

# GraPPI: A Retrieve-Divide-Solve GraphRAG Framework for Large-scale Protein-protein Interaction Exploration

Ziwen Li  
UCLA

Los Angeles, USA  
zil105@ucla.edu

Xiang ‘Anthony’ Chen  
UCLA

Los Angeles, USA  
xac@ucla.edu

Youngseung Jeon\*  
UCLA

Los Angeles, USA  
ysj@ucla.edu

## Abstract

Drug discovery (DD) has tremendously contributed to maintaining and improving public health. Hypothesizing that inhibiting protein misfolding can slow disease progression, researchers focus on target identification (Target ID) to find protein structures for drug binding. While Large Language Models (LLMs) and Retrieval-Augmented Generation (RAG) frameworks have accelerated drug discovery, integrating models into cohesive workflows remains challenging. We conducted a user study with drug discovery researchers to identify the applicability of LLMs and RAGs in Target ID. We identified two main findings: 1) an LLM should provide multiple Protein-Protein Interactions (PPIs) based on an initial protein and protein candidates that have a therapeutic impact; 2) the model must provide the PPI and relevant explanations for better understanding. Based on these observations, we identified three limitations in previous approaches for Target ID: 1) semantic ambiguity, 2) lack of explainability, and 3) short retrieval units. To address these issues, we propose *GraPPI*, a large-scale knowledge graph (KG)-based retrieve-divide-solve agent pipeline RAG framework to support large-scale PPI signaling pathway exploration in understanding therapeutic impacts by decomposing the analysis of entire PPI pathways into sub-tasks focused on the analysis of PPI edges<sup>1</sup>.

## 1 Introduction

The discovery of new drugs can potentially create treatments that save lives and enhance health outcomes globally. For example, penicillin, discovered in the early twentieth century, revolutionized bacterial infection treatments and saved countless lives (Drews, 2000). Based on the critical hypothesis that inhibiting and activating protein misfolding can slow disease progression, drug discovery

GraPPI framework to recommend PPIs with generated explanations

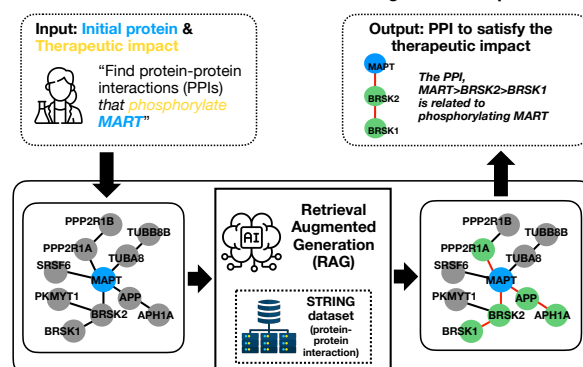


Figure 1: GraPPI for target identification (Target ID). Based on the two inputs- an initial protein and therapeutic impact on the initial protein- GraPPI recommends PPI pathways with explanations and retrieves text based on previous work.

(DD) researchers are focused on *Target identification* (Target ID), the process of elucidating the protein structures that drugs can bind to. As depicted in Figure 1, this process aims to identify protein-protein interactions (PPIs), which are protein pathways from an initial protein (IP) to protein candidates for the target protein (TP). The TP should have a therapeutic impact on the IP. Given that the number of protein candidates in the human body is several billion (Smith and Kelleher, 2013), Target ID is very time-consuming and expensive, requiring DD researchers to explore PPI candidates within the extensive protein space by scanning related literature for validation.

Recently, Large Language Models (LLMs) (OpenAI, 2023; Touvron et al., 2023; Devlin et al., 2019; AI, 2024) have become a trending generative model for understanding user intents (Li et al., 2023; Zhao et al., 2023; Kang et al., 2023) such as preferred therapeutic impacts and generating biomedical advice in natural language (Bender and Koller, 2020; Vaswani, 2017). However, they are prone to hallucinations that introduce unreliable results (Ji et al.,

<sup>1</sup>Code link: <https://github.com/AaronLi43/GraPPI>

\*Corresponding author.

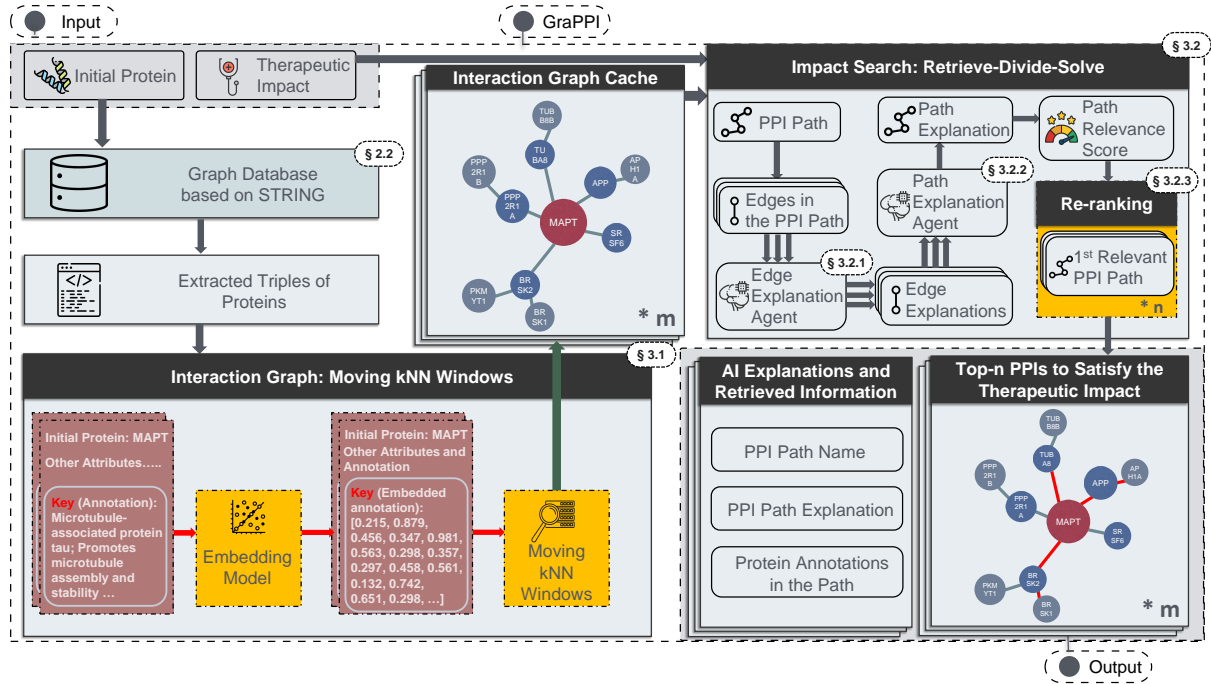


Figure 2: Overview of *GraPPI* framework. The input of the users to *GraPPI* is the name of the initial protein and the therapeutic impact query. The outputs are recommended PPIs with AI-generated explanations and retrieved information from the database.

