in two argumentative datasets. Our evaluations show that we obtain the highest number of connections from the static ConceptNet graph, however, they are often noisy due to unrelated intermediate nodes, and – even after replacements – still contain many unspecific relations. Our framework CO-NNECT, instead, combines relation classification and target prediction, leveraging the *high precision* of the former, and the *ability to perform forward chaining* of the latter, and obtain high-quality, meaningful and relevant knowledge paths that reveal implicit knowledge conveyed by the text, as shown in a profound manual evaluation experiment.

We believe that CO-NNECT is a useful framework which can be applied for different tasks, such as enriching texts with commonsense knowledge relations and paths, for dynamically enriching knowledge bases, or for building knowledge constraints for language generation. In Becker et al. (2021b) for example we inject single- and multihop commonsense knowledge paths predicted by CO-NNECT as constraints into language models and show that this improves the model’s ability of generating sentences that explicate implicit knowledge which connects sentences in texts. We furthermore believe that the paths established with CO-NNECT, which can provide explicitations of implicit knowledge, can be useful to enhance many other NLP downstream tasks, such as argument classification, stance detection, or commonsense reasoning.

Acknowledgements

This work has been funded by the DFG within the project ExpLAIN as part of the Priority Program “Robust Argumentation Machines” (SPP-1999). We thank our annotators for their contribution.

References

- Maria Becker, Katharina Korfhage, and Anette Frank. 2020. [Implicit Knowledge in Argumentative Texts: An Annotated Corpus](#). In *Proceedings of the Conference on Language Resources and Evaluation (LREC)*, pages 2316–2324, Online.
- Maria Becker, Katharina Korfhage, and Anette Frank. 2021a. COCO-EX: A Tool for Linking Concepts from Texts to ConceptNet. In *Proceedings of the Conference of the European Chapter of the Association for Computational Linguistics (EACL), Demo Papers*, Online.
- Maria Becker, Siting Liang, and Anette Frank. 2021b. Reconstructing Implicit Knowledge with Language Models. *Accepted at: Deep Learning Inside Out (DeeLIO): Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*.
- Maria Becker, Michael Staniek, Vivi Nastase, and Anette Frank. 2017. [Enriching Argumentative Texts with Implicit Knowledge](#). In *Applications of Natural Language to Data Bases (NLDB) - Natural Language Processing and Information Systems*, Lecture Notes in Computer Science, pages 84–96. Springer.
- Maria Becker, Michael Staniek, Vivi Nastase, and Anette Frank. 2019. [Assessing the Difficulty of Classifying ConceptNet Relations in a Multi-Label Classification Setting](#). In *Proceedings of RELATIONS - Workshop on Meaning Relations between Phrases and Sentences*, pages 1–15, Gothenburg, Sweden.
- Antoine Bosselut, Ronan Le Bras, , and Yejin Choi. 2021. Dynamic neuro-symbolic knowledge graph construction for zero-shot commonsense question answering. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*.
- Antoine Bosselut, Hannah Rashkin, Maarten Sap, Chaitanya Malaviya, Asli Celikyilmaz, and Yejin Choi. 2019. [COMET: Commonsense transformers for automatic knowledge graph construction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4762–4779, Florence, Italy. Association for Computational Linguistics.
- Chris Callison-Burch, Miles Osborne, and Philipp Koehn. 2006. [Re-evaluating the role of Bleu in machine translation research](#). In *11th Conference of the European Chapter of the Association for Computational Linguistics*, Trento, Italy. Association for Computational Linguistics.
- Ting-Yun Chang, Yang Liu, Karthik Gopalakrishnan, Behnam Hedayatnia, Pei Zhou, and Dilek Hakkani-Tur. 2020. [Incorporating commonsense knowledge graph in pretrained models for social commonsense tasks](#). In *Proceedings of Deep Learning Inside Out (DeeLIO): The First Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 74–79, Online. Association for Computational Linguistics.
- Zhijiang Guo, Yan Zhang, and Wei Lu. 2019. [Attention guided graph convolutional networks for relation extraction](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 241–251, Florence, Italy. Association for Computational Linguistics.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. [The argument reasoning comprehension task: Identification and reconstruction of implicit warrants](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1930–1940, New Orleans, Louisiana. Association for Computational Linguistics.
- Pavan Kapanipathi, Veronika Thost, Siva Patel, Spencer Whitehead, Ibrahim Abdelaziz, Avinash Balakrishnan, Maria Chang, Kshitij Fadnis, Chulaka Gunasekara, Bassem Makni, Nicholas Mattei, Kartik Talamadupula, and Achille Fokoue. 2020. [Infusing knowledge into the textual entailment task using graph convolutional networks](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:8074–8081.
- Xiang Li, Aynaz Taheri, Lifu Tu, and Kevin Gimpel. 2016. [Commonsense knowledge base completion](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1445–1455, Berlin, Germany. Association for Computational Linguistics.
- Bill Yuchen Lin, Xinyue Chen, Jamin Chen, and Xiang Ren. 2019. [KagNet: Knowledge-aware graph networks for commonsense reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2829–2839, Hong Kong, China. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. [Can a suit of armor conduct electricity? a new dataset for open book question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391, Brussels, Belgium. Association for Computational Linguistics.

