

BioPIE: A Biomedical Protocol Information Extraction Dataset for Experiment Understanding

Haofei Hou^{*,1,3}, Shunyi Zhao^{*,2,3}, Fanxu Meng^{*,1,3}, Kairui Yang^{1,3},
Lecheng Ruan^{✉,1,3}, Qining Wang^{✉,1,3}

¹School of Advanced Manufacturing and Robotics, Peking University, Beijing, China

²School of Integrated Circuits, Peking University, Beijing, China

³PKU-HEC Joint Lab for Industrial Intelligence, Beijing, China

*Equal Contribution ✉ ruanlecheng@ucla.edu, qiningwang@pku.edu.cn

Abstract

Understanding biomedical experiments provides a foundation for downstream tasks, *e.g.*, laboratory automation, and facilitates effective cross-disciplinary communication. Two challenges, High Information Density (HID) and Multi-Step Reasoning (MSR), pose unique difficulties for precise automatic experimental understanding. Extracting structured knowledge, *e.g.*, Knowledge Graphs (KGs), is an effective approach to address the HID and MSR. However, existing biomedical datasets for structured knowledge Information Extraction (IE) are limited to a general or coarse-grained level, hindering fine-grained experimental understanding. To address this gap, we introduce Biomedical Protocol Information Extraction Dataset (BioPIE), a dataset providing procedure-centric KGs that captures entities, actions, and relations at a scale sufficient for reasoning across biomedical protocols. We evaluate both supervised and LLM-based IE methods on BioPIE to verify its effectiveness, and implement a biomedical question answering system to provide a quantitative illustration of BioPIE’s effectiveness for downstream understanding tasks. The experimental results demonstrate improved understanding performance on both the HID and MSR question sets.

1 Introduction

Biomedical research spans diverse sub-fields and relies heavily on experimental workflows (Jin et al., 2019). These experiments involve multiple stages, including experiment design, execution, and result analysis. Due to the procedural complexity and wide-ranging domain coverage of biomedical experiments, automated understanding has significant potential for supporting cross-disciplinary collaboration (Rohrbach et al., 2022), laboratory automation, and Artificial Intelligence (AI)-assisted experimental planning (Steiner et al.,




Figure 1: **BioPIE enhances complex biomedical protocol understanding.** (A) The KGs in BioPIE provide fine-grained structural representations of experimental steps (*e.g.*, temperature, duration, and execution order), resulting in **high information density**, and enable **multi-step reasoning** by integrating sentence-level context with graph-structured knowledge. (B) Existing information extraction datasets involve a trade-off: general datasets lack biomedical knowledge, while domain-specific datasets may not generalize across diverse experiments.

2019; Mehr et al., 2020; Burger et al., 2020; Szymanski et al., 2023; Shi et al., 2026). However, the language describing these experiments is inherently complex and highly domain-specific, posing major challenges for automation (Frisoni et al., 2022; Shi et al., 2024).

Understanding biomedical experiments exhibits several characteristics. Experimental protocols often condense substantial operational detail into a single sentence (Li et al., 2018), *e.g.*, volumes, temperatures, timings, buffer compositions, and instrument settings, collectively termed **High Information Density (HID)**, which requires precise extraction and association with the corresponding operations. Moreover, protocols frequently feature chained steps, hierarchical subroutines, and im-

licitly distributed information (Shi et al., 2025), highlighting **Multi-Step Reasoning (MSR)**: answering a question often necessitates integrating information from multiple distinct steps or facts. Collectively, HID and MSR encompass key aspects of the complexity of biomedical experiment understanding.

Prior work has demonstrated that Information Extraction (IE) of structured knowledge can benefit natural language understanding (Fang et al., 2024). Knowledge Graphs (KGs) can structurally represent dense and heterogeneous parameters, such as entities, attributes, and their operational relations, thereby making them well-suited to capturing the HID and MSR structures inherent in biomedical experimental protocols. Therefore, a dataset that provides detailed and accurate biomedical experimental KGs may help natural language understanding of the biomedical protocols. However, there is a lack of datasets specifically designed for this purpose.

Existing IE datasets can be broadly categorized into two groups. Datasets covering general scientific information are broad in sub-fields (Nasar et al., 2021; Zhao et al., 2024). For example, SciERC and SciER annotate entities such as Method, Task, Metric, and Material, along with relations including Used-for, Part-of, Compare, and Evaluate (Luan et al., 2018; Zhang et al., 2024). However, the lack of domain-specific text makes it difficult for these datasets to fully represent the experimental reagents, materials, containers, and devices required in biomedical applications, as shown in Fig. 1(B).