2023; Zhang et al., 2023b). As a result, Retrieval-Augmented Generation (RAG) represents a promising solution by combining a constantly updated database with efficient information retrieval for more accurate and contextually-related responses (Lewis et al., 2020; Li and Huang, 2023; Khandelwal et al., 2019). After recognizing the potential to mitigate hallucination and improve the reliability in generative models (Martineau et al., 2023; Zhang et al., 2023a), RAG has been introduced to the biomedical-related realms (Lin et al., 2024; Wang et al., 2024; Li et al., 2024; Yang et al., 2024) where reliability and explainability matter especially for domain experts.

To identify the applicability of LLM- or RAG-based methods, we conducted user studies with five drug discovery researchers. We observed that an LLM should provide multiple Protein-Protein Interactions (PPIs) based on initial protein and protein candidates that align with the therapeutic impact. Additionally, the model must provide the PPI and relevant explanations for researchers' understanding. Based on these observation, we identified three main limitations in applying LLMs to *Target ID*.

**First**, previous works have semantic ambiguity issues where the subject and object shift (Ji et al., 2023). RAG may learn or predict incorrect relationships if it fails to interpret this accurately

and may struggle to recognize relation information from a corpus of drastically increasing biomedical literature before PPI signaling pathway analysis, due to noise or irrelevant information (Ji et al., 2023). Knowledge Graphs (KGs) are widely used to store structured information with well-assigned attributes to each entity. This can mitigate inaccurate data interpretation of the used dataset in the RAG database (Wang et al., 2023).

**Second**, the lack of explainability of LLM-based and RAG-based systems when reasoning and inference process hinders fact-checking and exacerbates over-reliance on AI recommendations (Si et al., 2023). RAG typically only provides retrieved information for the generators rather than to users. Domain experts cannot easily access to the related context to verify the truthfulness of their expertise.

**Third**, Short retrieval units in existing RAG frameworks restrict performance of tasks requiring extended context information, particularly in exploring pathways with multiple PPIs, while general KGs usually have shorter contexts (Sun et al., 2024; Luo et al., 2023). A limited context window leads to semantic incompleteness due to truncation of supporting information (Jiang et al., 2024).

To bridge these gaps, we propose *GraPPI*, with the key idea illustrated in Figure 2, a KG-based RAG framework that supports the exploration of

PPI signaling pathways. We incorporate all *Homo sapiens* PPI data from the STRING dataset to construct our KG in *GraPPI*. To guarantee data reliability in the used dataset, we consulted with DD domain experts, who confirmed full confidence in the STRING dataset. It utilizes fast  $k$ -nearest neighbor ( $k$ NN) search windows to extract relevant PPI context in the entire KG and breaks down the task of long PPI signaling pathways inference into parallel sub-tasks focused on PPI edges inference in the potential pathways using a retrieve-divide-solve style agent pipeline to **incorporate long contexts of protein annotations to PPI exploration to support fact-checking**. For LLM inference on single PPIs and entire PPI signaling pathways, we introduce co-designed Chain-of-Thought (CoT) with domain experts to help agents to understand single PPI analysis and multi-PPI pathway analysis.

In summary, the key contributions of our work are as follows:

- We propose *GraPPI*, a KG-based RAG framework in a retrieve-divide-solve agent style pipeline to break down the exploration of the entire PPI pathway into parallel analysis of sub-edges inside to enable accurate inference on long PPI contexts.
- We propose a hybrid retrieval strategy incorporating multiple fast  $k$ NN searches and LLM inference, which promotes efficient fact-checking based on contexts retrieved in large-scale KG for fact-checking and mitigate over-reliance on AI recommendations.
- To our best knowledge, we have constructed the largest KG covering all human PPI information for improved interpretation of entity relationships in a reliable dataset, consisting of 18,767 proteins and 2,955,220 PPIs from the STRING dataset.

## 2 User Study

In this appendix, we describe the details of our user study (Section 2.1), existing processes in Target ID (Section 2.2), and the challenges and applicability of RAG as potential solutions (Section 2.3).

### 2.1 Approach

We recruited five researchers specializing in drug discovery and conducted user studies in February 2024 to gain a deep understanding of their *Target ID* process. Their work experience spans 7 to 10

years (mean = 8.75, SD = 1.3). Each interview lasted approximately 60-90 minutes. Three authors attended all sessions, took notes during the discussions, and later consolidated and analyzed the notes in wrap-up meetings. Two participants, who were our collaborators, were not compensated, while the remaining three received a \$30 gift certificate after the study. Table 4 in the appendix shows participants' demographic details.

During the interviews, we asked about (1) their current practices regarding *Target ID* and (2) the applicability of LLMs in *Target ID*. After completing the interviews, we employed thematic analysis and iterative open coding (Clarke et al., 2015) to analyze the interview transcripts. Three researchers coded and analyzed the transcripts for emerging themes, and the findings were iteratively discussed among the co-authors until reaching a consensus.

### 2.2 Results

*Target identification (Target ID)* is introduced as a process to explore the protein space for PPI signaling pathways. PPI signaling pathways are paths of proteins that start from an initial protein (IP) that is therapeutically related to certain diseases to a target protein (TP). There are two conditions for TP. First, the target protein should have physical and functional interactions ranging from itself to the initial protein (C1). To identify these interactions, scientists input the initial protein on STRING<sup>2</sup>, and then they retrieve a protein interaction graph consisting of hundreds of proteins based on the initial protein. Before moving on to the next step, the scientists want to identify as many proteins as possible to increase the likelihood of finding the optimal target protein. Secondly, scientists look for possible therapeutic impacts to make an initial protein be inhibited or activated (C2). Scientists search for the therapeutic impacts of proteins on the interaction graph via Google Search. For example, MAPT is a key protein associated with Alzheimer's disease because excessive phosphorylation (activation) of MAPT promotes the disease. When researchers search for therapeutic impacts of proteins, they should verify whether a protein phosphorylates MAPT. As a result, scientists filter the proteins on the graph to retain only those with the desired impact, significantly reducing their number.