- Debjit Paul and Anette Frank. 2019. [Ranking and selecting multi-hop knowledge paths to better predict human needs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3671–3681, Minneapolis, Minnesota. Association for Computational Linguistics.
- Debjit Paul and Anette Frank. 2020. [Social commonsense reasoning with multi-head knowledge attention](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 2969–2980, Online. Association for Computational Linguistics.
- Debjit Paul, Juri Opitz, Maria Becker, Jonathan Kobbe, Graeme Hirst, and Anette Frank. 2020. [Argumentative Relation Classification with Background Knowledge](#). In *Proceedings of the International Conference on Computational Models of Argument (COMMA)*, pages 319–330, Online.
- Andreas Peldszus and Manfred Stede. 2016. An annotated corpus of argumentative microtexts. In *Argumentation and Reasoned Action: Proceedings of the 1st European Conference on Argumentation, Lisbon 2015 / Vol. 2*, pages 801–815, London. College Publications.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language Models are Unsupervised Multitask Learners](#).
- Itsumi Saito, Kyosuke Nishida, Hisako Asano, and Junji Tomita. 2018. [Commonsense knowledge base completion and generation](#). In *Proceedings of the 22nd Conference on Computational Natural Language Learning*, pages 141–150, Brussels, Belgium. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *EMC2: 5th Edition. Co-located with NeurIPS'19*.
- Robyn Speer, Joshua Chin, and Catherine Havasi. 2017. [Conceptnet 5.5: An open multilingual graph of general knowledge](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, pages 4444–4451. AAAI Press.
- Bayu Distiawan Trisedya, Gerhard Weikum, Jianzhong Qi, and Rui Zhang. 2019. [Neural relation extraction for knowledge base enrichment](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 229–240, Florence, Italy. Association for Computational Linguistics.
- Ekaterina Vylomova, Laura Rimell, Trevor Cohn, and Timothy Baldwin. 2016. [Take and took, gaggle and goose, book and read: Evaluating the utility of vector differences for lexical relation learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1671–1682, Berlin, Germany. Association for Computational Linguistics.
- Peifeng Wang, Nanyun Peng, Filip Ilievski, Pedro Szekely, and Xiang Ren. 2020. Connecting the dots: A knowledgeable path generator for commonsense question answering. *EMNLP Findings*, pages 4129–4140.
- Dirk Weissenborn, Tomas Kocisky, and Chris Dyer. 2018. Dynamic Integration of Background Knowledge in Neural NLU Systems. *ICLR 2018*.
- Shanchan Wu and Yifan He. 2019. [Enriching pre-trained language model with entity information for relation classification](#). In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM '19*, page 2361–2364, New York, NY, USA. Association for Computing Machinery.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *ICLR*.
- Xiaobin Zhang, Fucui Chen, and Ruiyang Huang. 2018. A Combination of RNN and CNN for Attention-based Relation Classification. In *Procedia Computer Science*, volume 131, pages 911 – 917.

APPENDIX

A Constructing the RANDOM Class for Training COREC-LM in an Open World Setting

Our downstream application task – finding connections between concepts – requires that our relation classifier also learns to detect that no direct relation holds between a given pair of concepts. We thus extend the data for training and testing COREC-LM with a RANDOM class which contains concept pairs that are not related which we add to CN-100k and CN-13 Instances for this class are generated similarly to [Vylomova et al. \(2016\)](#): 50% of them are opposite pairs which we obtain by switching the order of concept pairs within the same relation, and 50% are corrupt pairs, obtained by replacing one concept in a pair with a random concept from the same relation. Corrupt pairs ensure that COREC-LM learns to encode relation instances rather than simply learning properties of the word classes. We add these instances (2070 for training and 260 for development and testing, respectively) to CN-100k and CN-13 when training and evaluating in an open world setting.

	COREC-LM	COMET	CONNECT	CN Subgraphs
IKAT	ATLOCATION(25%) HASPROPERTY(20%) ISA(17%)	ISA(19%) HASA(18%) CAUSES(17%)	ATLOCATION(22%) ISA(17%) HASPROPERTY(16%)	RELATEDTO(71%) HASCONTEXT(8%) ISA(7%)
ARC	ATLOCATION(31%) ISA(18%) HASPROPERTY(14%)	ATLOCATION(22%) CAUSES(20%) HASA(18%)	ATLOCATION(27%) ISA(15%) HASPROPERTY(10%)	RELATEDTO(72%) HASCONTEXT(8%) ISA(7%)

Table 4: Most frequently used relations when constructing single and multihop knowledge paths using COMET, COREC-LM, their combination, and ranked subgraphs from CN.

B Relations Used for Constructing Single- and Multihop Paths

Table 4 lists the three most frequently used relations when constructing single and multihop knowledge paths using COMET, COREC-LM, their combination, and ranked subgraphs, respectively for the two datasets IKAT and ARC. The top three relations used by **COREC-LM** within both datasets are ATLOCATION, HASPROPERTY, and ISA. Interestingly, besides ISA and HASA, **COMET** frequently uses the only causal relation in the CN inventory CAUSES. In contrast to COREC-LM and COMET, the ranked CN-graphs are constructed using mostly the very general relation RELATEDTO, followed by the likewise vague relation HASCONTEXT. When excluding paths that contain RELATEDTO, only 2,551 connected concept pairs remain in IKAT and 6,858 in ARC.