

InvestmentHelper-AI

Progress Meeting 3

Emir Şahin Dilli
M. Onur Özdemir
Hakan Muluk

Text to Fivelets

- Naive RAG is not suitable for our project.
- Company news are very similar to each other.
- News of the same company with different dates are very similar to each other.

Text:

Bank of Canada's notification on 10.11.2024:

The Bank of Canada (BoC) announced that consumer inflation, which was 2.7% in June, had fallen to 1.6% in September. The statement noted that business and consumer inflation expectations have largely normalized because inflationary pressures are no longer broad-based. It also mentioned that inflation is expected to remain close to the target, with upward and downward pressures on inflation roughly balanced. The statement highlighted that inflation has returned to the 2% target and that the policy interest rate was reduced by 50 basis points, bringing it to 3.75%, to support growth and keep inflation near the midpoint of the 1% to 3% range.

Assistant


Here are the extracted fivelets from the provided text:

1. ("Bank of Canada (BoC)", "announced that consumer inflation", "fell from 2.7% in June to 1.6% in September", ["Bank of Canada"], ["01.06.2024", "30.09.2024"])
2. ("Bank of Canada (BoC)", "noted that business and consumer inflation expectations", "have largely normalized", ["Bank of Canada"], ["10.11.2024", "10.11.2024"])
3. ("Bank of Canada (BoC)", "mentioned that inflationary pressures", "are no longer broad-based", ["Bank of Canada"], ["10.11.2024", "10.11.2024"])
4. ("Bank of Canada (BoC)", "expects inflation", "to remain close to the target", ["Bank of Canada"], ["10.11.2024", "-"])
5. ("Bank of Canada (BoC)", "highlighted that inflation", "has returned to the 2% target", ["Bank of Canada"], ["10.11.2024", "-"])
6. ("Bank of Canada (BoC)", "reduced the policy interest rate by", "50 basis points to 3.75%", ["Bank of Canada"], ["10.11.2024", "10.11.2024"])
7. ("Policy interest rate reduction by BoC", "aims to support", "growth and keep inflation near the midpoint of the 1% to 3% range", ["Bank of Canada"], ["10.11.2024", "-"])

Neo4j Database

Entity-Relation-Entity triplets and the fivelet's corresponding document are embedded using the E-5-large-v2 model.

Vector indexes are created for faster processing.



Single-Hop User Query – Data Extraction

Related companies are extracted.

Start-end dates are extracted.

Ranking mode is extracted.

What new renewable energy investments has Enerjisa Enerji made?

Assistant

```
[ "Enerjisa Enerji", [ "-", "-" ], "late"
```

Retrieval of Related Texts

- First, filtering based on companies and dates are done.
- Semantic similarity scores are calculated with the following formula:

$$semanticScore = cos_sim(userQuery, tripletEmbedding) + 0.5 \cdot cos_sim(userQuery, documentEmbedding)$$

- After that, temporal score is calculated based on ranking mode:

$$\tau(qt, dt) = \frac{\alpha_{scale}}{qt - dt}$$

$$\tau(qt, dt) = \frac{\tau(qt, dt) - \mu_{\tau}}{\sigma_{\tau}} \times \sigma_s + \mu_s$$

- Lastly, semantic scores and temporal score are summed and retrieval is done based on that final score.

Retrieval Example

Fivelets retrieved for the user query “What is the relationship between YEOTK and SEP?” are below:

```
241 driver = GraphDatabase.driver(uri, auth=(username, password))
242
243 q = embed_text("What is the relationship between YEOTK and SEP?")
244
245 results = retrieve_and_rerank_with_temporal_in_neo4j(["YEOTK"], ["-", "-"], q, driver, 1, "late")
```

PROBLEMS OUTPUT TERMINAL PORTS DEBUG CONSOLE

zsh + v [] [] ... ^ X

```
{'fivelet': 2, 'entity1': 'YEOTK', 'entity2': 'on 15.09.2024', 'originalRelation': 'announced the partnership agreement', 'final_score': 2.7538870921695566, 'retrieval_score': 1.3674564361572266, 'normalized_temporal_score': 1.38643065601233, 'temporal_score': 0.012658227848101266}
```

```
{'fivelet': 4, 'entity1': 'YEOTK and SEP', 'entity2': 'on 13.08.2024', 'originalRelation': 'signed partnership agreement', 'final_score': 2.7483475455677175, 'retrieval_score': 1.4065340757369995, 'normalized_temporal_score': 1.341813469830718, 'temporal_score': 0.008928571428571428}
```

```
{'fivelet': 0, 'entity1': 'Contract for construction signed by YEOTK and SEP', 'entity2': 'EUR 65.8 million', 'originalRelation': 'is worth', 'final_score': 2.7425208389759064, 'retrieval_score': 1.3783987760543823, 'normalized_temporal_score': 1.364122062921524, 'temporal_score': 0.010793399638336348}
```

```
{'fivelet': 17, 'entity1': 'Investments by YEOTK and SEP', 'entity2': 'DEFIC Globe Enerji A.Ş. (Defic Globe)', 'originalRelation': 'are under subsidiaries of', 'final_score': 2.741777688264847, 'retrieval_score': 1.3776556253433228, 'normalized_temporal_score': 1.364122062921524, 'temporal_score': 0.010793399638336348}
```

```
{'fivelet': 3, 'entity1': 'YEOTK', 'entity2': 'Shanghai Electric Power Co Ltd (SEP) for joint investments', 'originalRelation': 'announced partnership agreement with', 'final_score': 2.7347606667911672, 'retrieval_score': 1.3929471969604492, 'normalized_temporal_score': 1.341813469830718, 'temporal_score': 0.008928571428571428}
```

```
{'fivelet': 16, 'entity1': 'DEFIC Globe Enerji A.Ş. (Defic Globe)', 'entity2': 'used for investments', 'originalRelation': 'is a subsidiary where YEOTK holds 51% stake', 'final_score': 2.7337313047016, 'retrieval_score': 1.34730064868927, 'normalized_temporal_score': 1.38643065601233, 'temporal_score': 0.012658227848101266}
```

