

---

# Is Complex Query Answering Really Complex?

---

Cosimo Gregucci<sup>1</sup> Bo Xiong<sup>2</sup> Daniel Hernández<sup>1</sup> Lorenzo Loconte<sup>3</sup>  
Pasquale Minervini<sup>3,4</sup> Steffen Staab<sup>1,5</sup> Antonio Vergari<sup>3</sup>

## Abstract

Complex query answering (CQA) on knowledge graphs (KGs) is gaining momentum as a challenging reasoning task. In this paper, we show that the current benchmarks for CQA might not be as *complex* as we think, as the way they are built distorts our perception of progress in this field. For example, we find that in these benchmarks, most queries (up to 98% for some query types) can be reduced to simpler problems, e.g. link prediction, where only one link needs to be predicted. The performance of state-of-the-art CQA models decreases significantly when such models are evaluated on queries that cannot be reduced to easier types. Thus, we propose a set of more challenging benchmarks composed of queries that *require* models to reason over multiple hops and better reflect the construction of real-world KGs. In a systematic empirical investigation, the new benchmarks show that current methods leave much to be desired from current CQA methods.

## 1. Introduction

A crucial challenge in AI and ML is learning to perform *complex reasoning*, i.e., solving tasks that involve several intermediate steps and sub-goals to be completed. Complex query answering (CQA; Hamilton et al., 2018; Zhang et al., 2021; Arakelyan et al., 2021; Zhu et al., 2022) emerged as a way to measure complex reasoning over external knowledge bases, encoded as knowledge graphs (KGs; Hogan et al., 2021). For instance, to answer the query:

“Which actor performed in a movie filmed in  
New York City and distributed on Blue Ray?”  $(q_1)$

---

<sup>1</sup>Shared supervision <sup>1</sup>Institute for Artificial Intelligence, University of Stuttgart, Germany <sup>2</sup>Stanford University <sup>3</sup>School of Informatics University of Edinburgh Edinburgh, UK <sup>4</sup>Miniml.AI <sup>5</sup>University of Southampton, UK. Correspondence to: Cosimo Gregucci <cosimo.gregucci@ki.uni-stuttgart.de>.

over a KG such as FreeBase (Bollacker et al., 2008), one would need to *first* intersect the set of movies found on *Blue Ray* and the ones shot in *New York City*, and *then* link these intermediate candidate answers to another entity, i.e., an *actor* participating in it. However, if some links are missing in the KG, not all relevant actor entities can be immediately reachable. To deal with the unavoidable incompleteness of real-world KGs, ML methods were developed to solve CQA in the presence of missing links. The current state-of-the-art (SoTA) for CQA are *neural query answering* models (Arakelyan et al., 2021; Zhu et al., 2022; Arakelyan et al., 2023; Ren et al., 2023; Galkin et al., 2024b), mapping queries and KGs (i.e., entities and relation names) into a unified latent space. Performance measured on de-facto-standard benchmarks such as FB15k237 (Toutanova & Chen, 2015) and NELL995 (Xiong et al., 2017) suggests impressive progress achieved in recent years on CQA on queries having different structures, and hence posing apparently different levels of difficulty to be answered.

The difficulty of a benchmark relates to the size and structural complexity<sup>1</sup> of its queries, and several query “types” have been proposed (Ren et al., 2020), each involving a different combination of logical operators—conjunctions, disjunctions, and requiring to traverse a number of missing links that generally increases with the number of logical conditions imposed (Figures 1 and 2). For instance, the query  $q_1$  is an example of a “*2ilp*” query type, since it comprises an intersection of two entity sets (*2i*) followed by a path<sup>2</sup> of length one (*lp*). In this paper, we argue that *the perception of progress on CQA benchmarks has been distorted by implicit assumptions in these benchmarks*. We start by noting how both in FB15k237 and NELL995, the vast majority of queries of a complex type simplify to one simpler type, thanks to the fact that to answer them, one can leverage links already appearing in the training data. Figure 1 shows an example of how answers to the query  $q_1$  on FB15k237— which should require predicting *three* missing links – require the same effort associated with simpler types involving fewer links. In fact, in FB15k237, the answer “K. Dunst” can be retrieved by predicting just one missing link

---

<sup>1</sup>Not to be confused with the computational complexity of query answering (Dalvi & Suciu, 2012).

<sup>2</sup>Also known as *projection* in related works (Ren et al., 2023).

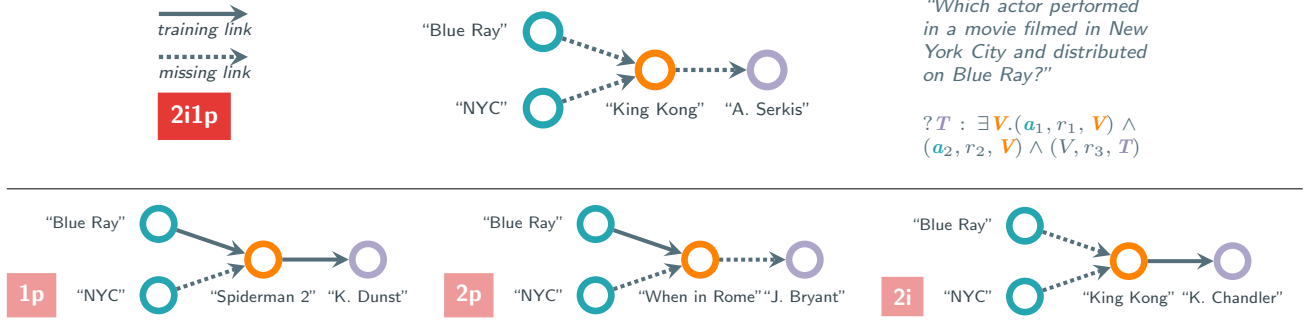


Figure 1. Query answers are not all equally hard when some links can be found in the training data as shown for the 2i1p query  $q_1$  and fragments of the KG FB15k237, where  $r_1 = \text{distributedVia}^{-1}$ ,  $r_2 = \text{locatedIn}^{-1}$  and  $r_3 = \text{performedIn}^{-1}$ . Its different answers can be obtained by traversing the training graph (continuous line) and predicting the missing links (dotted lines). (Top) Example answer that requires all links to be predicted. (Bottom) Example answers that require only a subset of the links to be predicted and that, therefore, can be reduced to the simpler types 1p, 2p, and 2i (see Sec. 2 and Figure 2).

when leveraging the training links. In this sense, a 2i1p query reduces to a "1p" query, i.e., the much simpler task of link prediction (see Sec. 2). Similarly, the answers "J. Bryant" and "K. Chandler" can be retrieved by predicting two links instead of three. Only the answer "A. Serkis" requires predicting three missing links, and thus follows the intuitive "hardness" associated to the type 2i1p, i.e., it is the hardest when no information is available. As we note here, there is a *spectrum of hardness*, predicting "K. Dunst" will be easier than answering "K. Chandler", which in turn is easier than saying "A. Serkis". However, current CQA benchmarks greatly overrepresent the easy query answers over the others – the "K. Dunst"-like answers are the vast majority – and this inflates performance as SoTA models are known to be prone to memorize training links. In this paper, we analyze this issue and propose a new set of benchmarks that cover the full spectrum of hardness in CQA.

**Contributions.** After revisiting the CQA task (Sec. 2) and highlighting how query types can simplify at test time in the presence of training links (Sec. 3), (C1) we show that up to 99.9% of the test queries can be reduced to simpler queries, with the majority of them (up to 98%) reduced to "one-step" link prediction problems (Sec. 4). This clearly inflates the reported performance. (C2) We re-evaluate previous SoTA approaches (Sec. 5), revealing that neural link predictors rely on memorized information from the training set. Furthermore, we show that the reported hardness of queries involving unions is only apparent. To better understand why performances are inflated, we introduce CQD-Hybrid, a hybrid CQA solver that combines classical graph matching with neural link predictors. (C3) We create new benchmarks for CQA (Sec. 6), FB15k237+H and NELL995+H from FB15k237 and NELL995, respectively, where we equally sample all intermediate query types one can consider by accounting for known links in the training data. To these, we build ICEWS18+H from the temporal KG

ICEWS18 (Boschee et al., 2015), where validation and test links are links that have been added to the KG *after* those in the training, to investigate more realistic sampling schemes to generate the train/validation/test splits. Lastly, we add the query types "4p" (a four-length path) and "4i" (a conjunction of four patterns), making them more challenging.

## 2. KGs and Complex Query Answering

**Knowledge graphs.** A KG is a graph-structured knowledge base where information is encoded as relationships between entities. More formally, a KG is a multi-relational graph  $G = (\mathcal{E}, \mathcal{R}, \mathcal{T})$ , where  $\mathcal{E}$  is a set of entities,  $\mathcal{R}$  is a set of relation names, and  $\mathcal{T} \subseteq \mathcal{E} \times \mathcal{R} \times \mathcal{E}$  is a set of links or *triples*, where each triple  $(s, p, o) \in \mathcal{T}$  represents a relationship of type  $p$  between a subject  $s$  and an object  $o$ . For instance, in a KG such as FreeBase, the fact that the movie "Spiderman 2" is distributed in Blue Ray is stored in the triple (Spiderman2, distributedVia, BlueRay) or equivalently (BlueRay, distributedVia<sup>-1</sup>, Spiderman2) where distributedVia<sup>-1</sup> is the inverse relation of distributedVia. Figure 1 shows examples of fragments of the KG FB15k237.

**The aim of CQA** is to retrieve a set of answers to a logical query  $q$  that poses conditions over entities and relation types in a KG. Following Arakelyan et al. (2021), we consider the problem of answering logical queries with a single *target* variable ( $T$ ), a set of constants including *anchor* entities ( $a_1, \dots, a_k \in \mathcal{E}$ ), given relation names ( $r_1, \dots, r_n \in \mathcal{R}$ ), and first-order logical operations that include conjunction  $\wedge$ , disjunction  $\vee$ , negation  $\neg$  and existential quantification  $\exists$ . Different queries are categorized into different *types* based on the structure of their corresponding logical sentence (Xiong et al., 2017). This taxonomy should imply that queries of the same type share the same "hardness", i.e., the level of difficulty to be answered, and different types correspond to tasks that map to more or less complex rea-




Figure 2. **Query structures we consider**, adapted from Ren & Leskovec (2020) and including *path* (**p**), *intersection* (**i**), *union* (**u**) structures. Sec. 2 for their logical formulation. See App. B for additional query structures including **negation**.

soning tasks. As we will show in the next sections, this is not the case, and to fully grasp the actual hardness to answer a query one needs to account for the “leakage” of intermediate training links. We now review the most common query types used in CQA. The simplest CQA task is *link prediction*, i.e., answering a query of the form:

$$?T : (a_1, r_1, T), \quad (1p)$$

that is, given an entity  $a_1$  (e.g., NYC) and a relation name  $r_1$  (e.g., `locatedIn`<sup>-1</sup>), find the entity that when substituted to  $T$  correctly matches the link in the KG (e.g., Spiderman2, see Figure 1). Instead of matching a single link, more complex queries involve matching *sub-graphs* in a KG (see Figure 2). Xiong et al. (2017) extend **1p** queries, and ask questions involving sequential *paths* made of two or three links, i.e.,

$$\exists V_1. (a_1, r_1, V_1) \wedge (V_1, r_2, T), \quad (2p)$$

$$\exists V_1, V_2. (a_1, r_1, V_1) \wedge (V_1, r_2, V_2) \wedge (V_2, r_3, T), \quad (3p)$$

where  $V_1, V_2$  denote variables to be grounded into entities appearing in the path. Moreover, multiple ground entities can participate in a conjunction, e.g., in queries such as:

$$?T : (a_1, r_1, T) \wedge (a_2, r_2, T), \quad (2i)$$

$$?T : (a_1, r_1, T) \wedge (a_2, r_2, T) \wedge (a_3, r_3, T), \quad (3i)$$

which represent the *intersection* of the target entity sets defined over two (**2i**) or three (**3i**) links. Path and intersection structures can be combined into more complex queries: for example, the natural language expression for query **q1** can be formalized as the formula

$$?T : \exists V_1. (a_1, r_1, V_1) \wedge (a_2, r_2, V_1) \wedge (V_1, r_3, T), \quad (2i1p)$$

involving one intersection followed by a length-one path.<sup>3</sup> See also the example in Figure 1. By inverting the order of operations, we obtain the query type “**1p2i**”:

$$?T : \exists V_1. (a_1, r_1, V_1) \wedge (V_1, r_2, T) \wedge (a_2, r_3, T). \quad (1p2i)$$

<sup>3</sup>In Ren et al. (2020) this type was referred to as “ip”. We explicitly mention the number of steps involved in a path or conjunction, as this is a factor of complexity. Analogously, “pi”, and “up” queries from Ren et al. (2020) are now **1p2i**, and **2u1p**.

Similarly to introducing conjunctions, we can consider disjunctions in queries, realizing the *union* query types which can be answered by matching one link or the other, such as

$$?T : (a_1, r_1, T) \vee (a_2, r_2, T), \quad (2u)$$

or by combining the union with the previous query types, e.g., obtained by combining **2u** and **1p**:

$$?T : \exists V_1. ((a_1, r_1, V_1) \vee (a_2, r_2, V_1)) \wedge (V_1, r_2, T). \quad (2u1p)$$

Note that despite their dissimilar syntaxes and the associated sub-graphs (Figure 2), the query type **2u** should be as difficult as **1p**, as to answer the first it suffices to match a single link correctly. Similarly, **2u1p** should be as complex as **2p**. The fact that **2u** and **2u1p** are reported to be harder to solve in practice than **1p** and **2p** (Ren et al., 2020) is due to the way standard benchmarks are built, which we discuss in Sec. 4, and the way in which CQA is evaluated, discussed next. App. B details further query types, including negation.

**Standard evaluation.** Given a KG  $G$  and a query  $q$ , answering  $q$  boils down to a graph matching problem (Hogan, 2020) if we assume that all the meaningful links are already in  $G$ . Instead, if  $G$  is incomplete, we will need to predict missing links while answering  $q$ . Many ML approaches to CQA, reviewed in Sec. 5, therefore assume a distribution over *possible* links (Loconte et al., 2023), thus requiring some form of *probabilistic reasoning*. To evaluate them, standard benchmarks such as FB15k237 and NELL995 artificially divide  $G$  into  $G_{\text{train}}$  and  $G_{\text{test}}$ , treating the triples in the latter as missing links. This splitting process is done uniformly at random, a procedural choice that is inherited from link prediction benchmarks (Ruffinelli et al., 2020) and that can alter the measured performance, as we discuss next. CQA is generally treated as a ranking task, counting how many true candidate answers to a query are ranked higher than non-answer ones. Let the rank of a **query answer (QA) pair**  $(q, t)$  be  $\text{rank}(q, t)$ , the performance for each query type is calculated as the mean reciprocal rank (MRR), i.e.:

$$|\mathcal{Q}|^{-1} \sum_{q \in \mathcal{Q}, t \in \mathcal{E}_q} |\mathcal{E}_q|^{-1} \text{rank}(q, t)^{-1}, \quad (1)$$

where  $\mathcal{Q}$  denotes the set of test queries of the considered type, and  $\mathcal{E}_q$  is the set of candidate answer entities for each

query  $q \in \mathcal{Q}$ . This average, across queries and answers, assumes that every QA pair having the same query type is equivalently hard, which is not the case. In fact, we show that certain QA pairs can be easier to predict if some links were already seen by the model during training (Figure 1) and that the distribution of the query answer pairs in the existing benchmarks is very skewed towards those involving a single missing link (Table 1). Thus, computing Equation (1) without understanding what the benchmark distribution is distorts the perception of performance gains.

### 3. What is the real “Hardness” of CQA in incomplete KGs?

As discussed in the previous section, the perceived complexity of a query is related to the graph structure associated to its query type (Figure 2): queries containing more hops/existentially-quantified variables are more challenging, e.g., a 3p query is harder than a 2p query. In this section, we give an alternative perspective on the difficulty of answering queries that take into account the information coming from the training data. We argue that predicting links that are truly missing, i.e., not accessible to a learned model, is actually what makes a query “hard”. To do so, we formalize the notion of a reasoning tree for a QA pair, and then define how we determine the practical hardness of a QA pair.