<sup>2</sup><https://sea.bkslab.org/>

### 2.3 Applicability

The space of possible targets is expansive, given that researchers estimate around 10,000 (Adkins et al., 2002) proteins in the human body. Scientists are constrained to exploring limited protein candidates, negatively impacting scientific discoveries. In our user study, the experts provided highly positive feedback on using LLMs with PPI graphs to explore the extensive protein space, citing their efficiency in identifying potential therapeutic target protein (TP) candidates. If LLMs can identify PPIs having desired TPs in a PPI graph based on an initial protein (IP) consisting of several PPIs, it demonstrates their strong potential to support scientific discovery effectively. However, they also emphasized the necessity of utilizing a confident dataset, such as the STRING dataset, that provides interactions for over 3 million proteins. Researchers also mentioned the importance of fact-checking with explanations for why the model recommends specific PPIs. The absence of scientific materials for results could diminish the quality of LLM-generated recommendations.

In summary, key findings for supporting Target ID are as follows:

- A RAG should provide a PPI composed of multiple proteins, with the IP and TP as start and end points, respectively, and ensure alignment with the desired therapeutic impact.
- A RAG should utilize datasets that scientists can trust, such as the STRING dataset.
- A RAG should provide the PPI alongside relevant explanations to enable fact-checking.

## 3 Preliminaries

In this section, we introduce the formulation of the problem and the construction of our medical graph database. We conducted a user study with domain experts in DD to understand their requirements for PPI exploration: a framework that can understand their initial protein and therapeutic impact then outputting recommended PPI pathways with explanations and retrieved contexts from reliable datasets for fact-checking. More details on formative study are provided in the Appendix.

### 3.1 Problem Formulation

Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  be an instance of a knowledge graph, where  $\mathcal{V}$  is the set of nodes and  $\mathcal{E}$  is the

set of edges. Node  $v \in \mathcal{V}$  and edge  $e \in \mathcal{E}$  represent protein and PPI in the PPI network. The goal of *GraPPI* is to provide potential PPI signaling pathways and offer AI-generated explanations and database-retrieved contextual information to support experts’ decision-making. The following equation can describe the input and output of *GraPPI*:

$$\text{GraPPI}(G, Q_{\text{users}}, IP) = \arg \max_{P \subseteq G} \left( S(P, Q_{\text{users}}, X_{\text{AI}}(P), I_{\text{DB}}(P)) \right) \quad (1)$$

In this equation,  $G$  represents the interaction graph derived from the database, while  $P$  represents the potential PPI signaling pathway.  $IP$  represents the initial protein.  $Q_{\text{users}}$  indicates the therapeutic impact from users.  $I_{\text{DB}}(P)$  represents the retrieved information about certain pathways from the database. The equation maximizes the relevance score  $S(P)$  for a pathway based on a given therapeutic impact  $Q_{\text{users}}$ .  $X_{\text{AI}}(P)$  (AI-generated explanation), and  $I_{\text{DB}}(P)$  (retrieved information from the database) over the pathway  $P$  are utilized as supplementary information for evaluation. We will assess the results of the recommended PPI candidates using quantitative metrics for semantic similarity literal alignment and subjective evaluation of domain experts in DD.

### 3.2 Medical Graph Database Construction

Dataset	#Entities	#Triples
Ours	18,767	2,955,220

Table 2: Basic statistics of our knowledge graphs.

In this section, we describe the process of constructing our KG database containing domain knowledge of all human PPIs from the STRING dataset (Szk-larczyk et al., 2023), through data collection and deployment. After domain experts in drug discovery confirmed the reliability of the STRING dataset, we collected information on all proteins within the *Homo sapiens* category, identifying and reorganizing representative features including *combined\_score*, *interaction\_type*, *protein*, *text annotation*, and *embedded-annotation vector*. The *text annotation* and *combined\_score* attributes are sourced from prior work, with their reliability validated by experts. Attribute definitions are summarized in Table 1. The embedding model used is Ope-



Notation	Description
<i>protein</i> (Node)	the name of a protein
<i>start_annotation</i> (Node)	the text to describe the property of the start protein in a PPI
<i>end_annotation</i> (Node)	the text to describe the property of the end protein in a PPI
<i>embedding</i> (Node)	the semantic embedding vector of the annotation text
<i>combined_score</i> (Edge)	a score to show the confidence of the PPI
<i>interaction_type</i> (Edge)	a text to indicate the type of the PPI
<i>path</i> (Path)	names of proteins along the PPI signaling pathway
<i>path_details</i> (Path)	generated explanations and retrieved annotations of a PPI pathway

Table 1: Attributes of the constructed knowledge graph (KG). (Node) indicates the attributes of nodes, (Edge) represents the attributes of edges, and (Path) represents the attributes of paths.

nAI’s *text-embedding-3-small*<sup>3</sup>. We employed a *neo4j* database dump to inject the reconstructed dataset and deployed it on Google Cloud Platform for cloud hosting. The resulting graph database contains 18,767 nodes and 2,955,220 edges, representing the largest PPI graph database for *Homo sapiens* proteins to date, with the most comprehensive PPI relationships.

## 4 GraPPI Framework

To add supporting information for LLMs to understand single PPI and multi-PPI signaling pathway analysis, we developed a KG-based RAG framework to support PPI pathway exploration. The framework contains two components: (1) **Interaction Graph: Moving  $k$ NN windows** to extract the relevant subgraphs from the KG to enable LLM inference over the large-scale KG, and (2) **Impact Search: A retrieve-divide-solve style agent pipeline** to understand PPI pathways on single PPI edge level and entire PPI pathway level context. The overview of *GraPPI* is illustrated in Figure 2.

### 4.1 Interaction Graph: Moving $k$ NN Graph Windows

After discussion with domain experts, we constructed the entities in KG:  $\langle head\ protein, combined\_score, tail\ protein \rangle$ . To start to generate *Interaction Graph* as the PPI context, *GraPPI* will receive the name of the initial protein from users and use Cypher statement to extract all the connected nodes (proteins) in the KG. Among all the connected protein candidates, *GraPPI* will implement a moving  $k$ -nearest neighbor ( $k$ NN) graph window strategy to form interaction graphs allow-

ing users to select potential PPIs and mitigate redundant inference in *Interaction Graph*. The pseudocode for moving  $k$ NN graph windows is shown in Algorithm 1. In this case, we use the FAISS open-source library for fast  $k$ NN search in high-dimensional dense representation space (Johnson et al., 2019).

