






# SLInterpreter: An Exploratory and Iterative Human-AI Collaborative System for GNN-based Synthetic Lethal Prediction

Haoran Jiang , Shaohan Shi , Shuhao Zhang , Jie Zheng  and Quan Li 

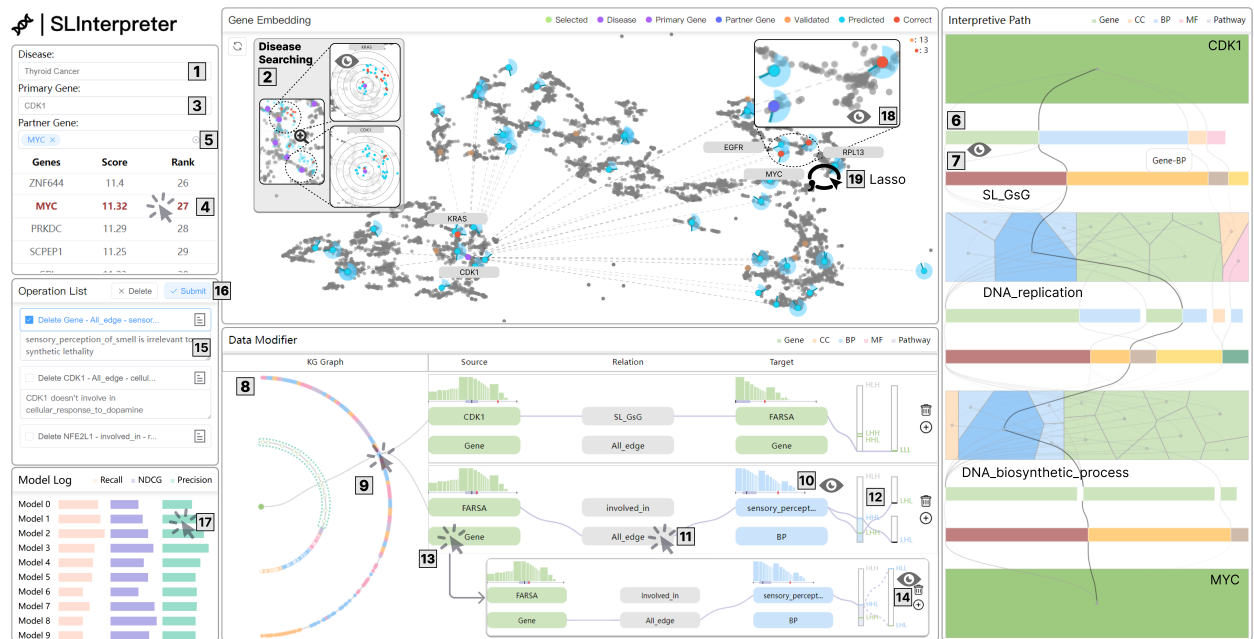


Fig. 1: In Case Study, the biologist’s analysis comprises: **1** Select the *Thyroid Cancer* in the *Disease search box*. **2** Observe the similar pattern between KRAS and CDK1 by zooming in on the *Radar* in the *Embedding View*. **3** Select CDK1 in the *Primary Gene search box*. **4** Select MYC, a correctly predicted partner gene of CDK1, in *Partner Gene Table*. **5** MYC tagged in *Partner Gene search box*. **6** The *Entity Flow Bar* indicates that gene-gene connections are no longer predominant. **7** The *Path Bar* displays a smaller ratio of SL\_GsG. **8** Narrow the investigation scope to CDK1 → Gene → BP by color encoding in the *Knowledge Graph*. **9** Select the irrational path in the *Knowledge Graph*. **10** Observe the histogram of *sensory\_perception\_of\_smell*, indicating that this entity often ends such metapaths. **11** Select L-H-L strategy by clicking on *All edge*. **12** Observe that the proportional relationship of L-H-L is not statistically significant within its parent set H-H-L. **13** Select H-H-L strategy by clicking on *Gene*. **14** Observe the proportional relationships of L-H-L and H-L-L, sub-strategies to H-H-L. **15** Add the formulated strategy to *Operation List* and enter memos. **16** Click  to retrain the model. **17** Select the retrained model to initiate an iterative exploration. **18** Identify the cluster of three similarly ranked genes, EGFR, MYC, and RPL13, using the *Rank Indicator*. **19** Lasso the three clustered genes.

**Abstract**—Synthetic Lethal (SL) relationships, though rare among the vast array of gene combinations, hold substantial promise for targeted cancer therapy. Despite advancements in AI model accuracy, there is still a significant need among domain experts for interpretive paths and mechanism explorations that align better with domain-specific knowledge, particularly due to the high costs of experimentation. To address this gap, we propose an iterative Human-AI collaborative framework with two key components: 1) *Human-Engaged Knowledge Graph Refinement based on Metapath Strategies*, which leverages insights from interpretive paths and domain expertise to refine the knowledge graph through metapath strategies with appropriate granularity. 2) *Cross-Granularity SL Interpretation Enhancement and Mechanism Analysis*, which aids experts in organizing and comparing predictions and interpretive paths across different granularities, uncovering new SL relationships, enhancing result interpretation, and elucidating potential mechanisms inferred by Graph Neural Network (GNN) models. These components cyclically optimize model predictions and mechanism explorations, enhancing expert involvement and intervention to build trust. Facilitated by *SLInterpreter*, this framework ensures that newly generated interpretive paths increasingly align with domain knowledge and adhere more closely to real-world biological principles through iterative Human-AI collaboration. We evaluate the framework’s efficacy through a case study and expert interviews.

**Index Terms**—Synthetic Lethality, Model Interpretability, Visual Analytics, Iterative Human-AI Collaboration.

## 1 INTRODUCTION

H. Jiang, S. Shi, S. Zhang, J. Zheng and Q. Li (corresponding author) are with School of Information Science and Technology, ShanghaiTech University, and Shanghai Engineering Research Center of Intelligent Vision and Imaging, China. E-mail: {hjianghr2023, shishh2023, zhangshh, zhengjie,

liquan}@shanghaitech.edu.cn.

Manuscript received 1 April 2024; accepted 15 July 2024. Date of Publication 21 October 2024; date of current version 18 July 2024. Digital Object Identifier: xx.xxxx/TVCG.201x.xxxxxx

Synthetic Lethal (SL) relationships among genes denote a specific type of gene interaction characterized by co-expression [21]. In this context, inhibiting one lethal gene alone does not affect cell survival, but simultaneous inhibition of both genes results in cell death [12], as depicted in Fig. 2. This concept is crucial for targeted cancer therapy, as targeting the SL partner gene of a known cancer-related gene allows for precise cancer cell elimination while sparing normal cells. Thus, it offers a viable strategy for inducing the demise of cancer cells when directly targeting certain genes is challenging.

Given the substantial medical relevance of SL pairs, screening them is crucial for clinical applications. Several experimental biology-based methods, such as RNA interference and genome editing with *CRISPR/Cas9* [41], have been introduced for comprehensive SL screening [59]. However, these traditional wet lab methods are resource and labor-intensive, time-consuming, and prone to off-target effects [17]. Additionally, validated SL pairs constitute less than 0.1% of all potential pairs due to the vast number of human gene combinations [4]. To overcome these challenges, computational methods [49], based on biological insights, have been suggested to speed up and enhance the precision of SL gene pair predictions. These methods span statistical approaches [9], network-based strategies [15], and classic machine learning (ML) [56], along with artificial intelligence (AI) techniques [36]. Notably, AI methods have demonstrated superior predictive accuracy. However, aside from achieving accurate predictions, biologists increasingly prioritize understanding the interpretation of SL relationship mechanisms implicit in these AI models. Nevertheless, grasping the biological mechanisms underlying the predicted results remains challenging [46]. Moreover, the current AI model tools generally lack sufficient interpretability, resulting in a skeptical stance among researchers toward the predictive outcomes in this field.




Fig. 2: SL gene pairs refer to a pair of genes where inhibiting one does not impact cancer cell survival, but inhibiting both leads to cell death.

Prior research [13, 50, 52, 64] has aimed to aid biologists in comprehending the behavior of prediction models, facilitating the exploration of mechanisms behind successful predictions. These approaches are often rooted in biological knowledge. For instance, *KG4SL* [52] enhanced SL prediction accuracy by directly incorporating various relevant biological factors into the Knowledge Graph (KG) and integrating them with Graph Neural Networks (GNN). Going a step further, *KR4SL* [64] developed an interpretable AI model for SL prediction. Expanding the scope, efforts have been made to support both domain experts and non-AI-experts in exploring and explaining model behaviors [13, 50].

Despite aiding target users in understanding parts of model behaviors to enhance interpretability, the methods mentioned above bring new challenges to Human-AI collaboration, which stem from two intertwined dimensions: the symbiotic interplay between *biologists' involvement* and the *application of AI*. On one hand, earlier research has mainly aimed at improving prediction accuracy by incorporating vast amounts of information into the model input [29]. However, in efforts to understand SL relationship mechanisms, the complex data used for model input significantly increases the complexity of processing and analysis. Given the intricacies of biological data, some gene or protein interactions might have subtle relevance to SL mechanisms. This subtlety potentially introduces new noise, influencing both prediction outcomes [67] and interpretive subgraphs [24]. Hence, a method *synergizing with domain expertise becomes crucial to effectively devise reasonable strategies for screening biological knowledge used to train AI models*. On the other hand, biologists seek to understand AI models and expect clear insights into predictions. This comprehension is vital for decision-making and selecting wet lab directions, allowing experts to trust or work with these models [16]. Within this collaborative process, biologists must carefully *consider the division of responsibilities between themselves and the AI models* while simultaneously evaluating their reliability. In the end, biologists bear

the accountability for prediction outcomes. Current approaches mainly provide AI-generated suggestions for biologists to make final decisions [64], with little focus on the *iterative synergy between AI and biologists*, such as effective means to rectify the misleading knowledge acquired by the model. Nevertheless, an iterative interaction [57] has the potential to significantly impact the quality and efficiency of the analytical process, which highlights the need for a more integrated and dynamic collaboration between biologists and AI systems.