Given a query  $q$  and every answer set of candidate answers  $t \in \mathcal{E}_q$ , we define the **reasoning tree** of each QA pair  $(q, t)$ , as the directed acyclic graph starting from the anchor entities of  $q$  to the target entity  $t$ , whose relational structure matches the query graph. Figure 1 provides examples of different reasoning trees for four different answers to the same query. There, we highlight whether a link belongs to  $G_{\text{train}}$  or not, i.e., it is a missing link. We assess the hardness of each answer  $t \in \mathcal{E}_q$ , by analyzing the composition of the reasoning tree required to predict  $t$ . As to answering a  $(q, t)$  pair, multiple reasoning trees are possible, so we select the one with the smallest number of missing links and the fewest number of hops. A  $(q, t)$  pair is **trivial** if the answer  $t$  can be entirely retrieved from  $G_{\text{train}}$ , i.e., there exists at least one reasoning tree where all the links in it are seen during training. This type of answer does not need probabilistic inference and hence is filtered out from current CQA benchmarks (Ren & Leskovec, 2020), which only consider non-trivial  $(q, t)$  pairs, which we call **inference-based** pairs. However, inference-based pairs do not need all links in their reasoning tree to be predicted as, by definition, it is sufficient that at least one link in the tree is missing. Therefore, we define a  $(q, t)$  pair to be a **partial-inference** pair if its reasoning tree contains at least one link in  $G_{\text{train}}$  and at least one link present in  $G_{\text{test}}$ . Alternatively, a  $(q, t)$  pair whose reasoning tree contains only links present in  $G_{\text{test}}$  is called a **full-inference** pair. *We claim that inference-*

*based queries encompass a spectrum of hardness, with full-inference queries being the hardest and queries that have only one truly missing link (1p) the easiest. However, current benchmarks flatten this complexity and highlight only performance on the easiest.*

To predict a partial-inference pair  $(q, t)$ , a ML model that has **explicit** access to  $G_{\text{train}}$  has to solve a simpler task to answer  $q$ , as only a subset of the links in the reasoning tree are missing and need to be predicted. As such, the  $(q, t)$  pair can be simplified to  $(q', t)$  pair, where the query  $q'$  is of an easier query type than  $q$ . Figure 1 shows one example of how a QA pairs of type 2i1p from FB15k237 are reduced in practice to the simpler types 1p, 2i and 2p. This is exploited by hybrid solvers such as QTO (Bai et al., 2023) and our CQD-Hybrid (Sec. 5). Note that this advantage applies also to ML models that have **implicit** access to  $G_{\text{train}}$ , e.g., by having memorized the triples during training, a common phenomenon for many neural link predictors (Nickel et al., 2014). We confirm this in Sec. 5.

### 4. How Many “Easy” Queries in Current CQA Benchmarks?

In this section, we systematically analyze the practical hardness of queries from the current CQA benchmarks and answer the following research question: **(RQ1)** *What is the proportion of QA pairs that can be classified as full-inference rather than partial-inference and how easy are the latter?* To this end, we consider the CQA benchmarks generated from FB15k237 (based on Freebase) and NELL995 (based on NELL systems (Carlson et al., 2010)) as they are the most used to evaluate SoTA methods for CQA (Ren et al., 2020; Ren & Leskovec, 2020; Arakelyan et al., 2021; Zhang et al., 2021; Zhu et al., 2022; Arakelyan et al., 2023).<sup>4</sup>

**Complex queries can be reduced to much simpler types.** We group the testing QA pairs into query types (Sec. 2), and we further split them based on whether they can be reduced to simpler types after observing the training links in  $G_{\text{train}}$  (Sec. 3). Table 1 shows that for both FB15k237 and NELL995 the vast majority of QA pairs can be reduced to simpler types. For FB15k237, 86.8% to 98.3% of QA pairs can be reduced to 1p queries, while only 0.1% to 4% require full inference. Similarly, for NELL995, 49.5% to 98.5% of QA pairs map to 1p queries, and only 0.1% to 6% to full inference. For instance, 96.7% of 2i1p QA pairs in FB15k237 can be reduced to 1p queries, 1.8% to 2p, and 1.4% to 2i. However, only 0.1% of these pairs cannot be reduced to any other QA pair, i.e., they require full inference in order to be predicted. The only exceptions to this trend are QA pairs where the query has a 1p or 2u structure which,

<sup>4</sup>We do not consider FB15k (Bordes et al., 2013) as it suffers from another form of data leakage (Toutanova & Chen, 2015).




**Table 1. The great majority of QA pairs are of the much easier partial-inference type (non-diagonal), rather than of the harder full-inference type (diagonal).** For each query type (as rows) we show the percentage of the QA pairs that can be reduced to an easier query type (as columns) for datasets FB15k237 and NELL995. Most of the complex queries can be reduced to simple link prediction queries (i.e., **1p**). We denote as ‘-’ those reductions that are not possible given the sub-graph structure induced by the query type.

	FB15k237									NELL995								
	1p	2p	3p	2i	3i	1p2i	2i1p	2u	2u1p	1p	2p	3p	2i	3i	1p2i	2i1p	2u	2u1p
1p	100	-	-	-	-	-	-	-	-	100	-	-	-	-	-	-	-	-
2p	98.1	1.9	-	-	-	-	-	-	-	97.6	2.4	-	-	-	-	-	-	-
3p	97.2	2.7	0.1	-	-	-	-	-	-	95.6	4.3	0.1	-	-	-	-	-	-
2i	96.0	-	-	4.0	-	-	-	-	-	94.0	-	-	6.0	-	-	-	-	-
3i	91.6	-	-	8.2	0.2	-	-	-	-	87.4	-	-	12.1	0.5	-	-	-	-
1p2i	86.8	1.0	-	12.0	-	0.2	-	-	-	49.5	0.6	-	49.0	-	0.9	-	-	-
2i1p	96.7	1.8	-	1.4	-	-	0.1	-	-	96.2	2.4	-	1.2	-	-	0.2	-	-
2u	0.0	-	-	-	-	-	-	100	-	0.0	-	-	-	-	-	-	100	-
2u1p	98.3	0.0	-	-	-	-	-	1.6	0.1	98.5	0.0	-	-	-	-	-	1.4	0.1

by definition, only require the prediction of a single link and therefore cannot be reduced by any other query type. We conduct the same analysis for queries that include negation, considering only the reasoning tree over non-negated links (see App. A.1 for details). Table A.1 presents results analogous to those in Table 1, showing that the vast majority of QA pairs including negation have training links in their reasoning tree. For example, 95.4% and 93.9% of **3in** QA pairs in FB15k237 and NELL995, respectively, have training links in their reasoning tree.

**Is this realistic?** This skewed distribution is the result of *implicit assumptions* behind how these benchmarks have been created. We argue they are arbitrary and not realistic. First, we note that the procedure to generate queries assumes that triples are *independent*, a simplifying assumption that does not hold in practice. Second, as triples are sampled independently, the quantity of QA pairs that can be reduced to **1p** depends on how many links are assumed to be missing. This is given by the ratio of  $|G_{\text{train}}|$  and  $|G_{\text{test}}|$ . Note that however the train/test split is arbitrary and has been inherited from the benchmarks for link prediction, where the task was rightfully to predict a single link.

**Non-existing links for union queries.** Finally, we discovered that there are *non-existing* links, i.e., links that are not in the original KG  $G$  and hence neither in  $G_{\text{train}}$  or  $G_{\text{test}}$ , in both FB15k237 and NELL995 in the reasoning trees of queries involving unions, i.e., the **2u** and **2u1p** types. These links violate our definitions of inference QA pairs, hence, we filter them out, and report only *filtered* QA pairs in Table 1. For example, for **2u** QA pairs, if either one of the two links in the reasoning tree does not exist in the graph, we remove the QA pair (see Figure 3 bottom for an example). It follows that only QA pairs where both links in the reasoning tree exist in the test graph are retained (see Figure 3 (Top) for an example). More crucially, these non-existing links can alter the performance of solvers for **2u** and **2u1p** types, as we discuss next.



**Figure 3. Some answers of union queries can be reached by a single missing link, while the other link does not exist**, as shown for this **2u** query and fragments of the KG FB15k237, where  $r_1 = \text{releasedIn}$ ,  $r_2 = \text{crewMemberOf}$ . (Top) Example of an answer that we retain in our benchmark, being reachable by two missing links. (Bottom) Example of answers we filter out, being only reachable by one missing link.

**Takeaway 1.** We discourage using current CQA benchmarks, as they essentially reflect only the performance of predicting answers where only one link is truly missing (for both hybrid and neural solvers).

## 5. What do SoTA Models Perceive as Hard?

In this section, we re-evaluate SoTA methods for CQA and answer to the following question: **(RQ2)** *How do they perform on QA pairs that are placed differently in the hardness spectrum?* To substantiate that memorizing training links helps in answering partial-inference queries but not full-inference ones (Sec. 3), we evaluate hybrid solvers who have access to  $G_{\text{train}}$ .

**Neural CQA methods.** Over the years, a significant number of neural models have been proposed for solving the CQA task, see App. C for an overview. Among purely-neural models, we consider five representative approaches

method	3p								2i1p							
	ovr	1p	2p	3p	2i	3i	1p2i	2i1p	ovr	1p	2p	3p	2i	3i	1p2i	2i1p
GNN-QE	11.8	12.0	4.4	4.8	-	-	-	-	18.9	19.2	8.2	-	6.2	-	-	3.4
ULTRAQ	8.9	9.0	4.6	4.4	-	-	-	-	18.6	18.9	8.5	-	5.1	-	-	8.1
ConE	11.0	11.0	5.4	2.6	-	-	-	-	14.0	13.8	9.7	-	12.8	-	-	5.6
CQD	7.8	7.8	3.4	1.8	-	-	-	-	20.7	21.0	10.5	-	6.7	-	-	7.6
CQD-Hybrid	11.0	11.1	3.6	4.4	-	-	-	-	24.0	24.6	10.0	-	4.6	-	-	7.5
QTO	<b>15.6</b>	15.8	4.5	<b>5.0</b>	-	-	-	-	<b>24.7</b>	25.3	10.3	-	8.6	-	-	<b>8.1</b>
CLMPT	11.3	11.3	6.1	4.0	-	-	-	-	19.3	19.1	18.7	-	7.9	-	-	<b>13.2</b>

1p queries are overrepresented in the existing benchmark (Table 1). Compare instead their values to those in the full-inference column (highlighted in orange) where all links in the reasoning tree of the QA pairs are missing: the latter are much lower. This applies also to intermediate partial-inference QA pairs (middle columns), entailing that the skewed distribution of these datasets (Table 1) greatly distorts the overall performance usually reported in previous papers. The same pattern appears for all query types in both FB15k237 and NELL995. See Tables A.4 and A.5 for the complete stratified analysis for all query types and models.

that significantly differ in their methodological designs and yield SoTA results compared to other models in their class: (1) Cone Embeddings (ConE; Zhang et al., 2021) is a geometry-based complex query answering model where logical conjunctions and disjunctions are represented as intersections and union of cones. (2) Graph Neural Network Query Executor (GNN-QE; Zhu et al., 2022) decomposes a complex first-order logic query into projections over fuzzy sets. (3) Conditional Logical Message Passing Transformer (CLMPT; Zhang et al., 2024) leverages existing KG embeddings to conduct one-hop inferences on atomic formulas in a message-passing scheme as in LMPNN (Wang et al., 2023), and extends it by using a transformer to aggregate messages incoming from neighbor nodes. (4) Continuous Query Decomposition (CQD; Arakelyan et al., 2021; 2023), reduces the CQA task to the problem of finding the most likely variable assignment, where the likelihood of each link (1p query) is assessed by a neural link predictor, and logical connectives are relaxed via fuzzy logic operators. (5) UltraQuery (ULTRAQ; Galkin et al., 2024b) is a foundation model for CQA inspired by Galkin et al. (2024a) where links and logical operations are represented by vocabulary-independent functions which can generalize to new entities and relation types in any KG.

**Hybrid solvers.** So far, we have assessed that current benchmarks are skewed towards easier query types (Table 1) and thus the perceived progress of current SoTA CQA methods boils down to their performance on 1p queries (Table 2). To have an undistorted view of this progress, in the next section, we will devise a benchmark that rebalances all partial-inference as well as full-inference queries types. That is, we will provide an equal number of QA pairs for each subtype a query can be reduced to. Now, we argue that in a real-world scenario one has to perform reasoning over *both* existing links *and* missing ones, hence discarding the information in  $G_{\text{train}}$  is wasteful. This, in addition to the fact that purely-neural SoTA CQA methods might implicitly exploit this aspect at test time if they are able to memorize well  $G_{\text{train}}$ , pushes us to evaluate also *hybrid* solvers.

Two remarkable examples are QTO (Bai et al., 2023) and FIT (Yin et al., 2024). They explicitly retrieve existing links from  $G_{\text{train}}$  when answering a query at test time. Inspired by them, and to evaluate our hypothesis, we create a light-weight hybrid solver, named *CQD-Hybrid*, that is a variant of CQD that uses a pre-trained link predictor to compute scores for the answer entities of 1p queries only, denoting the unnormalized probability of that single link to exist (Loconte et al., 2023). Then, assignments to existentially quantified variables in a complex query are greedily obtained by maximizing the combined score of the links, computed by replacing logical operators in the query with fuzzy ones (van Krieken et al., 2022). A complete assignments to the variables (and hence to the target variable  $T$ ), gives us an answer to the query. In our CQD-Hybrid, we assign the maximum score to those links that are observed in  $G_{\text{train}}$ . That is, training links will have a higher score than the one that the link predictor would output, hence effectively steering the mentioned greedy procedure at test time. This is the only minimal change we apply to CQD to test if its performance can depend on memorizing training triples. Note that this makes CQD-Hybrid different from QTO and FIT, which involve more sophisticated steps such as score calibration and a forward/backward update stage, that can boost performance. In our experiment we will only consider QTO, as for the set of queries we considered QTO and FIT are equivalent (Yin et al., 2024). We report hyperparameters and implementation details of CQD-Hybrid in App. F.1.

**All SoTA methods perform significantly worse on full-inference queries.** We re-evaluate the SoTA methods mentioned above and stratify their performances based on the type of partial- and full-inference queries. Where available, we reused the pre-trained models for their evaluation. Otherwise, we re-produced the SoTA results using the hyperparameters provided by the authors, listed in App. F.1.

Table 2 presents a portion of the results in terms of MRR for the testing QA pairs in the existing benchmarks FB15k237 and NELL995, along with the different QA pairs they can be reduced to when observing the training KG, grouped by

Table 2. SoTA performance drops significantly on full-inference QA pairs as shown for the MRR for 3p and 2i1p queries on FB15k237. Compare the overall MRR achieved on all available testing queries (column “ovr”) with that achieved on the 1p reductions (“1p” column, i.e., when only a single link is truly missing): they are extremely close. This is because

Table 3. Exploiting training links during inference boosts MRR on existing benchmarks, as shown by our CQD-Hybrid model compared to previous SoTA, for several query types. CQD-Hybrid always achieves higher MRR scores when compared to CQD, and outperforms more sophisticated non-hybrid baselines 9/14 times. Best values in bold and second best underlined.

	Model	2p	3p	2i	3i	1p2i	2i1p	2u1p
FB15k237	GNN-QE	<u>14.7</u>	<b>11.8</b>	<b>38.3</b>	<b>54.1</b>	<u>31.1</u>	18.9	9.7
	ULTRAQ	11.5	8.9	35.7	51.0	29.6	18.6	7.3
	CQD	13.2	7.8	35.0	48.5	27.5	<u>20.7</u>	<u>10.5</u>
	ConE	12.8	11.0	32.6	47.3	25.5	14.0	7.4
	CLMPT	13.9	<u>11.3</u>	37.5	51.9	28.4	19.3	<b>11.4</b>
	<b>CQD-Hybrid</b>	<b>15.0</b>	11.0	<u>37.6</u>	<u>52.7</u>	<b>31.2</b>	<b>24.0</b>	10.3
NELL995	GNN-QE	17.9	15.2	40.0	50.9	29.1	20.5	8.8
	ULTRAQ	11.2	9.7	36.3	47.7	25.1	15.6	8.4
	CQD	22.0	13.4	<u>42.2</u>	<u>51.8</u>	<u>31.5</u>	<u>25.8</u>	16.0
	ConE	16.0	13.8	39.6	50.2	26.1	17.6	11.1
	CLMPT	<u>22.1</u>	<b>18.2</b>	41.6	51.7	29.0	24.7	<u>16.1</u>
	<b>CQD-Hybrid</b>	<b>23.8</b>	<u>17.8</u>	<b>44.2</b>	<b>57.8</b>	<b>33.2</b>	<b>28.4</b>	<b>16.6</b>

each query type. There, the performance of all SoTA methods consistently drops for query answer pairs that have a higher number of missing links in their reasoning tree. Furthermore, Tables A.4 and A.5 compare the MRR computed on all the available queries originally used in the benchmarks (column “all”), and the scores on those query types that they can be reduced to. *There is a high similarity in MRR between the columns “all” and “1p” that is evidence that the good performances of CQA methods are explained by their link prediction performances*, as a very high percentage of queries in fact are reduced to **1p** queries (see Table 1). A similar conclusion can be drawn for queries including negation, as shown in Table A.2.

