

FEDERAL UNIVERSITY OF JUIZ DE FORA
EXACT SCIENCES INSTITUTE
BACHELOR'S DEGREE IN COMPUTER SCIENCE

A hereditary attentive question answering frameworks for knowledge bases

Rômulo Crispim de Mello

JUIZ DE FORA
02, 2025

A hereditary attentive question answering frameworks for knowledge bases

RÔMULO CHRISPIM DE MELLO

Federal University of Juiz de Fora
Exact Sciences Institute
Department of Computer Science
Bachelor's Degree in Computer Science

Supervisor: Jairo Francisco de Souza
Co-supervisor: Victor Strole de Andrade Menezes

JUIZ DE FORA

02, 2025

A HEREDITARY ATTENTIVE QUESTION ANSWERING
FRAMEWORKS FOR KNOWLEDGE BASES

Rômulo Chrispim de Mello

MONOGRAPHY SUBMITTED TO THE FACULTY OF THE EXACT SCIENCES INSTITUTE OF FEDERAL UNIVERSITY OF JUIZ DE FORA, AS AN INTEGRAL PART OF THE REQUIREMENTS NECESSARY TO OBTAIN THE BACHELOR'S DEGREE IN COMPUTER SCIENCE.

Approved by:

Jairo Francisco de Souza
PhD in Informatics, PUC-Rio

Victor Strole de Andrade Menezes
PhD in Systems and Computing Engineering, UFRJ

Jorão Gomes Jr.
Master in Computer Science, UFJF

Fabricio Martins Mendonça
PhD in Information Science, UFMG

JUIZ DE FORA

2025, 02 05

To my friends and brothers.

To the parents, for their support.

Resumo

Este trabalho descreve a criação de um framework para abordagens de C-KBQA (Answering Complex Questions over Conceptual Knowledge Bases), que é um problema de grande importância no campo da Ciência da Computação. O objetivo principal é desenvolver um sistema capaz de responder perguntas complexas que requerem o uso de conhecimento conceitual, como ontologias e taxonomias. O framework proposto tem como intuito superar os desafios enfrentados por soluções anteriores e oferecer vantagens significativas em relação a outras soluções existentes.

A pesquisa foi realizada na Universidade Federal de Juiz de Fora, com a orientação de Jairo Francisco de Souza e coorientação de Victor Strole de Andrade Menezes. O framework é composto por dois módulos principais: o módulo de pré-processamento e o módulo de geração de consultas. O módulo de pré-processamento é responsável por extrair informações relevantes de uma base de conhecimento e transformá-las em um formato adequado para a geração de consultas. Já o módulo de geração de consultas tem a função de gerar consultas com base nas perguntas do usuário e nas informações extraídas anteriormente.

O framework proposto traz consigo diversas vantagens em relação a outras soluções existentes. Ele é capaz de lidar com a ambiguidade e a incerteza presentes nas perguntas, além de ser capaz de lidar com perguntas complexas que requerem o uso de conhecimento conceitual. Adicionalmente, o framework também tem a capacidade de lidar com bases de conhecimento extensas e complexas.

Palavras-chave: C-KBQA, Framework, Natural Language Processing, Semantic parsing, Knowledge Base.

Abstract

O rápido crescimento dos dados online tornou a recuperação de informações relevantes uma tarefa desafiadora, impulsionando o surgimento de sistemas de Resposta a Perguntas Baseadas em Bases de Conhecimento (KBQA) que lidam com consultas complexas e multi-hop. Este trabalho expandido refina nossa pipeline anterior ao introduzir modelos estruturados simplificados, uma Hereditary Tree-LSTM (HTL) para classificação e análises mais abrangentes sobre reconhecimento de entidades, extração de propriedades e montagem de consultas SPARQL. Ampliamos o conjunto de dados LC-QUAD 2.1 com modelos padronizados e avaliamos uma pipeline flexível que integra DeepPavlov, Falcon, SpaCy, restrições de qualificadores e buscas reversas. Nossos experimentos mostram que o reconhecimento de entidades com múltiplas ferramentas supera métodos que utilizam apenas uma ferramenta, enquanto a extração de propriedades se beneficia de conjuntos de propriedades expandidos e estratégias de ranqueamento refinadas. A precisão global das consultas SPARQL atinge entre 70% e 80% para consultas de complexidade intermediária, mas permanece inferior em subconjuntos específicos do domínio. A sinergia proposta entre ferramentas de PLN e modelos simplificados refinados aumenta a cobertura para KBQA complexos, embora melhorias adicionais no tratamento morfológico e em embeddings especializados possam ser necessárias para lidar de forma abrangente com consultas multi-hop desafiadoras ou de nicho.

Keywords: C-KBQA, Framework, Natural Language Processing, Semantic parsing, Knowledge Base.

Acknowledgements

To everyone who has contributed to my growth, both as a researcher and as an individual, I would like to express my sincere gratitude.

I deeply appreciate the support and guidance of RNP (Rede Nacional de Ensino e Pesquisa) throughout this research. Your input was critical to the successful completion of this project.

I would also like to extend my sincere gratitude to the professors who shared their knowledge and experiences, enriching my understanding and shaping my research efforts. Your guidance has been invaluable and I am grateful for the opportunities you have given me to expand my horizons.

Furthermore, I would like to express my gratitude to my esteemed supervisor and co-supervisor, whose support and encouragement were crucial in navigating the complexities of this research. His experience, guidance and commitment to my development were truly inspiring.

Once again, I am immensely grateful to everyone who has contributed to my journey, and I look forward to continuing my quest for knowledge and making meaningful contributions to the field with your continued support.

Contents

List of Tables	6
List of Abbreviations	7
1 Introduction	8
1.1 Contextualization	10
1.2 Problem description	11
1.3 Justification	11
1.4 Objectives (General and Specific)	12
2 Related Work	14
2.1 Entity Recognition	14
2.2 Relation Extraction	15
2.3 Template Matching	16
2.4 Frameworks	16
3 Materials and Methods	19
3.1 Methodological Introduction	19
3.2 LC-QUAD 2.0	20
3.3 LC-QUAD 2.1	21
3.4 Creating Dummy Templates	23
3.5 Grouping Templates by Semantic Proximity	24
3.6 Hereditary Tree-LSTM (HTL)	26
3.7 Tools for Entity Extraction	28
3.8 Approach for Property Extraction	30
3.9 SPARQL Reverse Lookup for Coverage.	31
3.10 Pipeline Integration and Remaining Steps	32
4 Results and Discussion	35
4.1 Overall Results from Multiple Combinations	35
4.2 Entity Recognition in Each Group	36
4.3 Property Extraction in Each Group	37
4.4 Final SPARQL Completeness	38
4.5 Discussion and Future Directions	39
5 Conclusion	41
Bibliography	42

List of Tables

3.1	All 13 groups (0–12) of dummy templates in LC-QUAD 2.1, with representative examples	25
4.1	Entity Recognition Accuracy (%) by Group and Configuration	36
4.2	Property Extraction Accuracy (%) by Group and Configuration	38

List of Abbreviations

DCC	Departamento de Ciência da Computação
UFJF	Universidade Federal de Juiz de Fora
C-KBQA	Complex Knowledge Based Question Answering
KBQA	Knowledge Based Question Answering
POS	Part of speech
NER	Named Entity Recognition
QA	Question Answering

1 Introduction

The large volume of data available on the internet has made the task of finding relevant information even more challenging (JANG et al., 2017). In this context, new information retrieval methods have enabled more intelligent searches, taking into account the context of the search. Additionally, advances in Natural Language Processing (NLP) research have allowed a better understanding of the search context, enabling Knowledge-based Question-Answering (KBQA) frameworks to emerge as an efficient solution to meet this demand (LAN et al., 2021).

However, the complexity of the questions users ask remains a significant challenge for KBQA frameworks (QIN et al., 2020). Questions may involve multiple concepts, conditional or comparative reasoning, and temporal aspects, requiring the system to understand both the mentioned entities and the relationships among them. For instance, questions such as “Which films directed by Quentin Tarantino were nominated for an Oscar?” or “Is the Eiffel Tower taller than the Statue of Liberty?” demand a deep understanding of entities, properties, and comparisons (XIE et al., 2016). This complexity increases significantly when implicit context needs to be inferred, as in questions like “What is the country’s capital?” (HU et al., 2022).

Large language models (LLMs), such as GPT, have shown significant advances as Question-Answering (QA) systems, providing impressive results when answering natural language questions (DAULL et al., 2023). However, these models present limitations, especially in accessing information not included in their training data (TAN et al., 2023). On the other hand, KBQA frameworks are designed to access structured knowledge bases, such as knowledge graphs, enabling more efficient retrieval of updated and domain-specific information. Previous work has explored the use of semantic knowledge graphs for integrating corporate data, demonstrating the potential of these approaches in precise information retrieval (ROLIM et al., 2021).

Additionally, frameworks that evaluate the connectivity of entity pairs in knowledge bases have highlighted the importance of understanding relationships among en-

tities to construct accurate answers (JIMÉNEZ et al., 2021) and (JIMÉNEZ; LEME; CASANOVA, 2021). Such approaches emphasize the relevance of robust entity and property extraction systems, as used in KBQA frameworks. Moreover, methods that evaluate entity similarity in RDF knowledge bases have advanced the semantic understanding of queries, improving the retrieval of contextually relevant data (JIMÉNEZ; LEME; CASANOVA, 2022).