#### Algorithm 1 Moving $k$ NN Graph Windows

**Input:** : Initial protein  $p_{init}$ ,  $k$ NN constants  $K$ , Window indexes  $I$ , Graph size parameter  $M$   
**Output:** Expanded graph  $G$   
1: Initialize interaction graph  $G$  with  $p_{init}$  as the root node  
2: Set  $m \leftarrow 0$   
3: **while**  $m < M$  **do**  
4:   **for**  $protein \in G$  **do**  
5:     Determine the moving  $k$ NN window  $[i * k + 1, (i + 1) * k], i \in I, k \in K$   
6:     Retrieve  $k$  connected proteins using the moving  $k$ NN search of annotation embeddings  
7:     Add retrieved proteins and edges to  $G$   
8:     Increment  $m \leftarrow m + 1$   
9:   **end for**  
10: **end while**

### 4.2 Impact Search: Retrieve-Divide-Solve

In *Impact Search*, users can select the interaction graph that contains their preferred proteins for LLM inference based on therapeutic impact. After interaction graph selection, *Impact Search* will take the interaction graph and the therapeutic impact query as the input to perform the PPI pathway exploration using a retrieve-divide-solve style agent pipeline. Instead of directly incorporating long-context annotations in the long PPI signaling pathways, *Impact Search* first decomposes the analysis for the entire tasks into multiple subtasks, with each subtask retrieving the long annotations of two end proteins to explain each PPI in parallel for the whole of the PPI pathway recommendation. Next, we will walk through *Impact Search* in *GraPPI*

<sup>3</sup><https://platform.openai.com/docs/guides/embeddings>

step by step.

#### 4.2.1 Step I: Edge Explanation

Instead of directly analyzing the entire PPI pathway, which is very likely to have an extremely long PPI context, the edge explanation agent breaks down the entire path into a collection of edges with detailed information about proteins. Each edge will be analyzed with corresponding protein context and therapeutic impact to generate shorter and more concise explanations with less redundant information. The edge explanations will be the information source for path explanation and supporting context for fact-checking.

#### 4.2.2 Step II: Path Exploration

In Step II, the edge explanation agent will generate the explanation for the entire PPI signaling pathway. The generated path explanation and the relevance score will be stored as attributes of each path. After the generation, all the intermediate results from *k*NN search in *Interaction Graph* to Edge Explanation and Path exploration in *Impact Search* are accessible to domain experts to support their fact-checking.

#### 4.2.3 Step III: Re-rank

When generating path explanations, the path explanation agent is also responsible for evaluating the relevance of the pathway to the therapeutic impact using the demonstrated zero-shot ranking ability of LLMs (Hou et al., 2024; Zhao et al., 2023) since previous studies show the strong potential of transferring knowledge from LLMs as powerful recommendation models. Using the calculated relevance score, *GraPPI* will show the top *n* relevant pathways among all PPIs to users rather than using the ranking in the previous interaction graph. Users specify the ranking window length *n*.

## 5 Experiment

We will assess the performance of *GraPPI* in multiple configurations to answer the following experimental questions (EQs):

- **EQ1:** To what extent does *GraPPI* improve the accuracy of path explanations compared with baseline models?
- **EQ2:** How does the retrieve-divide-solve pipeline improve the accuracy and efficiency when scaling up interaction graphs?

- **EQ3:** How do the domain experts perceive the utility of AI-generated explanations and retrieved information in supporting their decision-making?

**EQ1, EQ2, and EQ3** will be explained in Section 5.4.1, 5.4.2, and 5.4.3, respectively.

### 5.1 Setup

For the evaluation, we accessed the performance of our framework in terms of two aspects: (1) accuracy to demonstrate how the generated contents align with the requirements of users and the background information, and (2) the scalability that evaluates the framework performance when *Interaction Graph* scales up. We conducted the experiments on both edge-level generated explanation and path-level explanation. For the comparative study and the ablation study, the graph size parameter was set to be 2, and the hyperparameters for *k*NN  $k_1$  and  $k_2$  to explore the nodes in depth 1 and depth 2 were set to be 10 and 2 respectively. *Impact Search* will recommend the top 10 PPI signaling pathway candidates for each interaction graph according to the relevance to the therapeutic impacts from users. For scalability analysis, we test the performance of *GraPPI* using 6 *Interaction Graph* with different sizes. The depths of those *Interaction Graph* are 2 and 3 while *k* value for *k*NN are 10, 15, and 20 respectively. The numbers of paths are set as follows: [40, 75, 76, 95, 113, 160]. The *Interaction Graph* in the experiment group is using *GraPPI* to generate path explanations based on the edge explanations. The *Interaction Graph* in the control group directly refers to raw annotation texts of each protein in the pathway to generate path explanations.

### 5.2 Baseline

We assess the performance of our framework using three state-of-the-art LLMs as base models: GPT-4o<sup>4</sup>, GPT-4-Turbo (Achiam et al., 2023), and GPT-4o-mini<sup>5</sup>. GPT-4o and GPT-4o-mini are the cost-effective variants of GPT-4-Turbo designed to balance performance and computational expense.

### 5.3 Tasks and Metrics

To evaluate the accuracy of the generated contents, we employed **BERTScore** (Zhang et al., 2019), **ROUGE-1**, and **ROUGE-L** (Lin, 2004)

<sup>4</sup><https://openai.com/index/hello-gpt-4o/>

<sup>5</sup><https://platform.openai.com/docs/models#gpt-4o-mini/>

Methods	metrics (Mean±Std)		
	BERTScore F1	ROUGE-1 F1	ROUGE-L F1
<b>GPT-4o-mini-Baseline</b>	86.76±2.13	24.30±8.54	21.33±6.83
<b>GPT-4o-mini-Zero shot w/ CoT</b>	88.55±1.99	26.38±8.13	30.96±11.29
<b>GPT-4o-mini-RAG w/o CoT</b>	87.53±1.61	27.80±11.79	22.51±6.26
<b>GPT-4o-mini-Ours</b>	<u>88.98±1.32</u>	<u>30.41±7.12</u>	<u>32.28±6.57</u>
<b>GPT-4-Turbo-Baseline</b>	88.00±2.00	32.46±11.17	24.40±9.14
<b>GPT-4-Turbo-Zero shot w/ CoT</b>	89.32±1.58	36.69±11.39	31.02±11.03
<b>GPT-4-Turbo-RAG w/o CoT</b>	89.95±0.94	38.70±7.04	31.93±6.61
<b>GPT-4-Turbo-Ours</b>	<b><u>90.02±0.93</u></b>	<u>42.19±5.53</u>	<b><u>37.47±6.96</u></b>
<b>GPT-4o-Baseline</b>	86.76±2.13	29.10±8.28	21.33±6.84
<b>GPT-4o-Zero shot w/ CoT</b>	87.64±2.28	34.07±12.25	30.76±10.75
<b>GPT-4o-RAG w/o CoT</b>	88.70±1.46	34.36±4.15	30.06±4.94
<b>GPT-4o-Ours</b>	<u>89.62±1.82</u>	<b><u>42.26±11.96</u></b>	<u>35.67±12.58</u>

Table 3: Performance comparison on different models and configurations. Underline means the best configuration in a given model and **Bold font** represents the top-performing methods in the evaluation. **Baseline** represents vanilla model setting without retrieved materials from our KG and a co-designed chain-of-thought (CoT) with domain experts. **Zero shot** means that models will do inference without the support of retrieved contexts from the KG. **Ours** includes both components to complete the pipeline.

to assess semantic alignment and literal overlapping. Since our framework aligns the generated content with both retrieved knowledge from our database and users’ intent, we combined input query with retrieved information and formed the input tuple of edge level as:  $\langle "query", "start\_protein", "end\_protein", "start\_annotation", "end\_annotation" \rangle$ . The input tuple of path level is:  $\langle "query", "path", "path\_details" \rangle$ . The definitions of those variables are shown in Table 1. The metrics are calculated between the input tuples of edges level along the recommended PPI pathway and the generated path explanation. Regarding scalability evaluation, we conducted scale variant evaluation to compare the performance of *GraPPI* under different scales of *Interaction Graph* shown in Figure 3.