Note that this happens consistently also for hybrid solvers such as QTO and our CQD-Hybrid: they have a harder time solving QA pairs with more than one missing link. While their performance was consistently surpassing purely-neural solvers, and CQD-Hybrid consistently boosted the performance of CQD, their poor performance on harder QA pairs is not reflected in the overall average score that is usually reported on papers, see Table 2.

Table 3 highlights that our CQD-hybrid is able to surpass existing non-hybrid SoTA, which is evidence that a pre-trained neural link prediction plus memorization of  $G_{\text{train}}$  achieves SoTA performances for FB15k237 and NELL995. The complete stratified comparison for all solvers is reported in Tables A.4 and A.5. However, in a few cases it happens that full-inference results are higher than some partial-inference; in such cases, the reason is to be found in the fact that full-inference QA pairs are very scarce (see Table 1).<sup>5</sup> *This motivates us to create a fully-balanced*

<sup>5</sup>We rule out that this is due to the influence of the number of existing entities bounding to existential quantified variables. For example we report such an analysis for CQD in Figures A.1

Figure 4. MRR of SoTA on the new benchmarks is significantly lower than the old ones as reported for FB15k237 and NELL995 (in blue) versus our new FB15k237+H and NELL995+H (in orange) for different query types and SoTA models.




Table 4. Queries of type 2u are less hard than previously thought in terms of MRR, with our filtered 2u queries having comparable performance to 1p, as discussed in Sec. 2, in contrast with the results of the 2u queries originally used (2u all).

method	FB15k237			NELL995		
	1p	2u all	2u filtered	1p	2u all	2u filtered
GNN-QE	42.8	16.2	<b>40.7</b>	53.6	15.4	34.8
ULTRAQ	40.6	13.2	33.6	38.9	10.2	21.3
CQD	<b>46.7</b>	<b>17.6</b>	32.7	<b>60.4</b>	<b>19.9</b>	35.9
ConE	41.8	14.9	29.9	60.0	14.9	28.2
CLMPT	45.3	14.8	32.8	58.1	18.6	<b>39.6</b>

*benchmark* in the next section.

**Non-existing links for union queries give a false sense of hardness.** For union queries, we filter query answer pairs, as discussed in Sec. 4, removing those with non-existing links in their reasoning trees. In Table 4, we show that if we do not do that, and consider also the pairs with non-existing links (denoted as “2u all”), MRR performance greatly drops. In this way, we reproduce the low performance for 2u queries that the original SoTA baselines reported in their papers. However, as Table 4 shows, this is just an artifact due to including non-filtered pairs while computing Equation (1). With our filtered pairs, instead, we recover the performance of 1p queries, as expected (Sec. 2). Similar conclusions can be drawn for other union queries as reported in Table A.3.

**Takeaway 2.** FB15k237 and NELL995 are not suitable to precisely assess CQA methods, as most of their QA pairs can be predicted by just predicting a single missing links. This results in highly inflated performance that distorts the perception of progress in the field and pushes the community to improve the performance only on the easiest QA pairs.

and D.2. Moreover, we found that no anchor entities nor relation names is predominant in both benchmarks, as shown in Table A.6.

## 6. New Benchmarks for Undistorted CQA Performance

In this section, we aim to answer the following question: **(RQ3)** *How can we construct a set of CQA benchmarks that let us measure the performance on each query type while not distorting the aggregated average?* To do so, we present a new set of CQA benchmarks based on the previously used FB15k237 and NELL995, as well as a novel one we design out of temporal knowledge graph ICEWS18 to question the current way to split the original KG  $G$  into  $G_{\text{train}}$  and  $G_{\text{test}}$  (see our discussion of how sampling triples independently is arbitrary in Sec. 4). Then, we evaluate current SoTA methods on these new benchmarks to establish a rigorous baseline for future works.

**Building new CQA benchmarks.** We generate our benchmarks to comprise both full-inference and partial-inference QA pairs, but we modify the algorithm of (Ren et al., 2020) to ensure that no training links are present in the reasoning trees of full-inference pairs, and for each query type, we build the benchmark such that the full spectrum of hardness appears with the same frequency in it. For example, for 3p queries, we sample 30,000 QA pairs, 10,000 for 3p that can be reduced to 1p, 10,000 for those reduceable to 2p, and 10,000 for full-inference QA pairs (3p QA pairs non-reduceable to any other type). Furthermore, we sample union queries 2u, 2u1p such that non-existing link are not present in their reasoning tree. Finally, to raise the bar of “complexity” in CQA, we introduce two query types that, in their full-inference versions, are harder than simpler types by design. Specifically, we design “4p” queries for a sequential path made of four links and “4i”, which represents the intersection of the target entity sets defined over four links. For their logical formulation, see App. E.

We do so for the query types shown in Figures 2 and B.1, and name the *harder* versions of FB15k237 and NELL995, FB15k237+H and NELL995+H. Only exceptions are 1p, 2in, 2nu1p those only being of type full-inference. See Apps. A.1 and B for details about queries including negation, and App. G.2 for details about query generation.

**Non-uniform at random splits.** Recall from Sec. 4 that FB15k237 and NELL995 have been created by assuming triple independence and by inheriting the train/test split from their link prediction counterparts. For FB15k237+H and NELL995+H, we removed the triple independence assumption as we sample the same number of QA pairs per type, but we keep the existing data splits of  $G_{\text{train}}$  and  $G_{\text{test}}$ , for retro-compatibility. However, to evaluate the impact of this artificial splitting process, we adopt a more realistic one for ICEWS18+H, where  $G_{\text{train}}$  contains only “past” links, and we might want to predict the future links contained in  $G_{\text{test}}$ . To this end, we leverage the temporal information in ICEWS18 by (1) ordering the links based on their times-

tamp; (2) removing the temporal information, thus obtaining regular triples; and (3) selecting the train set to be the first temporally-ordered 80% of triples, the valid the next 10%, and the remaining to be the test split. If the same fact appears with multiple timestamps, we retain only the link with the earliest timestamp. For detailed statistics about the splits, please refer to App. G.1. By doing so, we create a much more challenging benchmark even for easy 1p queries.

### How do SoTA methods fare on our new benchmarks?

For FB15k237+H and NELL995+H we re-used the pre-trained models used for the experiments in Sec. 4, while for the newly created ICEWS18+H we trained GNN-QE, CLMPT, ConE, and the link predictor used in CQD, CQD-Hybrid and QTO from scratch. As this is not necessary for ULTRAQ, being a zero-shot neural link prediction applicable to any KG, we re-used the checkpoint provided by the authors. Hyperparameters of all models are in App. F.2.

Figure 4 reports the MRR on the new benchmarks versus the old benchmarks, where the MRR of FB15k237 and NELL995 (in blue) is much higher than the one of FB15k237+H and NELL995+H (in orange) for all methods for all query types. Considering how MRR can drop for more than 30 points, e.g., for 3i queries, this highlights once more how distorted the perception of progress in the field can be.

Table 5 shows the results of the selected baselines on the new benchmarks, for all query types in Figures 2 and B.1. For stratified results on FB15k237+H, NELL995+H, ICEWS18+H see Tables D.3 to D.5, respectively. For the stratified analysis on queries including negation see Table D.6. **The surprising result from Table 5 is that no model obtains SoTA performance on most query type**, which was not the case for the old benchmarks, where QTO obtained SoTA performance for most query types (see Tables A.2, A.4 and A.5). This is further evidence that current SoTA methods mainly focused on improving the performance on the easiest QA pairs and likely overfit the benchmarks. Instead, with our new benchmarks, this will not be possible in the same way, as we compile the benchmark with the same amount of partial- and full-inference QA pairs per type.

Furthermore, we highlight how full-inference 2u queries have similar performance with 1p, similar to how full-inference 2u1p are expected to be as hard as 2p, as discussed in Sec. 3. Finally, we remark that ICEWS18+H is much harder than NELL995+H and FB15k237+H, across all query types, even 1p, thus raising the bar for neural link predictors as well. This confirms our hypothesis that the sampling method of the KG splits plays a big role in determining the hardness of the benchmark. For additional analysis on the new benchmarks, comprising the influence of the number of existing entities bounding to existential quantified variables, the most frequent relation name and anchor entity per query type, and statistical tests, see App. D.



**Table 5. There is no clear SoTA method for the new benchmarks.** As shown in Figure 4, the MRR on the new benchmarks is significantly lower than the old ones. For example, for 3i queries on FB15k237+H, QTO has an MRR of 10.1, while for FB15k237 it was 54.6. For the stratified analysis see Tables D.3 to D.6.

	Model	1p	2p	3p	2i	3i	1p2i	2i1p	2u	2u1p	4p	4i	2in	3in	2pi1pn	2nu1p	2in1p
FB15k237+H	GNN-QE	42.8	5.2	4.0	6.0	8.8	5.6	9.9	32.5	10.0	4.3	20.0	6.8	<b>6.5</b>	<b>3.7</b>	5.0	3.3
	ULTRAQ	40.6	4.5	3.5	5.2	7.2	5.3	10.1	29.4	8.3	3.8	16.4	5.3	5.5	2.6	3.7	2.2
	CQD	<b>46.7</b>	4.4	2.4	<b>11.3</b>	<b>12.8</b>	6.0	11.5	40.1	10.6	1.1	23.8	3.3	2.6	0.6	4.9	1.2
	CQD-Hybrid	<b>46.7</b>	4.8	3.1	6.0	8.6	5.5	12.9	<b>42.2</b>	12.0	2.4	17.4	4.7	1.6	1.0	3.2	1.3
	ConE	41.8	4.6	3.9	9.1	10.3	3.8	7.9	22.8	6.0	3.5	20.3	5.1	4.9	2.9	3.3	<b>3.6</b>
	QTO	<b>46.7</b>	4.9	3.7	8.7	10.1	<b>6.1</b>	13.5	30.6	11.2	3.9	20.2	<b>10.6</b>	3.1	2.0	<b>5.3</b>	1.5
	CLMPT	45.3	<b>5.3</b>	<b>4.7</b>	10.2	12.2	5.6	<b>14.9</b>	33.6	<b>14.2</b>	<b>4.5</b>	<b>24.0</b>	6.8	2.3	1.6	4.8	2.5
NELL995+H	GNN-QE	53.6	8.0	6.0	10.7	13.3	16.0	13.5	47.5	9.8	4.7	19.4	5.5	<b>6.4</b>	<b>5.8</b>	3.3	4.4
	ULTRAQ	38.9	6.1	4.1	7.9	10.2	15.8	9.3	28.1	9.5	4.2	15.6	4.5	5.9	4.3	2.7	3.6
	CQD	<b>60.4</b>	9.6	4.2	18.5	19.6	18.9	22.6	46.3	18.5	2.0	24.8	4.2	1.5	1.5	4.9	2.6
	CQD-Hybrid	<b>60.4</b>	9.0	6.1	12.1	14.4	17.4	21.2	46.4	19.3	3.5	20.4	5.1	1.2	1.4	4.3	2.4
	ConE	53.1	7.9	6.7	21.8	23.6	14.9	11.8	39.9	8.8	5.2	27.6	4.6	6.0	3.7	2.7	<b>6.4</b>
	QTO	60.3	9.8	<b>8.0</b>	14.6	15.8	<b>17.6</b>	21.1	49.1	18.9	7.0	20.9	<b>10.2</b>	2.3	3.1	<b>8.4</b>	2.4
	CLMPT	58.1	<b>10.1</b>	7.8	<b>22.7</b>	<b>25.0</b>	17.2	<b>24.4</b>	<b>50.0</b>	<b>22.0</b>	<b>7.2</b>	<b>29.1</b>	6.5	2.4	4.1	2.3	4.5
ICEWS18+H	GNN-QE	12.2	0.9	0.5	<b>16.1</b>	<b>26.5</b>	<b>19.1</b>	3.5	<b>7.6</b>	1.1	0.4	<b>34.0</b>	4.5	6.9	0.9	<b>3.5</b>	0.8
	ULTRAQ	6.3	1.2	1.2	7.0	11.7	8.8	1.3	3.3	1.2	0.8	15.9	2.3	4.8	<b>1.2</b>	2.2	<b>1.6</b>
	CQD	<b>16.6</b>	2.5	<b>1.5</b>	13.0	19.5	17.1	<b>6.7</b>	6.8	<b>5.9</b>	<b>1.1</b>	24.0	1.5	2.9	0.2	2.7	0.9
	CQD-Hybrid	<b>16.6</b>	<b>2.6</b>	<b>1.5</b>	15.0	25.5	17.5	5.8	6.8	5.6	0.9	33.2	1.7	4.0	0.3	2.2	1.1
	ConE	3.5	0.9	0.9	1.2	0.5	1.2	1.6	1.1	0.9	0.6	0.3	1.7	2.9	1.1	0.9	1.3
	QTO	<b>16.6</b>	<b>2.6</b>	1.4	15.7	25.0	18.4	6.2	6.7	4.9	<b>1.1</b>	31.5	<b>4.9</b>	<b>8.7</b>	<b>1.2</b>	3.0	0.9
	CLMPT	4.7	0.8	0.1	12.0	23.0	9.7	2.1	2.7	2.2	0.1	31.0	1.2	2.1	0.1	1.0	0.2

**Takeaway 3.** Our benchmarks are built with the same amount of partial-inference and full-inference QA pairs that depend on the query type, allowing to measure CQA performance while not distorting the aggregated performance across all QA pairs. Additionally, ICEWS18+H highlights that more realistic sampling strategies are challenging for current SoTA.

## 7. Conclusion

In this paper, we revisit CQA on KGs and reveal that the “good” performance of SoTA approaches predominantly comes from answers that can be reduced to easier types (Table 2), the vast majority of which boiling down to single link prediction (Table 1). We perform such an analysis by dissecting the arbitrary assumptions that were used to create these datasets and by highlighting how current neural and hybrid solvers can exploit (different) forms of triple memorization to make complex queries much easier. We confirm this by reporting their performance in a stratified analysis and by proposing our hybrid solver, CQD-Hybrid, which, while being a simple extension of an old method like CQD, can be very competitive against other SoTA models. This reveals that FB15k237 and NELL995 are not suitable to precisely assess the capability of CQA methods to answer complex queries and that most of their performance of

CQA on FB15k237 and NELL995 can be explained by their memorization ability.

We then created a set of new benchmarks FB15k237+H, NELL995+H and ICEWS18+H that comprise the same amount of full-inference and partial-inference QA pairs. The MRR performance of existing CQA greatly drops on our new benchmarks. We consider our new benchmarks to be a stepping stone towards benchmarking truly complex reasoning with ML models and advise all future works on CQA to use them. However, both the old and our new benchmarks only consider queries with bindings of single target variables. While this reflects an important class of queries on real-world KGs, many real-world queries require bindings to multiple target variables (i.e. answer tuples). We plan to extend our current study of the “real hardness” of CQA benchmarks to other popular settings in neural query answering, such as inductive scenarios (Galkin et al., 2022; Arun et al., 2025) where some entities and/or relations are unseen during the test stage. Moreover, our work could also be extended to queries requiring bindings to multiple target variables (Yin et al., 2023), and to queries having a DAG structure (He et al., 2025).

## Impact Statement

This work analyses current CQA benchmarks for knowledge graphs and shows that the way benchmarks are created

inflates perceived progress in the field. By demonstrating that most benchmark queries can be reduced to simpler problems and developing new, more balanced benchmarks, this work helps establish more accurate measures of model performance in CQA tasks. The findings have important implications for how the research community evaluates and develops CQA systems, potentially leading to more robust and capable models to better handle truly complex queries.