In this study, we propose a KBQA framework featuring a practical and flexible pipeline designed to handle complex knowledge-based questions. The pipeline is modularly structured, allowing for the replacement of entity recognition and property extraction models, as well as refinements in SPARQL queries using qualifier constraints. The flexibility of the proposed approach is inspired by studies demonstrating the effectiveness of customizable frameworks in different domains (JR et al., 2022). The proposed framework aims not only to improve the accuracy of responses but also to facilitate future adaptations to other languages and domains.

This paper is an extended and revised version of a previously published work (MELLO et al., 2024). We discuss the flexibility of KBQA frameworks and present new experiments to evaluate the application of the proposed techniques in more complex scenarios. The study includes enhancements in SPARQL query generation, with improved precision in entity and property recognition, as well as the integration of refinements in model adaptation. Our goal is to provide a practical and modular approach that can be tailored to different contexts and requirements, broadening the applicability of knowledge-based question-answering systems.

This paper is organized as follows: Section 2 presents related work, discussing existing approaches and their contributions to the field. Section 3 describes the materials and methods used in this study, including dataset preprocessing, template grouping by semantic proximity, dummy template creation, and the tools used for entity and property extraction. Section 4 presents and discusses the obtained results, highlighting the tool combinations that achieved the best performance. Section 5 concludes the work, discussing the implications of the results and suggesting future research directions.

1.1 Contextualization

KBQA (Knowledge-Based Question Answering) systems aim to extract information from knowledge bases to answer user questions. However, this task can become too complex depending on the question. The complexity of questions can exceed the capacity of existing systems, resulting in a low accuracy rate in responses (QIN et al., 2020).

The complexity of a question varies by the number of entities and properties and the relationships between them (XIE et al., 2016). For example, consider the question “What is the average temperature in Brazil during winter?”. In this question, the main entity is “Brazil”, and the related entities and properties are “average temperature” and “winter”. To answer this question, it is necessary to understand how these entities and properties relate to each other. Prior knowledge of the average temperature in Brazil and the climatic characteristics of winter in the country is required.

Another factor contributing to the complexity of questions is the need to understand the context in which the question is asked (HU et al., 2022). The same concept can have different meanings depending on the context in which it is used. For example, the word “java” can refer to both a programming language and a region in Indonesia. Therefore, if someone asks “What is java?”, the KBQA system must be able to determine the correct question based on the context in which it is asked.

Additionally, another factor that increases the complexity of questions is the dependence on additional information that was not mentioned in the question. For example, to answer the question “What is the capital of the country?”, it is necessary to know which country is being referred to. If this information is not provided, the system may have difficulty arriving at the correct answer.

Therefore, it is important that KBQA systems are capable of handling complex questions and considering multiple factors to provide accurate and reliable answers to users. Systems need to have mechanisms to understand the context in which questions are asked, which may involve the use of natural language processing (NLP) techniques and machine learning.

1.2 Problem description

Knowledge-Based Question Answering (KBQA) is a field of artificial intelligence that focuses on developing systems capable of answering natural language questions by leveraging knowledge from structured or unstructured sources.

KBQA approaches follow a set of information-processing steps. However, it is not easy to combine these steps. This work presents a KBQA approach to leverage available steps and algorithms, as well as some solutions created at LApIC. The goal is to study the available methods and build a system that can accurately answer questions based on knowledge stored in a knowledge base.

To achieve this goal, it is investigated various techniques such as natural language processing, semantic parsing, entity linking, and relation extraction. These techniques aim to extract meaning from the natural language question and map it to the entities and relations in the knowledge base.

By combining these techniques and algorithms, the aim of this work is to develop a robust and accurate KBQA system that can effectively answer a wide range of natural language questions.

1.3 Justification

The creation of a framework for Complex Knowledge Base Question Answering (KBQA) has several compelling justifications. One of the key benefits is the ability to establish well-defined pipelines, which facilitates code reuse and allows for more effective evaluation and experimentation with the solution.

By developing a structure with clear pipelines, different components and modules can be utilized and implemented. This modularity allows for the separation of steps, making it easier to understand, maintain, and extend the system. Furthermore, when individual parts of the pipeline are reusable, it reduces redundancy and promotes efficiency in development efforts. Users can perform specific tasks such as question analysis, entity linking, relation extraction, or answer generation without needing to go through the entire pipeline.

Additionally, a well-defined structure allows for more appropriate evaluation and experimentation of the KBQA solution. By dividing the KBQA pipeline into distinct stages, it becomes possible to evaluate the performance of each component separately. This granularity enables the identification of strengths and weaknesses in specific steps and facilitates targeted improvements.

Moreover, modularizing the structure promotes flexible experimentation with alternative techniques or algorithms. Researchers and developers can easily substitute or modify components within the pipeline, enabling the comparison of different approaches. This flexibility leads to a more thorough exploration of various strategies, resulting in improved performance and the ability to adapt to different knowledge domains or languages.

Overall, a framework for Complex KBQA provides several advantages. It promotes code reuse, accelerating development and reducing duplication efforts. It enables more effective evaluation and experimentation of individual components, leading to focused improvements and a better understanding of specific performance aspects. Additionally, it facilitates flexible experimentation with alternative techniques, ensuring the development of a robust KBQA solution that can adapt to different knowledge domains or languages in the future.

1.4 Objectives (General and Specific)

The objective of this work is to make ease the creation of C-KBQA approaches through the definition of a framework. The framework aims to help users extract information from questions in natural language and return a SPARQL query as the answer.

To achieve our overall objective, we have identified the following specific objectives:

- Improve models for Named Entity Recognition (NER) to identify entities in user queries. NER models will be evaluated based on their ability to correctly identify and classify entities with high accuracy.
- Enhance Relationship Linking models to link identified entities to the correct relationships in the knowledge base. Relationship linking models will be evaluated

based on their ability to accurately map entities to the appropriate relationships with high accuracy and retrieval.

- Develop a method for filter extraction to extract information from the question in natural language.
- Develop Entity and Relationship Ranking models to rank identified entities and relationships based on relevance. Classification models will be evaluated based on their ability to accurately classify entities and relationships according to their relevance to the query, as well as their effectiveness in improving the overall accuracy of the system.
- Develop template matching templates to match user queries with predefined templates. Model matching models will be evaluated based on their ability to correctly match user queries to relevant models with high accuracy and retrieval.
- Develop a Slot Filling algorithm to fill slots in predefined templates with identified entities.
- Enable the use of each step separately or all together for more flexibility. The system will allow users to choose whether they want to use each step separately or all together, depending on their specific needs.

This study aims to result in a KBQA framework that can be used to extract knowledge from data on a variety of questions. The framework should be flexible, allowing users to choose which steps to use and how to integrate the framework into their existing workflows. By evaluating each step of the framework and continually improving its accuracy and efficiency, we ensure that our KBQA system is a valuable tool for users looking to extract knowledge from data using natural language queries.

2 Related Work

Research in Knowledge-Based Question Answering (KBQA) has explored various approaches to enhance the precision and efficiency of these systems, particularly when dealing with complex questions. This section reviews significant contributions across three main areas: entity recognition, relation extraction, and template matching.

2.1 Entity Recognition

Entity recognition is a fundamental task in KBQA, involving identifying entities mentioned in the user’s query. One notable work in this area is TagMe (FERRAGINA; SCAIELLA, 2010), which introduces a technique for annotating short text fragments with relevant Wikipedia hyperlinks. The system leverages Wikipedia’s extensive collection of articles and anchor texts to provide informative and accurate annotations. The main technique used involves point identification, sense disambiguation, and annotation. Point identification analyzes the input text to identify potential points that can be linked to Wikipedia articles. Sense disambiguation selects each point’s most relevant Wikipedia page, considering context and statistical information. Finally, annotation adds hyperlinks to corresponding Wikipedia articles, allowing users to access additional information and context simply by clicking on the annotated points. TagMe demonstrates superior performance and speed compared to other systems, making it a valuable tool for entity recognition in short texts.

Another significant contribution in this field is Falcon (SAKOR et al., 2019), a rule-based approach for linking entities and relationships in Wikidata. Falcon employs core principles of English morphology, such as tokenization and N-gram tessellation, to link entity and relation surface forms in short sentences to Wikidata entries. This method includes a local knowledge base composed of DBpedia entities to enhance the recognition and linking process. Falcon provides a ranked list of entities and relations annotated with their Internationalized Resource Identifier (IRI) in Wikidata, aiding the NLP community

in entity and relation recognition. The approach outperforms existing baselines in entity linking tasks, demonstrating high F-score values and robustness across various datasets like QALD-9 (NGOMO, 2018) and LC-QuAD 2.0 (DUBEY et al., 2019).

2.2 Relation Extraction

Relation extraction is another crucial component of KBQA systems, focusing on identifying and linking relationships between entities within a query. SLING (MIHINDUKU-LASOORIYA et al., 2020) is a semantic analysis framework designed to link text relationships to knowledge bases accurately. The approach integrates multiple methods, including statistical Abstract Meaning Representation (AMR) mapping, distant supervision data generation, and various relation-linking modules. The statistical AMR mapping technique is pivotal in identifying relationships by normalizing syntactic variations between sentences and providing strong predicates. Distant supervision data generation creates training examples mapped to corresponding knowledge base relations, enhancing the system’s learning process. SLING leverages transformer-based architectures to encode AMR graphs and question text for relation linking, achieving state-of-the-art performance across datasets such as QALD-7 (USBECK et al., 2017), QALD-9 (NGOMO, 2018), and LC-QuAD 1.0 (TRIVEDI et al., 2017).