To tackle these challenges, we initiated a substantial partnership with six biologists and SL researchers, immersing ourselves in their daily research routines. Through a formative study involving these professionals, we conducted insightful interviews to discern their needs and expectations. These discussions not only provided an understanding of their requirements, but also yielded six design requirements. These requirements are structured into two tiers: *1) Human-Engaged Knowledge Graph Refinement Based on Metapath<sup>1</sup> Strategies* and *2) Cross-Granularity SL Interpretation Enhancement and Mechanism Analysis*. We propose an iterative framework of human-AI collaboration aligned with these design requirements. Initially, domain experts explore new SL pairs using predictions and interpretive paths<sup>2</sup> generated by a model trained on the entire data. During this process, irrelevant or incorrect paths that may introduce noise to the predictions are eliminated from the KG using appropriate metapath strategies. Subsequently, the model retrains, allowing domain experts to iteratively scrutinize predictions and interpretive paths, refining the KG. This iterative process continuously optimizes predictions and mechanism exploration, enhancing expert participation and intervention, ultimately leading to increased trust. Furthermore, mechanism exploration can identify more persuasive SL pairs for validation through wet lab experiments. To support this cycle, *SLInterpreter* is developed to help domain experts organize and compare interpretive paths across granularities, uncovering potential SL mechanisms and more persuasive SL pairs. We conducted a case study involving our collaborative experts, followed by interviews to gather their feedback. The findings notably underscored the approach's superiority in delivering informative domain insights, facilitating decision-making, and presenting an intuitive interface. In summary, the contributions of this study can be outlined as follows:

- We collaborated closely with biologists and SL researchers, conducting insightful interviews to uncover six essential design requirements for addressing SL mechanisms.
- We introduced the iterative framework *SLInterpreter*, tailored to meet these design prerequisites, empowering experts to explore SL pairs and mechanisms, refine the KG, and foster trust.
- We validated the effectiveness of *SLInterpreter* through a comprehensive case study and interviews, demonstrating its superior performance and value.

## 2 RELATED WORK

### 2.1 Computational Methods for Synthetic Lethal Prediction

Constrained by the limitations posed by wet lab experiments in screening SL pairs, numerous computational approaches have been introduced to infer SL pairs [49, 58]. These approaches fall into three main categories: *statistical-based*, *network-based*, and *ML-based* methods.

*Statistical-based* [6, 7, 9, 53] and *network-based* methods [23, 33] model existing SL data under specific assumptions. The interpretability of these methods mainly stems from their ability to reveal statistical or topological patterns associated with validated SL pairs. For instance, *DAISY* [18] employed three statistical methods for different cancer sets to infer SL interactions from genomic data. Kranthi et al. [22] used a graph-based approach with centrality measures, assuming that highly connected proteins are crucial and their knockout leads to lethality. However, these methods may not cover all data aspects, resulting in high prediction accuracy only for the aspects they cover, while accuracy may be low for other data [49]. Additionally, this can introduce biases

<sup>1</sup>A sequence of node types that guide a walk from the origin to the destination node, commonly applied in biomedical network analysis.

<sup>2</sup>Interpretive paths are explanatory links and nodes connecting the primary gene with the predicted gene through biological relationships.

and challenges in discovering new SL mechanisms. *ML-based* methods are designed to identify intricate patterns in large datasets that defy straightforward analysis, with fewer and more relaxed assumptions about the data [14], thus improving prediction accuracy for unobserved data. For instance, Wu et al. [56] developed a k-NN model, suggesting that gene pairs similar to validated SL pairs have higher SL potential. Lundberg S.M. [35] enhanced tree-based models' interpretability by measuring local feature interactions and using these insights to understand the model's global structure. However, most of these classic ML methods still depend on the manual extraction of gene features based on past knowledge, which can introduce human biases [49]. While deep learning methods reduce the need for manual feature engineering, they often lack transparency and function as opaque black boxes.

Given the limitations of the aforementioned methods, the KG-based GNN model leverages the extensive validated facts within the KG to learn gene representations and predict SL relationships with minimal or no explicit feature engineering [48]. The KG-based GNN model has demonstrated state-of-the-art predictive accuracy and provides interpretive paths with their weights through attentive aggregation. Consequently, we built our backend based on the KR4SL framework [64].

## 2.2 Graph Neural Network Interpretability

Interpreting GNN models is crucial, prompting many efforts to enhance their interpretability [36, 43, 45, 60, 61]. These efforts are primarily categorized into three levels of granularity [63]: 1) *Instance-Level Explanation*: This aims to elucidate the prediction for each instance by identifying significant substructures or subgraphs. 2) *Model-Level Explanation*: This focuses on comprehending the global decision rules captured by the GNN. 3) *Group-Level Explanation*: This provides explanations for a group of predicted instances.

Most GNN interpretability methods focus on *Model-Level Explanation* [2, 36, 45], aimed at helping data or AI experts understand the structures and behaviors of models or assist in debugging. Nevertheless, these methods often overlook the challenges faced by non-AI domain experts. To address this, some studies have sought to create more intuitive explanations for non-AI experts using *Instance-Level Explanation* [43, 62]. For instance, *GNNExplainer* [60] identifies critical subgraphs to elucidate GNN model predictions.

While *Instance-Level Explanation* methods enhance domain experts' understanding of a model's decisions for specific instances, they fall short for those concerned with potential mechanisms indicated by the model. To address this limitation, *DrugExplorer* [50] introduces *Group-Level Explanation*, which allows users to categorize drugs based on similar action mechanisms. Inspired by their approach, our study assists domain users in exploring potential mechanisms at the *Group-Level* by analyzing relationships between instances within interpretive paths throughout iterative interactions with AI. Notably, our target users are biologists researching SL. Unlike *DrugExplorer*, which focuses on drug-to-disease relationships, we emphasize gene-to-gene relationships. Furthermore, we enable users to iteratively refine AI predictions and discover new SL mechanisms through cross-granularity exploration.

Additionally, in GNN research, *CorGIE* [34] explores graph topology, node features, and embeddings interplay to decipher GNN operations. Building on this, *GNNLens* [20] provides a visual tool assisting ML experts in identifying potential error patterns, offering insights for optimizing model structures. However, these tools are unsuitable for non-AI domain experts to explore predictions and underlying mechanisms. Our work differs by focusing on assisting **domain experts** in a deeper exploration of the underlying mechanisms suggested by GNN models. Through iterative and cross-granularity visualization, we enable domain experts to rectify knowledge erroneously acquired by GNN models, facilitating a more profound understanding of predictions.

## 2.3 Human-AI Collaboration

*Human-AI collaboration* refers to the interactive cooperation between humans and AI to jointly address specific tasks, with responsibilities divided based on the strengths of each party [38, 44]. As AI performance continuously improves, the ways humans and AI collaborate are constantly evolving, with three main patterns proposed [44]. The first

pattern is *AI-assisted decision-making*, where AI provides necessary information and support for humans to make final decisions. Many studies have seamlessly integrated AI assistance into the workflow of domain experts [40, 54, 66]. The second pattern, *Human-in-the-loop*, involves incorporating expert intervention into the process of training AI models. This approach aims to train accurate models at minimal cost by integrating human knowledge and experience [8, 37, 38, 57, 68]. The third pattern, *joint action*, involves humans and AI working together towards a common objective, functioning as a single entity [1, 25–27, 38, 70].

Building upon existing research, our approach seamlessly integrates the expertise of biologists with the supportive role of AI. Biologists first explore new potential SL pairs through the predictions and interpretive paths generated by the model trained on entire data. Subsequently, we facilitate biologists' active involvement in a *human-in-the-loop* process, assisting them in refining the knowledge acquired by the model through appropriate metapath strategies, which results in interpretive paths that progressively reflect biologists' expertise and biological principles.

## 3 FORMATIVE STUDY

The formative study aimed to thoroughly understand and evaluate the challenges faced by SL researchers, their preferences and expectations regarding collaboration with AI, and their system design requirements. To achieve this, we conducted semi-structured interviews with four researchers (**E1-E4**) specializing in SL mechanisms and two clinical physicians (**E5-E6**) from a local university and a hospital (Mean Age = 47.5, SD = 46.6, four males and two females). Within this group, **E1-E2** specialize in screening SL pairs using the *CRISPR/Cas9* technique, focusing on thyroid cancer and glioma, respectively. **E3-E4** are engaged in anticancer drug research, collaborating with **E5-E6**. Each participant has extensive research experience in SL and possesses knowledge or experience in AI-assisted prediction of SL pairs. Through thematic analysis [5], we extracted valuable insights regarding the challenges faced by domain experts and summarized the design requirements for our approach. Each interview session lasted approximately 45 minutes.

### 3.1 Experts' Conventional Practices

Our experts emphasized their reliance on wet-lab experiments or statistical methods in the past for discovering and validating new SL pairs. Specifically, wet lab experiments require skilled researchers to culture relevant cells. Moreover, specific gene inhibitors are utilized, necessitating manual observation and recording of experimental data. For instance, in tumor experiments, it is crucial to observe and record proliferation-related experimental data. Some tumors require the assessment of their migration properties through methods like cell scratching, followed by cell washing or drug addition observation. If the experiment progresses smoothly, it typically involves about five researchers and takes approximately half a year. Considering potential issues such as unexpected cell death or microbial contamination during cell culture, the duration may extend further. Therefore, given the considerations of experimental costs, manpower, and equipment needs, most labs are limited to conducting one to two wet lab experiments at a time.

Given the high cost and limitations of wet lab experiments, our experts emphasized their need for a method that can predict SL pairs at a lower cost and with higher accuracy, while also providing reasonable explanations. To this end, they have attempted to use AI-assisted methods for experimental pair screening. However, in their attempts, they found that current AI methods lack sufficient interpretability and tools for mechanism exploration. Moreover, since the models can only provide one-way prediction results, applying their expertise to intervene in the AI prediction process proved to be difficult. The sole form of interaction experts have with the model is to accept or reject its predictions, which leads to a lack of trust in these models' predictions.

### 3.2 Experts' Concerns and Barriers

Domain experts initially voiced concerns over biological data accuracy. Several interviewees highlighted the importance of reliable and pertinent data in training AI models. They underscored that "...this basically decides how accurate the AI predictions are and how sensible the explanations it gives make sense." For instance, **E1** stated, "Picking the right data to train the model is really important, and where that data comes

from is definitely something to seriously think about.” E2 mentioned, “A worry we all share is that if the model is learning, the info we feed it better is true... When we’re looking into things, if we spot any issues with the data or conclusions, I hope there’s a way to quickly filter them out.” This reveals the first challenge: **C1. Reliability of biological data in AI training.** As experts analyze the model’s interpretive paths, having effective methods and interactive tools for screening data and addressing potential errors or misleading conclusions is vital.