## Author Contributions

AV, CG, and LL conceived the initial idea for which existing benchmarks were not hard enough for complex query answering. AV, CG, DH, and LL had the intuition that a simple hybrid solver could obtain SoTA results on the old benchmarks. BX, and CG found the issue with union queries. CG developed the necessary code, ran all the experiments, and plotted the figures, with the following exceptions: BX ran all the experiments involving ConE, LL contributed to the generation of the new benchmarks, specifically for the 4p and 4i queries, PM wrote the script for generating the KG splits of ICEWS18, AV drew Fig.1 and 2. BX, and CG wrote the initial draft of the paper, while AV shaped the story line of the paper and helped with the writing. All authors critically revised the paper. AV, PM, and SS supervised all the phases of the project and gave feedback.

## Acknowledgements

AV was supported by the “UNREAL: Unified Reasoning Layer for Trustworthy ML” project (EP/Y023838/1) selected by the ERC and funded by UKRI EPSRC. PM was partially funded by ELIAI (The Edinburgh Laboratory for Integrated Artificial Intelligence), EPSRC (grant no. EP/W002876/1), an industry grant from Cisco, and a donation from Accenture LLP. CG was funded by the CHIPS Joint Undertaking (JU) under grant agreement No. 101140087 (SMARTY), and by German Federal Ministry of Education and Research (BMBF) under the sub-project with the funding number 16MEE0444. CG acknowledge compute time on HoreKa HPC (NHR@KIT), funded by the BMBF and Baden-Württemberg’s MWK through the NHR program, with additional support from the DFG. BX, DH were supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1574 – 471687386.

## References

Arakelyan, E., Daza, D., Minervini, P., and Cochez, M. Complex query answering with neural link predictors. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021.

Arakelyan, E., Minervini, P., Daza, D., Cochez, M., and Augenstein, I. Adapting neural link predictors for data-efficient complex query answering. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023.

Arun, A., Kumar, S., Nayyeri, M., Xiong, B., Kumaraguru, P., Vergari, A., and Staab, S. Semma: A semantic aware knowledge graph foundation model, 2025. URL <https://arxiv.org/abs/2505.20422>.

Bai, Y., Lv, X., Li, J., and Hou, L. Answering complex logical queries on knowledge graphs via query computation tree optimization. In Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J. (eds.), *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pp. 1472–1491. PMLR, 2023.

Bollacker, K. D., Evans, C., Paritosh, P. K., Sturge, T., and Taylor, J. Freebase: a collaboratively created graph database for structuring human knowledge. In *SIGMOD Conference*, pp. 1247–1250. ACM, 2008.

Bordes, A., Usunier, N., García-Durán, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In Burges, C. J. C., Bottou, L., Ghahramani, Z., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*, pp. 2787–2795, 2013.

Boschee, E., Lautenschlager, J., O’Brien, S., Shellman, S., Starz, J., and Ward, M. ICEWS Coded Event Data, 2015.

Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Jr., E. R. H., and Mitchell, T. M. Toward an architecture for never-ending language learning. In Fox, M. and Poole, D. (eds.), *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*, pp. 1306–1313. AAAI Press, 2010. doi: 10.1609/AAAI.V24I1.7519.

Chen, X., Hu, Z., and Sun, Y. Fuzzy logic based logical query answering on knowledge graphs. In *AAAI*, pp. 3939–3948. AAAI Press, 2022.

Choudhary, N., Rao, N., Katariya, S., Subbian, K., and Reddy, C. K. Self-supervised hyperboloid representations from logical queries over knowledge graphs. In *WWW*, pp. 1373–1384. ACM / IW3C2, 2021.

- Dalvi, N. N. and Suciu, D. The dichotomy of probabilistic inference for unions of conjunctive queries. *J. ACM*, 59(6):30:1–30:87, 2012.
- Galkin, M., Zhu, Z., Ren, H., and Tang, J. Inductive logical query answering in knowledge graphs. *Advances in neural information processing systems*, 35:15230–15243, 2022.
- Galkin, M., Yuan, X., Mostafa, H., Tang, J., and Zhu, Z. Towards foundation models for knowledge graph reasoning. In *ICLR*. OpenReview.net, 2024a.
- Galkin, M., Zhou, J., Ribeiro, B. F., Tang, J., and Zhu, Z. Zero-shot logical query reasoning on any knowledge graph. *CoRR*, abs/2404.07198, 2024b.
- Guu, K., Miller, J., and Liang, P. Traversing knowledge graphs in vector space. In *EMNLP*, pp. 318–327. The Association for Computational Linguistics, 2015.
- Hamilton, W. L., Bajaj, P., Zitnik, M., Jurafsky, D., and Leskovec, J. Embedding logical queries on knowledge graphs. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 2030–2041, 2018.
- He, Y., Xiong, B., Hernández, D., Zhu, Y., Kharlamov, E., and Staab, S. Dage: Dag query answering via relational combinator with logical constraints. In *Proceedings of the ACM on Web Conference 2025*, pp. 2514–2529, 2025.
- Hogan, A. Sparql query language. *The Web of Data*, pp. 323–448, 2020.
- Hogan, A., Blomqvist, E., Cochez, M., d’Amato, C., Melo, G. D., Gutierrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., et al. Knowledge graphs. *ACM Computing Surveys (Csur)*, 54(4):1–37, 2021.
- Loconte, L., Di Mauro, N., Peharz, R., and Vergari, A. How to turn your knowledge graph embeddings into generative models. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 77713–77744. Curran Associates, Inc., 2023.
- Nickel, M., Jiang, X., and Tresp, V. Reducing the rank in relational factorization models by including observable patterns. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., and Weinberger, K. Q. (eds.), *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 1179–1187, 2014.
- Ren, H. and Leskovec, J. Beta embeddings for multi-hop logical reasoning in knowledge graphs. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- Ren, H., Hu, W., and Leskovec, J. Query2box: Reasoning over knowledge graphs in vector space using box embeddings. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020.
- Ren, H., Galkin, M., Cochez, M., Zhu, Z., and Leskovec, J. Neural graph reasoning: Complex logical query answering meets graph databases. *arXiv preprint arXiv:2303.14617*, 2023.
- Ruffinelli, D., Broscheit, S., and Gemulla, R. You CAN teach an old dog new tricks! on training knowledge graph embeddings. In *ICLR*. OpenReview.net, 2020.
- Sun, H., Arnold, A. O., Bedrax-Weiss, T., Pereira, F., and Cohen, W. W. Faithful embeddings for knowledge base queries. In *NeurIPS*, 2020.
- Toutanova, K. and Chen, D. Observed versus latent features for knowledge base and text inference. In *CVSC*, pp. 57–66. Association for Computational Linguistics, 2015.
- Trouillon, T., Dance, C. R., Gaussier, É., Welbl, J., Riedel, S., and Bouchard, G. Knowledge graph completion via complex tensor factorization. *J. Mach. Learn. Res.*, 18: 130:1–130:38, 2017.
- van Krieken, E., Acar, E., and van Harmelen, F. Analyzing differentiable fuzzy logic operators. *Artificial Intelligence*, 302:103602, 2022.
- Wang, Z., Song, Y., Wong, G. Y., and See, S. Logical message passing networks with one-hop inference on atomic formulas. In *ICLR*, 2023.
- Xiong, W., Hoang, T., and Wang, W. Y. Deeppath: A reinforcement learning method for knowledge graph reasoning. In Palmer, M., Hwa, R., and Riedel, S. (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pp. 564–573. Association for Computational Linguistics, 2017. doi: 10.18653/V1/D17-1060.
- Yin, H., Wang, Z., Fei, W., and Song, Y.  $EFO_k$ -cqa: Towards knowledge graph complex query answering beyond set operation. *arXiv preprint arXiv:2307.13701*, 2023.

Yin, H., Wang, Z., and Song, Y. Rethinking complex queries on knowledge graphs with neural link predictors. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024.

Zhang, C., Peng, Z., Zheng, J., and Ma, Q. Conditional logical message passing transformer for complex query answering. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4119–4130, 2024.

Zhang, Z., Wang, J., Chen, J., Ji, S., and Wu, F. Cone: Cone embeddings for multi-hop reasoning over knowledge graphs. In *NeurIPS*, pp. 19172–19183, 2021.

Zhu, Z., Galkin, M., Zhang, Z., and Tang, J. Neural-symbolic models for logical queries on knowledge graphs. In *ICML*, volume 162 of *Proceedings of Machine Learning Research*, pp. 27454–27478. PMLR, 2022.



## A. Additional analysis on the old benchmarks

### A.1. Influence of existing links on queries involving negation

In Table A.1 we show an analysis for queries involving negation. For such analysis, we split the reasoning tree, defined in Sec. 3, into two subparts, namely *positive reasoning tree*, composed by the non-negated triples, and the *negative reasoning tree*, only composed by the negated triples. In particular, in the same spirit of Table 1, by only looking at the positive reasoning tree, we report the percentage of QA pairs that can be reduced to easier types (partial-inference) and the one that cannot be reduced to a simpler type (full-inference). Furthermore, we argue that also the negative reasoning tree contains training triples, but how this propagates to performance is less clear than the positive part, as each method treats negation differently. Table A.1 shows that in both FB15k237 and NELL995, queries can be reduced to easier types, revealing that our analysis also extends to negated queries. The only exceptions are *2in*, *2nu1p* where the positive reasoning tree consists of a single link, see App. B. Consequently, no reduction to partial inference is possible (see Sec. 3).

Table A.1. **Most negated queries have training links in the positive reasoning tree of the QA pairs.** Partial-inference and full-inference refer to the **positive** reasoning tree. '-' refers to reductions that are not possible.

	FB15k237		NELL995	
	partial-inference	full-inference	partial-inference	full-inference
2in	-	100	-	100
3in	95.4	4.6	93.9	6.1
2pi1pn	98.4	1.6	97.8	2.2
2nu1p	-	100	-	100
2in1p	97.4	2.6	95.6	4.4

### A.2. Additional comparisons

Table A.5 and Table A.4 show the full results of FB15k237 and NELL995, including the results comparing all methods, all query types, and the subtypes they are reduced to.

Moreover, A.3 show the full results of the union queries, including the old overall results, which include QA pairs having non-existing links in their reasoning tree (see Figure 3), and the “new overall”, in which those QA pairs are filtered out. In particular, for “2u1p” queries, we observe that for FB15k237, the full-inference queries have a higher MRR than the overall, while for NELL995, we find the MRR results on full-inference QA pairs to be lower than the new overall results as expected. We attribute the results on FB15k237 to the presence of bias in the QA pairs due to their scarcity in the benchmarks. We argue that for the same reason, 2u1p performances are not comparable to 2p. However, we will show that for the full-inference QA pairs in the new benchmark, our claim is confirmed (see Tables D.3 to D.5).

Moreover, for query types including negation, results in Table A.2 show the same pattern of Table 2, i.e. the MRR of the full-inference QA pairs is much lower than the one of the partial-inference QA pairs and of the overall.

**Performance analysis of QTO** The stratified MRR performance of QTO (Bai et al., 2023) for the old benchmarks have been included in Tables A.4 and A.5. For a fair comparison, we remark that we used the same link predictor for CQD, CQD-Hybrid, and QTO. Details and hyperameters are available in App. F.1. Our analysis reveals that, similarly to the other baselines, QTO performance drops when evaluated on the full-inference QA pairs only (diagonal), w.r.t overall results. Moreover, even when QTO is the SoTA on a certain query type, most of the time it is not so on the portion of full-inference QA pairs only, showing that improving performance on the partial-inference QA pairs not necessarily results on improvements over the full-inference ones.

Moreover, results on NELL995, shown in Table A.4, reveal that CQD-Hybrid outperforms QTO for some query types of NELL995, i.e. *3i*, *1p2i*, *2i1p*, suggesting that even by only setting a score=1 to the training triples it is possible to obtain SoTA results on the old benchmarks. We remark that, while both in CQD-Hybrid and QTO a score=1 is set to the training links, QTO is much more sophisticated than CQD-Hybrid, as they 1) calibrated the scores with a heuristics, 2) store a matrix  $|V| \times |V|$  for each relation name containing the score for every possible triples, 3) have a forward/backward mechanism in the reasoning.

Table A.2. Even for queries including negation the MRR results are highly influenced by the presence of existing links in the positive reasoning tree of the QA pairs. For example, for 3in queries of FB15k237, the overall MRR of QTO is 16.1, while its MRR on the portion of full-inference QA pairs is only 2.1. Hence, the overall result is highly influence by the vast number of partial inference QA pairs in the benchmark (see Table A.1)

Query type	method	FB15k237			NELL995		
		overall	partial-inf	full-inf	overall	partial-inf	full-inf
2in	GNN-QE	6.8	-	6.8	5.5	-	5.5
	ULTRAQ	5.3	-	5.3	4.5	-	4.5
	CQD	3.3	-	3.3	4.2	-	4.2
	CQD-hybrid	4.7	-	4.7	5.1	-	5.1
	QTO	<b>10.6</b>	-	<b>10.6</b>	<b>10.2</b>	-	<b>10.2</b>
	ConE	5.1	-	5.1	4.6	-	4.6
	CLMPT	6.8	-	6.8	6.5	-	6.5
3in	GNN-QE	11.3	11.7	1.1	9.0	9.5	0.6
	ULTRAQ	9.6	9.9	0.6	8.2	8.6	0.7
	CQD	7.9	8.0	<b>2.7</b>	6.5	6.7	1.1
	CQD-hybrid	10.4	10.7	1.1	8.3	8.8	0.7
	QTO	<b>16.1</b>	16.7	2.1	<b>14.2</b>	14.9	1.2
	ConE	8.0	8.3	1.7	6.5	6.8	<b>1.3</b>
	CLMPT	13.1	11.1	1.7	7.9	6.9	1.2
2pilpn	GNN-QE	4.8	4.8	1.8	3.7	3.7	1.5
	ULTRAQ	3.6	3.6	1.5	3.2	3.2	1.7
	CQD	2.0	2.0	0.6	2.2	2.2	1.4
	CQD-hybrid	3.1	3.1	1.1	3.1	3.2	0.8
	QTO	<b>8.9</b>	9.0	<b>2.5</b>	<b>7.4</b>	7.5	<b>2.7</b>
	ConE	4.1	4.1	1.7	3.0	3.1	1.6
	CLMPT	4.6	4.0	1.8	3.7	3.2	1.7
2nu1p	GNN-QE	5.0	-	5.0	3.3	-	3.3
	ULTRAQ	3.7	-	3.7	2.7	-	2.7
	CQD	4.9	-	4.9	4.9	-	4.9
	CQD-hybrid	3.2	-	3.2	4.3	-	4.3
	QTO	<b>5.3</b>	-	<b>5.3</b>	<b>8.4</b>	-	<b>8.4</b>
	ConE	3.3	-	3.3	2.7	-	2.7
	CLMPT	4.8	-	4.8	4.5	-	4.5
2in1p	GNN-QE	5.2	5.3	1.7	5.4	5.5	2.7
	ULTRAQ	3.8	3.8	1.4	4.7	4.8	1.9
	CQD	3.6	3.6	1.4	5.5	5.5	2.6
	CQD-hybrid	4.1	4.1	1.5	6.6	6.7	2.5
	QTO	<b>7.9</b>	8.1	1.6	7.4	7.6	2.0
	ConE	6.3	6.4	<b>3.0</b>	7.5	7.6	4.0
	CLMPT	<b>7.9</b>	6.1	<b>3.0</b>	<b>11.6</b>	7.8	<b>4.2</b>

### A.3. Influence of intermediate existing entities on the results

For query types having intermediate variables, i.e., **2p**, **3p**, **2i1p**, **1p2i**, and **2u1p**, we analyzed how the number of existing intermediate entities matching the existentially quantified variables could influence the results. We refer to this number as *cardinality of existing entities*. For example, in Figure 1 the entities “When in Rome” and “Spiderman 2” are existing intermediate entities matching the existential quantified variable  $V$  for the query  $q_1$ . The intuition is that while some of these entities simplify the prediction of the answers of the query (see Sec. 3), some others might be misleading for the prediction. Hence, if there is a high number of existing entities in the path, it would be harder for the model to select only the existing entities needed to predict the next link. For this reason, we claim that a high value of cardinality of existing entities would make a QA pair *harder*.

To support our claim, we analyze the percentage of QA pairs having different values of cardinality, shown in Figure A.2, and their MRR, using CQD (Arakelyan et al., 2021), shown in Figure A.1. Figure A.1 shows that CQD results are highly influenced by the value of cardinality, having decreasing performance at the increase of the cardinality. Note that by matching this plot with the one in Figure A.2, one can understand if results on a specific query type are highly influenced by

Table A.3. **Full inference 2u QA pairs show higher performance than non-filtered, while full-inference 2u1p pairs show lower or comparable performance.** Results on union queries for existing benchmarks and comparison between old and new overalls. ‘-’ refers to reductions that are not possible, while ‘/’ to reductions that are possible but that are not present in the data.