The other category comprises datasets specifically designed for biomedical research (Arsenyan et al., 2024; Peng et al., 2024). These biomedical datasets primarily focus on entities such as proteins, chemicals, drugs, and diseases (Kringelum et al., 2016; Krallinger et al., 2021), with relation types covering the molecular and pharmacological interactions (Herrero-Zazo et al., 2013; Zhang et al., 2022; Luo et al., 2022; Lai et al., 2024). Nevertheless, existing biomedical datasets typically provide a relatively coarse-grained scheme centered on molecular and pharmacological entities, lacking the procedural granularity required for biomedical experiment understanding. Together, these observations highlight the need for a dataset that encompasses diverse, cross-disciplinary experimental protocols with sufficient procedural detail to support biomedical experiment understand-

ing (Perera et al., 2020). To the best of our knowledge, such a dataset does not currently exist.

In this paper, we introduce Biomedical Protocol Information Extraction Dataset (BioPIE), a new dataset specifically designed to support biomedical experiment understanding. BioPIE provides clearly defined biomedical experimental protocols and corresponding KGs, with broad cross-disciplinary coverage of experimental entities, actions, and procedural relations. It is constructed to support generalizable reasoning and machine understanding of biomedical protocols, thereby improving biomedical experimental understanding, which is illustrated in Fig. 1(A).

The contributions of this paper are as follows: (1) We construct BioPIE, an IE dataset for understanding complex biomedical experiments, including both HID and MSR aspects; (2) We systematically evaluate different IE algorithms on BioPIE, including both supervised models and Large Language Models (LLMs) under different settings; and (3) We develop a Question Answering (QA) system to provide a quantitative illustration that BioPIE effectively supports biomedical experiment understanding, particularly for HID and MSR.

2 BioPIE Dataset

2.1 Data Scheme

To support robust knowledge extraction from biomedical protocols, we design an annotation scheme that captures the essential operational structure while avoiding overly domain-specific categories. The scheme is fine-grained for procedural reasoning yet generalizable across biomedical experimental workflows rather than being tied to a narrow subdomain, as shown in Fig. 2(A), which is different from Fig. 1(B). The detailed design and development process of ontologies can be found in Appx. A.1.

BioPIE comprises 34 entity types encompassing actions, materials, laboratory instruments, biological samples, and key experimental parameters (*e.g.*, time, temperature, and force). BioPIE not only provides general definitions for actions and processes, but also focuses on operational elements that recur across diverse biomedical procedures, *e.g.*, centrifugation forces, incubation temperatures, buffers, and consumables. Following standard practice in scientific IE (Stenetorp et al., 2012; Zhang et al., 2024), annotators allow




Figure 2: Overview of BioPIE. (A) An annotated plasmid DNA preparation protocol illustrating how domain-specific experimental descriptions are mapped into a shared, procedure-centric representation of actions, materials, parameters, and procedural dependencies, thereby supporting consistent annotation across diverse biomedical workflows. (B) Distributions of the 34 entity types and 21 relation types in the dataset. (C) Definitions and examples of representative entity and relation labels, with examples drawn from different protocols to demonstrate their applicability across biomedical procedures.

nested spans when necessary for relation attachment.

We define 21 relation types to describe how experimental entities interact within a protocol, including action-object relations (*is_object_of*), action-parameter relations (*have_parameter*), resource-usage relations (*use_device*), structural relations (*contain*), and procedural logic (*next_step*). These relations capture the diverse aspects of human instructions in experimental protocols without relying on domain-specific biological semantics, making them suitable for heterogeneous protocols. Some detailed explanations of representative entity and relation labels are provided in Fig. 2(C).

Our balanced design allows the KGs to support understanding while maintaining broad applicability across diverse biomedical experiments. First, our scheme is procedure-centric and focuses on operational details such as actions, materials, and parameters. At the same time, it deliberately avoids narrowly specialized biomedical categories, enabling consistent annotation across cell culture protocols, microscopy workflows, sequencing preparations, biomaterial fabrication, and other experimental contexts. BioPIE provides a comprehensive representation of biomedical workflows, achieved through these detailed, fine-grained local representations. Detailed exam-

ples can be found in Appx. A.1.4.