In terms of collaboration with AI, several interviewees expressed concerns about the current pattern. The prevalent human-AI collaboration approach in SL involves domain experts passively accepting prediction results. Subsequently, experts simply choose a few gene pairs for initial experimental validation, typically those not yet investigated or related to their current research direction. This collaboration mode entirely overlooks the active involvement of domain experts in the prediction task, leading to researchers distrusting the conclusions. As mentioned by E4, “I simply get the results the model spits out, and when I see some obvious mistakes, there’s no way for me to fix [them]. It just keeps eroding my trust in AI.” This highlights the second challenge: **C2. Lack of intervention in AI prediction process.**

Interviewees also expressed concerns in decision-making, potential SL mechanism exploration, and continuous collaboration with AI. Specifically, determining whether to accept and validate AI’s prediction results through experiments poses a significant challenge. As highlighted by E3, “We’re talking about over 20,000 human genes here, and doing experiments for all those gene pairs? That’s just not economically achievable.” Similarly, E2 pointed out, “...So, AI may predict a bunch of gene pairs, but it takes us at least half a year to check out a few pairs at the same time. So, we really hope the predictions are highly accurate and AI can give us a good explanation of how they came up [with the predictions].” This reveals the third challenge: **C3. High trial-and-error costs.** Wet lab experiments are time-consuming and resource-intensive. Researchers cannot experimentally validate all predictions, and careful consideration is needed in selecting the next research targets. For AI predictions, there is a need for more abundant and reasonable evidence to demonstrate their experimental value.

Furthermore, even if the model provides a plausible explanation for individual gene pairs, domain experts express greater concern about extracting underlying mechanisms and identifying commonalities from multiple prediction results. For example, E2 pointed out, “The prediction results contain too many possible interpretive paths. Although these paths are sorted by scores, we still need to focus on at least 10 or so paths, and once too many entities are set up, it puts a very heavy burden on us to explore possible mechanisms”, highlighting the challenge: **C4. Limitations in SL mechanism exploring.** While domain experts find the commonalities among these explanations intriguing, current methods remain cumbersome and insufficient in assisting experts to effectively compare and analyze these commonalities.

While some SL experts acknowledge the accuracy and efficiency of AI predictions, researchers still harbor widespread concerns regarding sustaining ongoing collaboration with AI. E5 noted, “We’ve got some AI-related projects going on, but they’re mostly just starting out... Honestly, AI models are still kind of new to us compared to the usual methods, so we stick to what we know best - traditional research methods using stats or genomics.” Similarly, E6 commented, “Yeah, I totally agree that we need [tools] to fix AI when it learns wrong stuff. In the past, we didn’t have such support, and we’re not used to working with interactive models like that... So, if those tools come along, we’ve gotta make sure they actually work by checking how they affect the model’s performance,” highlighting the final challenge: **C5. Lack of experience in iterative collaboration with AI.**

### 3.3 Experts’ Needs and Expectations

After conducting interviews with experts, we’ve compiled a set of requirements aimed at effectively addressing the discussed challenges. Our approach is tailored to seamlessly integrate with conventional practices, allowing experts to grasp prediction model behaviors and delve into the mechanisms behind successful SL predictions. To delve deeper, we systematically tackle these requirements in two

parts: **[KG Refinement]**Human-Engaged Knowledge Graph Refinement Based on Metapath Strategies and **[SL Interpretation]**Cross-Granularity SL Interpretation Enhancement and Mechanism Analysis. These parts operate in an **[Iterative Cycle]**, continuously refining model predictions and exploring mechanisms, thereby enhancing expert involvement and intervention, and bolstering trust.

**[KG Refinement]**Initially, researchers underscored the urgency of conducting a rapid and comprehensive exploration of training data and interpretive paths, which is essential for gaining understanding and familiarity with both the model’s input data and its performance. To tackle challenges concerning the reliability of biological data **[C1]**, providing a concise summary of training data and interpretive paths is essential. This forms the basis for our first design requirement: **DR1. Offer a clear summary of training data, prediction results, and interpretive paths.** Furthermore, addressing the urgent need expressed by several researchers to quickly identify erroneous or noisy data in the training set through the error-prone interpretive paths indicated by the model, our second design requirement emerges: **DR2. Detect relevant training data corresponding to erroneous or noisy paths.**

**[KG Refinement]**Addressing the second challenge **[C2]**, which concerns researchers’ skepticism towards AI predictions due to inadequate intervention, we introduce **DR3. Provide appropriate interventions to filter data corresponding to interpretive paths.** This requirement was consistently emphasized during semi-structured interviews, highlighting researchers’ worries about the negative impact of insufficient intervention on AI predictions. To alleviate these concerns, our approach offers a comprehensive visual analysis of potentially erroneous or noisy data, along with interactive features allowing users to correct any identified issues promptly.

**[SL Interpretation]**To address the third challenge **[C3]**, it is proposed that the system should not only enhance the accuracy of predicting gene pairs through proper data input but also offer a clear and sufficient representation of interpretive paths, considering the high cost of trial and error. This entails the ability to vividly present information about both the genes and selected entities, along with the attributes of edge relationships. The emphasis lies on minimizing visual clutter and ensuring the smooth presentation of both topological and textual information. During semi-structured interviews, E1 stressed that “only when the interpretive paths are presented clearly and intuitively, and align well with our domain expertise, will the predicted results be convincingly justified for wet lab experiments validation.” This underscores the importance of our fourth design requirement: **DR4. Provide detailed displays of paths and entities for interpretive paths.**

**[SL Interpretation]**Regarding the fourth challenge, identified by researchers in identifying common potential mechanisms from AI predictions **[C4]**, the system must clearly present and analyze similarities across various interpretive paths, which is crucial as it can significantly aid researchers in comprehending existing mechanisms and uncovering new ones. Acknowledging the complexity researchers face in exploring common potential mechanisms from numerous paths, the design must balance cognitive load with information presentation. According to E1, “Exploring potential mechanisms in SL is highly complex, and current AI struggles to directly accomplish theoretical derivations.” Additionally, as noted by E2, “However, identifying commonalities in AI predictions can suggest fruitful directions, provided these commonalities are easily extracted and summarized. Comparing them individually in common AI-assisted systems with limited attention is impractical and can narrow our focus.” This leads us to our fifth design requirement: **DR5. Comparative analysis and mechanism exploration of different prediction results and interpretive paths.**

**[Iterative Cycle]**In response to **[C5]**, we introduce the last design requirement: **DR6. Facilitating iterative Human-AI interaction with feedback and recordings.** This requirement is crafted to support domain experts in overcoming unfamiliarity and uncertainty during their collaboration with AI systems. By ensuring a seamless exchange of feedback and recording interactions, we empower domain experts to grasp the repercussions of their interventions on both datasets and model performance through ongoing engagement. The essence of this requirement lies in capturing users’ actions comprehensively, thus

enabling them to retrace their steps when faced with unsatisfactory intervention outcomes or when model metrics fall short of expectations. This iterative process fosters domain experts’ familiarity with AI collaboration, allowing them to refine model performance progressively. Echoing the sentiments of E5, “During the early stages of AI interaction, keeping up with iterative exploration is key. You know, being able to jot down all those steps you’re taking? It’s like having a handy tool that lets you go back and forth, compare stuff, and double-check what you’ve done. This way, you can make sure your interventions are on point and your analysis is as sharp as it can be.”

## 4 SLINTERPRETER

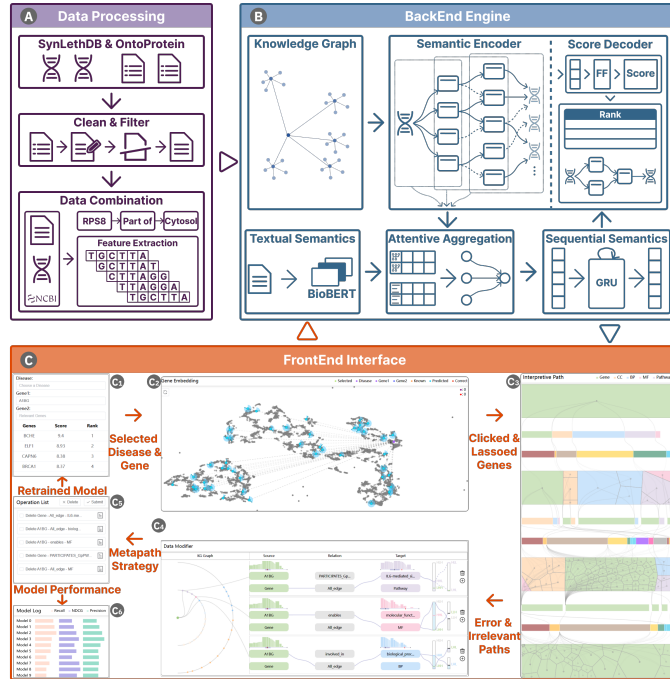


Fig. 3: *SLInterpreter* comprises **A** data processing module, **B** backend engine, and **C** frontend interface, facilitating an iterative workflow.

Our analysis approach, implemented through *SLInterpreter*, consists of three main components: **A** data processing module, **B** backend engine, and **C** frontend interface (Fig. 3). The data processing module prepares data for the KG and extracts gene features. The backend engine trains the GNN model with the KG and gene features, predicting SL pairs and interpretive paths. The frontend then provides multiple interconnected views for iterative and cross-granularity SL analysis.

### 4.1 Data Description and Processing

We provide detailed information regarding the data utilized in the backend engine, including both *Gene Feature* and *Knowledge Graph*.

**Gene Feature.** The gene feature data includes both gene sequence  $\mathcal{S}$  and gene description  $\mathcal{D}$  sourced from the National Center for Biotechnology Information (NCBI) database [10], recognized as an authoritative gene database in the biological domain, ensuring data reliability and professionalism. Each gene sequence consists of approximately 19,000 base pairs (ACGT). However, directly converting these sequences into numerical form and performing operations like dimensionality reduction would entail significant computational costs due to the extensive number of genes and the relatively lengthy gene sequences for each. To address this challenge, we utilize the *k-mer* method [11], a common method of DNA feature extraction in the biological field, to identify local patterns in gene sequences  $\mathcal{S}$ . Meanwhile, we utilize *BioBERT* [28] to process functional and descriptive gene information  $\mathcal{D}$ .

**Knowledge Graph.** The biological data used to build the Knowledge Graph (KG) is sourced from two main repositories (Tab. 1): the SL relationship graph *SynLethDB* [52] and the biomedical knowledge dataset *ProteinKG25* [65]. These datasets include validated gene SL

interactions and external knowledge about gene functions, such as pathways and biological processes. The resulting KG comprises five types of entities: *Gene*, *Pathway (PW)*, and three types of *Gene Ontology (GO)*: *Biological Process (BP)*, *Molecular Function (MF)*, and *Cellular Component (CC)*. Additionally, it incorporates multiple relationship types, primarily indicating gene associations with specific paths or annotations by GO. Entities are connected by these relationships to form triplets, which are the fundamental components of KG structures, expressing how two entities are linked (e.g., *RPS8—Part of—Cytosol*). Collectively, these triplets constitute the final KG.