Another innovative approach in relation extraction is presented by (ROSSIELLO et al., 2021), which proposes a sequence-to-sequence model enhanced with structured data from the target knowledge base. This model generates a sequence of relations based on the input question text, enriched by an entity-linking system that queries the knowledge base to retrieve candidate relations. The model’s decoder then uses the enriched input representation to generate a structured sequence of argument-relation pairs, considering the contextual information and candidate relations. This approach significantly improves relation-linking performance in question-answering systems, demonstrating notable enhancements over existing methods.

2.3 Template Matching

Template matching involves identifying patterns in user questions and matching them to predefined templates, facilitating the generation of structured and well-formatted responses. In (DILEEP et al., 2021), the authors explore machine learning models and pre-processing techniques to classify natural language questions into appropriate templates using the LC-QUAD 2.0 dataset. They train classifiers such as XGBoost and Random Forest, utilizing Part-of-Speech (POS) tagging and FastText for preprocessing. POS tagging assigns grammatical tags to words, helping the system understand the syntactic structure of questions, while FastText captures the semantic meaning of words through embeddings. The combination of XGBoost and POS+FastText preprocessing achieves superior accuracy in classifying questions into relevant templates, showcasing the effectiveness of template-based question answering.

Another noteworthy study is presented by (GOMES et al., 2022), which introduces a hereditary attention mechanism combined with template matching to enhance semantic extraction from questions. This approach categorizes complex questions into answer templates, leveraging hierarchical structures within the questions. The hereditary attention mechanism operates bottom-up, where each neural network cell inherits attention from another cell, capturing and prioritizing the most relevant information at different levels of the question’s structure. This method improves the robustness and accuracy of KBQA systems, providing a reliable technique for answering complex questions from knowledge bases.

2.4 Frameworks

Recent advancements in KBQA framework have highlighted the effectiveness of integrating various methodologies to improve performance in complex queries. Two works in this area are the RnG-KBQA and ReTraCk framework, which present approaches to enhance the accuracy and generalization capabilities of KBQA framework.

The RnG-KBQA system, developed by (YE et al., 2021), addresses the limitations of traditional KBQA systems that struggle with unseen knowledge base (KB)

schema items. RnG-KBQA combines a ranking-based approach with a generation model to enhance coverage and generalization. The system first ranks candidate logical forms based on the question. Then, it uses a generation model conditioned on the question and top-ranked candidates to create the final logical form. This dual approach significantly improves performance, achieving state-of-the-art results on the GRAILQA and WEBQSP datasets, with notable improvements in zero-shot generalization.

However, RnG-KBQA also presents some limitations. The complexity of combining ranking and generation can lead to longer processing times and increased computational resource requirements. Additionally, the effectiveness of the generation model may be limited by the quality of the training data, affecting the system’s ability to handle ambiguous or poorly formulated queries.

The ReTraCk, developed by (CHEN et al., 2021), introduces a flexible framework that integrates multiple stages for entity recognition, relation extraction, and ranking. ReTraCk emphasizes the importance of a modular design, allowing for easy integration of different models and techniques at each stage. The system has shown remarkable performance in various benchmarks, demonstrating its adaptability and efficiency in handling diverse KBQA tasks. The approach focuses on iterative refinement and ranking to enhance the accuracy of the generated answers, contributing to the development of a more robust KBQA framework.

Despite its advantages, ReTraCk also faces challenges. Integrating and adjusting multiple models and techniques can increase the system’s complexity, making maintenance and updates more difficult. Additionally, since ReTraCk relies on multiple processing stages, any error in one of these stages can compromise the accuracy of the final answer. The modular approach, while flexible, can result in inconsistencies when different modules are not perfectly aligned. Moreover, ReTraCk is a resource-intensive system requiring considerable computational resources, which can be a barrier to deployment in production environments with limited resources.

These works underscore the importance of combining ranking and generation techniques to overcome the limitations of the traditional KBQA framework, paving the way for more accurate and generalizable solutions. By integrating advanced NLP and

machine learning models, both RnG-KBQA and ReTraCk contribute significantly to the field, offering valuable insights and methodologies for future research in KBQA. However, the identified limitations of these systems highlight the ongoing need for refinement and innovation to improve the efficiency and effectiveness of KBQA systems.

This work shares the same principles of these frameworks, such as flexibility and their use for different datasets. However, our approach stands out by focusing on complex queries, which often require the integration of multiple pieces of information, reasoning over data, and understanding nuanced context. The solution allows new entity recognition and property extraction models if needed. We implemented methods to retrieve all properties related to an entity and developed a system to dynamically fill placeholders in templates with the correct values extracted from the knowledge base. These methods ensure the pipeline can easily adapt to new requirements and improvements. Finally, the system’s structured design can decompose complex queries, identify relevant entities and relationships, and construct precise SPARQL queries using entity recognition, relation extraction, template matching, ranking, and slot-filling methods.

3 Materials and Methods

This section provides a more detailed view of the study’s methodological underpinnings, describing the choice and preparation of datasets, as well as the processes of template grouping and the creation of **dummy templates** with placeholders. In this expanded version, the objective is not only to describe *what* was done but also to explain *why* each step was essential for handling complex questions in a Knowledge-Based Question Answering (KBQA) environment.

3.1 Methodological Introduction

The primary motivation for our C-KBQA framework is to ensure that complex questions—those involving multiple entities, temporal filtering, conditional counting, and chained relationships—can be consistently converted into coherent SPARQL queries. Such questions often require more than simply identifying one entity; they frequently include elements such as “What is the birth date of the director of a certain film, as long as the film received an award after 2000?” or “Which players were transferred to a different team before a specific year?” These examples reveal the variety of multi-hop operations and precise filters that must be managed. With this in mind, we looked for a dataset that both captured these scenarios of high complexity and allowed for fair performance assessment.

We focused on the *Large-scale Complex Question Answering Dataset (LC-QUAD) 2.0*, subsequently refined into *LC-QUAD 2.1*. In the following sections, we outline the key characteristics of each dataset, emphasizing the inconsistencies that led to LC-QUAD 2.1 and explaining how we organize and structure SPARQL templates for such complexity. Although the first subsection addresses why these data were chosen, subsequent subsections explain the semantic grouping of templates and the introduction of **dummy templates**—placeholders that simplify the mapping from questions to SPARQL queries.

3.2 LC-QUAD 2.0

LC-QUAD 2.0 (DUBEY et al., 2019) was created to push beyond the limitations of earlier QA datasets by significantly expanding both the number of queries and the breadth of question types. Unlike prior resources that often featured only a few thousand simple questions, LC-QUAD 2.0 contains over 30,000 English-language questions addressing complex, multi-hop structures and advanced SPARQL features such as numeric and textual filters, COUNT operations, and temporal qualifiers. This magnitude of content was crowd-sourced via the Amazon Mechanical Turk platform, ensuring broader linguistic variety and the inclusion of paraphrased versions for many questions.

A central goal in developing LC-QUAD 2.0 was to promote more realistic testing of how QA systems handle multi-faceted queries. Rather than restricting questions to a single triple pattern or straightforward relationships, the authors introduced scenarios requiring two or more interconnected facts, time-based conditions (for instance, an event that must occur before or after a given year), and textual constraints (involving `contains` or `STRSTARTS`). These additions compel systems to integrate multiple knowledge graph edges, interpret qualifiers (such as an event’s date or location), and cope with more diverse language variations. To facilitate this diversity, the dataset was constructed around both Wikidata and DBpedia 2018, allowing queries to tap into distinct but overlapping knowledge graphs and highlight potential differences in how systems map question text to each resource’s ontology.

Another distinctive contribution is the deliberate inclusion of paraphrased questions. The dataset’s creators recognized that QA systems trained exclusively on a single wording of a query might overfit to specific linguistic cues. By asking Mechanical Turk workers to restate or paraphrase each question, LC-QUAD 2.0 ensures that machine-learning approaches cannot rely solely on narrow syntactic patterns. This requirement to handle varying formulations of the same intent is especially relevant for advanced queries where multi-hop reasoning or qualifiers are combined. Systems that succeed on LC-QUAD 2.0, therefore, must prove robust against rephrasings and synonyms.

Although LC-QUAD 2.0 represented a major leap in complexity and quantity over its predecessor LC-QUAD 1.0, certain practical issues emerged. The large-scale

crowdsourcing methodology led to occasional inconsistencies between question text and the assigned SPARQL query, along with instances of overlap or duplication in the data. Some queries proved too short or vague to give the system meaningful clues, while others veered toward extreme length or unwieldy lexical constructions. These challenges made it harder to ascertain whether errors arose from a model’s limitations or from mismatched or duplicated entries. Nonetheless, the essential complexity the dataset introduced—covering two-intent questions, qualifier usage, date-based restrictions, and textual substring operations—marked a turning point in the field. No earlier dataset had offered such an extensive set of natural language queries with corresponding SPARQL templates across DBpedia and Wikidata.

Researchers seeking to build or evaluate complex QA systems found multiple uses for LC-QUAD 2.0. One common application involved entity-linking or predicate-linking tasks, where the dataset’s short, sometimes paraphrased questions forced algorithms to handle ambiguous references and partial textual matches. Others employed LC-QUAD 2.0 for training end-to-end question-answering pipelines that generate SPARQL queries. The variety of advanced operations present in the dataset—a reflection of the depth of Wikidata’s and DBpedia’s schemas—stimulated the development of machine-learning approaches that handle multi-hop patterns and complicated string filters more gracefully.

