Software Engineering for Business Information Systems
School of Computation, Information and Technology
Technical University of Munich

Exploring LLMs, Knowledge Graphs, and Their Joint Use in Multi-Domain Claim Verification

Master Thesis Proposal

Ivana Hacajová

Software Engineering for Business Information Systems
School of Computation, Information and Technology
Technical University of Munich

November 27rd, 2023

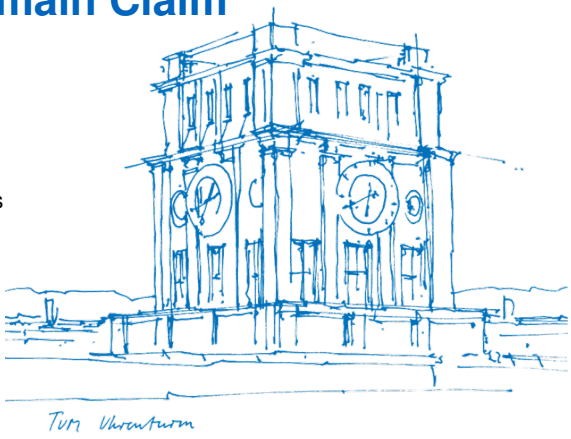

Outline

About me

Academic background

- Bachelor's Degree in **Knowledge Engineering**, Faculty of Information Technology, Czech Technical University in Prague, Czech Republic
- Master's of **Data Engineering and Analytics**, School of Computation, Information and Technology, Technical University of Munich, Germany

Other studies

RMIT Australia (2018), NCKU Taiwan (2019), CTU (2020), NTU Taiwan (2023)

About me

Relevant Courses at TUM

- Introduction to Deep Learning
- Natural Language Processing
- Seminar in Computational Social Science
- Advanced Praktikum in NLP
- Guided Research in Legal Tech

Work Experience

- Teaching Assistant for Foundations of Data Engineering, TUM
- BigData Consultant, Profinit, Czech Republic

Outline

Claim Verification
Also known as Fact Checking

- Online, but also offline space is filled with misinformation, whether intentional or unintentional
- Such claims are of various nature, for example medical (COVID-19 most notably), climate change or just very general
- To tackle this problem, initiatives for manual fact checking exist
- Due to the large volume of such claims to be verified, automated claim verification is needed

Automated Claim Verification

- In general consists of 3 steps: **document retrieval**, **sentence retrieval**, **claim verification**
- As a source for "truth" is usually used Wikipedia, scientifc articles, Google Search, ...
- Traditionally BERT-like models are used to find evidence sentences and then decide if the claim is true or false
- Nowadays, the research aims towards the use of LLMs

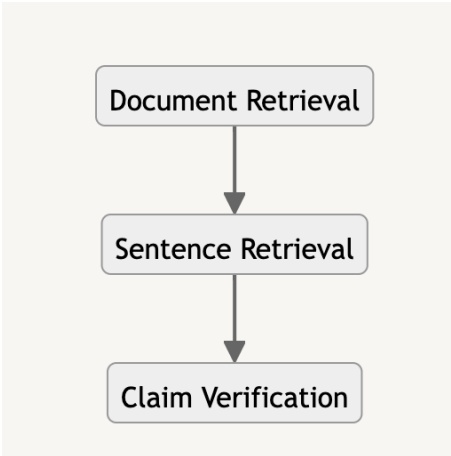

Outline

Research Questions

1. How can the use of LLMs help claim verification?

2. Does leveraging knowledge from knowledge graphs and structured reasoning improve performance?

3. How do different domains compare in this task?

Outline

Datasets

Fact Checking Datasets

- **HOVER** - Claims requiring multiple source documents to be verified, no specific domain
- **CLIMATE-FEVER** - Climate related claims
- **HealthFC** - Medical claims
- **FactKG** - Claims verifiable using the Wikipedia's knowledge graph

Knowledge Basis

- **DBpedia** - Information from Wikipedia organised into a knowledge graph
- **UMLS** - Medical terms and concepts organised into a knowledge graph

Dataset Examples
HOVER

- Claim: Before I Go to Sleep stars an Australian actress, producer and occasional singer.
- Label: SUPPORTS

- Supporting facts: ["key": "Before I Go to Sleep (film)", "value": 0 , "key": "Before I Go to Sleep (film)", "value": 1 , "key": "Nicole Kidman", "value": 0]
- Hops: 2

Dataset Examples
CLIMATE-FEVER

- Claim: Global warming is driving polar bears toward extinction
- Label: SUPPORTS

- Evidences: [{ "evidence_id": "Extinction risk from global warming:170", "evidence_label": 2, # "NOT_ENOUGH_INFO" "article": "Extinction risk from global warming", "evidence": "R̈ecent Research Shows Human Activity Driving Earth Towards Global Extinction Event:", "entropy": 0.6931471805599453, "votes": ["SUPPORTS", "NOT_ENOUGH_INFO", null, null, null] },

Dataset Examples
HealthFC

- Claim: Can regular intake of vitamin C prevent colds?
- Label: Refuted

- Document: The recommendation to take high-dose vitamin C at the first signs of a cold cannot be confirmed by studies. If cough, sniffing or sore throat are already present, vitamin C does not seem to have any detectable effect. The daily requirement for the vitamin is about 100 milligrams, with the recommendations slightly fluctuating [2,3]. This amount is contained in an apple, half a pepper or two tomatoes [4]. (...)