Single-Hop Examples

Single-Hop questions require only one chunk of text while generating an answer.

Multi-Hop questions require different text chunks to be combined for answering.

User: I heard that YEOTK has made an investment in the energy sector, which energy type it is?

Chatbot: YEOTK has made an investment in the solar power sector, specifically in solar power plants in Romania.

○ (base) hakanmuluk@Hakan-MacBook-Pro-2 rag % █

User: What was the inflation rate in Canada in June?

Chatbot: The inflation rate in Canada in June was 2.7%.

○ (base) hakanmuluk@Hakan-MacBook-Pro-2 rag % █

Another Database

MongoDB is used for storing chats, messages within chats, users etc.

A user might have multiple chat instances with the Chatbot.

Each chat stored belongs to a user.

Each message belongs to a chat.

Relation order: message → chat → user

Collection Name	Documents	Logical Data Size	Avg Document Size
chats	6	378B	63B
messages	24	11.66KB	498B
users	4	246B	62B

FastAPI

The backend endpoints for relevant functionalities is implemented as a FastAPI.

```
app = FastAPI()
```

The pseudo-schema of object stored inside the database is also indicated here. Example:

```
class Message(BaseModel):  
    sender: str  
    message: str  
    timestamp: datetime
```

Currently, we have all the endpoints needed for creating users, creating chats, sending messages, displaying messages to user, rephrasing queries for follow-up questions and so on.

We can test the skeleton of our project with the existing endpoints.

Follow-up Questions

User might ask follow-up questions regarding previous conversation with the LLM.

Example:

Initial Question to LLM: *Who founded Apple?*

Llama: The founders of Apple Inc. were:

1. **Steve Jobs** (1955-2011)
2. **Steve Wozniak** (born 1950), also known as "Woz"
3. **Ronald Wayne** (born 1934)

However, it's worth noting that Ronald Wayne sold his share of Apple back to the other two founders just a few months after its inception in 1976, so he is often not considered one of the key figures in the company's success.

Steve Jobs and Steve Wozniak are widely credited with creating Apple Computer, Inc. (later renamed Apple Inc.) and revolutionizing the personal computer industry with their innovative designs, such as the Apple I and Apple II.

You: How were they related before founding the company?

Rephrased Query: What was the relationship between Steve Jobs, Steve Wozniak, and Ronald Wayne before they founded Apple?

Follow-up Questions

The previous conversation and the user query is processed together, and a single self-contained query is generated.

Then, this new query is used in the LLM call.

The extent of previous information that needs referencing before rephrasing the user query is a hyperparameter of the pipeline.

That is, look at 2 previous messages, or 3, 4 ...n.

Follow-up Questions

A structured conversation (date and sender of each message is indicated), with the user query is sent to a LLM.

```
[2024-12-01 20:48:40] You: Where was Einstein born?  
[2024-12-01 20:48:47] Llama: Albert Einstein was born on March 14, 1879, in Ulm, King  
[2024-12-01 20:49:43] You: Is there any other known scientist from where he was born?  
[2024-12-01 20:49:53] Llama: The only well known scientist I could verify is Max Planck  
quantum theory.
```

The LLM is generates a new query that incorporates necessary information referencing the conversation given, that is fully self-contained.

This newly generated query is returned for the original use.

Main Page



Indicators

Tools

Watchlist

Alerts





Hourly ▾

17

19:20

20:20

21:20

22:20

23:20

5

01:20

02:20

03:20

04:20

05:20

06:20

167.50
165.00
162.50
160.00
157.50
155.00
152.50
150.00
147.50
145.00
142.50
140.00
137.50
135.00
132.50
130.00
127.50
125.00
122.50
120.00
117.50
115.00
112.50
110.00
107.50
105.00
102.50

Chats



Chat 1



Chat 2



Chat 1

Hello! How can I assist you?

Chatbot

This is a placeholder response!

Chatbot

What is the current situation of AKBANK

User

Type your message...