Table 1: Data Description.

	#Entities	#Relations	#Triples
SL Graph	9746	1	35374
ProteinKG25	42547	32	361245
<b>Knowledge Graph</b>	42547	33	396619

### 4.2 Backend Engine

Predicting SL gene pairs can be seen as akin to link prediction. Following the methodology described in *KR4SL* [64], we employ a GNN-based model. This model predicts potential SL relationships among gene pairs lacking direct connections by utilizing established relationships to generate prediction paths and allocate weights to these paths, reflecting the strength of the predicted SL relationship.

**Training Data Construction.** According to *KR4SL* [64], we begin by constructing a directed heterogeneous KG by integrating the well-established SL graph *SynLethDB* [52] with the comprehensive biological dataset *ProteinKG25* [65]. Notably, merging these datasets does not increase the number of entities since all human genes in *SynLethDB* are already included in *ProteinKG25*. This augmented graph serves as the input for the GNN model, enabling it to capture a rich set of relational patterns and interactions among genes. For transparency, we have cited the public databases used in this study within our open-source project<sup>3</sup>. Additionally, in the context of *SLInterpreter*, paths that are irrelevant or less related to SL are iteratively pruned from the KG, refining the input for the model.

**Predicting SL Pairs.** Following the encoder-decoder architecture in *KR4SL*, we utilize the heterogeneous KG structure to identify potential SL pairs by tracing relational paths. This approach integrates structural semantics from the graph with textual semantics extracted from gene descriptions. Additionally, it employs attentive aggregation among triples with the same target, enhancing semantic flow through *GRU modules* at each layer and accounting for the paths’ weights. In the final encoder layer, candidate SL partners are pinpointed and evaluated by a *score decoder*. The decoder ranks candidates by their potential to form an SL relationship with the target gene, selecting the top-50 as predicted SL partners, a commonly used metric in such link prediction tasks [64]. Simultaneously, the model outputs the 3-hop path containing the predicted link as an interpretive path along with its score.

The triplets in the dataset are divided into 80% training, 10% validation, and 10% test sets. Performance evaluation focused on precision for the top-50 candidates in the test set. The experiments, detailed in Table 1 in the Appendix, highlighted the model’s efficacy, achieving a precision rate of 53% for the top-50 candidates. In comparison, using the same data, the average precision of other similar models [51, 52, 69] was 16.29%, with the second-best model, *NSF4SL* [51], achieving 34.7%, which demonstrates the effectiveness of this model in detecting key gene interactions, providing valuable insights for cancer research.

### 4.3 Frontend Interface

In collaboration with domain experts, we developed a frontend interface for exploring biological and genetic data for SL analysis, featuring: 1) *Session View ([KG Refinement])*: Includes disease and gene search boxes, a log of modifier operations, and model performance overview. 2) *Embedding View ([SL Analysis])*: Provides an overview of prediction results on feature reduction. 3) *Interpretation View ([SL Analysis])*: Displays cross-granularity interpretation

<sup>3</sup><https://github.com/jianghr-shanghaitech/SLInterpreter-Demo>

paths, and 4) *Modifier View* ([**KG Refinement**]): Shows KG representations for gene and modifier selection in metapath strategies. After setting the metapath strategies, users can start a new exploration cycle based on a retrained model ([**Iterative Cycle**]).

#### 4.3.1 Session View

The *Session View* (Fig. 3-C<sub>1</sub> & C<sub>5</sub> & C<sub>6</sub>) offers a platform to interact with the model’s prediction outcomes and to review previous logs and model performance. This view comprises three subviews: the *Panel*, the *Operation List*, and the *Model Log*.

**Panel.** Within the *Panel* (Fig. 3-C<sub>1</sub>), three search boxes are designed for selecting prediction results at varying levels. The *Disease search box* (Fig. 1-1) showcases genes linked to a specific disease, while the *Primary Gene search box* (Fig. 1-3) enables direct identification of genes of interest through auto-complete. Upon selecting a *Primary Gene*, *Partner Gene Table* (Fig. 1-4) below the search boxes displays the top-50 *Partner Genes* from the prediction results in ascending order, detailing *Name*, *Score*, and *Rank*. Correctly predicted *Partner Genes* are highlighted in brown font. We verify the accuracy of predicted *Partner Genes* using ground truth data from *SynLethDB*, sourced from previous wet lab experiments or relevant literature [52]. Moreover, users can engage with the table by selecting gene rows, which generates tags in the *Partner Gene search box* (Fig. 1-5). Concurrently, the *Embedding View* and *Interpretation View* update all prediction results related to the selected *Partner Gene*. **Operation List.** The *Operation List* (Fig. 3-C<sub>5</sub>) records user interactions with the KG in chronological order, featuring collapsible text boxes beneath each action for users to elaborate on their operations and corresponding notes (**DR6**). Users can select metapaths using checkboxes. Undesired metapaths can be removed by clicking  Delete. Once the selection of metapaths is finalized and  is clicked, The backend automatically deletes the corresponding paths, retrains the model, and displays its performance in the *Model Log*. **Model Log.** The *Model Log* (Fig. 3-C<sub>6</sub>) displays grouped bars detailing model performance indicators, including *Recall*, *NDCG*, and *Precision*. Indicators are color-coded with bar length representing values. The vertical arrangement allows intuitive performance comparison across models. Users can access models by clicking the corresponding bars.

#### 4.3.2 Embedding View

The *Embedding View* (Fig. 3-C<sub>2</sub>) offers users an overview of clustering based on gene features, along with model prediction results and validated SL partners (**DR1**), enabling exploration of SL pairing patterns within comparable gene groups (**DR5**). To fully leverage gene attributes, including sequences  $\mathbb{S}$  and descriptions  $\mathbb{D}$ , we use concatenation fusion methods to merge the processed data. We then apply *UMAP* [39] to generate the embedding result due to its efficiency in compressing high-dimensional data into a lower-dimensional representation. Our experiments (Appendix Fig. 1&2&3) showed that *UMAP* outperformed *PCA* [55] and *t-SNE* [47] in preserving global data structure and execution speed. *PCA* often overlooks crucial biological variations [19], while *t-SNE* primarily reveals local structures without effectively showing similarities between clusters [31]. Both *PCA* and *t-SNE* also perform slowly with large datasets. Although *UMAP* is non-linear, meaning distances in the *Embedding View* are not linear, it effectively preserves local distances over global distances, enhancing clustering within local regions, which makes it suitable for our focus on the correspondence between local clusters. Consequently, the dashed lines connecting nodes illustrate pairing relationships between genes rather than distance comparisons, helping users identify patterns where genes from one cluster frequently link to genes from another cluster.

**Visual Design.** The *Embedding View* enhances disease exploration through intuitive visual cues. Upon selecting a disease from the *Disease search box*, the *Embedding View* highlights associated genes with distinct  $\bullet$  points. The surrounding *Radars* directly project the relative locations of *Partner Genes* onto concentric circles centered on the *Primary Gene*. The top-50 predicted *Partner Genes* are highlighted in  $\bullet$  and correctly predicted *Partner Genes* are shown in  $\bullet$  (Fig. 4-A). Users can employ semantic zooming to manage overlapping points or observe radar patterns more clearly. Hovering over a highlighted point

reveals an enlarged *Radar* (Fig. 4-B), facilitating the identification of tightly clustered *Partner Genes* with common pairing patterns.




Fig. 4: Glyph Designs for *Embedding View*. (A) is the overview of *Radar* in disease searching. (B) is the enlarged *Radar*. (C) is the *Rank Indicator*.

Once a *Primary Gene* is identified, users can select it in the *Primary Gene search box* for further investigation. Upon selecting, the corresponding *Primary Gene* remains highlighted, while other genes revert to the default color. Simultaneously, the *Embedding View* highlights top-50 predicted *Partner Genes*  $\bullet$  with dashed lines connecting to the *Primary Gene*, validated *Partner Genes* of the *Primary Gene*  $\bullet$ , and correctly predicted *Partner Genes*  $\bullet$ . Each top-50 predicted *Partner Gene* is encircled by *Rank Indicator* (Fig. 4-C), an arc indicating its rank of the *Partner Gene*. The arc starts at the top and lengthens clockwise as the rank increases, with a bar representing the end of the arc, helping to discover clusters of similar ranked genes at a glance. When tags are created in the *Partner search box*, corresponding points are highlighted  $\bullet$ , and the dashed line turns solid.

**Interaction.** When no *Disease* or *Primary Gene* is selected, users can lasso all genes to explore patterns among genes with similar features, with the selected genes highlighted  $\bullet$ . Once a *Disease* is chosen, lassoing among disease-related genes helps explore patterns within genes linked to the same disease. Similarly, when a *Primary Gene* is selected, lassoing among the top-50 *Partner Genes* enables the exploration among *Partner Genes* belonging to the same *Primary Gene*.

#### 4.3.3 Interpretation View

The *Interpretation View* (Fig. 3-C<sub>3</sub>) offers users an across-granularity analysis of paths: spanning from comprehensive cross-gene analyses (**DR5**) to intricate single-path analyses (**DR4**). To enhance the clarity of metapath content and flow, we break them down based on path hierarchy into three primary elements: 1) the name, type, and proportion of the current hierarchical node; 2) the flow from the current hierarchical node type to the category of the subsequent hierarchical node, and 3) the names of the paths linking to the next hierarchical node.

**Visual Design.** To effectively illustrate the name, type, and proportion of entities on each layer while mitigating excessive overlapping of connecting lines, we employ Voronoi Treemaps [3]. This method uses the weight of different entities within interpretive paths as their value, dividing space into irregular polygons, optimizing display space utilization, and offering clear differentiation between various entities. Clicking on an entity’s polygon will highlight other entities connected to it. Consequently, entities belonging to the same category, such as genes or biological processes (BP), are assigned the same color and grouped together within the same area. There are a total of five categories of entities: *Gene*, *Cellular Component* (CC), *Biological Process* (BP), *Molecular Function* (MF), and *Pathway*, as well as 33 types of path relationships, which are represented using consistent colors.