## 5.4 Experimental Results

### 5.4.1 Ablation Study

**EQ1: Accuracy:** In this experiment, we conducted the accuracy evaluation of *GraPPI* under a given data scale. As shown in Table 3, *GraPPI* outperforms not only the baseline but all other configurations with all three base models in terms of semantic similarity (BERTScore) and literal alignment (ROUGE-1 and ROUGE-L). Four configurations are designed: Ours, RAG w/o Chain-of-Thought (CoT), Zero shot w/ CoT, and Zero shot w/o CoT

(Baseline). To ensure consistency, we implemented those configurations on the same initial protein from domain experts for sub-graph retrieval, and we adjusted LLM inference settings for the ablation study. One interesting thing we observed while manually checking the generated explanations is that although GPT-4-Turbo has the best quantitative performance among all three base models, its explanations may prioritize protein annotation details over user-specified therapeutic impacts. GPT-4-Turbo generates explanations covering more divergent topics, while GPT-4o focuses more on the users’ query to give more concise explanations.

### 5.4.2 Scale Variant Evaluation

**EQ2: Scalability and Efficiency:** We evaluate the performance of *GraPPI* using GPT-4o under interaction graphs of different scales. As shown in Figure 3, by comparing all four metrics in the six scales, significant differences exist between the two data groups for all metrics except ROUGE-1 and ROUGE-L when the number of pathways reaches 160. In this experiment, six *Interaction Graph* containing different numbers of pathways are used to test *GraPPI*’s scalability. For each *Interaction Graph*, *GraPPI* recommends 10 potential PPI pathway candidates based on their relevance to user queries. We employed paired t-tests to determine whether *GraPPI*’s introduction leads to significant performance differences. While no statistically sig-

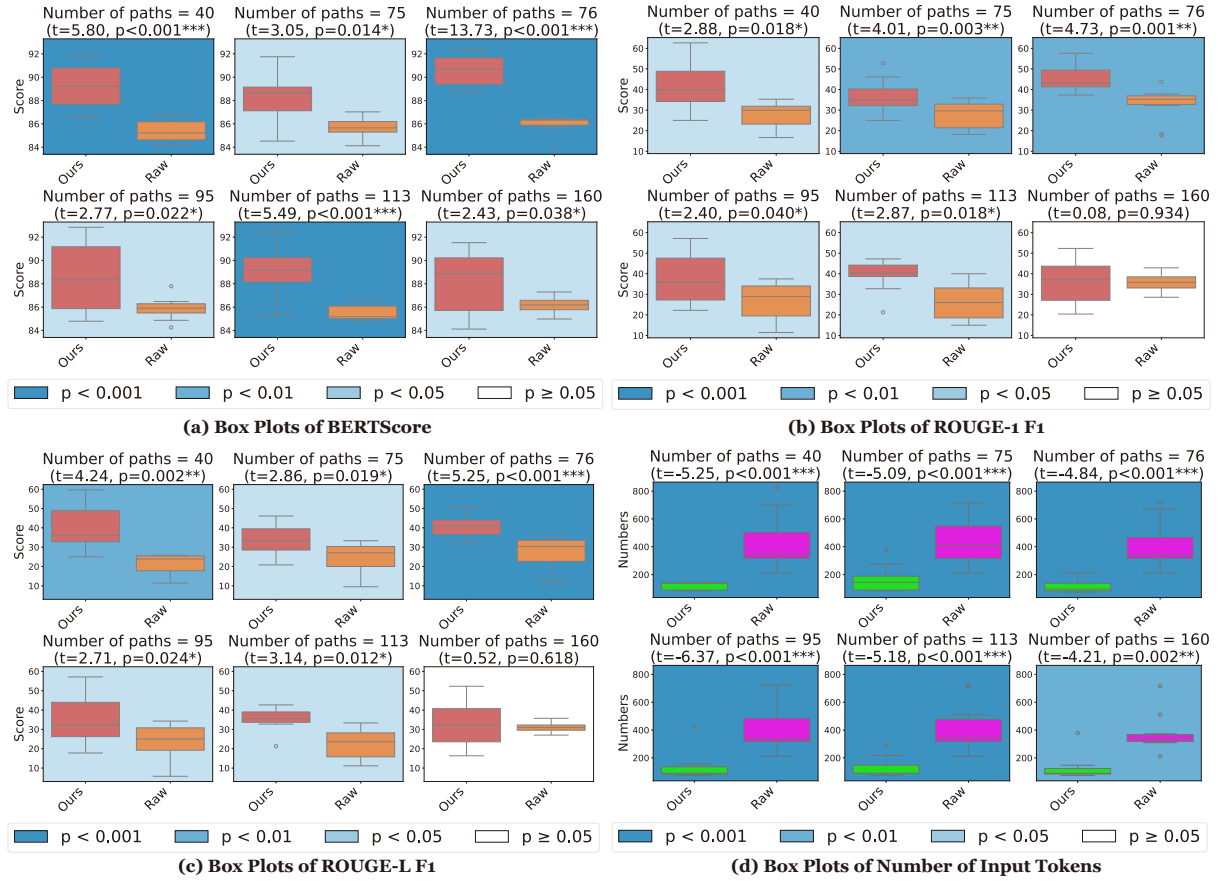


Figure 3: Box plots of results under different graph sizes. The blue color in the background indicates the level of difference between the two groups. Darker blue represents a more significant difference. For plots (a), (b), and (c), the red and orange boxes represent different accuracies of *GraPPI* and the system directly using protein annotations, respectively. For plot (d), the green and purple boxes indicate the number of input tokens they have. **Raw** refers to the methods using raw annotations texts as contexts while **Ours** utilizes the edge explanations with more concise representation of biomedical context.

nificant differences were observed between groups for ROUGE-1 and ROUGE-L at 160 paths, the mean scores for *GraPPI* still surpassed those of the Raw group, indicating superior literal alignment. Overall, *GraPPI* achieves better semantic similarity and lexical alignment while using significantly fewer tokens. These results highlight *GraPPI*'s potential as an accurate and token-efficient computational tool for exploring PPI pathways in large-scale datasets.