3.3 LC-QUAD 2.1

LC-QUAD 2.1 (Gomes Jr. et al., 2021) emerged as a refined and standardized update to the large-scale LC-QUAD 2.0 dataset. Building upon the original collection of more than 30,000 questions, its creators undertook a comprehensive review to eliminate duplicate entries, remove malformed or excessively short items, and address inconsistencies between question text and the associated SPARQL queries (GOMES et al., 2022). The resulting dataset preserves the core complexity of multi-hop reasoning, numeric and textual filters, and advanced operators, yet offers a more uniform set of mappings from natural language to formal query structures. By mitigating issues like duplicated questions and mismatched references, LC-QUAD 2.1 reduces the possibility of QA systems failing for reasons unrelated to genuine semantic or lexical challenges.

A defining enhancement introduced in LC-QUAD 2.1 was the concept of “dummy” elements within each SPARQL query. Rather than embedding fixed QIDs or numeric values in the query, placeholders such as `DUMMY_S` (for subjects), `DUMMY_P` (for properties), `DUMMY_O` (for objects), and `DUMMY_F` (for numeric or string filters) replaced the original references. This strategy separates the question’s broad logical shape—for example, the number of triple patterns or the presence of filters—from its specific data-level details. Consequently, the pipeline can be split into two distinct phases. In the first phase, the system classifies a user query to determine which structural “template” it matches. In the second, the recognized placeholders are filled with the relevant entity IDs, property IDs, or numeric/string constraints discovered by entity- and property-linking modules.

This two-phase design streamlines experimentation in several ways. It explicitly distinguishes structural misinterpretation (choosing the wrong SPARQL “shape”) from errors in entity linking or property mapping (supplying the wrong QIDs or thresholds). If the pipeline is already confident in the structural template, any subsequent errors become easier to localize and correct, often without having to retrain the entire model. Moreover, dummy placeholders allow incremental improvements in entity recognition or property detection to be introduced without disrupting the dataset’s underlying template library. As a result, a flawed final SPARQL no longer conflates poor structural analysis with simple linking missteps, making system failures more traceable to specific components. This modularity in query building has proven particularly valuable when researchers experiment with new entity-linking strategies or domain-specific expansions.

By isolating each question’s logical skeleton and its data-level references, LC-QUAD 2.1 presents a cleaner testbed for KBQA systems. Whether a query demands two or three hops, textual substring checks, or date-qualifier logic, the dummy format ensures the system’s classification step accurately pins down which query form applies, while the linking step concentrates on populating placeholders with the correct QIDs or numeric constraints. In this sense, LC-QUAD 2.1 retains the multifaceted complexity of LC-QUAD 2.0 but provides a more stable platform for diagnosing the root cause of system errors and advancing the field of complex KBQA.

3.4 Creating Dummy Templates

A central refinement introduced by LC-QUAD 2.1, in comparison to LC-QUAD 2.0, was the decision to separate the fixed logical form of each SPARQL query from the specific QIDs, property labels, and numeric or textual filters that appeared in it. In LC-QUAD 2.0, each query was fully specified and directly referenced entities such as `wd:Q3772` (Quentin Tarantino) and properties like `wdt:P57` (directed by). However, for handling complex queries in a more flexible, generalizable fashion, LC-QUAD 2.1 introduced placeholders in place of these explicit identifiers. This step is referred to as “dummy-fication.”

During dummy-fication, each original SPARQL query was rewritten with placeholders (`DUMMY_S` for subjects, `DUMMY_P` for predicates, `DUMMY_O` for objects, and `DUMMY_F` for numeric or string filters). For example, a fully specified query involving two properties and a numeric cutoff would be converted to a structure that indicates the same number of triple patterns and filters, but no longer binds them to a specific QID or integer. Through this process, the dataset preserves only the “shape” of each query (such as whether it is `ASK` or `SELECT`, how many hops it contains, or how many comparison filters are included) without tying those shapes to fixed IDs or values.

Performing this dummy creation before any classification or grouping steps has practical benefits. It clarifies that two queries referencing different QIDs but using the same underlying pattern (for instance, “Which films directed by X were nominated for Y after year Z?”) share a single structural form. Consequently, the pipeline can focus on the logic of the query—how many relationships or filters it contains—rather than on the specific details of each entity or property. Table ?? illustrates a few dummy templates, highlighting how the “raw” references get replaced by placeholders while preserving the SPARQL skeleton.

By ensuring every query adheres to this placeholder-based format, the dataset becomes more amenable to subsequent pipeline tasks such as grouping queries by complexity or training a model that classifies questions according to their structure.

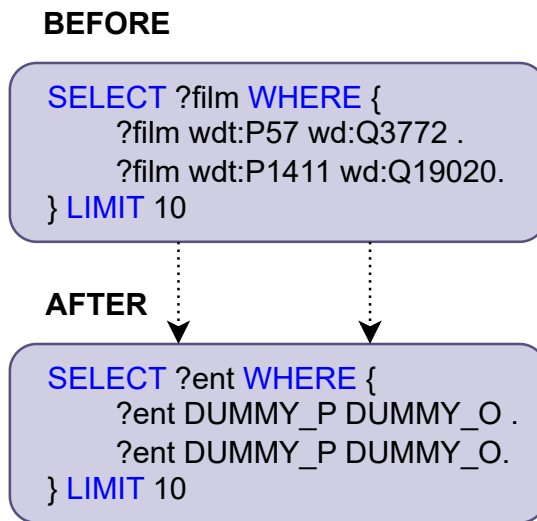


Figure 3.1: Illustration of how raw SPARQL becomes a dummy template with placeholders

3.5 Grouping Templates by Semantic Proximity

After converting all LC-QUAD 2.0 SPARQL queries into *dummy templates* (i.e., replacing explicit entity IDs, property IDs, and numeric/string values with placeholders), the next essential step in LC-QUAD 2.1 was to cluster these templates according to their structural and semantic features. This process yielded a total of **13 groups** (indexed 0 through 12), covering **29 distinct dummy templates**. Each group comprises one or more templates that share a common logic or level of complexity, such as whether the query is **ASK** or **SELECT**, if it requires multi-hop reasoning, or whether it applies numeric/string filters. Table 3.1 summarizes the groups, noting the number of templates in each, a brief description, and a single *Dummy SPARQL example* from that group.

By grouping dummy templates that share common structures, LC-QUAD 2.1 provides a concise map of how query complexity progresses. A single **ASK** plus numeric filter might belong to Group 0, a two-hop **SELECT** might reside in Group 4 or 5, and textual label checks fall under Groups 7, 9, or 11. This organization allows the pipeline to narrow down the relevant set of template “shapes” when faced with a user question that, for instance, demands multi-hop plus a textual filter.

Table 3.1: All 13 groups (0–12) of dummy templates in LC-QUAD 2.1, with representative examples

Group ID	#Templates	Description	Dummy SPARQL Example
0	4	Basic ASK queries with at most one numeric filter (yes/no checks)	0.2: ASK WHERE { DUMMY_S DUMMY_P ?obj FILTER(?obj = DUMMY_F) }
1	1	ASK queries referencing two relationships	1.0: ASK WHERE { DUMMY_S DUMMY_P DUMMY_O . DUMMY_S DUMMY_P DUMMY_O }
2	2	Single-hop SELECT retrieval, typically one relationship	2.2: SELECT DISTINCT ?answer WHERE { DUMMY_S DUMMY_P ?answer }
3	2	“How many...” questions using COUNT (single-hop)	3.1: SELECT (COUNT(?sub) AS ?value) { ?sub DUMMY_P DUMMY_O }
4	5	Multi-hop SELECT, often 2 steps with an intermediate variable	4.2: SELECT ?answer WHERE { DUMMY_S DUMMY_P ?answer . ?answer DUMMY_P DUMMY_O }
5	2	Multi-hop SELECT with extra reference to the same subject/object	5.1: SELECT ?obj WHERE { DUMMY_S DUMMY_P ?s . ?s DUMMY_P ?obj . ?s DUMMY_P DUMMY_O }
6	2	Multi-hop plus ORDER BY and LIMIT	6.1: SELECT ?ent WHERE { ?ent DUMMY_P DUMMY_O . ?ent DUMMY_P ?obj . ?ent DUMMY_P DUMMY_O } ORDER BY ASC(?obj) LIMIT 10
7	4	SELECT queries with string/date-based filters	7.2: SELECT ?value WHERE { DUMMY_S DUMMY_P ?s . ?s DUMMY_P ?x FILTER(contains(?x, 'DUMMY_F')) . ?s DUMMY_P ?value }
8	1	SELECT retrieving multiple answers in one step	8.0: SELECT ?ans_1 ?ans_2 WHERE { DUMMY_S DUMMY_P ?ans_1 . DUMMY_S DUMMY_P ?ans_2 }
9	2	Label-based constraints, partial string checks, lang(...)	9.2: SELECT DISTINCT ?subj ?subj_label WHERE { ?subj DUMMY_P DUMMY_O . ?subj DUMMY_P ?subj_label . FILTER(STRSTARTS(lcase(?subj_label), 'DUMMY_F')) . FILTER(lang(?subj_label)='DUMMY_F') } LIMIT 10
10	1	SELECT retrieving multiple columns (?value1, ?obj) in one hop	10.0: SELECT ?value1 ?obj WHERE { DUMMY_S DUMMY_P ?s . ?s DUMMY_P ?obj . ?s DUMMY_P ?value1 }
11	2	Repeated references plus textual filters for labels	11.2: SELECT DISTINCT ?subj ?subj_label WHERE { ?subj DUMMY_P DUMMY_O . ?subj DUMMY_P DUMMY_O . ?subj DUMMY_P ?subj_label . FILTER(STRSTARTS(lcase(?subj_label), 'DUMMY_F')) . FILTER(lang(?subj_label)='DUMMY_F') } LIMIT 10
12	1	Multi-value SELECT retrieving two results from the same intermediate node	12.0: SELECT ?value1 ?value2 WHERE { DUMMY_S DUMMY_P ?s . ?s DUMMY_P DUMMY_O . ?s DUMMY_P ?value1 . ?s DUMMY_P ?value2 }

Although the table above highlights the 13 main groups, each group can contain multiple related templates, reflecting minor variations in operators, filters, or multi-hop steps. For instance, Group 0 encloses 4 templates (IDs 0.1–0.4), whereas Group 4 includes 5 (4.1–4.5). This arrangement ensures that if new SPARQL forms arise—e.g., a specialized subquery or advanced date operator—one can extend or refine an existing group or define a new one, keeping the classification consistent.