The *Entity Flow Bar* (Fig. 5-7) and the *Path Bar* (Fig. 5-4) are both presented as horizontally arranged stacked bars. The *Entity Flow Bar* is positioned beneath the *Entity Treemap* (Fig. 5-6) and corresponds to the categories of entities depicted in the *Entity Treemap* above. It illustrates the proportion of the flow from a specific category of entity to the succeeding layer’s entity category. The *Entity Flow Bars* for distinct entity categories operate independently, yet entities of the same category are color-coded identically to streamline user comprehension. Flow lines traverse through the designated areas defined by the *Entity Flow Bar*, effectively distinguishing the connections. On the other hand, the *Path Bar* is a comprehensive stacked bar including every type of relation observed, along with its relative proportion. By analyzing these proportions, users can discern the frequency of each relation type. Moreover, connections stemming from the *Entity Flow Bar* are linked to the *Path Bars* representing different edges based on their

corresponding relation type, which illustrates the connection between the source entities and the relations.

**Design Alternatives.** In the iterative design process, we explored two alternatives. The first alternative (Fig. 5-**A**) includes the *Sankey Flow* (Fig. 5-**1**) illustrating entity flow, and the *Sankey Bar* (Fig. 5-**2**) showing entity proportion within the interpretation path. In the Packchart diagram, *Entity Nodes* (Fig. 5-**3**) display entity type and proportion, while the *Path bar* (Fig. 5-**4**) shows path type and proportion with connecting lines. This design maintains consistent colors for entity type and accurately represents entity percentage and weight, but confuses hierarchical correspondence between the Sankey and Packchart diagrams, requiring users to toggle between the two. Moreover, the use of nodes in the Packchart is spatially inefficient. In the second alternative (Fig. 5-**B**), we integrate both diagrams by merging *Entity Nodes* and *Sankey Bar* into the *Entity Matrix* (Fig. 5-**5**), where each matrix’s size indicates the entity’s weight. Retaining *Sankey Flow* and *Path Bar*, this design results in considerable overlap in entities’ vertical coordinates, leading to disorganized connections and compromised visibility. Furthermore, the presentation of the hierarchy remains confusing. In the final design (Fig. 5-**C**), *Entity Treemap* (Fig. 5-**6**) eliminates coordinate overlap with irregular shapes, and *Entity Flow Bar* (Fig. 5-**7**) represents the proportion of different entity types, replacing *Sankey Flow*. This design displays hierarchical paths without overlap, ensuring clear and uninterrupted layer connections within an integrated view.




Fig. 5: Design alternatives for *Interpretation View*. **A** is an alternative based on Sankey and Packchart. **B** is an alternative based on Sankey and node matrix. **C** is the final design based on Voronoi Treemap.

#### 4.3.4 Modifier View

The *Modifier View* (Fig. 3-**C4**) incorporates both the *Knowledge Graph* (Fig. 1-**8**) and *Metapath Modifier* (Fig. 6) functionalities to aid users in exploring the KG (**DR1**), pinpointing pertinent training data (**DR2**), and devising human-engaged strategic refinement of metapaths to filter out irrelevant or noisy data within the dataset (**DR3**).

**Visual Design of the KG.** Upon selecting a *Primary Gene* in the *Panel*, the KG displays entities originating from this gene. Inspired by the design of *EgoComp* [32], we used an ego-network layout to show the KG structure. Entities of the same layer are placed on circular arcs with equal radius, preferentially grouping similar entities. The colors of entities in the KG correspond to those in the *Interpretation View*, ensuring design consistency. To manage the exponential increase in entities, we display up to 2-hop connections to balance path length and entity count. When users hover over different entities, the paths leading to that entity and all related entities along the path are highlighted, with their names displayed. This interactive feature allows users to leverage their expertise to evaluate if entities across various paths are biologically related to SL or have plausible links to the *Primary Gene*. Automating selection based on interactions with the *Interpretation View* was also considered, but expert feedback showed it could lead users to select many irrelevant paths or paths not present in the KG. Therefore, manual selection was chosen for simplicity and to minimize errors.

**Concept of the Granularity Level Flow.** To assist users in formulating detailed metapath strategies and deriving insights, we introduced the *Granularity Level Flow* (Fig. 6-**D**). This flow illustrates the transition of metapath granularity from **high (H)** to **low (L)**, moving left to right, and displays the hierarchy of metapath strategies across granularity levels. In this flow, **H-H-H** represents the metapath with the highest granularity, **1-Low Granularity** represents the metapath with one low granularity entity, and **2-Low Granularity** represents the metapath with two low granularity entity. **L-L-L** denotes specific triplets, with a count

of 1. Users begin formulating their metapath strategies by selecting a specific entity from the KG, with the default granularity being **L-L-L**. By converting a low-granularity entity to high-granularity, more paths are included in the formulated metapath strategies, resulting in the removal of more paths from the KG due to the expanded selection range. Therefore, user selections progress from right to left, gradually increasing in granularity. This linear progression helps users determine the appropriate level of granularity, minimizing the risk of accidental or erroneous deletions caused by sudden changes in granularity.

**Visual Design of the Metapath Modifier.** When a specific entity is selected, the *Metapath Modifier* generates a horizontal metapath display box on the right, linked with the entity. This display box, shaped by usage habits from expert interviews, is intended to show information in a table format. Initially, the display box showcases the selected path (Fig. 6-**B1**), with each entity’s higher granularity elements (Fig. 6-**B2**) displayed beneath it. Higher granularity elements denote wider biological categories, each potentially encompassing several entities of lower granularity. For example, beneath the entity *A1BG*, the metapath element *Gene* is displayed. Users can modify metapath strategies by clicking corresponding entity granularities (Fig. 6-**B1** & **B2**). A frequency histogram (Fig. 6-**A**) is presented above the source (Fig. 6-**1**) and target entities (Fig. 6-**2**), indicating the distribution of path counts that begin or end with each type of entity. Given the wide range and long-tail distribution of entity frequencies, we applied a log transformation to the x-axis to enhance the comparison of the data, resulting in varying bar widths. Below the histogram, a horizontal box plot (Fig. 6-**3**) shows the central tendency and dispersion of frequency data. The position of the current entity is marked red, helping users grasp its relative placement within the dataset. To provide detailed information, hovering over a bar in the histogram reveals its interval and corresponding value. To further illustrate the relative proportions of metapath strategies with varying granularity, we have designed two vertical bars (i.e., *Primary Bar* and *Secondary Bar*) (Fig. 6-**4** & **5**) situated in the *Path-Level Summary* (Fig. 6-**C**). These bars are designed in a manner akin to the *Granularity Level Flow* (Fig. 6-**D**).

**Primary Bar.** The left bar (Fig. 6-**4**), designated as the *Primary Bar*, of which the height corresponds to the number of metapaths with the highest granularity, such as *Gene—All\_edge—Pathway* ( **H-H-H** ). Since each entity offers two granularities to choose from—low (Fig. 6-**B1**) or high (Fig. 6-**B2**), the *Primary Bar* reflects this distinction. Within the *Primary Bar*, three sub-segments (Fig. 6-**6**) depict the relative proportions of metapath strategies with exactly one low granularity entity (i.e., **L-H-H**, **H-L-H** and **H-H-L**) selected within the metapath strategies of highest granularity (i.e., **H-H-H**), with colors indicating the type of low granularity entity.

**Secondary Bar.** The height of the right bar (Fig. 6-**5**), termed the *Secondary Bar*, corresponds to the number of metapaths involving one low granularity selected entity (i.e., **L-H-H**, **H-L-H** or **H-H-L**). Its sub-segment heights depict the relative proportions of metapath strategies involving two selected low granularity entities within the metapaths of one selected low granularity entity. There are two cases regarding the display of the *Secondary Bar*: 1) If the selected metapath strategy includes exactly one entity of low granularity, the *Secondary Bar* is displayed as a complete bar. Its height represents the number of paths included in the selected metapath strategy, while the sub-segments’ height represents the proportion of the two sub-strategies of the selected metapath strategy within it. 2) If the selected metapath strategy includes exactly two entities of low granularity, the *2-Low Granularity* metapath strategies have two parents strategies, such as **L-L-H** among **H-L-H** and **L-H-H**. In this case, the *Secondary Bar* is divided into the top (Fig. 6-**9**) and bottom (Fig. 6-**10**) bars to represent the two proportions, (e.g., **L-L-H** / **H-L-H** and **L-L-H** / **L-H-H**).

**Connecting Rules of Primary and Secondary Bars for Granularity Selection.** Based on the number of low granularity entities selected in the strategy, the *Primary Bar* and the *Secondary Bar* utilize two connecting rules to represent corresponding information: **1) In cases where exactly one entity of lower granularity is selected**, such as *Gene—All\_edge—IL6-mediated\_si...* ( **H-H-L** ) (Fig. 6-**C1**), the target




Fig. 6: Metapath Modifier design for Modifier View. (A) Design for the Entity-level Summary. (B) Design for the Metapath Strategy. (C) Design for the Path-level Summary. (D) Granularity Level Flow for the Path-level Summary, H represents High Granularity and L represents Low Granularity.

entity (Fig. 6-2) links to the corresponding sub-segment (Fig. 6-6) on the *Primary Bar* (Fig. 6-4). This linkage indicates the strategy’s proportion within the highest granularity strategy ( H-H-H ). Two lines (Fig. 6-7) extend from the *Primary Bar* to the *Secondary Bar*, delineating the relative proportions of two lower granularity strategies based on the selected entity within the current metapath strategy ( H-H-L ). Since these two lower granularity strategies are not truly selected, they are linked in dash lines. In this example, the L represents the shared low granularity entity, making the 2-Low Granularity strategies on the *Secondary Bar* subsets of the current strategy. 2) **In cases where two entities of lower granularity are selected**, such as *A1BG—enables—MF* ( L-L-H ) depicted in Fig. 6-C, the strategies at this granularity are subsets of two higher-level strategies ( H-L-H and L-H-H ), each involving 1-Low Granularity entity. Consequently, the proportions of these two parent strategies within the highest granularity are simultaneously displayed in the *Primary Bar* and connected to the target entity (Fig. 6-8). In this scenario, the *Secondary Bar* (Fig. 6-9 & 10) shows the proportion of the current metapath selection strategy ( L-L-H ) within two higher granularity strategies ( H-L-H & L-H-H ).