### 5.4.3 Case Study

**EQ3: Explainability:** We further conducted a case study to evaluate the quality of responses generated by *GraPPI*. Users reviewed the generated content in *Interaction Graph* and *Impact Search* to assess *GraPPI*'s contribution to exploring novel PPI pathways. We interviewed users about their current *Target ID* practices and the applicability of LLMs in *Target ID* to validate the generated content. Given

an initial protein and therapeutic impact, *GraPPI* produced 5 PPI pathways (containing 9 PPI edges) with explanations and retrieved descriptive texts for pathway validation. Users leveraged these outputs to explore new PPI pathways. All generated explanations were deemed coherent, and users could integrate them with prior work context to make decisions. Baseline model outputs were less convincing due to the lack of fact-checking support, with some requiring manual calibration. *GraPPI* offers a novel approach to PPI pathway exploration, providing biomedical explanations and retrieved information to mitigate over-reliance on AI. Case study details and sample results are provided in Figure 4 in the Appendix.



## 6 Related Works

### 6.1 LLM-based and RAG-based systems in PPI prediction

Large Language Models (LLMs) have demonstrated significant potential in understanding and generating natural language responses (Wang et al., 2024; Li et al., 2023; Kang et al., 2023). Inspired by LLMs, pioneering work has focused on building protein language models (PLMs), which are pre-trained on large-scale protein sequences (Hsu et al., 2022; Elnaggar et al., 2021; Li and Huang, 2023). PLMs capture more accurate protein features by representing sequences as high-dimensional embedding vectors. Previous studies have leveraged PLMs to enhance performance in downstream tasks such as protein structure prediction (Lin et al., 2023) and PPI prediction (Jin et al., 2024). However, while PLMs exhibit strong capabilities in biomedical tasks, they remain prone to hallucinations—generating text that is nonsensical or unfaithful to source content (Ji et al., 2023; Zhang et al., 2023b). To address this, Retrieval-Augmented Generation (RAG) has been introduced to mitigate hallucinations and improve trustworthiness by retrieving contextual information from external databases (Khandelwal et al., 2019; Lewis et al., 2020). RAG enhances the reliability and transparency of biomedical LLM-based systems (Yang et al., 2024).

### 6.2 Explainable AI

Heated discussions have centered on explainable AI (XAI), which generates actionable explanations for AI outputs instead of treating LLM-based systems as black boxes (Wiegrefe and Marasović, 2021; Schuff et al., 2022; Lamm et al., 2021). Transparent LLM-based systems that provide explanations and retrieve materials from existing research enhance human verification of AI-generated content (Vasconcelos et al., 2023), assist in decision-making (Fok and Weld, 2023), calibrate user trust (Bussone et al., 2015), and reduce over-reliance on AI (Das et al., 2022; Zhang et al., 2020; Bansal et al., 2021). Thus, providing reliable background information is critical to mitigating over-reliance on AI and enabling experts to calibrate trust through explanations and retrieved context (Si et al., 2023).

## 7 Conclusion

In this paper, through a carefully designed user study with drug discovery researchers, we observed the importance of retrieving PPIs with the therapeutic impact they target in *Target ID* (i.e., find PPIs to inhibit MARK4). Therefore, we proposed *GraPPI*, which is a large-scale KG-based retrieve-divide-solve style agent pipeline RAG framework to support PPI pathways exploration in understanding therapeutic impacts with two key components: 1) moving  $k$ NN windows for sub-graph extraction, and 2) a retrieve-divide-solve style agent pipeline to incorporate long protein annotations for inference. Through extensive experiments, we demonstrate that: (1) *accuracy*: *GraPPI* consistently outperforms baseline methods in semantic and lexical alignment; (2) *scalability and efficiency*: *GraPPI* achieves comparable results using significantly fewer tokens than raw annotation-based systems, and (3) *explainability*: *GraPPI* generates more reliable explanations and retrieved information, validated by domain experts.

## 8 Limitations

In this paper, we developed a KG-based RAG framework to support PPI signaling pathway exploration. However, our study has several limitations.

**First**, while domain experts in DD expressed full confidence in the STRING dataset during interviews, the dataset does not comprehensively cover all known PPIs, as biomedical research continues to discover new interactions, and the STRING dataset requires constant updates. During the case study, one expert noted that protein *TUBB3* interacts with *MAPT* via phosphorylation based on their expertise, but this interaction is absent in STRING. Integrating features such as PPI prediction models or allowing researchers to input custom findings could enhance dataset coverage.

**Second**, the case study utilized only two initial proteins validated by DD experts. While these proteins were well-understood, they do not represent the full diversity of protein interactions, leaving potential edge cases unexamined. Future work should incorporate a broader range of initial proteins to rigorously assess *GraPPI*'s applicability.

**Third**, while the retrieve-solve-merge pipeline improves PPI exploration at moderate scales, its benefits diminish as interaction graphs grow larger. This is due to the increasing length of edge explanations approaching LLM context window limits.