Clustering these dummy templates by semantic proximity underpins the entire pipeline’s ability to locate the correct SPARQL skeleton quickly. Once a user query is recognized as, say, a multi-hop SELECT with date-based filtering, the system knows exactly which group (and subsequently which specific template) to match. This strategy significantly reduces the complexity of forming valid SPARQL queries, since the pipeline

only needs to fill placeholders (`DUMMY_S`, `DUMMY_P`, `DUMMY_O`, `DUMMY_F`) with the correct QIDs or numeric thresholds after the structure is identified.

3.6 Hereditary Tree-LSTM (HTL)

After defining the overall strategy for grouping SPARQL templates and refining the dataset into dummy structures, a crucial question arises: how can the system determine which template best matches a user’s complex question, especially when multiple relational hops or filters may be involved? A purely rule-based or semantic-parsing approach might try to construct logical forms from scratch, which can be computationally expensive and prone to errors if the user’s query contains multiple constraints or unusual syntactic cues. To address this challenge, we introduced the **Hereditary Tree-LSTM (HTL)**, drawing on the concepts put forth in (GOMES et al., 2022), as an efficient and robust means of classifying complex natural language questions into their corresponding dummy templates.

A standard Tree-LSTM departs from the linear token-processing found in traditional LSTM architectures by operating directly on a syntactic tree, such as one produced by a dependency parser. Each node in this tree represents a token, while edges reflect syntactic or semantic relationships between words, like subject–verb or adjective–noun links. By summing the hidden states of child nodes, the Tree-LSTM captures the hierarchical structure of language in a way that a linear model may overlook. This structural advantage becomes critical in questions that reference multiple entities, date thresholds, or chained comparisons, since relevant words are not always linearly adjacent in the sentence.

The HTL approach extends this by incorporating what is termed a “hereditary attention” mechanism. Rather than having a single attention layer that may overlook subtle interactions, each subtree computes local attention signals that highlight tokens or phrases deemed most significant. These attention signals then pass upward in a bottom-up fashion. By the time the hidden representations accumulate at the root node of the tree, the model has integrated key features from lower-level phrases—references to comparative operators such as “greater than” or date markers like “after 2000” or relational patterns suggesting multi-hop logic. This richer, aggregated representation at the root node is then

passed to a classification layer, typically a softmax, to decide which dummy template is structurally most suitable for the user’s question.

The “hereditary” aspect of HTL is especially useful in questions where relevant pieces of information are scattered in different branches of the sentence. In a linear LSTM, it might be difficult to keep track of these scattered clues without heavily engineered gating or attention modules. By contrast, the Tree-LSTM natively deals with linguistic hierarchies, and the hereditary attention ensures that essential fragments are not lost as intermediate child nodes pass their states to the parent node. This approach proves especially powerful for complex queries that combine numeric or temporal conditions with multi-hop relationships, since each child node can measure local significance and transmit that forward.

In practical terms, the decision process with HTL involves two major phases. The first is the construction of the tree representation. The question is parsed using a syntactic or dependency parser, producing a tree whose nodes correspond to tokens. Each node is associated with an embedding (for instance, via pretrained word embeddings or contextual representations) and a dependency relation to its parent. The second phase is the bottom-up accumulation of hidden states and attentional weights. Each node computes, for each child, how relevant it is to the overarching meaning. These local attentions are aggregated, culminating in a top-level vector. A final classification stage maps this top-level vector to a dummy template ID.

Several advantages emerge from this design. One advantage is that HTL substantially reduces the need to enumerate all possible logical forms or to rely on purely text-based classification methods that might fail to appreciate the compositional nature of the sentence. Another advantage lies in providing better interpretability: by inspecting the attention distributions at each node, one can see which parts of the question the Tree-LSTM deems most crucial. This can be invaluable when explaining why a given template was chosen over alternatives. A potential limitation is that it requires reliable dependency parsing; errors in the parse tree can propagate upward and cause misclassification. Nevertheless, for well-formed sentences such as those curated in LC-QUAD 2.1, this risk is mitigated by the dataset’s improved consistency.

Once the HTL selects the best matching template, the process of constructing the final SPARQL query is greatly simplified. The question no longer needs an entirely new logical form; it only requires filling the dummy slots (`DUMMY_S`, `DUMMY_P`, `DUMMY_O`, `DUMMY_F`) with the right entities, properties, or filter values. This separation between structure classification and slot filling yields both computational efficiency and clarity: the system pinpoints the structural skeleton first, then calls upon entity-extraction modules to fill in details. In the context of LC-QUAD 2.1, whose queries already come in dummy form, HTL stands out as a stable and powerful means of bridging the gap between the question’s linguistic complexity and the final query’s logical architecture.

3.7 Tools for Entity Extraction

Once the HTL indicates which dummy template structure should apply to a user’s question, the pipeline still needs to identify the exact entities mentioned so it can fill placeholders (e.g., `DUMMY_S`, `DUMMY_O`) with the right QIDs. Although entity linking seems straightforward when questions reference well-known concepts such as “Quentin Tarantino” or “Oscar,” it can be more difficult with ambiguous or lesser-known items, domain-specific abbreviations, and nonstandard phrasing. To address these diverse scenarios, we employ three tools in parallel: DeepPavlov, Falcon 2.0, and spaCy. Each tool focuses on detecting named entities and linking them to potential Wikidata identifiers, but they do so using different methodologies, thereby reducing the chance that one tool’s blind spots will lead to missed or incorrect matches.

DeepPavlov. It provides a wide range of pretrained NLP models, including pipelines for both named entity recognition (NER) and entity linking. When a user query is processed by DeepPavlov, the system first segments the text to identify candidate entity boundaries. For instance, it can discern that “New York City” is a single entity mention, rather than erroneously splitting “New York” and “City.” Once the candidate mentions are extracted, DeepPavlov’s linking model uses embedding-based similarity to map each mention to potential Wikidata entries. This stage can struggle with short or rare mentions that do not exhibit strong embedding matches to any known label or alias in the knowledge base.

However, DeepPavlov’s flexibility in customizing models or training them on domain-specific text is advantageous if the pipeline needs to adapt to specialized vocabularies (e.g., medical or technical domains).

Falcon 2.0. Adopts a rule-driven approach to entity and relation linking, emphasizing tokenization, morphology, and N-Gram tiling. It is particularly adept at capturing short compound forms, multiword expressions, or property-like phrases that may not appear in standard NER outputs. For example, it can detect “operating income” or “manager position” as meaningful references to possible relations or specialized entity phrases. Once Falcon identifies these candidate tokens, it queries a local index of Wikidata labels and aliases and ranks the retrieved entities by textual similarity. Although Falcon 2.0 can sometimes over-prioritize close string matches in ambiguous cases (e.g., multiple items sharing a term like “Ford”), its ability to identify less common or domain-specific phrases complements the embedding-based coverage gaps that DeepPavlov might leave.

spaCy. Offers a fast, comprehensive NLP pipeline that performs tokenization, part-of-speech (POS) tagging, dependency parsing, and built-in NER. We use spaCy not just for its NER component but also to parse each query into a structured dependency tree, which supports the HTL-based classification and helps interpret multiword sequences. spaCy’s default entity linker, while somewhat generic, often succeeds on widely used entity names or standard lexical items. Its speed and practical approach to common named entities make it a useful baseline for everyday queries, and its token and syntactic boundaries allow the other two tools to coordinate over consistent phrase segments. In many cases, spaCy can rapidly resolve simpler queries or confirm well-known entities so that the pipeline can focus more time on ambiguous or domain-specific references.

In practice, all three tools run concurrently on each user query, producing separate lists of entity mentions and their top candidate Wikidata links. We then merge these lists using a set of heuristic rules that check for overlaps or disagreements. If two or more tools converge on the same entity mention, confidence in that match increases significantly. Conversely, if DeepPavlov proposes a mention that Falcon 2.0 and spaCy do not see (or strongly disagree with), the pipeline may attempt partial substring matching or check the

local neighborhood of tokens to see if the mention is plausible. In ambiguous cases, we also consider how frequently a proposed entity ID appears in Wikidata and whether it aligns with the query’s overall domain or syntactic clues from spaCy’s dependency parse. This ensures that a single tool’s shortcoming—like missing a name variant or failing to parse compound terms—does not cause the pipeline to discard a valid candidate.