**Design Takeaway.** Through observation and interaction with the metapath display box, users can thoroughly assess the proportions and metapaths with specific granularity and derive insights from them. For instance, *A1BG—PARTICIPATES\_GpPW—IL6-MEDIATED\_...* (Fig. 6-B), where its H-H-L (Fig. 6-C) subset occupies a smaller proportion within H-H-H, and the H-L-L subset occupies a larger proportion within H-H-L, user can discern that despite the fewer H-H-L paths, the majority of paths between *Gene* (H) and *IL6-MEDIATED\_...* (L) are connected by relation *PARTICIPATES\_GpPW* (L). However, some paths are connected by other relations, like *involved\_in* (L). Gaining this insight, if user determines that *Gene* should not be connected to *IL6-MEDIATED\_...* through *PARTICIPATES\_GpPW*, he can opt for the H-L-L metapath instead of H-H-L, thereby ensuring the exclusion of any undesired paths from deletion.

## 5 EVALUATION

To evaluate the *SLInterpreter*’s effectiveness, we conducted a case study with E1-E3, who had participated in our formative study. Following the Co-discovery Learning Protocol [30], one author briefly guided the system’s operation while E1 operated the system and E2 and E3 engaged in discussions. We then conducted expert interviews to gather insights and feedback on their user experience. The entire process lasted about 2 hours: 30 minutes for familiarizing participants with the system, followed by two 50-minute rounds of iterative exploration.

### 5.1 Case Study

**Investigating CDK1 and Identifying Problem.** E1, a thyroid cancer specialist, initiated by selecting thyroid cancer in the *Disease search box* (Fig. 1-1). Within the *Embedding View*, he identified a region (Fig. 1-2) where accurately predicted *Partner Genes* are closely clustered. Employing semantic zoom, E1 noticed a cluster of KRAS predictions in the upper right area. Given the extensive study of KRAS in SL and its role in the RAS/MAPK signaling pathway regulating cells, the experts held a solid understanding of KRAS. Observing that correctly predicted *Partner Genes* of CDK1 also cluster in the upper right, similar to KRAS, E1 selected CDK1 in the *Primary Gene search box* (Fig. 1-3). However, he found the interpretive paths for CDK1 predominantly linked

by gene, making these paths less convincing. Therefore, E1 shifted focus to different path types. By clicking on the non-gene entities (Fig. 7-1), E1 discovered that these paths primarily lead to MYC and RPL13. Among the non-gene entities, *sensory\_perception\_of\_smell* has the highest weight in BP, which E2 deemed irrelevant to SL.

### Analysing the Interpretive Path of CDK1→MYC

Delving deeper into the embedding of *CDK1*→*MYC*, E1 noted three correctly predicted genes: *KRAS*, ranked 13<sup>th</sup>, *MYC* at 27<sup>th</sup>, and *RPL13* at 33<sup>rd</sup>. After excluding higher-ranked results that were predicted solely based on SL relationships and had lower biological significance, E1 found the remaining ranks sufficiently high to merit further exploration. Notably, *CDK1* and *KRAS* are close in the embedding, as are *MYC* and *RPL13*. This proximity suggests a similarity in their features. Given their relevance to thyroid cancer, the clustering of *CDK1* and *KRAS* is intriguing to E1. He emphasized that targeted therapy for RAS gene mutations is a key focus in cancer treatment, with extensive research ongoing. Consequently, E1 was interested in how the model explains the predictions for the other two genes. Drawing from his previous observations, E1 proceeded by selecting the higher-ranked *MYC* in the *Partner Gene Table* (Fig. 1-4), and the corresponding tag is generated (Fig. 1-5). *MYC* is a critical oncogene encoding the transcription factor *MYC* protein, pivotal in cell growth, division, and metabolism processes. Analyzing the *CDK1*→*MYC* interpretive path, E1 observed that the *Entity Treemap*’s middle layers primarily consist of *Genes*, *BP*, and *CC*. Contrary to common patterns observed in *CDK1*’s pairing with *Partner Genes*, the *Entity Flow Bar* (Fig. 1-6) indicates that gene-gene connections are not predominant here. Additionally, the *Path Bar* (Fig. 1-7) displays a smaller ratio of *SL\_GsG*, which indicates validated SL partner relationship, compared to those belonging to BP and CC. The linkage between *CDK1* and *MYC* mainly involves *Gene*→*BP* paths, such as *CDK1—involved\_in—DNA\_replication—has\_part—DNA\_biosynthetic\_process—involved\_in\_inv—MYC*. E1 found this path consistent with his domain knowledge, particularly as *MYC* over-expression, a transcription factor driving cell growth, may enhance the reliance on *CDK1* for DNA biosynthesis.

**Refining the KG.** From the initial exploration, E2 identified an implausible path *CDK1—SL\_GsG—FARSA—involved\_in—sensory\_perception\_of\_smell—involved\_in\_inv—MYC* and similar paths inconsistent with his expertise. To investigate potential data-level issues causing these irrational predictions, he employed the *Modifier View*. E1 first narrowed the investigation scope to *CDK1*→*Gene*→*BP* in the *Knowledge Graph*. Since *CDK1* belongs to the gene category, he began exploring from the dashed green areas (Fig. 1-8). He noticed that the path of *Gene*→*sensory\_perception\_of\_smell* indeed appears in the *Knowledge Graph* (Fig. 1-9), E1 then refined metapath strategies in the *Metapath Modifier*. He first noticed the histogram of *sensory\_perception\_of\_smell* (Fig. 1-10), which indicates that among all entities in its high granularity category of *Biological Processes* (BP),




Fig. 7: 1 E1 clicked on the non-gene entities during investigating CDK1.



this entity frequently appears as the endpoint of such metapaths (*Gene*—*All edge*—*BP*). Based on his expertise, **E1** determined that any path involving *sensory\_perception\_of\_smell* is irrelevant to SL. Therefore, **E1** selected the L-H-L strategy by clicking on *All edge* (Fig. 1-11).

**E3** noted the relatively low height of the primary bar, indicating that the proportional relationship of *FARSA*—*All edge*—*sensory\_perception\_of\_smell* (L-H-L) within its parent set, *Gene*—*All edge*—*sensory\_perception\_of\_smell* (H-H-L) (Fig. 1-12), was not statistically significant. This indicates that many other *Genes* are also connected to *sensory\_perception\_of\_smell*, which are all irrelevant to SL. Consequently, **E1** selected the H-H-L metapath strategy by clicking on *Gene* (Fig. 1-13). Upon doing so, **E1** observed the proportional relationships of the L-H-L and H-L-L sub-strategies to H-H-L (Fig. 1-14). Specifically, there is only one occurrence of L-H-L, while H-L-L dominates the entire bar. This suggests that almost all *Genes* are associated with *sensory\_perception\_of\_smell* through the *involved\_in* relationship. Simultaneously, *FARSA* represents only a fraction of *Gene* connected to *sensory\_perception\_of\_smell*, even when considering all possible edge types. These findings aligned with the experts' expectations. Consequently, **E1** added this strategy to the *Operation List* and entered "*sensory\_perception\_of\_smell is irrelevant to synthetic lethality*" in the text box (Fig. 1-15). Subsequently, **E1** clicked  (Fig. 1-16) to retrain the model. Upon receiving the retrained model and its performance metrics from the backend, **E1** observed from the *Model Log* that the new model's performance was slightly improved compared to the original model. The accuracy increased by 0.88%, an 1.66% improvement over the original rate, while the deleted path data accounted for only 0.13% of the total data.

### Exploring New Potential SL Partner of CDK1

**E1** proceeded by selecting the retrained model (Fig. 1-17) to initiate an iterative exploration. This time, he directly targeted CDK1 in the *Primary Gene search box*. Upon examining the *Interpretation View* of CDK1, **E2** noticed that the newly generated paths no longer include *sensory\_perception\_of\_smell*. Furthermore, MYC and RPL13 remained accurately identified, maintaining their rankings at 23<sup>rd</sup> and 29<sup>th</sup>. In the *Embedding View*, **E1** noted that RPL13 and EGFR are close to MYC. Additionally, through the *Rank Indicator* (Fig. 1-18), **E3** noticed that these three genes share similar rankings. Notably, EGFR, ranked 31<sup>st</sup>, indicates a potential but unidentified SL relationship with CDK1. To investigate potential patterns among them, **E1** conducted a lasso operation (Fig. 1-19). Among these genes' interpretive paths, **E3** observed a high proportion for the pattern *Gene*→*BP*→*BP*→*Gene*. By highlighting EGFR's interpretive paths (Fig. 8-1), **E3** observed striking similarities to the patterns of the other two correctly predicted genes. At this point, **E3** remarked, "If the model could give us clear paths like this for every prediction, and we could compare them with other known patterns, then I'd feel more confident in trusting those predictions and putting resources into checking them."

## 5.2 Expert Interview

**System Design and Feedback.** Biologists confirmed the effectiveness of *SLInterpreter* in a human-AI collaboration that enhances model interpretability and explores SL mechanisms. **E1** remarked, "The design of the Session View fits well with my habits of using biological databases; we can start exploring from a specific disease of interest." **E2** praised the *Embedding View* for quickly indicating potential common patterns and appreciated the *Rank Indicator*, which helps determine the ranks of predicted genes. **E3** was pleased with the system's comprehensive and intuitive features that aid in removing irrelevant or erroneous data, especially given the high number of data proven to be erroneous or biased. He suggested the system could further improve by automatically col-




Fig. 8: 1 **E1** highlighted EGFR's interpretive paths to reveal patterns similar to MYC and RPL13.

lecting and displaying suspicious data. Additionally, **E1** and **E3** noted that traditional AI-assisted methods required sequential reviewing of a large number of materials based on confidence ranking alone, typically taking **two to three days** to filter 50 predicted results. This lengthy duration was due to the extensive time spent retrieving potential associations between predicted SL pairs, as the information was scattered and lacked targeted retrieval, often resulting in unexpected delays. The discontinuity of this review process also reduced their confidence in the results. In contrast, using *SLInterpreter* made their exploration more targeted, enabling them to review 50 predicted results within **five to eight hours**, considering the varying time required for retraining the model. *SLInterpreter* eliminated the need for further retrieval of obviously erroneous prediction paths, and the continuity and intervenability of the process increased their confidence in the results.

**Learning Curve.** The experts noted that the system's design matches their usage habits and the interaction is intuitive. However, there is a certain learning curve associated with the *Metapath Modifier* in the *Modifier View* and the whole iterative cycle. The entire system requires some time to understand and master (about 20 minutes), and learning through specific cases can make the learning curve more gradual. Nevertheless, the experts still believe that *SLInterpreter* can help them gain more insights from their explorations and have greater confidence in the potential SL pairs they discover.