Thus, *GraPPI* is not a universal solution for arbitrarily large graphs; users must judiciously set the *k*NN window size to balance exploration scope and explanation quality. This approach not only generates smaller interaction graphs for analysis but also prioritizes shorter PPI pathways, which are clinically preferable due to reduced side effects and more efficient signaling transmission.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Joshua N Adkins, Susan M Varnum, Kenneth J Auberry, Ronald J Moore, Nicolas H Angell, Richard D Smith, David L Springer, and Joel G Pounds. 2002. Toward a human blood serum proteome: analysis by multi-dimensional separation coupled with mass spectrometry. *Molecular & Cellular Proteomics*, 1(12):947–955.
- Google AI. 2024. Gemini: Google’s large language models. <https://gemini.google.com/>. Accessed: 2024-08-31.
- Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel Weld. 2021. Does the whole exceed its parts? the effect of ai explanations on complementary team performance. In *Proceedings of the 2021 CHI conference on human factors in computing systems*, pages 1–16.
- Emily M Bender and Alexander Koller. 2020. Climbing towards nlu: On meaning, form, and understanding in the age of data. In *Proceedings of the 58th annual meeting of the association for computational linguistics*, pages 5185–5198.
- Adrian Bussone, Simone Stumpf, and Dymrna O’Sullivan. 2015. The role of explanations on trust and reliance in clinical decision support systems. In *2015 international conference on healthcare informatics*, pages 160–169. IEEE.
- Victoria Clarke, Virginia Braun, and Nikki Hayfield. 2015. Thematic analysis. *Qualitative Psychology: A Practical Guide to Research Methods*, pages 222–248.
- Anubrata Das, Chitrang Gupta, Venelin Kovatchev, Matthew Lease, and Junyi Jessy Li. 2022. Prototex: Explaining model decisions with prototype tensors. *arXiv preprint arXiv:2204.05426*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. *Bert: Pre-training of deep bidirectional transformers for language understanding*. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Jurgen Drews. 2000. Drug discovery: a historical perspective. *science*, 287(5460):1960–1964.
- Ahmed Elnaggar, Michael Heinzinger, Christian Dal-lago, Ghalia Rehawi, Yu Wang, Llion Jones, Tom Gibbs, Tamas Feher, Christoph Angerer, Martin Steinegger, et al. 2021. Prottrans: Toward understanding the language of life through self-supervised learning. *IEEE transactions on pattern analysis and machine intelligence*, 44(10):7112–7127.
- Raymond Fok and Daniel S Weld. 2023. In search of verifiability: Explanations rarely enable complementary performance in ai-advised decision making. *AI Magazine*.
- Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*, pages 364–381. Springer.
- Chloe Hsu, Robert Verkuil, Jason Liu, Zeming Lin, Brian Hie, Tom Sercu, Adam Lerer, and Alexander Rives. 2022. Learning inverse folding from millions of predicted structures. In *International conference on machine learning*, pages 8946–8970. PMLR.
- Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. *ACM Computing Surveys*, 55(12):1–38.
- Ziyan Jiang, Xueguang Ma, and Wenhui Chen. 2024. Longrag: Enhancing retrieval-augmented generation with long-context llms. *arXiv preprint arXiv:2406.15319*.
- Mingyu Jin, Haochen Xue, Zhenting Wang, Boming Kang, Ruosong Ye, Kaixiong Zhou, Mengnan Du, and Yongfeng Zhang. 2024. Prollm: Protein chain-of-thoughts enhanced llm for protein-protein interaction prediction. *bioRxiv*, pages 2024–04.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data*, 7(3):535–547.
- Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. 2023. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv preprint arXiv:2305.06474*.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2019. Generalization through memorization: Nearest neighbor language models. *arXiv preprint arXiv:1911.00172*.

- Matthew Lamm, Jennimaria Palomaki, Chris Alberti, Daniel Andor, Eunsol Choi, Livio Baldini Soares, and Michael Collins. 2021. Qed: A framework and dataset for explanations in question answering. *Transactions of the Association for computational Linguistics*, 9:790–806.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Jinming Li, Wentao Zhang, Tian Wang, Guanglei Xiong, Alan Lu, and Gerard Medioni. 2023. Gpt4rec: A generative framework for personalized recommendation and user interests interpretation. *arXiv preprint arXiv:2304.03879*.
- Mingchen Li and Lifu Huang. 2023. Understand the dynamic world: An end-to-end knowledge informed framework for open domain entity state tracking. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 842–851.
- Mingchen Li, Halil Kilicoglu, Hua Xu, and Rui Zhang. 2024. Biomedrag: A retrieval augmented large language model for biomedicine. *arXiv preprint arXiv:2405.00465*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Xinyi Lin, Gelei Deng, Yuekang Li, Jingquan Ge, Joshua Wing Kei Ho, and Yi Liu. 2024. Generag: Enhancing large language models with gene-related task by retrieval-augmented generation. *bioRxiv*, pages 2024–06.
- Zeming Lin, Halil Akin, Roshan Rao, Brian Hie, Zhongkai Zhu, Wenting Lu, Nikita Smetanin, Robert Verkuil, Ori Kabeli, Yaniv Shmueli, et al. 2023. Evolutionary-scale prediction of atomic-level protein structure with a language model. *Science*, 379(6637):1123–1130.
- Linhao Luo, Yuan-Fang Li, Gholamreza Haffari, and Shirui Pan. 2023. Reasoning on graphs: Faithful and interpretable large language model reasoning. *arXiv preprint arXiv:2310.01061*.
- Kim Martineau, AI Explainable, and AI Generative. 2023. What is retrieval-augmented generation? *IBM Research Blog*, 22.
- OpenAI. 2023. Chatgpt: Optimizing language models for dialogue. <https://openai.com/chatgpt>. Accessed: 2024-08-31.
- Hendrik Schuff, Alon Jacovi, Heike Adel, Yoav Goldberg, and Ngoc Thang Vu. 2022. Human interpretation of saliency-based explanation over text. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, pages 611–636.
- Chenglei Si, Navita Goyal, Sherry Tongshuang Wu, Chen Zhao, Shi Feng, Hal Daumé III, and Jordan Boyd-Graber. 2023. Large language models help humans verify truthfulness—except when they are convincingly wrong. *arXiv preprint arXiv:2310.12558*.
- Lloyd M Smith and Neil L Kelleher. 2013. Proteoform: a single term describing protein complexity. *Nature methods*, 10(3):186–187.
- Lei Sun, Zhengwei Tao, Youdi Li, and Hiroshi Arakawa. 2024. Oda: Observation-driven agent for integrating llms and knowledge graphs. *arXiv preprint arXiv:2404.07677*.
- Damian Szklarczyk, Rebecca Kirsch, Mikaela Koutrouli, Katerina Nastou, Farrokh Mehryary, Radja Hachilif, Annika L Gable, Tao Fang, Nadezhda T Doncheva, Sampo Pyysalo, et al. 2023. The string database in 2023: protein–protein association networks and functional enrichment analyses for any sequenced genome of interest. *Nucleic acids research*, 51(D1):D638–D646.
- Hugo Touvron et al. 2023. Llama: Open and efficient foundation language models. <https://github.com/facebookresearch/llama>. Accessed: 2024-08-31.
- Helena Vasconcelos, Matthew Jörke, Madeleine Grunden-McLaughlin, Tobias Gerstenberg, Michael S Bernstein, and Ranjay Krishna. 2023. Explanations can reduce overreliance on ai systems during decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1):1–38.
- A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*.
- Chengrui Wang, Qingqing Long, Xiao Meng, Xunxin Cai, Chengjun Wu, Zhen Meng, Xuezhi Wang, and Yuanchun Zhou. 2024. Biorag: A rag-llm framework for biological question reasoning. *arXiv preprint arXiv:2408.01107*.
- Jianing Wang, Qiushi Sun, Xiang Li, and Ming Gao. 2023. Boosting language models reasoning with chain-of-knowledge prompting. *arXiv preprint arXiv:2306.06427*.
- Sarah Wiegrefe and Ana Marasović. 2021. Teach me to explain: A review of datasets for explainable natural language processing. *arXiv preprint arXiv:2102.12060*.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. 2024. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *ACM Transactions on Knowledge Discovery from Data*, 18(6):1–32.
- Shuo Zhang, Liangming Pan, Junzhou Zhao, and William Yang Wang. 2023a. The knowledge alignment problem: Bridging human and external knowledge for large language models. *arXiv preprint arXiv:2305.13669*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023b. Siren’s song in the ai ocean: a survey on hallucination in large language models. *arXiv preprint arXiv:2309.01219*.