A recurring challenge arises when user questions involve abbreviations (e.g., “NYC” instead of “New York City”), short acronyms, or domain-specific synonyms rarely seen in ordinary embeddings. Such references often result in low matching scores for embedding-based approaches like DeepPavlov. Falcon 2.0’s rule-based expansions can partially help, but only if the relevant aliases appear in its index. spaCy may detect the tokens but lack a strong domain linker. In these cases, the pipeline checks morphological expansions, or potential substring matches across all candidate sets.

Advantages of a Multi-Tool Strategy. By combining these three approaches, the pipeline gains resilience against the varied weaknesses inherent in any single method. DeepPavlov’s embedding-based linking can handle well-known or phrasally consistent names; Falcon 2.0 captures tricky morphological variants and property-like expressions; spaCy quickly resolves standard NER segments and yields useful syntactic boundaries for the question. The cost is higher runtime, but for complex knowledge-based QA, the reduced risk of missing critical entities often outweighs that overhead. Moreover, as new domain modules or updated embeddings become available, they can be slotted into this parallel pipeline, building on the collaboration that already exists. While no single entity pipeline is flawless, coordinating their results through conflict-resolution heuristics can push the overall entity-linking accuracy far beyond what we could achieve with only one or two of these tools, particularly on the broad and sometimes peculiar set of queries found in LC-QUAD 2.1.

3.8 Approach for Property Extraction

Whereas entity extraction resolves DUMMY_S (subject) and DUMMY_O (object) placeholders, another key question is how to find the properties—i.e., the DUMMY_P slots—that define

the relationships or attributes mentioned in the user’s query. Sometimes the query text is explicit (e.g., “*directed by*”), making detection relatively easy. More often, especially in multi-hop questions, relational clues may be implicit (“*starring in a film*” might map to “*cast member*” in Wikidata). Disambiguation is equally vital here, since Wikidata hosts thousands of properties with labels that may overlap or sound similar.

Initial Property Hints from Extraction Tools. Falcon 2.0 is often the most direct source for early property hints, since it is designed to locate both entities and relations in short text. When the user’s question includes a phrase like “*nominated for an Oscar*”, Falcon 2.0 might produce “wdt:P1411” (nominated for) as a candidate. DeepPavlov’s main emphasis is on entity linking, though it can occasionally flag property-like tokens. spaCy itself does not inherently propose property IDs, but can help segment the question to see where relational phrases appear.

Difficulties in Multi-Hop or Implicit Relations. Challenges multiply if the question calls for more than one property (e.g., “*Which politician’s spouse held office before 2000?*”) or if the user’s wording does not align neatly with Wikidata property labels. The HTL’s template classification might reveal that the question structure involves two relationships plus a temporal filter, but it does not specify which properties to use. If the user’s text yields only partial lexical clues, the pipeline may need to guess among multiple property candidates. This is where we incorporate both textual cues and knowledge-graph lookups to increase accuracy.

3.9 SPARQL Reverse Lookup for Coverage.

If a recognized entity appears as a subject, we can query Wikidata for all properties where is the subject of a statement. We do the same if seems to be an object. This “reverse lookup” step is crucial in capturing less obvious properties not explicitly suggested by the user’s phrasing or by initial text-based extraction. However, for prominent entities with hundreds of statements, retrieving so many property IDs can become overwhelming. Thus, we filter and rank them based on semantic similarity to the question text and any

relevant contexts extracted by spaCy’s dependency analysis.

Ranking Property Candidates. Because the reverse lookup or textual hints may produce multiple possibilities, we apply a final **ranking** step to each candidate property. We compare the property’s label or description to key words in the question, using word-embedding vectors (for instance, from FastText). If the user’s query references “*acted in movie*” a property whose label includes “*cast member*” might score well. Conversely, a property about “*producer*” or “*distributor*” might rank lower. Only the top few property candidates (commonly the top 5) progress to the final pipeline stage, preventing exponential growth in potential SPARQL combinations.

Qualifiers and Constraints. Even when the correct property is identified, Wikidata can embed qualifiers specifying contextual constraints, such as date ranges or roles. If the question implies these qualifiers (e.g., “*during the 1990s*”), ignoring them can yield an incomplete or incorrect final query. For relevant properties, the pipeline checks whether qualifiers exist that match the question’s temporal or contextual conditions. Such qualifiers might be introduced as additional statements or filter clauses in the final query, ensuring alignment with the user’s multi-layered request.

3.10 Pipeline Integration and Remaining Steps

As depicted in Figure 3.2, by the time the pipeline has identified a suitable dummy template through the HTL and consolidated a set of entity and property matches from the multi-tool extraction and ranking modules, it can proceed to fill each placeholder (DUMMY_S, DUMMY_P, DUMMY_O, and optionally DUMMY_F) with the selected QIDs, property IDs, and any numeric or string constraints. When qualifier constraints have been detected, these too must be integrated into the final template structure to capture time-sensitive or contextual details.

Instead of executing this fully populated SPARQL query on a live endpoint, the system compares the filled template directly to the gold-standard SPARQL provided by LC-QUAD 2.1. This approach allows us to confirm whether the discovered entities,

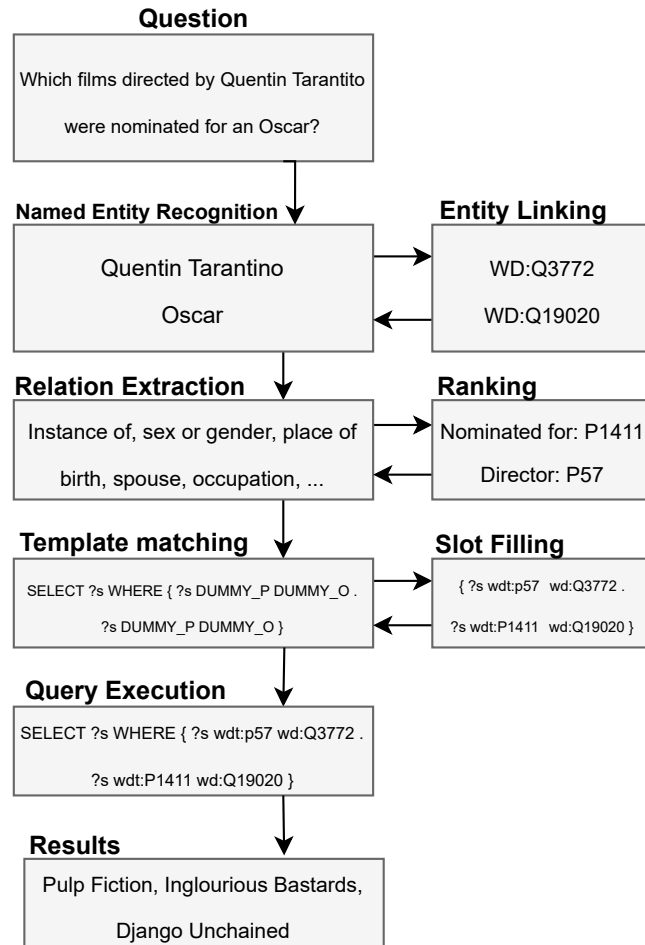


Figure 3.2: Fluxogram illustrating how the pipeline integrates template selection (HTL), entity recognition, property extraction, and final alignment with LC-QUAD 2.1’s gold SPARQL.

properties, and constraints align precisely with those in the question’s reference query—an indication that the pipeline arrived at the correct combination of placeholders and slot values for that template. We measure performance at multiple stages (Correct Entities, Correct Properties, and Correct SPARQL) without relying on any actual query execution. A misplaced entity disrupts `DUMMY_S`, a property mismatch impacts `DUMMY_P`, and broader structural errors may point to shortcomings in the HTL-based template selection itself.

By leveraging the LC-QUAD 2.1 gold SPARQL queries, we do not need to issue real-time queries against a knowledge base to determine accuracy. Consistency with the reference indicates that the pipeline successfully reproduced the question’s logic—covering entities, properties, filters, and qualifiers. Where mismatches arise, more targeted refinements can be made: for instance, repeated minor property errors may suggest improve-

ments in the ranking module or reverse lookups; systematic placeholder mismatches might point to gaps in entity linking or qualifier integration.

This design not only streamlines our evaluation but also preserves a modular structure. Each subsystem—classification, entity extraction, property selection, or template refinement—can be updated or replaced, while the gold SPARQL references in LC-QUAD 2.1 remain a stable benchmark for correctness. As a result, we can continuously refine or expand the pipeline (through new extraction models, advanced ranking heuristics, or expanded dummy templates) while maintaining a consistent means of verifying overall alignment with the intended SPARQL.

4 Results and Discussion

This section provides a more detailed analysis of how each phase of our KBQA pipeline—entity recognition, property extraction, and final SPARQL query alignment—performs under multiple experimental setups. We tested configurations ranging from single-tool approaches (e.g., **DP**: DeepPavlov alone) to pairwise integrations (e.g., **DP+F**: DeepPavlov + Falcon), and finally a three-tool combination with additional refinements (qualifier constraints and reverse property lookups). In the following tables, the columns indicate the entity or property accuracy (percentage of queries where *all* relevant entities or properties were correctly identified) and the final SPARQL completeness (the percentage of queries that fully matched the gold-standard structure in LC-QUAD 2.1).

4.1 Overall Results from Multiple Combinations

We begin by noting that simpler pipelines using a single tool (**DP**, Falcon, or SpaCy) often handle the more elementary queries in LC-QUAD 2.1 (e.g., Groups 0, 1, or 2) reasonably well, as these queries typically reference only one or two entities without complex filters or multi-hop reasoning. However, once queries introduce multiple relationships or require unusual constraints (especially in Groups 4, 6, 9, and 11), single-tool approaches generally fall short. Pairwise integrations improve coverage, but the best results consistently emerge from the full three-tool setup (**DP+F+S+Q**).