Query type	method	FB15k237					NELL995				
		overall	overall (new)	1p	2u	2u1p	overall	overall (new)	1p	2u	2u1p
2u	GNN-QE	16.2	<b>40.7</b>	/	<b>40.7</b>	-	15.4	34.8	/	34.8	-
	ULTRAQ	13.2	<u>33.6</u>	/	<u>33.6</u>	-	10.2	21.3	/	21.3	-
	CQD	<b>17.6</b>	32.7	/	32.7	-	19.9	35.9	/	35.9	-
	ConE	14.3	29.9	/	29.9	-	14.9	28.2	/	28.2	-
2u1p	GNN-QE	<b>13.4</b>	9.7	8.7	50.9	13.1	8.8	8.8	6.3	53.1	2.0
	ULTRAQ	10.2	7.3	6.6	33.2	<b>15.0</b>	8.4	7.0	6.6	37.1	6.5
	CQD	11.3	10.5	10.4	13.7	<u>14.7</u>	16.0	14.6	14.2	51.8	9.5
	ConE	10.6	7.4	7.0	20.3	11.9	11.1	7.0	9.4	43.2	4.0

a high proportion of QA pairs having the same cardinality. For example, the MRR of 1p2i partial-inference QA pairs in FB15k237 is highly influenced by the presence of about 30% of QA pairs having cardinality of “1”.

#### A.4. Imbalances of relation names and anchor nodes

Next, we check if there is some imbalances in both relation names and anchors. Note that if a relation name or an anchor entity is present multiple times in a query, this is counted only once. As shown in Table A.6, for FB15k237 and NELL995 we find that in most cases there is no predominant relation name nor anchor entity. However, there are exceptions, as for example, for anchor entities of 3i queries of FB15k237, where the anchor node “USA” is present in 30.1% of them, and for 2u queries, where it is present in 22.5%. We also notice that there is a predominance of USA as an anchor entity across the vast majority of query types, which is most likely given by the vast presence of this entity in the knowledge graph.

Figure A.1. **Higher cardinality of intermediate existing entities leads to lower MRR.** Influence of the number of existing intermediate variables on the MRR for datasets FB15k237 and NELL995, using CQD.




Figure A.2. **Proportion of QA pairs having different cardinality of intermediate existing entities in the old benchmarks.**






Table A.4. The best model on all available NELL995 queries (column “overall”) and the best on the full inference queries (diagonal) is different for 6/7 query types, excluding 1p and 2u, being already of type full-inference. Comparison of MRR scores for SoTA methods for the different subtypes NELL995 queries can be reduced to. Best in bold.

Query type	method	overall	1p	2p	3p	2i	3i	1p2i	2i1p	2u	2u1p
1p	GNN-QE	53.6	53.6	-	-	-	-	-	-	-	-
	ULTRA-query	38.9	38.9	-	-	-	-	-	-	-	-
	ConE	60.0	60.0	-	-	-	-	-	-	-	-
	CQD	<b>60.4</b>	<b>60.4</b>	-	-	-	-	-	-	-	-
	CQD-hybrid	<b>60.4</b>	<b>60.4</b>	-	-	-	-	-	-	-	-
	QTO	60.3	60.3	-	-	-	-	-	-	-	-
	CLMPT	58.1	58.1	-	-	-	-	-	-	-	-
2p	GNN-QE	17.9	18.2	6.1	-	-	-	-	-	-	-
	ULTRA-query	11.2	11.5	4.6	-	-	-	-	-	-	-
	ConE	16.0	16.3	7.2	-	-	-	-	-	-	-
	CQD	22.0	22.3	7.6	-	-	-	-	-	-	-
	CQD-hybrid	23.8	24.2	6.2	-	-	-	-	-	-	-
	QTO	<b>24.7</b>	25.1	7.3	-	-	-	-	-	-	-
	CLMPT	22.1	22.3	<b>8.9</b>	-	-	-	-	-	-	-
3p	GNN-QE	15.2	15.1	8.1	3.4	-	-	-	-	-	-
	ULTRA-query	9.7	9.8	5.1	4.1	-	-	-	-	-	-
	ConE	13.8	13.8	8.2	6.5	-	-	-	-	-	-
	CQD	13.4	12.8	7.8	<b>8.5</b>	-	-	-	-	-	-
	CQD-hybrid	17.8	17.2	9.9	8.1	-	-	-	-	-	-
	QTO	<b>22.6</b>	22.3	10.9	7.7	-	-	-	-	-	-
	CLMPT	18.3	17.7	10.7	8.2	-	-	-	-	-	-
2i	GNN-QE	40.0	41.4	-	-	3.6	-	-	-	-	-
	ULTRA-query	36.3	37.5	-	-	3.2	-	-	-	-	-
	ConE	39.6	40.6	-	-	8.0	-	-	-	-	-
	CQD	42.2	43.3	-	-	6.9	-	-	-	-	-
	CQD-hybrid	<b>44.2</b>	45.7	-	-	6.7	-	-	-	-	-
	QTO	<b>44.2</b>	45.7	-	-	5.1	-	-	-	-	-
	CLMPT	41.6	42.5	-	-	<b>8.9</b>	-	-	-	-	-
3i	GNN-QE	50.9	53.6	-	-	10.9	2.2	-	-	-	-
	ULTRA-query	47.7	50.0	-	-	11.3	1.5	-	-	-	-
	ConE	50.2	51.9	-	-	18.3	4.8	-	-	-	-
	CQD	51.8	53.9	-	-	14.4	3.6	-	-	-	-
	CQD-hybrid	<b>57.8</b>	61.6	-	-	10.7	2.5	-	-	-	-
	QTO	56.2	60.0	-	-	12.4	2.4	-	-	-	-
	CLMPT	51.7	53.4	-	-	20.3	<b>5.3</b>	-	-	-	-
1p2i	GNN-QE	29.1	42.7	21.4	-	12.0	-	10.2	-	-	-
	ULTRA-query	25.1	39.0	25.6	-	8.2	-	7.0	-	-	-
	ConE	26.1	38.5	25.2	-	10.5	-	8.6	-	-	-
	CQD	31.5	44.0	25.7	-	14.5	-	<b>12.6</b>	-	-	-
	CQD-hybrid	<b>33.2</b>	48.7	24.2	-	13.7	-	9.4	-	-	-
	QTO	32.8	47.4	22.0	-	14.5	-	9.8	-	-	-
	CLMPT	29.0	39.3	15.3	-	21.8	-	12.1	-	-	-
2i1p	GNN-QE	20.5	20.8	16.3	-	14.6	-	-	20.7	-	-
	ULTRA-query	15.6	16.5	9.5	-	9.6	-	-	11.4	-	-
	ConE	17.6	17.7	16.4	-	25.7	-	-	19.5	-	-
	CQD	25.8	25.9	21.2	-	26.3	-	-	23.6	-	-
	CQD-hybrid	<b>28.4</b>	28.9	20.4	-	22.8	-	-	23.3	-	-
	QTO	28.2	28.5	20.2	-	24.0	-	-	24.4	-	-
	CLMPT	24.7	24.6	26.5	-	26.7	-	-	<b>25.2</b>	-	-
2u	GNN-QE	34.8	-	-	-	-	-	-	-	34.8	-
	ULTRA-query	21.3	-	-	-	-	-	-	-	21.3	-
	ConE	28.2	-	-	-	-	-	-	-	28.2	-
	CQD	35.9	-	-	-	-	-	-	-	35.9	-
	CQD-hybrid	35.9	-	-	-	-	-	-	-	35.9	-
	QTO	37.6	-	-	-	-	-	-	-	37.6	-
	CLMPT	<b>39.6</b>	-	-	-	-	-	-	-	<b>39.6</b>	-
2u1p	GNN-QE	8.8	6.3	-	-	-	-	-	-	53.1	2.0
	ULTRA-query	8.4	6.6	-	-	-	-	-	-	37.1	6.5
	ConE	7.0	9.4	-	-	-	-	-	-	43.2	4.0
	CQD	14.6	14.2	-	-	-	-	-	-	51.8	9.5
	CQD-hybrid	15.0	14.5	-	-	-	-	-	-	53.5	9.6
	QTO	16.4	15.8	-	-	-	-	-	-	61.2	9.6
	CLMPT	<b>17.7</b>	17.4	-	-	-	-	-	-	37.5	<b>10.2</b>

Table A.5. The best model on all available FB15k237 queries (column “overall”) and the best on the full inference queries (diagonal) is different for 5/7 query types, excluding 1p and 2u, those (q,a) pairs being all of type full-inference. Comparison of MRR scores for SoTA methods for the different subtypes FB15k237 queries can be reduced to. Best in bold.

Query type	method	overall	1p	2p	3p	2i	3i	1p2i	2i1p	2u	2u1p
1p	GNN-QE	42.8	42.8	-	-	-	-	-	-	-	-
	ULTRA-query	40.6	40.6	-	-	-	-	-	-	-	-
	ConE	41.8	41.8	-	-	-	-	-	-	-	-
	CQD	46.7	46.7	-	-	-	-	-	-	-	-
	CQD-hybrid	46.7	46.7	-	-	-	-	-	-	-	-
	QTO	<b>46.7</b>	<b>46.7</b>	-	-	-	-	-	-	-	-
	CLMPT	45.3	45.3	-	-	-	-	-	-	-	-
2p	GNN-QE	14.7	14.8	4.7	-	-	-	-	-	-	-
	ULTRA-query	11.5	11.5	4.2	-	-	-	-	-	-	-
	ConE	12.8	12.8	<b>5.2</b>	-	-	-	-	-	-	-
	CQD	13.2	13.3	3.5	-	-	-	-	-	-	-
	CQD-hybrid	15.0	15.2	3.5	-	-	-	-	-	-	-
	QTO	<b>16.6</b>	16.7	4.0	-	-	-	-	-	-	-
	CLMPT	13.9	13.9	5.1	-	-	-	-	-	-	-
3p	GNN-QE	11.8	12.0	4.4	4.8	-	-	-	-	-	-
	ULTRA-query	8.9	9.0	4.6	4.4	-	-	-	-	-	-
	ConE	11.0	11.0	5.4	2.6	-	-	-	-	-	-
	CQD	7.8	7.8	3.4	1.8	-	-	-	-	-	-
	CQD-hybrid	11.0	11.1	3.6	4.4	-	-	-	-	-	-
	QTO	<b>15.6</b>	15.8	4.5	<b>5.0</b>	-	-	-	-	-	-
	CLMPT	11.3	11.3	6.1	4.0	-	-	-	-	-	-
2i	GNN-QE	38.3	39.3	-	-	4.0	-	-	-	-	-
	ULTRA-query	35.7	36.7	-	-	3.4	-	-	-	-	-
	ConE	32.6	33.3	-	-	5.5	-	-	-	-	-
	CQD	35.0	35.8	-	-	<b>7.3</b>	-	-	-	-	-
	CQD-hybrid	37.6	38.7	-	-	3.8	-	-	-	-	-
	QTO	<b>39.7</b>	40.8	-	-	5.7	-	-	-	-	-
	CLMPT	37.5	38.4	-	-	<b>7.3</b>	-	-	-	-	-
3i	GNN-QE	54.1	56.0	-	-	10.9	5.2	-	-	-	-
	ULTRA-query	51.0	52.9	-	-	9.7	4.3	-	-	-	-
	ConE	47.3	48.3	-	-	15.6	4.0	-	-	-	-
	CQD	48.5	49.6	-	-	17.0	<b>6.0</b>	-	-	-	-
	CQD-hybrid	52.7	54.8	-	-	9.9	4.6	-	-	-	-
	QTO	<b>54.6</b>	56.4	-	-	15.4	5.4	-	-	-	-
	CLMPT	51.9	53.1	-	-	17.7	5.9	-	-	-	-
1p2i	GNN-QE	31.1	32.8	15.5	-	5.7	-	9.1	-	-	-
	ULTRA-query	29.6	31.5	19.2	-	4.2	-	7.4	-	-	-
	ConE	25.5	26.6	13.9	-	5.4	-	<b>9.7</b>	-	-	-
	CQD	27.5	28.7	13.4	-	7.1	-	9.0	-	-	-
	CQD-hybrid	31.2	33.4	14.8	-	4.8	-	7.0	-	-	-
	QTO	<b>33.8</b>	35.9	15.8	-	6.2	-	7.3	-	-	-
	CLMPT	28.4	30.0	16.2	-	5.9	-	8.3	-	-	-
2i1p	GNN-QE	18.9	19.2	8.2	-	6.2	-	-	3.4	-	-
	ULTRA-query	18.6	18.9	8.5	-	5.1	-	-	8.1	-	-
	ConE	14.0	13.8	9.7	-	12.8	-	-	5.6	-	-
	CQD	20.7	21.0	10.5	-	6.7	-	-	7.6	-	-
	CQD-hybrid	24.0	24.6	10.0	-	4.6	-	-	7.5	-	-
	QTO	<b>24.7</b>	25.3	10.3	-	8.6	-	-	8.1	-	-
	CLMPT	19.3	19.1	18.7	-	7.9	-	-	<b>13.2</b>	-	-
2u	GNN-QE	<b>40.7</b>	-	-	-	-	-	-	-	<b>40.7</b>	-
	ULTRA-query	33.6	-	-	-	-	-	-	-	33.6	-
	ConE	29.9	-	-	-	-	-	-	-	29.9	-
	CQD	32.7	-	-	-	-	-	-	-	32.7	-
	CQD-hybrid	32.7	-	-	-	-	-	-	-	32.7	-
	QTO	37.0	-	-	-	-	-	-	-	37.0	-
	CLMPT	32.8	-	-	-	-	-	-	-	32.8	-
2u1p	GNN-QE	9.7	8.7	-	-	-	-	-	-	50.9	13.1
	ULTRA-query	7.3	6.6	-	-	-	-	-	-	33.2	15.0
	ConE	7.4	7.0	-	-	-	-	-	-	20.3	11.9
	CQD	10.5	10.4	-	-	-	-	-	-	13.7	14.7
	CQD-hybrid	10.3	10.0	-	-	-	-	-	-	18.5	13.6
	QTO	12.2	11.6	-	-	-	-	-	-	35.3	11.3
	CLMPT	<b>12.7</b>	12.5	-	-	-	-	-	-	10.6	<b>27.7</b>

Table A.6. **There is no significant unbalancement in anchor nor relation names for each query type** . Most frequent relation name and anchor per query type in the old benchmarks. Highest value in bold.

Query	relation name	frequency(%)	anchor entity	frequency(%)
FB15k237				
1p	film_release_region	7.5	USA	0.7
2p	location/location/contains	8.6	USA	4.0
3p	location/location/contains	11.6	USA	3.3
2i	people/person/profession	16.4	USA	16.4
3i	people/person/gender	18.7	USA	<b>30.1</b>
pi	people/person/profession	13.7	USA	14.1
ip	people/person/profession	8.6	US dollars	5.7
2u	film/film_crew_role	14.6	USA	22.5
2u1p	taxonomy_entry/taxonomy	9.9	/m/08mbj5d	5.4
2in	people/person/profession	15.7	USA	15.2
3in	common/webpage/category	<b>20.7</b>	USA	28.5
2pi1pn	people/person/profession	12.4	USA	11.7
2nu1p	people/person/profession	17.1	USA	14.5
2in1p	location/contains	13.4	USA	14.3
NELL995				
1p	concept:atdate	9.4	concept_stateorprovince_california	0.5
2p	concept:mutualproxyfor	8.7	concept_lake_new	1.5
3p	concept:mutualproxyfor	15.9	concept_book_new	2.5
2i	concept:atdate	22.5	concept_book_new	7.6
3i	concept:atdate	21.6	concept_company_pbs	12.0
pi	concept:proxyfor	11.4	concept_book_new	6.1
ip	concept:subpartof	9.4	concept_athlete_sinorice_moss	3.8
2u	concept:atdate	10.8	concept_company_pbs	4.9
2u1p	concept:subpartof	11.1	concept_athlete_sinorice_moss	3.3
2in	concept:atdate	20.9	concept_book_new	6.8
3in	concept:atdate	<b>45.8</b>	concept_book_new	<b>14.2</b>
2pi1pn	concept:proxyfor	14.9	concept_book_new	5.7
2nu1p	concept:proxyfor	21.9	concept_book_new	8.0
2in1p	concept:atdate	15.0	concept_lake_new	6.1




Figure B.1. **Query structures including negation**, adapted from Ren & Leskovec (2020), where “2in1p” queries are referred as “inp”. We explicitly mention the number of steps involved in a path or conjunction, as this is a factor of complexity. Analogously, “pin”, “pni” queries from Ren & Leskovec (2020) are now “2pi1p”, and “2nu1p”. See App. B for their logical formulation.