Yunfeng Zhang, Q Vera Liao, and Rachel KE Bellamy. 2020. Effect of confidence and explanation on accuracy and trust calibration in ai-assisted decision making. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*, pages 295–305.

Zihuai Zhao, Wenqi Fan, Jiatong Li, Yunqing Liu, Xiaowei Mei, Yiqi Wang, Zhen Wen, Fei Wang, Xiangyu Zhao, Jiliang Tang, et al. 2023. Recommender systems in the era of large language models (llms). *arXiv preprint arXiv:2307.02046*.

## **A Appendix**



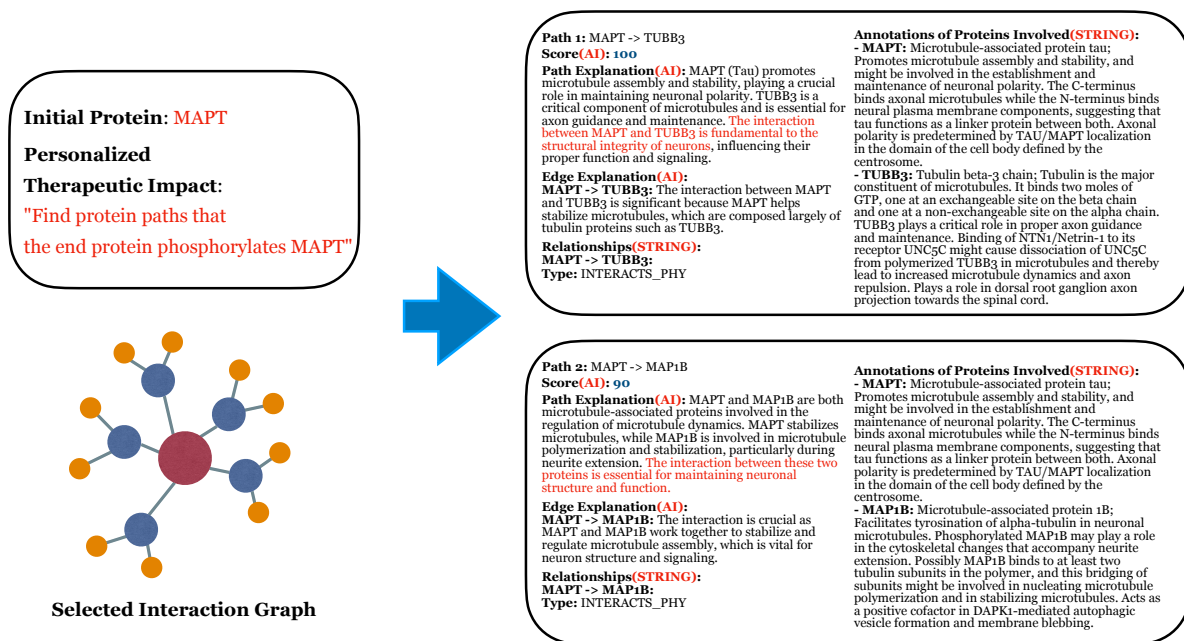


Figure 4: Part of the results of the case study showing the input and output content regarding certain recommended PPI signaling pathways.

Table 4: Demographic distribution of the five user study participants. The table includes details about gender, age range, profession, and work experience.

Demographic group		User Study	
		N	%
Gender	Female	1	25%
	Male	4	75%
Age Range	26-35	3	60%
	36-45	2	40%
Profession	Post-Doctor	2	40%
	Researcher	3	60%
Work Experience	7-9 years	2	40%
	over 10 years	3	60%

### **Prompt Template of Edge Explanations:**

*As a protein-protein interaction expert, evaluate the relevance of the interaction between {start\_protein} and {end\_protein} to the query: "{query}".*

*You should follow those steps for analysis:*

**<Steps>**

- 1. Use the name of proteins along the path to figure out the directional graph.  
Consider the direction of the interaction and always and only consider the impact of the direct last protein. The effect propagates backward to the start protein.*
- 2. The previous protein should be able to directly interact with the current protein according to the information provided. If not, the interaction is not relevant.*
- 3. Explain potential biological processes or mechanisms before recommending it. If you can not find any explicit reasons to indicate that they are relevant, the path will be irrelevant.*
- 4. The proteins on the ends of one edge must be able to directly interact with each other.  
If not, the edge does not make any sense and the relevance is zero.*
- 5. The effect propagates from the end protein to the start protein.*

**</Steps>**

*Protein information:*

*{start\_protein}: {start\_annotation}*

*{end\_protein}: {end\_annotation}*

*Explain why this interaction is relevant or not relevant to the query using less than 35 words.  
Try to explain the edge informatively using simple and short sentence structure to improve the readability.*

### **Prompt Template of Path Explanations:**

*As a protein-protein interaction expert, choose the following protein interaction paths to satisfy the query: "{query}"*

*You should follow those steps for analysis:*

**<Steps>**

- 1. Use the name of proteins along the path to figure out the directional graph.  
Protein A -> Protein B -> Protein C means that protein c can influence protein b and protein b can influence protein a. Any other directions are not relevant.*
- 2. The path should satisfy the requirements in the query. The more likely the requirements get satisfied, the higher the relevance score is.*
- 3. The previous protein should be able to directly interact with the current protein according to the information provided. If not, the interaction is not relevant.*
- 4. Explain potential biological processes or mechanisms before recommending it. If you can not find any explicit reasons to indicate that they are relevant, the path will be irrelevant.*
- 5. Provide a relevance score from 0 to 100, where 0 is not relevant at all and 100 is highly relevant. This score will be used to rank the paths according to their relevance of the explanation to the query.*
- 6. The proteins on the ends of one edge must be able to directly interact with each other. If not, the path does not make any sense and the relevance is zero.*
- 7. The effect propagates from end protein to start protein.*

**</Steps>**

*Path: {path}*

*Details of each step: {path\_details}*

*Discuss any potential biological processes or mechanisms that this path might represent and explain why you recommend this path, using less than 80 words.*

*Your explanation should be informative and accessible, as if you are explaining it to a fellow researcher. Your response MUST be in the following JSON format:*

```
{{  
  "explanation": "Your explanation here",  
  "relevance_score": "0-100"  
}}
```

Figure 5: Prompt Templates of Edge Explanation and Path Explanations