Groups are identified in LC-QUAD 2.1 to represent different structural SPARQL patterns (see Section 3.5 for details). For instance, Group 0 includes basic **ASK** queries, while Group 4 covers two-hop **SELECT** queries, and Group 9 focuses on label-based textual constraints. When reading the tables below, it helps to remember that:

- Lower group IDs (0–3) usually represent simpler queries like **ASK** or single-hop **SELECT** with minimal filters.
- Mid-range group IDs (4–7) reflect multi-hop queries or more intricate operations,

such as date-based filters.

- Higher group IDs (8–12) handle either multi-value retrieval, label-based textual searches, or repeated references, which can be trickier for standard extraction tools.

4.2 Entity Recognition in Each Group

Table 4.1 reports the entity recognition accuracy across 13 groups, comparing six different tool configurations. The metric indicates the proportion of queries in which *all* entities were correctly identified. If even one entity was missed or mislabeled, the entire query is considered incorrect at the entity level.

Table 4.1: Entity Recognition Accuracy (%) by Group and Configuration

Grupo	DP	DP+F	S	DP+S	F+S	DP+F+S
0	28.50	46.97	62.53	65.96	80.74	82.32
1	49.72	64.25	68.16	76.54	83.24	84.92
2	36.97	54.51	66.27	69.68	76.66	76.66
3	34.18	44.30	75.32	79.11	94.30	94.30
4	25.96	39.37	52.93	57.31	63.44	65.01
5	50.52	71.28	62.47	70.44	79.04	81.76
6	7.10	13.66	36.07	37.70	42.08	42.62
7	75.22	87.61	87.61	92.92	94.69	96.76
8	46.25	68.75	75.00	83.75	96.25	97.50
9	0.00	2.26	46.79	46.79	53.96	53.96
10	80.25	88.89	79.01	92.59	93.83	95.06
11	17.72	24.68	48.73	50.63	54.43	55.06
12	53.38	70.27	66.22	76.35	81.76	84.46

Key Observations. Groups 0, 1, 2, and 3 often reference single-hop or simpler relations, so entity coverage easily exceeds 60–70% in most multi-tool setups. In particular, Group 3 sees an especially large jump in accuracy when using more than one tool (e.g., from 44.30% in DP+F to 94.30% in F+S), which suggests that the counting queries often rely on subtle cues for identifying the subject entity. By contrast, Groups 6 and 9 stand out

as notably difficult, with the triple configuration (**DP+F+S+Q**) only reaching about 42.62% (Group 6) and 53.96% (Group 9). These groups typically feature domain-specific synonyms (Group 6) or label-based constraints (Group 9) that may not map cleanly onto each tool’s internal dictionary or require morphological expansions not adequately covered by the default models.

Discussion. Entity recognition emerges as the pipeline’s strongest step overall, but certain specialized or ambiguous queries remain a challenge. The collaboration among DeepPavlov (flexible entity recognition), Falcon (rule-based linking for short tokens), and SpaCy (robust parsing) lifts the average entity accuracy considerably. However, Groups like 6 (often multi-hop with ordering) and 11 (repeated references and textual filters) reveal persistent issues when the question’s phrasing deviates from standard synonyms or tokens.

4.3 Property Extraction in Each Group

Table 4.2 presents the property extraction accuracy for each group. Like entity recognition, the metric indicates the proportion of queries in which *all* required properties were correctly detected and linked.

Key Observations. Whereas entity accuracy often climbs above 70–80%, property extraction remains below 60% in most groups. Groups 1 and 8 do exceed 80% property correctness under the best setup, suggesting that straightforward “double-check” relationships (Group 1) or multi-answer queries (Group 8) can be resolved if the lexical overlap is reasonably direct. Still, advanced queries in Groups 6, 9, and 11 again stand out as problematic. Partial coverage improvements (e.g., from 0.75% to 7.17% in Group 9) underscore that text-based property references remain an Achilles’ heel when synonyms or domain-specific phrases are not captured by default extraction modules.

Reverse Lookups and Qualifiers. Leveraging reverse SPARQL lookups helps broaden the candidate property set, sometimes boosting coverage by 10–25% in multi-hop or filter-

Table 4.2: Property Extraction Accuracy (%) by Group and Configuration

Grupo	DP	DP+F	S	DP+S	F+S	DP+F+S+Q
0	28.75	36.74	28.75	51.76	36.74	67.09
1	73.33	75.00	73.33	85.00	75.00	86.67
2	19.25	31.35	19.25	33.39	31.35	52.47
3	31.01	34.81	31.01	65.19	34.81	74.68
4	28.09	38.68	28.09	42.08	38.68	54.92
5	45.18	51.68	45.18	49.58	51.68	55.97
6	10.47	12.40	10.47	18.99	12.40	24.42
7	39.68	46.90	39.68	43.22	46.90	51.18
8	50.63	63.75	50.63	71.88	63.75	84.38
9	0.75	0.75	0.75	5.66	0.75	7.17
10	51.85	59.26	51.85	55.56	59.26	66.67
11	17.72	18.99	17.72	29.11	18.99	32.91
12	31.08	41.44	31.08	37.84	41.44	50.90

laden queries (Groups 2, 4, 5, 12). Qualifier constraints further refine the pipeline’s ability to handle temporal or contextual edges (Groups 3, 7), preventing incomplete placeholders. Nonetheless, if a question’s phrasing diverges from known synonyms—or the question includes nested multi-hop logic—misalignments remain common.

4.4 Final SPARQL Completeness

Definition of Completeness. The final metric measures how many queries matched the gold SPARQL *exactly* in terms of entities, properties, filters, and structure. Table ?? (in Section 4) already showcased the best setup’s completeness values, which vary widely depending on how many pieces must align. Missing a single property or entity is enough to spoil the final match.

Group-Level Patterns. Groups 3 (counting) and 8 (multi-answer) reach above 70% completeness in the best-case scenario, reflecting that once the pipeline identifies the correct subject or object references, it can fill placeholders consistently. More challenging sets, like Groups 4, 5, 7, 9, 11, 12, rarely exceed 20% completeness, indicating that partial success in entity or property extraction does not yield a final SPARQL alignment.

In Group 6, no queries are fully correct under any setup, reinforcing the complexity of queries with ordering, multiple relationships, or ambiguous references.

Filtering Outliers. If we remove subsets of queries with extremely low coverage (e.g., Group 6 or Group 9), the final completeness among remaining groups rises significantly. This highlights how certain domain-specific or morphological details hamper the pipeline’s property recognition enough to break the entire SPARQL chain. Integrating more advanced morphological expansions and property disambiguation heuristics might improve these outlier groups in the future.

4.5 Discussion and Future Directions

Overall, these experiments confirm that combining multiple extraction tools (**DP+F+S+Q**) markedly improves entity and property coverage across most LC-QUAD 2.1 groups, compared to single-tool or pairwise settings. Nevertheless, advanced queries in Groups 6, 9, and 11, among others, expose how lexical variations, rarely used properties, and nested filters can still confound the pipeline. Specific directions for improvement include:

- **Domain-Specific Embeddings.** Incorporating custom embeddings or fine-tuning on domain text could capture synonyms or acronyms that default models miss, especially in Groups with specialized phrasing (e.g., 6 or 9).
- **Enhanced Ranking.** Dedicating more robust property-ranking heuristics or neural ranking methods might help handle borderline property matches, which hamper the final SPARQL alignment.
- **Morphological/Tokenization Improvements.** Groups that embed entities in multiword expressions (e.g., 11) might benefit from specialized tokenization logic and morphological checks.
- **Handling Nested or Rare Properties.** Introducing an expanded knowledge base of synonyms or advanced “reverse” lookups beyond standard properties could reduce the gap in multi-hop or heavily constrained questions.

Despite lingering gaps, the pipeline’s success in many medium-complexity queries (like Groups 3 and 8) underscores the effectiveness of dummy templates, synergy among multiple extraction modules, reverse lookups, and qualifier constraints. By systematically identifying each group’s structural logic, the pipeline reliably matches user questions to the correct SPARQL form, as long as the underlying entity and property references can be recognized.

5 Conclusion

This paper introduced a comprehensive and modular pipeline for Knowledge Base Question Answering (KBQA) that addresses significant shortcomings in existing entity- and property-extraction tools. Our framework employs multiple strategies, such as reverse SPARQL lookups and qualifier constraints, to broaden the set of potential predicates beyond what standard modules (DeepPavlov, Falcon, and SpaCy) typically capture. Although these off-the-shelf tools handle many straightforward cases, they often overlook essential properties in multi-hop or context-laden queries, leading to incomplete SPARQL alignments. By merging refined extraction processes with structured dummy templates and dynamic slot filling, we demonstrated a notable improvement in covering complex relationships, even when questions impose constraints or reference specialized knowledge. These enhancements highlight the importance of effectively retrieving not just entities but also less prominent predicates, a task crucial for assembling accurate and thorough SPARQL queries.

Despite these advances, there remain questions that involve rare or domain-specific properties, subtle morphological variations, or multiple layers of filtering and temporal qualifiers. Our results indicate that while the pipeline substantially raises recall in many scenarios, certain specialized or niche aspects still pose difficulties. Future research could focus on embedding-based expansions to better handle domain-specific synonyms, more refined morphological handling for short or compound tokens, and improved ranking models that weigh contextual information more robustly. The pipeline’s modular design, anchored by interchangeable extraction and ranking components, should facilitate iterative improvements in these directions. Ultimately, this work underscores the need to look beyond entity identification and toward a richer, more adaptable strategy for retrieving properties, which is key for ensuring that KBQA systems can accurately address the intricacies and breadth of real-world queries.