## B. Query types including negation

In Figure B.1 we include the query types including the negation operator involved in the analysis described in App. A.1. Those queries are obtained by adding a negation to some of the query structures described in Sec. 2. For example, by negating a triple pattern involved in the intersection of a 2i, 3i, 2i1p queries, we obtain respectively:

$$?T : (a_1, r_1, T) \wedge \neg(a_2, r_2, T), \quad (2in)$$

$$?T : (a_1, r_1, T) \wedge (a_2, r_2, T) \wedge \neg(a_3, r_3, T), \quad (3in)$$

$$?T : \exists V_1. ((a_1, r_1, V_1) \wedge \neg(a_2, r_2, V_1) \wedge (V_1, r_3, T)), \quad (2in1p)$$

Moreover, by placing the negation in different triple patterns of the query type 1p2i, we obtain two query types:

$$?T : \exists V_1. ((a_1, r_1, V_1) \wedge (V_1, r_2, T) \wedge \neg(a_2, r_3, T)), \quad (2pi1p)$$

where the negation is placed in the triple pattern directly involved in the intersection, and

$$?T : \exists V_1. (\neg((a_1, r_1, V_1) \wedge (V_1, r_2, T)) \wedge (a_2, r_3, T)), \quad (n2pi1p)$$

where the negation is placed on the path query 2p. Using De Morgan laws, this query can also be rewritten as:

$$?T : \exists V_1. (\neg(a_1, r_1, V_1) \vee \neg(V_1, r_2, T)) \wedge (a_2, r_3, T), \quad (2nu1p)$$

Note that the visual representation of 2nu1p provided by (Ren & Leskovec, 2020) was misleading, as in practice the negation is applied on the whole 2p query, rather than just on one of its links.

## C. More related works

Guu et al. (2015) propose compositional training for embedding methods to predict answers for path queries. GQE (Hamilton et al., 2018) learns a geometric intersection operator to answer conjunctive queries in embedding space; this approach was later extended by Query2Box (Ren et al., 2020), BetaE (Ren & Leskovec, 2020), and GNN-QE (Zhu et al., 2022). FuzzQE (Chen et al., 2022) improves embedding methods with t-norm fuzzy logic, which satisfies the axiomatic system of classical logic. Some recent works such as HypE (Choudhary et al., 2021) and ConE (Zhang et al., 2021) use geometric interpretations of entity and relation embeddings to achieve desired properties for the logical operators. Other solutions to CQA combine neural methods with symbolic algorithms. For example, EmQL (Sun et al., 2020) ensembles an embedding model and a count-min sketch, and is able to find logically entailed answers, while CQD (Arakelyan et al., 2021; 2023) extends a pretrained knowledge graph embedding model to infer answers for complex queries.



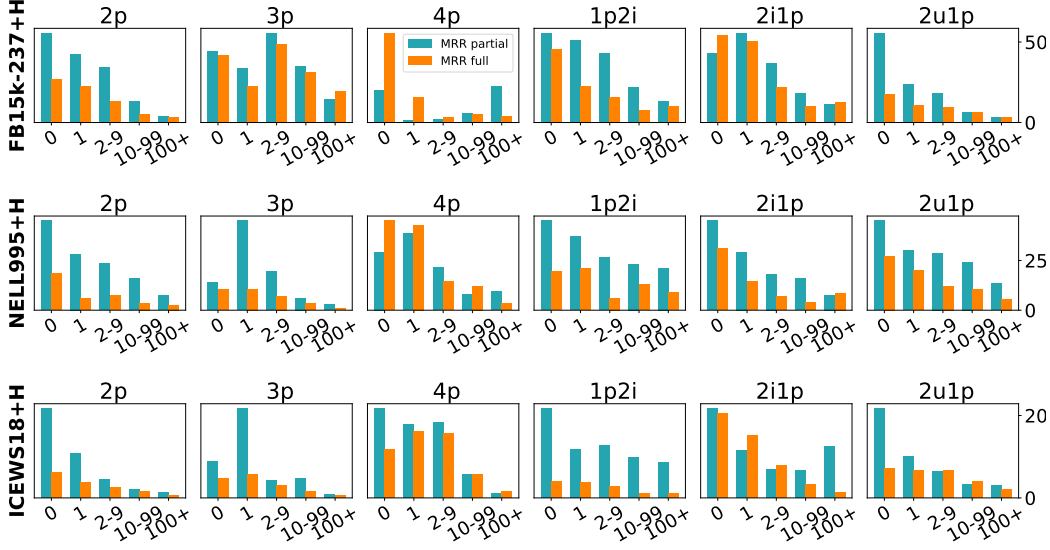


Figure D.1. **Higher cardinality of intermediate existing entities leads to lower MRR.** Influence of the number of existing intermediate variables on the MRR for datasets FB15k237+H, NELL995+H, ICEWS18+H, using CQD.

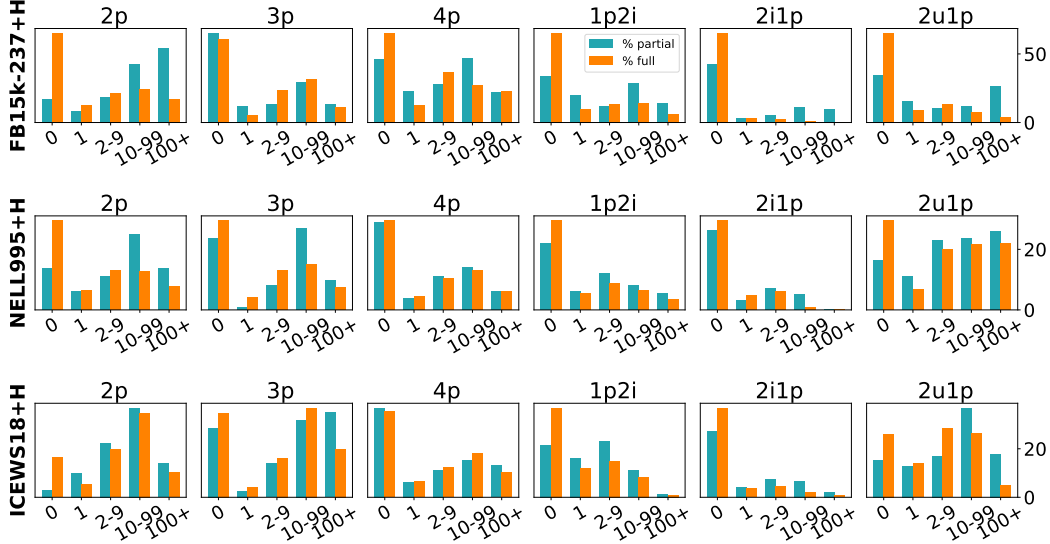


Figure D.2. **Proportion of QA pairs having different cardinality of intermediate existing entities in new benchmarks.**

## D. Additional Analysis on the new benchmarks

### D.1. Influence of intermediate existing entities on the results

Similarly to what was done for the old benchmarks in App. A, we analyze the proportion and influence (on the MRR) of the number of intermediate existing entities for CQD, on the new benchmarks FB15k237+H, NELL995+H, ICEWS18+H. Figure D.1 show a very similar result to the one in Figure A.1, showing a decreasing MRR at the increase of the value of cardinality for all benchmarks. Moreover, for reference, we also report the proportion of the different QA pairs for each cardinality category in Figure D.2.

## Is Complex Query Answering Really Complex?

Query	relation name	frequency(%)	anchor entity	frequency(%)
FB15k237+H				
1p	film_release_region	7.5	USA	0.7
2p	/people/person/nationality	15.3	/m/08mbj5d/m/08mbj5d	3.6
3p	/people/person/profession	14.2	/m/092ys-y	5.4
4p	/dataworld/gardening_hint/split.to	20.0	/m/0dq_5	5.5
2i	/people/person/profession	17.2	/m/09c7w0	15.5
3i	/people/person/profession	19.6	/m/09c7w0	19.7
4i	/people/person/profession	19.2	/m/09c7w0	20.0
pi	/people/person/nationality	19.7	/m/09c7w0	14.9
ip	/people/person/gender	18.5	/m/048_p	8.9
2in	people/person/profession	15.7	USA	15.2
3in	/people/person/gender	20.0	/m/09c7w0	20.0
2pi1pn	/people/person/nationality	18.4	/m/08mbj5d	13.3
2nu1p	people/person/profession	17.1	USA	14.5
2in1p	/people/person/profession	13.6	/m/08mbj5d	12.2
2u	/people/person/profession	15.6	/m/09c7w0	12.1
2u1p	/common/topic/webpage./common/webpage/category	19.4	/m/0kc9f	14.5
NELL995+H				
1p	concept:atdate	9.4	concept_stateorprovince_california	0.5
2p	concept:proxyfor	17.9	concept_sportsteam_ncaa_mens_midwest_regionals	6.4
3p	concept:proxyfor	20.0	concept_geopoliticallocation_national	5.9
4p	concept:atlocation	20.0	concept_website_new_york_times	6.5
2i	concept:atdate	20.0	concept_dateliteral_n2006	5.4
3i	concept:atdate	20.0	concept_sportsleague_mlb	5.7
4i	concept:atdate	20.0	concept_date_n2004	7.9
pi	concept:proxyfor	15.6	concept_book_new	7.8
ip	concept:atdate	19.4	concept_lake_new	9.4
2in	concept:atdate	20.9	concept_book_new	6.8
3in	concept:atdate	20.0	concept_sportsteamposition_center	15.5
2pi1pn	concept:proxyfor	20.0	concept_lake_new	15.2
2nu1p	concept:proxyfor	21.9	concept_book_new	8.0
2in1p	concept:atdate	20.0	concept_lake_new	11.2
2u	concept:atdate	20.0	concept_lake_new	2.7
2u1p	concept:proxyfor	19.8	concept_lake_new	11.0
ICEWS18+H				
1p	Make statement	18.8	United States	3.0
2p	Make statement	20.0	Russia	6.7
3p	Consult	20.0	United States	3.4
4p	Make statement	20.0	Yukio Edano	4.4
2i	Make statement	20.0	Citizien(India)	7.0
3i	Make statement	20.0	Citizien(India)	4.2
4i	Make statement	20.0	United States	4.5
pi	Make statement	20.0	Unspecified Actor	9.7
ip	Make statement	20.0	Russia	6.6
2in	Make statement	19.9	Unspecified Actor	17.1
3in	Make statement	20.0	Citizien(India)	20.0
2pi1pn	Make statement	20.0	United States	11.2
2nu1p	Make statement	19.9	Unspecified Actor	16.46
2in1p	Consult	20.0	United States	10.5
2u	Consult	19.9	China	3.5
2u1p	Make statement	19.9	United States	4.9

Table D.1. The new benchmarks are generated such that anchors and relation names cannot appear more than 20% in each query type. Most frequent relation name and anchor per query type in the new benchmarks.

**Imbalances of relation names and anchor nodes** When creating the new benchmarks FB15k237+H, NELL995+H, ICEWS18+H, to make sure that no anchor entities nor relation name was predominant in each query type, we set a maximum of 20% of QA pairs having the same anchor entity or relation name across all benchmarks and query type, as shown in Table D.1

**Stratified analysis on new benchmarks** For reference, we also provide the stratified analysis on the new benchmarks in Tables A.4 and A.5, which shows that 1) the full-inference QA pairs have always a lower MRR than the overall, and 2) the overall MRR is much lower than the 1p-reductions only, as for the new benchmark the proportion of QA pairs is equally distributed.

The same considerations also apply to queries including negation, as shown in Table D.6.

### D.2. Statistical tests

As an illustrative case, we analyze 2p queries from FB15k237+H to assess whether the observed differences in model performance are statistically significant, despite similar MRR in Table 5. We conduct an all-vs-all paired Mann–Whitney U test on the reciprocal ranks across all models for each 2p QA pair. The results in Table D.2 reveal strong statistical differences in most pairwise comparisons, indicating that models rank correct answers differently even when aggregate

### Is Complex Query Answering Really Complex?

Model	CLMPT	CQD	CQD-h	ConE	GNNQE	QTO	ULTRA
CLMPT	—	3.32e-63	1.82e-80	3.92e-01	8.28e-71	1.06e-188	8.55e-35
CQD	3.32e-63	—	1.48e-04	1.94e-52	1.72e-01	5.27e-38	2.31e-06
CQD-h	1.82e-80	1.48e-04	—	8.88e-75	5.08e-03	8.41e-15	6.91e-16
ConE	3.92e-01	1.94e-52	8.88e-75	—	6.65e-60	1.38e-165	7.04e-28
GNNQE	8.28e-71	1.72e-01	5.08e-03	6.65e-60	—	4.79e-30	2.32e-09
QTO	1.06e-188	5.27e-38	8.41e-15	1.38e-165	4.79e-30	—	2.22e-70
ULTRA	8.55e-35	2.31e-06	6.91e-16	7.04e-28	2.32e-09	2.22e-70	—

Table D.2. Pairwise Mann–Whitney U test  $p$ -values on reciprocal ranks of 2p queries in FB15k237+H. Lower values indicate stronger evidence of a significant difference between model performances.

metrics are close. However, a few model comparisons—such as ConE vs. CLMPT and CQD vs. GNNQE—exhibit statistical similarity, suggesting comparable ranking behavior on 2p queries. To examine whether the similarity is limited to specific query types, we apply the same test across all QA pairs of FB15k237+H from all query types. The resulting  $p$ -values are effectively zero for all comparisons, confirming that significant performance differences are widespread, though limited cases of similarity may occur between some models on some specific query types.