## 6 DISCUSSION AND LIMITATION

**Lessons Learned.** Our collaboration with biologists has provided valuable insights. First, SL-related biological data is often extensive and high-dimensional, making rule-based filtering insufficient and necessitating biologists' expertise. However, exhaustive evaluations by biologists are unsustainable, highlighting the importance of leveraging AI or visualization tools to focus their assessments. Second, we recognized the need to refine metapath granularity levels. While biological knowledge categorization is intricate, adding granularity levels for certain entity types can rapidly expand metapath strategies. Thus, balancing granularity levels with user burden is crucial for effectiveness.

**Generalization and Scalability.** The *SLInterpreter* enhances interpretability and facilitates cross-granularity exploration for AI models through interpretive paths. Validated by biologists, it is applicable to various network predictive tasks and models across multiple domains, including drug prediction, social network forecasting, and recommendation systems. To ensure scalability, features like semantic zooming and highlighting have been implemented to prevent visual clutter. Additionally, for larger entity sets in knowledge graphs, stacking or grouping nodes can further enhance spatial information capacity.

**Limitations.** This study presents several limitations. First, the dataset has limited biological relationships. Despite SL spans various domains like biomolecular studies and genomics, resulting explanations remain somewhat broad and constrained. Second, the model primarily predicts new SL pairs via existing pairs, which may have limited value due to relatively few connected paths via biological information. Third, experts note that SL gene pairs are not exclusively limited to pairs [42]. Although rare, confirmed instances of SL relationships involving multiple genes exist, expanding prediction from two genes to more, complicating their connections. This highlights the need for comprehensive systems and visualizations for complex SL networks. Additionally, due to extended cycles of wet lab experiments, adequate quantitative experiments were not feasible within the available time.

## 7 CONCLUSION AND FUTURE WORK

This study introduces an iterative Human-AI collaboration framework using *SLInterpreter*, aimed at 1) *Human-Engaged Knowledge Graph Refinement based on Metapath Strategies* and 2) *Cross-Granularity SL Interpretation Enhancement and Mechanism Analysis* for domain experts. A case study and expert interviews demonstrate *SLInterpreter*'s ability to efficiently discover new SL pairs, provide substantial interpretability, and offer effective methods for commonality exploration. Future plans include extending the analysis from single paths to complex networks involving multiple genes, further integrating additional data sources, and expanding the system to other GNN prediction tasks.

## ACKNOWLEDGMENTS

We would like to express our gratitude to our domain experts and the anonymous reviewers for their insightful comments. This work is funded by grants from the National Natural Science Foundation of China (No. 62372298), the Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), and the Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University), Ministry of Education. A part of the experiments of this work were supported by the Core Facility Platform of Computer Science and Communication, SIST, ShanghaiTech University.

## REFERENCES

- [1] Z. Ashktorab, Q. V. Liao, C. Dugan, J. Johnson, Q. Pan, W. Zhang, S. Kumaravel, and M. Campbell. Human-ai collaboration in a cooperative game setting: Measuring social perception and outcomes. *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–20, 2020. doi: 10.1145/3415167 3
- [2] F. Baldassarre and H. Azizpour. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686*, 2019. doi: 10.48550/arXiv.1905.1368 3
- [3] M. Balzer and O. Deussen. Voronoi treemaps. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 49–56, 2005. doi: 10.1109/INFVIS.2005.1532128 6
- [4] R. Cai, X. Chen, Y. Fang, M. Wu, and Y. Hao. Dual-dropout graph convolutional network for predicting synthetic lethality in human cancers. *Bioinformatics*, 36(16):4458–4465, 2020. doi: 10.1093/bioinformatics/btaa211 2
- [5] V. Clarke and V. Braun. Thematic analysis. *The journal of positive psychology*, 12(3):297–298, 2017. doi: 10.1080/17439760.2016.1262613 3
- [6] N. Conde-Pueyo, A. Munteanu, R. V. Solé, and C. Rodríguez-Caso. Human synthetic lethal inference as potential anti-cancer target gene detection. *BMC Systems Biology*, 3(1):1–15, 2009. doi: 10.1186/1752-0509-3-116 2
- [7] X. Deng, S. Das, K. Valdez, K. Camphausen, and U. Shankavaram. SI-biodp: multi-cancer interactive tool for prediction of synthetic lethality and response to cancer treatment. *Cancers*, 11(11):1682, 2019. doi: 10.3390/cancers11111682 2
- [8] J. J. Dudley and P. O. Kristensson. A review of user interface design for interactive machine learning. *ACM Transactions on Interactive Intelligent Systems (TiIS)*, 8(2):1–37, 2018. doi: 10.1145/3185517 3
- [9] X. Feng, N. Arang, D. C. Rigracciolo, J. S. Lee, H. Yeerna, Z. Wang, S. Lubrano, A. Kishore, J. A. Pachter, G. M. König, et al. A platform of synthetic lethal gene interaction networks reveals that the gnaq uveal melanoma oncogene controls the hippo pathway through fak. *Cancer cell*, 35(3):457–472, 2019. doi: 10.1016/j.ccell.2019.01.009 2
- [10] L. Y. Geer, A. Marchler-Bauer, R. C. Geer, L. Han, J. He, S. He, C. Liu, W. Shi, and S. H. Bryant. The ncbi biosystems database. *Nucleic acids research*, 38(suppl\_1):D492–D496, 2010. doi: 10.1093/nar/gkp858 5
- [11] M. Ghandi, D. Lee, M. Mohammad-Noori, and M. A. Beer. Enhanced regulatory sequence prediction using gapped k-mer features. *PLOS Computational Biology*, 10(7):1–15, 07 2014. doi: 10.1371/journal.pcbi.1003711 5
- [12] C. Gourley, J. Balmaña, J. A. Lederemann, V. Serra, R. Dent, S. Loibl, E. Pujade-Lauraine, and S. J. Boulton. Moving from poly (adp-ribose) polymerase inhibition to targeting dna repair and dna damage response in cancer therapy. *Journal of Clinical Oncology*, 37(25):2257–2269, 2019. doi: 10.1200/JCO.18.02050 2
- [13] J. Huang, A. Mishra, B. C. Kwon, and C. Bryan. Conceptexplainer: Interactive explanation for deep neural networks from a concept perspective. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):831–841, 2023. doi: 10.1109/TVCG.2022.3209384 2
- [14] H. Ij. Statistics versus machine learning. *Nat Methods*, 15(4):233, 2018. doi: 10.1038/nmeth.4642 3
- [15] A. Jacunski, S. J. Dixon, and N. P. Tatonetti. Connectivity homology enables inter-species network models of synthetic lethality. *PLoS computational biology*, 11(10):e1004506, 2015. doi: 10.1371/journal.pcbi.1004506 2
- [16] J. Jakubik, J. Schöffner, V. Hoge, M. Vössing, and N. Kühn. An empirical evaluation of predicted outcomes as explanations in human-ai decision-making. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 353–368. Springer, 2022. doi: 10.48550/arXiv.2208.04181 2
- [17] H. Jariyal, F. Weinberg, A. Achreja, D. Nagarath, and A. Srivastava. Synthetic lethality: a step forward for personalized medicine in cancer. *Drug Discovery Today*, 25(2):305–320, 2020. doi: 10.1016/j.drudis.2019.11.014 2
- [18] L. Jerby-Arnon, N. Pfetzer, Y. Y. Waldman, L. McGarry, D. James, E. Shanks, B. Seashore-Ludlow, A. Weinstock, T. Geiger, P. A. Clemons, et al. Predicting cancer-specific vulnerability via data-driven detection of synthetic lethality. *Cell*, 158(5):1199–1209, 2014. doi: 10.1016/j.cell.2014.07.027 2
- [19] D. Jiang, C. Tang, and A. Zhang. Cluster analysis for gene expression data: a survey. *IEEE Transactions on Knowledge and Data Engineering*, 16(11):1370–1386, 2004. doi: 10.1109/TKDE.2004.68 6
- [20] Z. Jin, Y. Wang, Q. Wang, Y. Ming, T. Ma, and H. Qu. Gnnlens: A visual analytics approach for prediction error diagnosis of graph neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 2022. doi: 10.1109/TVCG.2022.3148107 3
- [21] W. G. Kaelin Jr. The concept of synthetic lethality in the context of anticancer therapy. *Nature reviews cancer*, 5(9):689–698, 2005. doi: 10.1038/nrc1691 2
- [22] T. Kranthi, S. Rao, and P. Manimaran. Identification of synthetic lethal pairs in biological systems through network information centrality. *Molecular bioSystems*, 9(8):2163–2167, 2013. doi: 10.1039/c3mb25589a 2
- [23] A. A. Ku, H.-M. Hu, X. Zhao, K. N. Shah, S. Kongara, D. Wu, F. McCormick, A. Balmain, and S. Bandyopadhyay. Integration of multiple biological contexts reveals principles of synthetic lethality that affect reproducibility. *Nature communications*, 11(1):2375, 2020. doi: 10.1038/s41467-020-16078-y 2
- [24] M. Lai, G. Chen, H. Yang, J. Yang, Z. Jiang, M. Wu, and J. Zheng. Predicting synthetic lethality in human cancers via multi-graph ensemble neural network. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pp. 1731–1734, 2021. doi: 10.1109/EMBC46164.2021.9630716 2
- [25] V. Lai, S. Carton, R. Bhatnagar, Q. V. Liao, Y. Zhang, and C. Tan. Human-ai collaboration via conditional delegation: A case study of content moderation. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–18, 2022. doi: 10.1145/3491102.3501999 3
- [26] V. Lai, H. Liu, and C. Tan. " why is 'chicago' deceptive?" towards building model-driven tutorials for humans. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020. doi: 10.1145/3313831.3376873 3
- [27] V. Lai and C. Tan. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of the conference on fairness, accountability, and transparency*, pp. 29–38, 2019. doi: 10.1145/3287560.3287590 3
- [28] J. Lee, W. Yoon, S. Kim, D. Kim, S. Kim, C. H. So, and J. Kang. Biobert: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240, 2020. doi: 10.1093/bioinformatics/btz682 5
- [29] J. Li, D. Tang, X. Lu, F. Sun, K. Jiang, and J. Ruan. Latent space feature representation on multiple biological network for synthetic lethality interaction prediction. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pp. 1236–1241, 2023. doi: 10.1109/BIBM58861.2023.10385727 2
- [30] K. H. Lim, L. M. Ward, and I. Benbasat. An empirical study of computer system learning: Comparison of co-discovery and self-discovery methods. *Information Systems Research*, 8(3):254–272, 1997. doi: 10.1287/isre.8.3.254 8
- [31] G. C. Linderman and S. Steinerberger. Clustering with t-sne, provably. *SIAM journal on mathematics of data science*, 1(2):313–332, 2019. doi: 10.1137/18M1216134 6
- [32] D. Liu, F. Guo, B. Deng, H. Qu, and Y. Wu. Egocomp: A node-link-based technique for visual comparison of ego-networks. *Information Visualization*, 16, 09 2016. doi: 10.1177/14738716166667632 7
- [33] L. Liu, X. Chen, C. Hu, D. Zhang, Z. Shao, Q. Jin, J. Yang, H. Xie, B. Liu, M. Hu, et al. Synthetic lethality-based identification of targets for anticancer drugs in the human signaling network. *Scientific Reports*, 8(1):8440, 2018. doi: 10.1038/s41598-018-26783-w 2
- [34] Z. Liu, Y. Wang, J. Bernard, and T. Munzner. Visualizing graph neural networks with corgie: Corresponding a graph to its embedding. *IEEE Transactions on Visualization and Computer Graphics*, 28(6):2500–2516, 2022. doi: 10.1109/TVCG.2022.3148197 3