Bibliography

CHEN, S.; LIU, Q.; YU, Z.; LIN, C.-Y.; LOU, J.-G.; JIANG, F. Retrack: A flexible and efficient framework for knowledge base question answering. In: *Proceedings of the 59th annual meeting of the association for computational linguistics and the 11th international joint conference on natural language processing: system demonstrations*. [S.l.: s.n.], 2021. p. 325–336.

DAULL, X.; BELLOT, P.; BRUNO, E.; MARTIN, V.; MURISASCO, E. Complex qa and language models hybrid architectures, survey. *arXiv preprint arXiv:2302.09051*, 2023.

DILEEP, A. K.; MISHRA, A.; MEHTA, R.; UPPAL, S.; CHAKRABORTY, J.; BANSAL, S. K. Template-based question answering analysis on the lc-quad2.0 dataset. In: *2021 IEEE 15th International Conference on Semantic Computing (ICSC)*. [S.l.: s.n.], 2021. p. 443–448.

DUBEY, M.; BANERJEE, D.; ABDELKAWI, A.; LEHMANN, J. Lc-quad 2.0: A large dataset for complex question answering over wikidata and dbpedia. In: SPRINGER. *International Semantic Web Conference*. [S.l.], 2019. p. 69–78.

FERRAGINA, P.; SCAIELLA, U. Tagme: On-the-fly annotation of short text fragments (by wikipedia entities). In: *Proceedings of the 19th ACM International Conference on Information and Knowledge Management*. New York, NY, USA: Association for Computing Machinery, 2010. (CIKM '10), p. 1625–1628. ISBN 9781450300995. Disponível em: <https://doi.org/10.1145/1871437.1871689>.

GOMES, J.; de Mello, R. C.; STRÖELE, V.; de Souza, J. F. A hereditary attentive template-based approach for complex knowledge base question answering systems. *Expert Systems with Applications*, v. 205, p. 117725, 2022. ISSN 0957-4174. Disponível em: <https://www.sciencedirect.com/science/article/pii/S0957417422010089>.

Gomes Jr., J.; MELLO, R. C. de; STRÖELE, V.; SOUZA, J. F. de. *LC-QuAD 2.1*. Zenodo, 2021. Disponível em: <https://doi.org/10.5281/zenodo.5508297>.

HU, X.; WU, X.; SHU, Y.; QU, Y. Logical form generation via multi-task learning for complex question answering over knowledge bases. In: *Proceedings of the 29th International Conference on Computational Linguistics*. Gyeongju, Republic of Korea: International Committee on Computational Linguistics, 2022. p. 1687–1696. Disponível em: <https://aclanthology.org/2022.coling-1.145>.

JANG, H.; OH, Y.; JIN, S.; JUNG, H.; KONG, H.; LEE, D.; JEON, D.; KIM, W. Kbqa: Constructing structured query graph from keyword query for semantic search. In: *Proceedings of the International Conference on Electronic Commerce*. New York, NY, USA: Association for Computing Machinery, 2017. (ICEC '17). ISBN 9781450353120. Disponível em: <https://doi.org/10.1145/3154943.3154955>.

JIMÉNEZ, G.; LEME, P. P.; CASANOVA, T. I. N. A framework to compute entity relatedness in large rdf knowledge bases. *Journal of Information and Data Management*, v. 13, n. 2, Sep. 2022. Disponível em: <https://journals-sol.sbc.org.br/index.php/jidm/article/view/2435>.

JIMÉNEZ, J.; LEME, L.; CASANOVA, M. Coepinkb: A framework to understand the connectivity of entity pairs in knowledge bases. In: *Anais do XLVIII Seminário Integrado de Software e Hardware*. Porto Alegre, RS, Brasil: SBC, 2021. p. 97–105. ISSN 2595-6205. Disponível em: <https://sol.sbc.org.br/index.php/semish/article/view/15811>.

JIMÉNEZ, J. G.; LEME, L. P. P.; IZQUIERDO, Y. T.; NEVES, A. B.; CASANOVA, M. A distributed framework to investigate the entity relatedness problem in large rdf knowledge bases. In: *Anais do XXXVI Simpósio Brasileiro de Bancos de Dados*. Porto Alegre, RS, Brasil: SBC, 2021. p. 121–132. ISSN 2763-8979. Disponível em: <https://sol.sbc.org.br/index.php/sbbd/article/view/17871>.

JR, J. G.; MELLO, R. C. de; STRÖELE, V.; SOUZA, J. F. de. A study of approaches to answering complex questions over knowledge bases. *Knowledge and Information Systems*, Springer, v. 64, n. 11, p. 2849–2881, 2022.

LAN, Y.; HE, G.; JIANG, J.; JIANG, J.; ZHAO, W. X.; WEN, J. A survey on complex knowledge base question answering: Methods, challenges and solutions. *CoRR*, abs/2105.11644, 2021. Disponível em: <https://arxiv.org/abs/2105.11644>.

MELLO, R.; JR., J. G.; SOUZA, J.; STRÖELE, V. Constructing a kbqa framework: Design and implementation. In: *Proceedings of the 30th Brazilian Symposium on Multimedia and the Web*. Porto Alegre, RS, Brasil: SBC, 2024. p. 89–97. ISSN 0000-0000. Disponível em: <https://sol.sbc.org.br/index.php/webmedia/article/view/30301>.

MIHINDUKULASOORIYA, N.; ROSSIELLO, G.; KAPANIPATHI, P.; ABDELAZIZ, I.; RAVISHANKAR, S.; YU, M.; GLIOZZO, A.; ROUKOS, S.; GRAY, A. G. Leveraging semantic parsing for relation linking over knowledge bases. *CoRR*, abs/2009.07726, 2020. Disponível em: <https://arxiv.org/abs/2009.07726>.

NGOMO, N. 9th challenge on question answering over linked data (qald-9). *language*, v. 7, n. 1, p. 58–64, 2018.

QIN, K.; WANG, Y.; LI, C.; GUNARATNA, K.; JIN, H.; PAVLU, V.; ASLAM, J. A. A complex KBQA system using multiple reasoning paths. *CoRR*, abs/2005.10970, 2020. Disponível em: <https://arxiv.org/abs/2005.10970>.

ROLIM, T.; AVILA, C. V.; JUNIOR, N. A.; COSTA, F.; MARIANO, R.; CALIXTO, T.; VIDAL, V. M. Kg-e: Um grafo de conhecimento semântico baseado na integração de dados de empresas e sancionados. In: *Anais do IX Workshop de Computação Aplicada em Governo Eletrônico*. Porto Alegre, RS, Brasil: SBC, 2021. p. 155–166. ISSN 2763-8723. Disponível em: <https://sol.sbc.org.br/index.php/wcge/article/view/15985>.

ROSSIELLO, G.; MIHINDUKULASOORIYA, N.; ABDELAZIZ, I.; BORNEA, M.; GLIOZZO, A.; NASEEM, T.; KAPANIPATHI, P. Generative relation linking for question answering over knowledge bases. In: HOTHO, A.; BLOMQUIST, E.; DIETZE, S.; FOKOUE, A.; DING, Y.; BARNAGHI, P.; HALLER, A.; DRAGONI, M.; ALANI, H. (Ed.). *The Semantic Web – ISWC 2021*. Cham: Springer International Publishing, 2021. p. 321–337. ISBN 978-3-030-88361-4.

SAKOR, A.; SINGH, K.; PATEL, A.; VIDAL, M. FALCON 2.0: An entity and relation linking tool over wikidata. *CoRR*, abs/1912.11270, 2019. Disponível em: <http://arxiv.org/abs/1912.11270>.

TAN, Y.; MIN, D.; LI, Y.; LI, W.; HU, N.; CHEN, Y.; QI, G. Evaluation of chatgpt as a question answering system for answering complex questions. *arXiv preprint arXiv:2303.07992*, 2023.

TRIVEDI, P.; MAHESHWARI, G.; DUBEY, M.; LEHMANN, J. Lc-quad: A corpus for complex question answering over knowledge graphs. In: SPRINGER. *International Semantic Web Conference*. Cham: Springer International Publishing, 2017. p. 210–218.

USBECK, R.; NGOMO, A.-C. N.; HAARMANN, B.; KRITHARA, A.; RÖDER, M.; NAPOLITANO, G. 7th open challenge on question answering over linked data (qald-7). In: SPRINGER. *Semantic Web Challenges: 4th SemWebEval Challenge at ESWC 2017, Portoroz, Slovenia, May 28-June 1, 2017, Revised Selected Papers*. [S.l.], 2017. p. 59–69.

XIE, Z.; ZENG, Z.; ZHOU, G.; HE, T. Knowledge base question answering based on deep learning models. In: SPRINGER. *Natural Language Understanding and Intelligent Applications: 5th CCF Conference on Natural Language Processing and Chinese Computing, NLPCC 2016, and 24th International Conference on Computer Processing of Oriental Languages, ICCPOL 2016, Kunming, China, December 2–6, 2016, Proceedings 24*. [S.l.], 2016. p. 300–311.

YE, X.; YAVUZ, S.; HASHIMOTO, K.; ZHOU, Y.; XIONG, C. Rng-kbqa: Generation augmented iterative ranking for knowledge base question answering. *arXiv preprint arXiv:2109.08678*, 2021.