Table D.3. Stratified analysis on FB15k237+H

Query type	method	overall	1p	2p	3p	4p	2i	3i	4i	1p2i	2i1p	2u	2u1p
1p	GNN-QE	42.8	42.8	-	-	-	-	-	-	-	-	-	-
	ULTRA-query	40.6	40.6	-	-	-	-	-	-	-	-	-	-
	ConE	41.8	41.8	-	-	-	-	-	-	-	-	-	-
	CQD	46.7	46.7	-	-	-	-	-	-	-	-	-	-
	CQD-hybrid	46.7	46.7	-	-	-	-	-	-	-	-	-	-
	QTO	46.7	46.7	-	-	-	-	-	-	-	-	-	-
	CLMPT	45.3	45.3	-	-	-	-	-	-	-	-	-	-
2p	GNN-QE	5.2	11.8	4.4	-	-	-	-	-	-	-	-	-
	ULTRA-query	4.5	9.5	3.9	-	-	-	-	-	-	-	-	-
	ConE	4.6	9.9	4.0	-	-	-	-	-	-	-	-	-
	CQD	4.5	7.4	4.1	-	-	-	-	-	-	-	-	-
	CQD-hybrid	4.8	15.1	3.5	-	-	-	-	-	-	-	-	-
	QTO	4.9	14.4	3.8	-	-	-	-	-	-	-	-	-
	CLMPT	5.3	6.8	5.1	-	-	-	-	-	-	-	-	-
3p	GNN-QE	4.0	16.2	4.6	3.6	-	-	-	-	-	-	-	-
	ULTRA-query	3.5	14.0	3.3	3.2	-	-	-	-	-	-	-	-
	ConE	3.9	10.1	5.0	3.6	-	-	-	-	-	-	-	-
	CQD	2.4	3.7	2.2	2.4	-	-	-	-	-	-	-	-
	CQD-hybrid	3.1	10.7	3.8	2.7	-	-	-	-	-	-	-	-
	QTO	3.7	21.4	5.8	3.0	-	-	-	-	-	-	-	-
	CLMPT	4.7	8.3	4.1	4.5	-	-	-	-	-	-	-	-
4p	GNN-QE	4.3	16.1	14.1	4.8	3.8	-	-	-	-	-	-	-
	ULTRA-query	3.8	14.4	12.3	4.6	3.3	-	-	-	-	-	-	-
	ConE	3.5	12.0	8.8	4.2	3.2	-	-	-	-	-	-	-
	CQD	1.1	0.3	0.2	0.7	1.1	-	-	-	-	-	-	-
	CQD-hybrid	2.4	9.8	10.5	4.2	2.0	-	-	-	-	-	-	-
	QTO	3.9	18.2	21.7	5.5	3.2	-	-	-	-	-	-	-
	CLMPT	4.5	11.9	13.6	3.9	4.1	-	-	-	-	-	-	-
2i	GNN-QE	6.0	20.4	-	-	5.4	-	-	-	-	-	-	-
	ULTRA-query	5.2	20.7	-	-	4.6	-	-	-	-	-	-	-
	ConE	9.1	15.5	-	-	8.8	-	-	-	-	-	-	-
	CQD	11.3	18.4	-	-	10.8	-	-	-	-	-	-	-
	CQD-hybrid	6.0	22.0	-	-	5.4	-	-	-	-	-	-	-
	QTO	8.7	24.4	-	-	8.0	-	-	-	-	-	-	-
	CLMPT	10.2	18.4	-	-	9.7	-	-	-	-	-	-	-
3i	GNN-QE	8.8	32.2	-	-	-	2.5	7.5	-	-	-	-	-
	ULTRA-query	7.2	32.4	-	-	-	2.4	5.9	-	-	-	-	-
	ConE	10.3	25.0	-	-	-	6.9	9.0	-	-	-	-	-
	CQD	12.8	30.5	-	-	-	8.5	11.5	-	-	-	-	-
	CQD-hybrid	8.6	38.2	-	-	-	2.7	7.1	-	-	-	-	-
	QTO	10.1	40.0	-	-	-	6.5	8.2	-	-	-	-	-
	CLMPT	12.2	31.8	-	-	-	7.6	10.7	-	-	-	-	-
4i	GNN-QE	19.6	47.9	-	-	-	18.3	16.5	15.8	-	-	-	-
	ULTRA-query	16.4	48.9	-	-	-	17.5	13.2	12.9	-	-	-	-
	ConE	20.3	36.9	-	-	-	24.9	18.2	15.8	-	-	-	-
	CQD	23.8	45.1	-	-	-	27.0	20.9	19.8	-	-	-	-
	CQD-hybrid	17.4	58.3	-	-	-	17.2	13.2	14.6	-	-	-	-
	QTO	20.2	60.2	-	-	-	25.1	16.3	15.6	-	-	-	-
	CLMPT	24.0	40.5	-	-	-	26.7	21.9	18.7	-	-	-	-
1p2i	GNN-QE	5.6	28.1	6.3	-	-	3.0	-	-	3.5	-	-	-
	ULTRA-query	5.3	31.4	7.6	-	-	1.9	-	-	3.1	-	-	-
	ConE	3.8	16.7	5.1	-	-	2.5	-	-	2.6	-	-	-
	CQD	6.0	20.1	7.2	-	-	3.5	-	-	4.2	-	-	-
	CQD-hybrid	5.5	33.2	7.7	-	-	2.6	-	-	3.0	-	-	-
	QTO	6.1	32.0	8.1	-	-	3.8	-	-	3.6	-	-	-
	CLMPT	5.6	18.1	5.5	-	-	26.5	-	-	4.3	-	-	-
2i1p	GNN-QE	9.9	13.5	6.1	-	-	16.8	-	-	-	7.2	-	-
	ULTRA-query	10.1	13.8	5.7	-	-	18.4	-	-	-	7.2	-	-
	ConE	7.9	10.8	4.8	-	-	10.4	-	-	-	7.2	-	-
	CQD	11.5	12.2	6.4	-	-	21.7	-	-	-	10.5	-	-
	CQD-hybrid	12.9	23.0	6.0	-	-	24.7	-	-	-	8.2	-	-
	QTO	13.5	23.5	6.6	-	-	27.3	-	-	-	9.0	-	-
	CLMPT	14.9	10.5	9.4	-	-	11.0	-	-	-	14.2	-	-
2u	GNN-QE	32.4	-	-	-	-	-	-	-	-	-	32.4	-
	ULTRA-query	29.4	-	-	-	-	-	-	-	-	-	29.4	-
	ConE	22.8	-	-	-	-	-	-	-	-	-	22.8	-
	CQD	40.0	-	-	-	-	-	-	-	-	-	40.0	-
	CQD-hybrid	42.2	-	-	-	-	-	-	-	-	-	42.2	-
	QTO	30.6	-	-	-	-	-	-	-	-	-	30.6	-
	CLMPT	33.6	-	-	-	-	-	-	-	-	-	33.6	-
2u1p	GNN-QE	10.0	7.6	-	-	-	-	-	-	-	-	54.0	6.5
	ULTRA-query	8.3	7.4	-	-	-	-	-	-	-	-	42.2	4.8
	ConE	6.0	6.0	-	-	-	-	-	-	-	-	17.1	4.7
	CQD	10.6	9.2	-	-	-	-	-	-	-	-	13.9	8.4
	CQD-hybrid	12.0	10.5	-	-	-	-	-	-	-	-	44.0	8.1
	QTO	11.2	10.0	-	-	-	-	-	-	-	-	53.3	6.7
	CLMPT	14.2	8.3	-	-	-	-	-	-	-	-	9.1	13.4



Table D.4. Stratified analysis on NELL995+H

Query type	method	overall	1p	2p	3p	4p	2i	3i	4i	1p2i	2i1p	2u	2u1p
1p	GNN-QE	53.6	53.6	-	-	-	-	-	-	-	-	-	-
	ULTRA-query	38.9	38.9	-	-	-	-	-	-	-	-	-	-
	ConE	60.0	60.0	-	-	-	-	-	-	-	-	-	-
	CQD	60.4	60.4	-	-	-	-	-	-	-	-	-	-
	CQD-hybrid	60.4	60.4	-	-	-	-	-	-	-	-	-	-
	QTO	60.3	60.3	-	-	-	-	-	-	-	-	-	-
	CLMPT	58.1	58.1	-	-	-	-	-	-	-	-	-	-
2p	GNN-QE	8.0	21.6	6.3	-	-	-	-	-	-	-	-	-
	ULTRA-query	6.1	16.2	5.0	-	-	-	-	-	-	-	-	-
	ConE	7.9	15.3	7.1	-	-	-	-	-	-	-	-	-
	CQD	9.6	21.7	8.0	-	-	-	-	-	-	-	-	-
	CQD-hybrid	9.0	28.4	6.6	-	-	-	-	-	-	-	-	-
	QTO	9.8	25.9	7.7	-	-	-	-	-	-	-	-	-
	CLMPT	10.1	17.3	9.0	-	-	-	-	-	-	-	-	-
3p	GNN-QE	6.0	33.3	14.1	3.9	-	-	-	-	-	-	-	-
	ULTRA-query	4.1	14.7	9.4	3.1	-	-	-	-	-	-	-	-
	ConE	6.7	21.0	12.2	5.5	-	-	-	-	-	-	-	-
	CQD	4.2	6.2	5.6	3.7	-	-	-	-	-	-	-	-
	CQD-hybrid	6.1	22.2	10.9	4.4	-	-	-	-	-	-	-	-
	QTO	8.0	42.1	15.3	5.2	-	-	-	-	-	-	-	-
	CLMPT	7.8	20.1	10.6	6.4	-	-	-	-	-	-	-	-
4p	GNN-QE	4.7	4.9	25.0	11.9	2.8	-	-	-	-	-	-	-
	ULTRA-query	4.2	30.5	17.1	11.5	2.9	-	-	-	-	-	-	-
	ConE	5.2	27.3	17.3	7.5	4.5	-	-	-	-	-	-	-
	CQD	2.0	2.7	1.4	0.9	2.0	-	-	-	-	-	-	-
	CQD-hybrid	3.5	34.9	19.3	8.1	2.1	-	-	-	-	-	-	-
	QTO	7.0	54.4	29.1	12.4	4.9	-	-	-	-	-	-	-
	CLMPT	7.2	30.0	18.7	7.6	5.8	-	-	-	-	-	-	-
2i	GNN-QE	10.7	24.5	-	-	9.8	-	-	-	-	-	-	-
	ULTRA-query	7.9	18.7	-	-	7.3	-	-	-	-	-	-	-
	ConE	21.8	23.3	-	-	21.1	-	-	-	-	-	-	-
	CQD	18.5	26.3	-	-	17.6	-	-	-	-	-	-	-
	CQD-hybrid	12.1	28.7	-	-	11.0	-	-	-	-	-	-	-
	QTO	14.6	30.5	-	-	13.5	-	-	-	-	-	-	-
	CLMPT	22.7	27.1	-	-	21.9	-	-	-	-	-	-	-
3i	GNN-QE	13.3	34.7	-	-	-	11.7	10.6	-	-	-	-	-
	ULTRA-query	10.2	26.7	-	-	-	11.4	7.1	-	-	-	-	-
	ConE	23.6	31.8	-	-	-	22.5	21.5	-	-	-	-	-
	CQD	19.6	34.6	-	-	-	17.4	17.1	-	-	-	-	-
	CQD-hybrid	14.4	43.3	-	-	-	11.7	11.8	-	-	-	-	-
	QTO	15.8	46.2	-	-	-	14.1	12.8	-	-	-	-	-
	CLMPT	25.0	36.3	-	-	-	23.4	22.2	-	-	-	-	-
4i	GNN-QE	19.4	51.9	-	-	-	24.6	14.3	13.5	-	-	-	-
	ULTRA-query	15.6	42.6	-	-	-	24.9	12.2	8.2	-	-	-	-
	ConE	27.6	47.9	-	-	-	35.0	21.8	23.8	-	-	-	-
	CQD	24.8	52.9	-	-	-	30.6	17.7	20.8	-	-	-	-
	CQD-hybrid	20.4	65.1	-	-	-	24.9	14.2	14.5	-	-	-	-
	QTO	20.9	68.3	-	-	-	26.8	13.8	15.4	-	-	-	-
	CLMPT	29.1	50.5	-	-	-	36.9	22.4	24.6	-	-	-	-
1p2i	GNN-QE	16.0	55.9	19.5	-	-	9.1	-	-	7.3	-	-	-
	ULTRA-query	15.8	55.0	23.0	-	-	5.5	-	-	5.9	-	-	-
	ConE	14.9	32.6	18.9	-	-	10.7	-	-	7.0	-	-	-
	CQD	18.9	54.3	22.1	-	-	12.0	-	-	10.0	-	-	-
	CQD-hybrid	17.4	64.7	19.0	-	-	12.4	-	-	8.1	-	-	-
	QTO	17.6	58.4	20.6	-	-	13.4	-	-	7.9	-	-	-
	CLMPT	17.2	37.9	18.2	-	-	13.1	-	-	10.6	-	-	-
2i1p	GNN-QE	13.5	26.2	12.9	-	-	17.0	-	-	-	8.2	-	-
	ULTRA-query	9.3	22.2	7.6	-	-	22.3	-	-	-	6.0	-	-
	ConE	11.8	22.6	11.1	-	-	23.2	-	-	-	8.4	-	-
	CQD	22.6	29.9	20.1	-	-	36.9	-	-	-	16.1	-	-
	CQD-hybrid	21.2	35.8	18.9	-	-	34.5	-	-	-	13.7	-	-
	QTO	21.1	36.3	19.2	-	-	36.0	-	-	-	13.4	-	-
	CLMPT	24.4	25.9	21.8	-	-	35.2	-	-	-	20.0	-	-
2u	GNN-QE	47.5	-	-	-	-	-	-	-	-	-	47.5	-
	ULTRA-query	28.1	-	-	-	-	-	-	-	-	-	28.1	-
	ConE	39.9	-	-	-	-	-	-	-	-	-	39.9	-
	CQD	46.3	-	-	-	-	-	-	-	-	-	46.3	-
	CQD-hybrid	46.4	-	-	-	-	-	-	-	-	-	46.4	-
	QTO	49.1	-	-	-	-	-	-	-	-	-	49.1	-
	CLMPT	50.0	-	-	-	-	-	-	-	-	-	50.0	-
2u1p	GNN-QE	9.8	11.4	-	-	-	-	-	-	-	-	42.9	5.2
	ULTRA-query	9.5	12.3	-	-	-	-	-	-	-	-	36.9	5.1
	ConE	8.8	10.3	-	-	-	-	-	-	-	-	21.9	6.3
	CQD	18.5	18.8	-	-	-	-	-	-	-	-	53.5	13.4
	CQD-hybrid	19.3	20.5	-	-	-	-	-	-	-	-	63.7	13.1
	QTO	18.9	21.1	-	-	-	-	-	-	-	-	58.2	11.7
	CLMPT	22.0	17.6	-	-	-	-	-	-	-	-	29.7	19.1

Table D.5. Stratified analysis on ICEWS18+H

Query type	method	overall	1p	2p	3p	4p	2i	3i	4i	1p2i	2i1p	2u	2u1p
1p	GNN-QE	12.2	12.2	-	-	-	-	-	-	-	-	-	-
	ULTRA-query	6.3	6.3	-	-	-	-	-	-	-	-	-	-
	ConE	3.5	3.5	-	-	-	-	-	-	-	-	-	-
	CQD	16.6	16.6	-	-	-	-	-	-	-	-	-	-
	CQD-hybrid	16.6	16.6	-	-	-	-	-	-	-	-	-	-
	QTO	16.6	16.6	-	-	-	-	-	-	-	-	-	-
	CLMPT	4.7	4.7	-	-	-	-	-	-	-	-	-	-
2p	GNN-QE	0.9	2.2	0.8	-	-	-	-	-	-	-	-	-
	ULTRA-query	1.2	2.1	1.1	-	-	-	-	-	-	-	-	-
	ConE	0.9	1.9	0.9	-	-	-	-	-	-	-	-	-
	CQD	2.6	5.4	2.4	-	-	-	-	-	-	-	-	-
	CQD-hybrid	2.6	10.0	2.0	-	-	-	-	-	-	-	-	-
	QTO	2.6	8.4	2.2	-	-	-	-	-	-	-	-	-
	CLMPT	0.8	0.3	0.8	-	-	-	-	-	-	-	-	-
3p	GNN-QE	0.5	1.0	1.2	0.4	-	-	-	-	-	-	-	-
	ULTRA-query	1.2	2.1	1.5	1.2	-	-	-	-	-	-	-	-
	ConE	0.9	1.5	1.4	0.9	-	-	-	-	-	-	-	-
	CQD	1.5	2.6	3.2	1.4	-	-	-	-	-	-	-	-
	CQD-hybrid	1.5	9.8	4.4	1.2	-	-	-	-	-	-	-	-
	QTO	1.4	6.2	3.7	1.2	-	-	-	-	-	-	-	-
	CLMPT	0.1	0.2	0.1	0.1	-	-	-	-	-	-	-	-
4p	GNN-QE	0.4	1.6	1.4	1.0	0.3	-	-	-	-	-	-	-
	ULTRA-query	0.8	2.9	1.9	0.9	0.7	-	-	-	-	-	-	-
	ConE	0.6	1.8	1.8	0.9	0.6	-	-	-	-	-	-	-
	CQD	1.1	1.6	1.5	1.4	1.1	-	-	-	-	-	-	-
	CQD-hybrid	0.9	31.4	39.3	20.1	0.7	-	-	-	-	-	-	-
	QTO	1.1	14.3	5.4	2.5	0.7	-	-	-	-	-	-	-
	CLMPT	0.1	0.1	0.1	0.1	0.1	-	-	-	-	-	-	-
2i	GNN-QE	16.1	30.0	-	-	2.4	-	-	-	-	-	-	-
	ULTRA-query	7.0	12.3	-	-	2.1	-	-	-	-	-	-	-
	ConE	1.2	1.9	-	-	0.8	-	-	-	-	-	-	-
	CQD	13.0	22.1	-	-	4.3	-	-	-	-	-	-	-
	CQD-hybrid	15.0	27.3	-	-	3.3	-	-	-	-	-	-	-
	QTO	15.7	28.3	-	-	3.4	-	-	-	-	-	-	-
	CLMPT	12.0	20.5	-	-	2.3	-	-	-	-	-	-	-
3i	GNN-QE	26.5	58.5	-	-	-	19.3	2.0	-	-	-	-	-
	ULTRA-query	11.7	25.9	-	-	-	8.0	1.6	-	-	-	-	-
	ConE	0.5	1.2	-	-	-	0.6	0.2	-	-	-	-	-
	CQD	19.5	42.3	-	-	-	13.4	3.2	-	-	-	-	-
	CQD-hybrid	25.6	58.5	-	-	-	15.8	2.7	-	-	-	-	-
	QTO	24.7	56.7	-	-	-	16.3	2.5	-	-	-	-	-
	CLMPT	23.0	44.0	-	-	-	21.4	2.9	-	-	-	-	-
4i	GNN-QE	34.0	75.0	-	-	-	44.9	13.9	1.6	-	-	-	-
	ULTRA-query	15.9	37.1	-	-	-	19.1	6.0	1.0	-	-	-	-
	ConE	0.3	0.9	-	-	-	0.4	0.2	0.1	-	-	-	-
	CQD	24.0	54.2	-	-	-	28.8	10.2	2.3	-	-	-	-
	CQD-hybrid	33.2	78.9	-	-	-	39.4	11.8	2.1	-	-	-	-
	QTO	31.5	73.6	-	-	-	38.0	11.9	18.6	-	-	-	-
	CLMPT	31.0	60.3	-	-	-	41.5	18.4	2.7	-	-	-	-
1p2i	GNN-QE	19.1	62.5	30.1	-	-	2.5	-	-	2.3	-	-	-
	ULTRA-query	8.8	19.6	13.2	-	-	28	-	-	2.2	-	-	-
	ConE	1.2	2.6	1.5	-	-	1.0	-	-	0.8	-	-	-
	CQD	17.1	38.5	25.8	-	-	50.1	-	-	4.4	-	-	-
	CQD-hybrid	17.5	66.9	26.2	-	-	4.1	-	-	3.5	-	-	-
	QTO	18.4	59.3	28.0	-	-	4.4	-	-	3.6	-	-	-
	CLMPT	9.7	8.5	15.1	-	-	1.0	-	-	2.3	-	-	-
2i1p	GNN-QE	3.5	4.5	3.0	-	-	5.7	-	-	-	3.4	-	-
	ULTRA-query	1.3	2.3	1.2	-	-	1.6	-	-	-	1.2	-	-
	ConE	1.6	2.4	1.6	-	-	2.4	-	-	-	1.3	-	-
	CQD	6.9	6.1	6.0	-	-	11.4	-	-	-	6.8	-	-
	CQD-hybrid	5.8	9.2	4.6	-	-	14.8	-	-	-	5.0	-	-
	QTO	6.2	10.3	5.5	-	-	13.1	-	-	-	5.1	-	-
	CLMPT	2.1	0.5	1.6	-	-	0.4	-	-	-	2.3	-	-
2u	GNN-QE	7.6	-	-	-	-	-	-	-	-	-	7.6	-
	ULTRA-query	3.3	-	-	-	-	-	-	-	-	-	3.3	-
	ConE	1.1	-	-	-	-	-	-	-	-	-	1.1	-
	CQD	6.8	-	-	-	-	-	-	-	-	-	6.8	-
	CQD-hybrid	6.8	-	-	-	-	-	-	-	-	-	6.8	-
	QTO	6.7	-	-	-	-	-	-	-	-	-	6.7	-
	CLMPT	2.7	-	-	-	-	-	-	-	-	-	2.7	-
2u1p	GNN-QE	1.1	0.8	-	-	-	-	-	-	-	-	4.7	0.8
	ULTRA-query	1.2	1.5	-	-	-	-	-	-	-	-	2.0	1.0
	ConE	0.9	1.0	-	-	-	-	-	-	-	-	1.3	0.8
	CQD	5.9	3.2	-	-	-	-	-	-	-	-	11.6	5.6
	CQD-hybrid	5.6	2.7	-	-	-	-	-	-	-	-	17.8	4.5
	QTO	4.9	3.3	-	-	-	-	-	-	-	-	17.8	3.7
	CLMPT	2.2	0.3	-	-	-	-	-	-	-	-	0.3	2.5