Dataset Examples
FactKG

- Claim: Adam McQuaid weighed 94.8024 kilograms and is from Pleasant Springs, Wisconsin.
- Label: False

- Entity_set: ['Adam_McQuaid', '"94802.4"', 'Pleasant_Springs,_Wisconsin']
- Evidence': 'Adam_McQuaid': [['weight'], ['placeOfBirth']], '"94802.4"': [[' weight']], 'Pleasant_Springs,_Wisconsin': [[' placeOfBirth']]

QACheck
Base Approach

- Iterative and structured approach
- Uses LLMs to generate questions leading to claim's verification
- This workflow provides good explainability

Room for Experiments

- Unreliable source of knowledge - LLMs
- Has not been tested on domain specific datasets

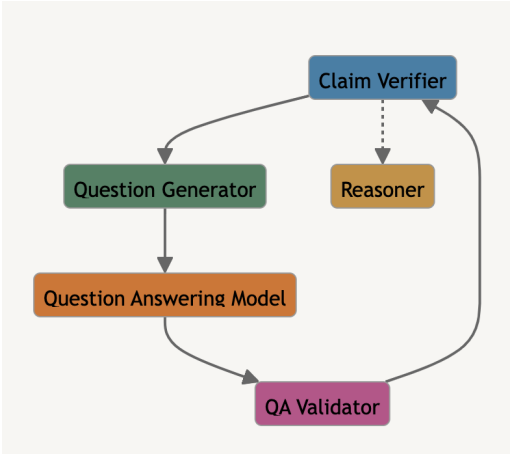

QACheck

Prompt Examples - Question Generator Module

Prompts for the initial question generation

```
Claim = Superdrag and Collective Soul are
both rock bands.
To verify the above claim, we can
first ask a simple question:
Question = Is Superdrag a rock band?

<10 demonstrations in total>
--------
Claim = [[CLAIM]]
To verify the above claim, we can
first ask a simple question:
Question =
```

Prompts for the follow-up question generation

```
Claim = Superdrag and Collective Soul are
both rock bands.
We already know the following:
Question 1 = Is Superdrag a rock band?
Answer 1 = Yes
To verify the claim, what is the
next question we need to know the
answer to?
Question 2 = Is Collective Soul a rock band?

<10 demonstrations in total>
--------
Claim = [[CLAIM]]
We already know the following:
[[QA_CONTEXTS]]
To verify the claim, what is the
next question we need to know the
answer to?
Question [[Q_INDEX]] =
```

FOLK
First-OrderLogic-Guided Knowledge-Grounded

- Uses LLMs
- Claims are decomposed into first-order-logic clauses consisting of predicates
- Simple questions are then generated from the predicates
- Each question is answered by googling it
- Gathered evidence is then used for the final verdict

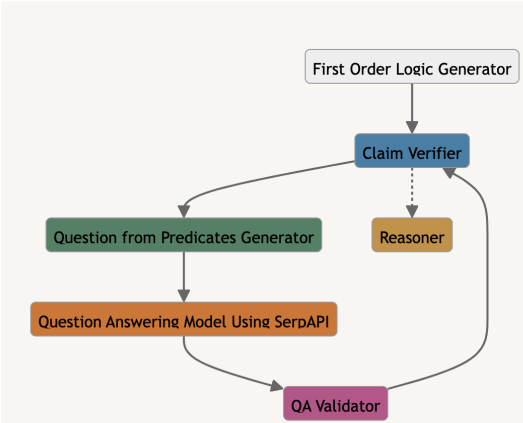

FOLK

Prompt Examples - Claim Decomposition

```
You are given a problem description and a claim. The task is to:
1) define all the predicates in the claim
2) parse the predicates into followup questions
3) answer the followup questions

Claim: Howard University Hospital and Providence Hospital are both located in Washington, D.C.
>>>>>>
Predicates:
Location(Howard Hospital, Washington D.C.) ::: Verify Howard University Hospital is located in Washington, D.C.
Location(Providence Hospital, Washington D.C.) ::: Verify Providence Hospital is located in Washington, D.C.

Followup Question: Where is Howard Hospital located?
Followup Question: Where is Providence Hospital located?
------
Claim: An IndyCar race driver drove a Formula 1 car designed by Peter McCool during the 2007 Formula One season.
>>>>>>
Predicates:
Designed(Peter McCool, a Formula 1 car) ::: Verify a Formula 1 car was designed by Peter McCool during the 2007 Formula One
     season.
Drive(An IndyCar race driver, a Formula 1 car) ::: Verify an IndyCar driver drove a Formula 1 car.

Followup Question: Which Formula 1 car was designed by Peter McCool during the 2007 Formula One season?
Followup Question: Did an IndyCar driver drove a Formula 1 car designed by Peter McCool during the 2007 Formula One season?
------
```

