



HAL
open science

CASIMIR : un Corpus d'Articles Scientifiques Intégrant les ModIfications et Révisions des auteurs

Léane Jourdan, Richard Dufour, Nicolas Hernandez, Florian Boudin

► **To cite this version:**

Léane Jourdan, Richard Dufour, Nicolas Hernandez, Florian Boudin. CASIMIR : un Corpus d'Articles Scientifiques Intégrant les ModIfications et Révisions des auteurs. Atelier sur l'Analyse et la Recherche de Textes Scientifiques, CORIA-TALN 2023, Jun 2023, Paris, France. hal-04122594

HAL Id: hal-04122594

<https://hal.science/hal-04122594v1>

Submitted on 8 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



CASIMIR: un Corpus d'Articles Scientifiques Intégrant les Modifications et Révisions des auteurs

Léane Jourdan, Florian Boudin, Richard Dufour, Nicolas Hernandez
{prénom.nom}@univ-nantes.fr

Objectifs

- Nouveau corpus d'articles scientifiques avec révisions
- Un alignement des modifications entre les versions des articles
- Enrichit avec les métadonnées des articles et les relectures par les pairs

Exemple de révisions

Abstract
Knowledge graphs are often used to store common sense information that is useful for various tasks. However, the extraction of contextually-relevant knowledge is an unsolved problem, and current approaches are relatively simple. Here we introduce a triple selection method based on a ranking model and find that it improves question answering accuracy over existing methods. We additionally investigate methods to ensure that extracted triples form a connected graph. Graph connectivity is important for model interpretability, as paths are frequently used to understand reasoning from question to answer. We make our code and data available at <https://github.com/anonymous>.

1 Introduction
For models to be able to reason about situations that arise in everyday life, they must have access to contextually appropriate common sense information. This information is commonly stored as a large set of facts from which the model must identify a relevant subset. One approach to structuring these facts is as a knowledge graph. Here, nodes represent high-level concepts, and relationships are expressed via typed edges joining two nodes. Each edge type represents a different kind of conceptual relationship between concepts. The contextually-relevant subgraphs ("schema graphs") that are extracted from these graphs are often encoded using neural models, which are trained for tasks including question answering or natural language inference.

Prior work has focused on different ways to encode this information, including using it as input directly to a transformer or using a graph neural network (GNN) (Feng et al., 2020; Yasunaga et al., 2021). However, the question of how to identify useful information has been under-explored, particularly in work that uses GNN encoders. The

simple retrieval methods that are used could limit performance on tasks if contextually-important information is not retrieved.

In this paper we explore methods to construct high-quality schema graphs containing contextually-relevant information. We approach this as a ranking task across triples in a knowledge graph, and select the highest scoring for the schema graph. However, simply using the most relevant triples as input for a GNN is insufficient, as the resulting subgraph is likely to have low connectivity. This is problematic for two reasons. First, it inherently limits the power of the GNN, as nodes in different graph components will not be updated with information from each other. Second, paths of reasoning through the schema graphs are often used as explanations for model behaviour (Feng et al., 2020; Yasunaga et al., 2021; Wang et al., 2020). If the graph consists of multiple separate components, this becomes impossible.

The issue of graph disconnectedness is compounded when certain nodes are required to be included. For example, in question answering, concepts mentioned in the question and in a candidate answer should be identified and included. A path starting from a question concept can then be evaluated for plausibility of reasoning as it progresses towards an answer concept. We therefore also apply a graph algorithm to ensure that the schema graph is connected, taking into account the identified edges and desired nodes, and use an embedding-based method to identify concepts mentioned.

Our contributions are summarised as follows:

- Apply a ranking model to identify common sense triples that are relevant to some context.
- Identify and thoroughly investigate methods to ensure schema graph connectivity.
- Compare existing lexical approaches to entity linking to a simple embedding-based method.

Abstract
Knowledge graphs are often used to store common sense information that is useful for various tasks. However, the extraction of contextually-relevant knowledge is an unsolved problem, and current approaches are relatively simple. Here we introduce a triple selection method based on a ranking model and find that it improves question answering accuracy over existing methods. We additionally investigate methods to ensure that extracted triples form a connected graph. Graph connectivity is important for model interpretability, as paths are frequently used as explanations for the reasoning that connects question and answer.

1 Introduction
For models to be able to reason about situations that arise in everyday life, they must have access to contextually appropriate common sense information. This information is commonly stored as a large set of facts from which the model must identify a relevant subset. One approach to structuring these facts is as a knowledge graph. Here, nodes represent high-level concepts, and typed edges represent different kinds of relationship between concepts. In practice, a subset of facts that are thought to be contextually relevant are extracted from the graph, as using all facts in each instance is unnecessary, noisy, and computationally expensive.

Prior work has focused on different ways to encode these facts, including by inputting them into a graph neural network (GNN) or into a transformer (Feng et al., 2020; Yasunaga et al., 2021). However, the question of how to identify useful information has been under-explored, particularly in work that uses GNN encoders. If contextually important information is not retrieved then performance could be dramatically reduced, a potential result of the use of overly simplistic retrieval methods.

In this paper we explore methods to extract high-quality subgraphs containing contextually relevant information.¹ We approach this as a ranking task across triples in a knowledge graph, and propose two methods that use the scores to extract a subgraph. The first is a weighted pathfinding approach which extends prior work (Lin et al., 2019), while the second builds a minimum spanning tree that includes the highest-ranked triples (figure 1). Both approaches ensure that all or most nodes in the subgraph are reachable from each other, which is important for two reasons. First, it means that the GNN can update node embeddings with information from most other nodes, which would not be possible if the graph were disconnected. Second, it allows paths of reasoning to be extracted from the subgraph, which are often used as explanations for model behaviour (Feng et al., 2020; Wang et al., 2020; Yasunaga et al., 2021).

There are also situations when specific concepts need to be included in order for a subgraph to be of high enough quality. For example, in question answering, a full explanation must include one

¹We call these "relevant subgraphs" or "extracted subgraphs", noting that others use "schema graphs" (Lin et al., 2019).

Processus de collecte

- 1 Collecte des événements (ateliers, conférences, etc) sur Open Review
- 2 Collecte des articles : métadonnées, relectures et PDF disponibles des versions des articles
- 3 Filtrage des articles ayant une seule version (89,33% des articles conservés soit 97,46% des PDF)
- 4 Conversion des PDF vers XML (outil Grobid)

Articles

les PDF des versions des articles

Fichiers de correspondance

- entre les versions finales et antérieures des articles
- entre les relectures et les articles

Contenu

Metadonnées des articles

- dates
- auteurs
- mots-clés
- événement
- identifiants


Relectures

- commentaires
- notes
- décisions
- dates


Description

Contient 730 événements et 118 415 PDF pour 26 355 articles
Domaines : apprentissage auto, robotique, TAL, vision, etc.

Distribution des articles suivant leur nombre d'anciennes versions



Distribution des articles suivant leur nombre de relectures



Et ensuite?

- Améliorer la conversion des PDF (choix de l'outil, traitement des figures, tables)
- Aligner les versions paragraphe à paragraphe et phrase à phrase puis extraire les révisions
- Annoter les documents selon une taxonomie de révisions à définir (ex: clarté, grammaire, style)
- Exploitation pour la mise en place d'outils d'aide à l'écriture