Table D.6. Also for queries involving negation, there is no clear SoTA method for the new benchmarks. The MRR on the new benchmarks is significantly lower than the old ones. For example, for 3in queries on FB15k237+H, QTO has an MRR of only 3.1, while for FB15k237 it was 16.1.

Query type	method	FB15k237+H			NELL995+H			ICEWS18+H		
		ovr	part-inf	full-inf	ovr	par-inf	full-inf	ovr	par-inf	full-inf
2in	GNN-QE	6.8	-	6.8	5.5	-	5.5	4.5	-	4.5
	ULTRAQ	5.3	-	5.3	4.5	-	4.5	2.3	-	2.3
	CQD	3.3	-	3.3	4.2	-	4.2	1.5	-	1.5
	CQD-hybrid	4.7	-	4.7	5.1	-	5.1	1.7	-	1.7
	QTO	<b>10.6</b>	-	<b>10.6</b>	<b>10.2</b>	-	<b>10.2</b>	<b>4.9</b>	-	<b>4.9</b>
	ConE	5.1	-	5.1	4.6	-	4.6	1.7	-	1.7
	CLMPT	6.8	-	6.8	6.5	-	6.5	1.2	-	1.2
3in	GNN-QE	<b>6.5</b>	9.0	1.0	<b>6.4</b>	7.2	1.0	6.9	12.6	1.4
	ULTRAQ	5.5	7.5	0.6	5.9	6.2	0.9	4.8	7.1	2.4
	CQD	2.6	5.9	<b>2.5</b>	1.5	6.0	1.3	2.9	4.0	1.6
	CQD-hybrid	1.6	8.4	1.2	1.2	7.9	0.9	4.0	6.5	1.3
	QTO	3.1	13.9	2.3	2.3	11.8	<b>1.8</b>	<b>8.7</b>	14.5	<b>2.8</b>
	ConE	4.9	6.4	2.0	6.0	5.8	<b>2.8</b>	2.9	3.3	2.0
	CLMPT	2.3	9.3	1.7	2.4	5.7	2.1	2.1	3.4	0.9
2pi1pn	GNN-QE	<b>3.7</b>	4.0	1.4	<b>5.8</b>	7.0	1.6	0.9	1.0	0.5
	ULTRAQ	2.6	2.8	1.3	4.3	5.2	1.7	<b>1.2</b>	1.3	0.9
	CQD	0.6	1.3	0.5	1.5	3.2	1.4	0.2	0.7	0.1
	CQD-hybrid	1.0	2.7	0.9	1.4	5.9	1.1	0.3	1.0	0.2
	QTO	2.0	7.3	<b>1.7</b>	3.1	9.4	<b>2.6</b>	<b>1.2</b>	3.4	<b>1.0</b>
	ConE	2.9	2.7	1.5	3.7	4.9	2.0	1.1	1.3	<b>1.0</b>
	CLMPT	1.6	2.9	1.6	4.1	4.3	2.1	0.1	0.1	0.1
2nu1p	GNN-QE	<b>5.0</b>	-	<b>5.0</b>	3.3	-	3.3	<b>3.5</b>	-	<b>3.5</b>
	ULTRAQ	3.7	-	3.7	2.7	-	2.7	2.2	-	2.2
	CQD	4.9	-	4.9	4.9	-	4.9	2.7	-	2.7
	CQD-hybrid	3.2	-	3.2	4.3	-	4.3	2.2	-	2.2
	QTO	5.3	-	5.3	<b>8.4</b>	-	<b>8.4</b>	3.0	-	3.0
	ConE	3.3	-	3.3	2.7	-	2.7	0.9	-	0.9
	CLMPT	4.8	-	4.8	2.3	-	2.3	1.0	-	1.0
2in1p	GNN-QE	3.3	2.4	1.4	4.4	4.6	2.1	0.8	1.0	0.5
	ULTRAQ	2.2	2.2	1.2	3.6	3.7	1.6	<b>1.6</b>	2.2	<b>1.4</b>
	CQD	1.2	1.3	1.2	2.6	3.5	2.6	0.9	0.2	0.9
	CQD-hybrid	1.3	1.1	1.2	2.4	4.1	2.3	1.1	1.9	1.0
	QTO	1.5	5.4	1.3	2.4	7.9	1.9	0.9	2.5	0.9
	ConE	<b>3.6</b>	2.8	2.2	<b>6.4</b>	7.0	<b>3.9</b>	1.3	2.3	1.2
	CLMPT	2.5	3.4	<b>2.4</b>	4.5	6.3	3.8	0.2	0.1	0.2

## Is Complex Query Answering Really Complex?

Dataset	2p		3p		2i		3i		1p2i		2i1p		2u		2u1p		2in		3in		2pi1pn		2nu1p		2in1p	
	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn
FB15k237	512	p	8	p	128	p	128	p	256	p	64	p	512	m	512	m	2	p	2	p	512	prod	512	prod	512	prod
NELL995	512	p	2	p	2	p	2	p	256	p	256	p	512	m	512	m	2	p	2	p	512	prod	512	prod	512	prod

Table F.1. CQD hyperparameters old benchmarks. tn = tnorm. p=prod, m=min. Note that for 1p neither the k nor the tnorm are needed.

### E. Queries with four hops

Logical formulation of the 4p and 4i queries we introduced in the benchmarks FB15k237+H, NELL995+H, ICEWS18+H

$$\begin{aligned} ?T : \exists V_1, V_2, V_3. (a_1, r_1, V_1) \wedge (V_1, r_2, V_2) \wedge \\ (V_2, r_3, V_3) \wedge (V_3, r_4, T), \end{aligned} \quad (4p)$$

$$\begin{aligned} ?T : (a_1, r_1, T) \wedge (a_2, r_2, T) \wedge \\ (a_3, r_3, T) \wedge (a_4, r_4, T). \end{aligned} \quad (4i)$$

### F. Hyperparameters

In this section, we detail the hyperparameters used for each dataset and model. Old and new benchmarks, the generation scripts, and the implementation of CQD-hybrid are included in our official repo.<sup>6</sup>

#### F.1. Old benchmarks

**GNN-QE** We did not tune hyperparameters, but re-used the ones provided in the official repo.<sup>7</sup>

**ULTRA** We did not tune hyperparameters, but re-used the ones provided in the official repo.<sup>8</sup>

**CQD** As mentioned in Sec. 5, we re-used the pre-trained link predictor (Trouillon et al., 2017) provided by the authors.<sup>9</sup> However, we tuned CQD-specific hyperparameters, namely the CQD beam “k”, ranging from [2,512] and the t-norm type being “prod” or “min”. In Table F.1 we provide the hyperparameter selection for the old benchmarks FB15k237 and NELL995. Also note that we normalize scores with min-max normalization. For negation, we always use standard negation.

**QTO** We re-used the hyperparameters provided in the official repo.<sup>10</sup> For a fair comparison with CQD, CQD-hybrid, we used the same pre-trained link predictor.

**CLMPT** We did not tune hyperparameters but re-used the ones provided in the official repo.<sup>11</sup>

**ConE** We did not tune hyperparameters but re-used the ones provided in the official repo.<sup>12</sup>

**CQD-Hybrid** For CQD-Hybrid, to make the comparison with CQD fair, we re-used the hyperparameters found for CQD and fixed an upper bound value for the CQD beam “k” to 512, even when there are more existing entities matching the existentially quantified variables, to match the upper bound of the “k” used for CQD. Additionally, the scores of the pre-trained link predictor are normalized between [0, 0.9] using min-max normalization, and a score of 1 is assigned to the existing triples.

<sup>6</sup><https://github.com/april-tools/is-cqa-complex>

<sup>7</sup><https://github.com/DeepGraphLearning/GNN-QE/tree/master/config>

<sup>8</sup><https://github.com/DeepGraphLearning/ULTRA/tree/main/config/ultraquery>

<sup>9</sup><https://github.com/Blidge/KGReasoning/>

<sup>10</sup><https://github.com/bys0318/QTO/tree/main>

<sup>11</sup><https://github.com/qianlima-lab/clmpt>

<sup>12</sup><https://github.com/MIRALab-USTC/QE-ConE/blob/main/scripts.sh>

## Is Complex Query Answering Really Complex?

Dataset	2p		3p		2i		3i		1p2i		2i1p		2u		2u1p		4i		4p		2in		3in		2pi1pn		2nu1p		2in1p	
	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn	k	tn
ICEWS18+H	32	p	2	p	2	p	2	p	512	p	256	p	2	m	2	m	2	p	2	p	2	p	2	p	512	prod	512	prod	512	prod

Table F.2. CQD hyperparameters ICEWS18+H benchmark. tn = tnorm. p=prod, m=min

### F.2. New benchmarks

For FB15k237+H and NELL995+H, we re-used the same models trained for the old benchmarks, and using the same hyperparameters presented in App. F.1. The only exception is the hyperparameters for CQD and CQD-hybrid of 4p and 4i queries that were not tuned for the old benchmarks. We used k=2 and tnorm=prod for both query types in all benchmarks.

Instead, for the new benchmark ICEWS18+H we trained every model. The used hyperparameters for each model are presented in the following:

**GNN-QE** For GNN-QE, we tuned the following hyperparameters: (1) *batchsize*, with values 8 or 48, and *concat hidden* being True or False, while the rest are the same used for the old benchmarks and do not change across benchmarks. For ICEWS18+H the best hyperparameters are “batchsize=48” and “concat hidden=True”.

**ULTRAQ** Being a zero-shot neural link predictor, we re-used the same checkpoint provided in the official repo, as for F.1.

**CQD** We train ComplEx (Trouillon et al., 2017) link predictor with hyperparameters “regweight” 0.1 or 0.01, and batch size 1000 or 2000, with the best being, respectively 0.1 and 1000. Moreover, in F.2 are shown the hyperparameters for CQD on the new benchmark ICEWS18+H.

**QTO** We re-used the link predictor used in CQD and CQD-Hybrid, and for QTO-specific hyperparameters, we used the same of FB15k237.

**CQD-Hybrid** For CQD-Hybrid, as mentioned above, we re-used the hyperparameters used for CQD. Note that for 4p queries, we fixed the “k” upper bound to 64 for memory constraints.

**ConE** We re-used the same hyperparameters of NELL995 for ICEWS18+H.

**CLMPT** For CLMPT, we tuned the following hyperparameters: (1) learning rate, with values in [1e-5, 5e-2, 5e-3, 5e-4, 5e-5, 5e-6], (2) temp, with values in [0.1, 0.2] We found the best hyperparameters to match the same found for FB15k237. Hence, we re-used them for ICEWS18+H.

## G. Statistics

### G.1. KG splits statistics

The statistics, i.e., number of entities, relation names, training/validation/test links, of the knowledge graphs used in this paper are shown in Table G.1.

Table G.1. Statistics of knowledge graphs used to generate complex queries

Dataset	Entities	Relation Names	Training Links	Validation Links	Test Links	Total Links
FB15k237	14,505	237	272,115	17,526	20,438	310,079
NELL995	63,361	200	114,213	14,324	14,267	142,804
ICEWS18+CQ ( <i>new</i> )	20,840	250	213,304	25,048	24,689	263,041

The KG splits of NELL995 and FB15k237 are the same of the one used in (Ren & Leskovec, 2020; Ren et al., 2020).

## G.2. New benchmark statistics

### G.2.1. NELL995+H AND FB15k237+H

**Training queries** were not changed from the one proposed by [Ren & Leskovec \(2020\)](#); [Ren et al. \(2020\)](#).

**Validation and Test** queries are generated to ensure a balanced distribution of hardness across the sub-query types to which a given query can be reduced. As described in Sec. 6, we generate queries such that each sub-query type contains exactly 10,000 QA pairs, along with an additional 10,000 pairs corresponding to full-inference queries (i.e., non-reducible cases). To achieve this, we adapt the algorithm from ([Ren et al., 2020](#)), sampling queries until the required number of QA pairs for each sub-query type is reached. For instance, for 2p queries, we continue sampling until we obtain 10,000 QA pairs reducible to 1p and 10,000 full-inference queries. When a query is sampled, all its corresponding answers are included in the benchmark until the target of 10,000 QA pairs per subtype is met. If incorporating all answers would exceed this limit, a random subset is selected to ensure exactly 10,000 pairs. Similarly, once a sub-query type reaches the 10,000-pair threshold, any additional answers belonging to that category are discarded.

Only exceptions are **1p**, **2in**, and **2nu1p**, which being already of full-inference type we kept the same as in ([Ren & Leskovec, 2020](#); [Ren et al., 2020](#)).

### G.2.2. ICEWS18+H

**Training queries** were generated using the same strategy used in ([Ren & Leskovec, 2020](#); [Ren et al., 2020](#)), namely by (1) creating **1p** queries by considering every triple in the training set, (2) retrieve the number  $n$  of generated **1p** queries, (3) generate  $n$  queries for all training query types.

**Validation and test queries** were generated using the same strategy detailed in App. G.2.1. Moreover, since we generated this benchmark from scratch, we also had to generate **1p** (which following ([Ren & Leskovec, 2020](#); [Ren et al., 2020](#)), were generated for each triple in the validation and test split), and **2in**, **2nu1p** queries.