- [35] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S.-I. Lee. From local explanations to global understanding with explainable ai for trees. *Nature machine intelligence*, 2(1):56–67, 2020. doi: 10.1038/s42256-019-0138-9 3
- [36] D. Luo, W. Cheng, D. Xu, W. Yu, B. Zong, H. Chen, and X. Zhang. Parameterized explainer for graph neural network. *Advances in neural information processing systems*, 33:19620–19631, 2020. doi: 10.48550/arXiv.2011.04573 2, 3
- [37] S. Ma, M. Sun, and X. Ma. Modeling adaptive expression of robot learning engagement and exploring its effects on human teachers. *ACM Transactions on Computer-Human Interaction*, 30(5):1–48, 2023. doi: 10.1145/3571813 3
- [38] M. Mackeprang, C. Müller-Birn, and M. T. Stauss. Discovering the sweet spot of human-computer configurations: A case study in information extraction. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW):1–30, 2019. doi: 10.1145/3359297 3
- [39] L. McInnes and J. Healy. Umap: Uniform manifold approximation and projection for dimension reduction. *ArXiv*, abs/1802.03426, 2018. doi: 10.21105/joss.00861 6
- [40] B. J. Oakes, M. Famelis, and H. Sahraoui. Building domain-specific machine learning workflows: A conceptual framework for the state-of-the-practice. *ACM Trans. Softw. Eng. Methodol.*, 2023. doi: 10.1145/3638243 3
- [41] N. J. O’Neil, M. L. Bailey, and P. Hieter. Synthetic lethality and cancer. *Nature Reviews Genetics*, 18(10):613–623, 2017. doi: 10.1038/nrg.2017.47 2
- [42] C. J. Ryan, L. P. S. Devakumar, S. J. Pettitt, and C. J. Lord. Complex synthetic lethality in cancer. *Nature Genetics*, 55(12):2039–2048, 2023. doi: 10.1038/s41588-023-01557-x 9
- [43] C. Shan, Y. Shen, Y. Zhang, X. Li, and D. Li. Reinforcement learning enhanced explainer for graph neural networks. *Advances in Neural Information Processing Systems*, 34:22523–22533, 2021. 3
- [44] C. Shi, Y. Hu, S. Wang, S. Ma, C. Zheng, X. Ma, and Q. Luo. Retrolens: A human-ai collaborative system for multi-step retrosynthetic route planning. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI ’23, article no. 770, 20 pages. Association for Computing Machinery, New York, NY, USA, 2023. doi: 10.1145/3544548.3581469 3
- [45] I. Spinelli, S. Scardapane, and A. Uncini. A meta-learning approach for training explainable graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2022. doi: 10.1109/TNNLS.2022.3171398 3
- [46] W. Topatana, S. Juengpanich, S. Li, J. Cao, J. Hu, J. Lee, K. Suliyanto, D. Ma, B. Zhang, M. Chen, et al. Advances in synthetic lethality for cancer therapy: cellular mechanism and clinical translation. *Journal of hematology & oncology*, 13:1–22, 2020. doi: 10.1186/s13045-020-00956-5 2
- [47] L. van der Maaten and G. E. Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605, 2008. 6
- [48] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Lio, and Y. Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. doi: 10.48550/arXiv.1710.10903 3
- [49] J. Wang, Q. Zhang, J. cai Han, Y. Zhao, C. Zhao, B. Yan, C. Dai, L. Wu, Y. Wen, Y. Zhang, D. Leng, Z. Wang, X. Yang, S. He, and X. Bo. Computational methods, databases and tools for synthetic lethality prediction. *Briefings in Bioinformatics*, 23, 2022. doi: 10.1093/bib/bbac106 2, 3
- [50] Q. Wang, K. Huang, P. Chandak, M. Zitnik, and N. Gehlenborg. Extending the nested model for user-centric xai: A design study on gnn-based drug repurposing. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):1266–1276, 2023. doi: 10.1109/TVCG.2022.3209435 2, 3
- [51] S. Wang, Y. Feng, X. Liu, Y. Liu, M. Wu, and J. Zheng. Nsf4sl: negative-sample-free contrastive learning for ranking synthetic lethal partner genes in human cancers. *Bioinformatics (Oxford, England)*, 38:ii13–ii19, 09 2022. doi: 10.1093/bioinformatics/btac462 5
- [52] S. Wang, F. Xu, Y. Li, J. Wang, K. Zhang, Y. Liu, M. Wu, and J. Zheng. Kg4sl: knowledge graph neural network for synthetic lethality prediction in human cancers. *Bioinformatics*, 37(Supplement\_1):i418–i425, 2021. doi: 10.1093/bioinformatics/btab271 2, 5, 6
- [53] X. Wang and R. Simon. Identification of potential synthetic lethal genes to p53 using a computational biology approach. *BMC medical genomics*, 6(1):1–10, 2013. doi: 10.1186/1755-8794-6-30 2
- [54] X. Wang and M. Yin. Effects of explanations in ai-assisted decision making: Principles and comparisons. *ACM Trans. Interact. Intell. Syst.*, 12(4), article no. 27, 36 pages, nov 2022. doi: 10.1145/3519266 3
- [55] S. Wold, K. Esbensen, and P. Geladi. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52, 1987. doi: 10.1016/0169-7439(87)80084-9 6
- [56] L.-L. Wu, Y.-Q. Wen, X.-X. Yang, B.-W. Yan, S. He, and X.-C. Bo. Synthetic lethal interactions prediction based on multiple similarity measures fusion. *Journal of Computer Science and Technology*, 36:261–275, 2021. doi: 10.1007/s11390-021-0866-2 2, 3
- [57] X. Wu, L. Xiao, Y. Sun, J. Zhang, T. Ma, and L. He. A survey of human-in-the-loop for machine learning. *Future Generation Computer Systems*, 135:364–381, 2022. doi: 10.1016/j.future.2022.05.014 2, 3
- [58] C. Yang, Y. Guo, R. Qian, Y. Huang, L. Zhang, J. Wang, X. Huang, Z. Liu, W. Qin, C. Wang, et al. Mapping the landscape of synthetic lethal interactions in liver cancer. *Theranostics*, 11(18):9038, 2021. doi: 10.7150/thno.63416 2
- [59] W. Yang, J. Soares, P. Greninger, E. J. Edelman, H. Lightfoot, S. Forbes, N. Bindal, D. Beare, J. A. Smith, I. R. Thompson, et al. Genomics of drug sensitivity in cancer (gdsc): a resource for therapeutic biomarker discovery in cancer cells. *Nucleic acids research*, 41(D1):D955–D961, 2012. doi: 10.1093/nar/gks1111 2
- [60] Z. Ying, D. Bourgeois, J. You, M. Zitnik, and J. Leskovec. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems*, 32, 2019. doi: 10.48550/arXiv.1903.03894 3
- [61] X. Yu, D. Liang, and Q. Li. Improved graphsvx for gnn explanations based on cross entropy. In *2023 4th International Conference on Electronic Communication and Artificial Intelligence (ICECAI)*, pp. 147–152. IEEE, 2023. doi: 10.1109/ICECAI58670.2023.10176786 3
- [62] H. Yuan, H. Yu, J. Wang, K. Li, and S. Ji. On explainability of graph neural networks via subgraph explorations. In *International conference on machine learning*, pp. 12241–12252. PMLR, 2021. doi: 10.48550/arXiv.2102.05152 3
- [63] J. Zhang, D. Luo, and H. Wei. Mixupexplainer: Generalizing explanations for graph neural networks with data augmentation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3286–3296, 2023. doi: 10.1145/3580305.3599435 3
- [64] K. Zhang, M. Wu, Y. Liu, Y. Feng, and J. Zheng. Kr4sl: knowledge graph reasoning for explainable prediction of synthetic lethality. *Bioinformatics*, 39(Supplement\_1):i158–i167, 2023. doi: 10.1093/bioinformatics/btad261 2, 3, 5
- [65] N. Zhang, Z. Bi, X. Liang, S. Cheng, H. Hong, S. Deng, J. Lian, Q. Zhang, and H. Chen. Ontoprotein: Protein pretraining with gene ontology embedding. *arXiv preprint arXiv:2201.11147*, 2022. doi: 10.48550/arXiv.2201.11147 5
- [66] Y. Zhang, Q. V. Liao, and R. K. E. Bellamy. 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, FAT\* ’20*, 11 pages, p. 295–305. Association for Computing Machinery, New York, NY, USA, 2020. doi: 10.1145/3351095.3372852 3
- [67] Y. Zhang, Q. Yao, L. Yue, X. Wu, Z. Zhang, Z. Lin, and Y. Zheng. Emerging drug interaction prediction enabled by a flow-based graph neural network with biomedical network. *Nature Computational Science*, 3(12):1023–1033, 2023. doi: 10.1038/s43588-023-00558-4 2
- [68] Z. Zhao, P. Xu, C. Scheidegger, and L. Ren. Human-in-the-loop extraction of interpretable concepts in deep learning models. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):780–790, 2022. doi: 10.1109/TVCG.2021.3114837 3
- [69] Y. Zhu, Y. Zhou, Y. Liu, X. Wang, and J. Li. SLGNN: synthetic lethality prediction in human cancers based on factor-aware knowledge graph neural network. *Bioinformatics*, 39(2):btad015, 01 2023. doi: 10.1093/bioinformatics/btad015 5
- [70] O. Zuckerman, V. S. Press, E. Barda, B. Megidish, and H. Erel. Tangible collaboration: A human-centered approach for sharing control with an actuated-interface. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2022. doi: 10.1145/3491102.3517449 3