KG-GPT
Reasoning on Knowledge Graphs

- Uses LLMs
- Claims are decomposed into simple claims
- From each claim, two entities are extracted and then they are retrieved from the graph with relations they share
- Most fitting relations to the subclaim is added to the evidence for the final verdict

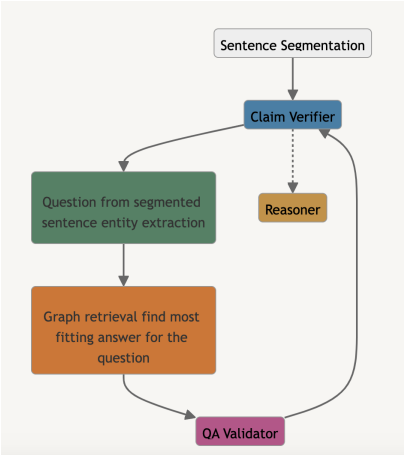

KG-GPT

Prompt Examples - Sentence Segmentation and Inference

Examples)
Sentence A: Ahmad Kadhim Assad's club is Al-Zawra'a SC.
Entity set: ['Ahmad_Kadhim_Assad' ## "Al-Zawra'a_SC"]
–>Divided:
1. Ahmad Kadhim Assad's club is Al-Zawra'a SC., Entity set: ['Ahmad_Kadhim_Assad' ## "Al-Zawra'a_SC"]

...

Sentence L: An academic journal with code IJPHDE is also Acta Math. Hungar.
Entity set: ["Acta Math. Hungar." ## "IJPHDE"]
–>Divided:
1. An academic journal is with code IJPHDE., Entity set: ['academic journal' ## "IJPHDE"]
2. An academic journal is also Acta Math. Hungar., Entity set: ['academic journal' ## "Acta Math. Hungar."]

Claim A: Ahmad Kadhim Assad's club is Al-Zawra'a SC.
Evidence set: [['Ahamad_Kadhim', 'clubs', "Al-Zawra'a SC"]]
Answer: True, based on the evidence set, Ahmad Kadhim Assad's club is Al-Zawra'a SC.

...

Claim L: The place, designed by Huseyin Butuner and Hilmi Guner, is located in a country, where the leader is Paul Nurse.
Evidence set: [["Baku_Turkish_Martyrs'_Memorial", 'designer', "Hüseyin Bütüner and Hilmi Güner"], ["Baku_Turkish_Martyrs'_Memorial", 'location', 'Azerbaijan']]
Answer: False, there is no evidence for Paul Nurse.

Outline

Timeline

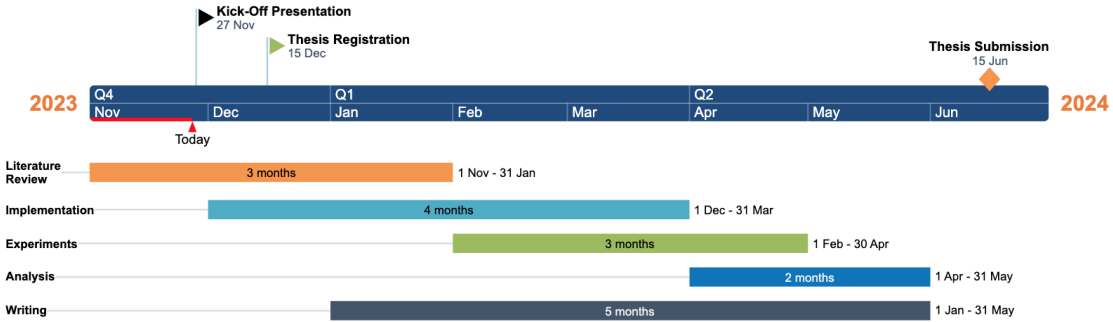
Kick-Off Presentation
27 Nov
Thesis Registration
15 Dec
Thesis Submission
15 Jun
2023
Q4
Q1
Q2
2024
Nov
Dec
Jan
Feb
Mar
Apr
May
Jun
Today
Literature Review
3 months
1 Nov - 31 Jan
Implementation
4 months
1 Dec - 31 Mar
Experiments
3 months
1 Feb - 30 Apr
Analysis
2 months
1 Apr - 31 May
Writing
5 months
1 Jan - 31 May