2.2 Data Collection and Processing

We first collect protocols from peer-reviewed biology research articles and use Qwen-Max to clean and normalize them into stepwise imperative sentences. From these protocols, we select 344 sub-protocols covering four major domains: genetic manipulation, molecular interactions, physiological signal detection, and regenerative medicine, from the complete collection as our in-domain (ID) set, covering a broad range of common biomedical workflows. To construct an out-of-domain (OOD) set, we select the other 45 sub-protocols whose biomedical sub-fields are not represented in the ID set. These OOD protocols cover distinct experiment types such as plant-based biopharmaceuticals, functional imaging, and epidemiological analysis. We also analyze the detailed distribution of data sources (Appx. A.2), and the impact of this normalization of protocols qualitatively and quantitatively (Appx. A.3). The analyses demonstrate the normalization has no significant impact on the information within experimental workflows.

2.3 Data Annotation

Two annotators with graduate-level backgrounds in computer science and biomedical research are

	ACE2005	SciERC	ChemProt	BioPIE
Entity Types	7	6	3	34
Relation Types	6	7	11	21
Entities	38287	8089	17340	10982
Relations	7070	4716	10065	8848
Sentences	10372	2679	7552	1916
Relations/Sent.	0.68	1.76	1.33	4.62

Table 1: **Comparison of BioPIE and three datasets.**

recruited. The annotators receive training before starting the task. The dataset was annotated through an independent parallel annotation process by two annotators who had no access to each other’s results during the task. A lead annotator with the same background then reconciled all discrepancies according to annotation guideline to produce the final gold-standard labels. For all protocols, we compute inter-annotator agreement using Cohen’s kappa (Davies and Fleiss, 1982). The kappa score for entity annotation is 79.20% and for relation annotation is 68.26%, achieving a level of consistency comparable to that reported in existing literature (Luan et al., 2018; Zhang et al., 2024).

2.4 Data Statistics

As shown in Tab. 1, BioPIE contains over 10.9k entities and 8.8k relations, which are comparable in scale to existing IE datasets (Doddington et al., 2004; Kringelum et al., 2016; Luan et al., 2018). BioPIE exhibits substantially higher relational density than prior datasets, averaging 4.6 relations per sentence compared to 0.7 (Doddington et al., 2004)–1.7 (Luan et al., 2018) in existing corpora. This reflects the inherently structured and interaction-rich nature of biomedical protocols. High relational density directly indicates information density but does not necessarily imply multi-step logical chain complexity. We randomly split the dataset into the training, development, and ID test sets using a 7:1:2 ratio. The additional protocols are used as the OOD test set. ID evaluation assesses model reasoning on seen objects and protocols under similar distributions. Conversely, OOD protocols cover entirely unseen biomedical fields and experiment types, serving as a true bottleneck to test cross-domain generalization. Fig. 2(B) presents the detailed distribution of each entity and relation type.

3 BioPIE Benchmarking

3.1 Information Extraction Baselines

We consider two commonly used types of IE methods: supervised models, which exhibit strong IE performance on specific tasks; and LLMs, which are pretrained on broad-coverage corpora and provide more general IE capabilities (Chang et al., 2024; Naveed et al., 2025). For both supervised models and LLMs, we investigate two architectural frameworks: a pipeline framework, which performs Named-Entity Extraction (NER) and Relation Extraction (RE) separately, and a joint Entity and Relation Extraction (ERE) framework, which performs NER and RE jointly.

With the above baseline selection criteria, we select two State-Of-The-Art (SOTA) supervised models as baselines: PL-Marker (Ye et al., 2022), which adopts a span-based representation strategy within a pipeline framework, and HGERE (Yan et al., 2023), which introduces a joint ERE framework. Considering the zero-shot, few-shot, and Low Rank Adaptation (LoRA) (Hu et al., 2022) settings of LLMs, we combine each setting with the pipeline and joint frameworks, resulting in six LLM-based configurations in total. Under both zero-shot and few-shot settings, we evaluate GPT-5, Claude-4.5-Opus, Llama-4, and Qwen-max. Under the LoRA setting, we evaluate Llama-3-8B and Qwen-3-8B.

3.2 Information Extraction Evaluation Details

For supervised baselines, we use *scibert* (Beltagy et al., 2019) and *biobert* (Lee et al., 2020) as the encoder. In the few-shot setting for LLMs, we employ a sentence retriever to select the most similar training examples as in-context demonstrations (Dong et al., 2024). For each task, we retrieve up to 20 candidate demonstrations and select the number of demonstrations that yield the highest Rel+ score on the validation set. We use OpenAI’s *text-embedding-3-large* model as the retriever backbone. The instruction part of our prompt is adapted from ChatIE (Wei et al., 2023), and we additionally provide component label definitions to improve clarity and model understanding (Zhang et al., 2024). The complete prompt can be found in Appx. A.4. During our experiments, the random seed is set to zero.

Given an input protocol D with sentences $S = \{s_1, s_2, \dots, s_N\}$, we define IE independently

	In-domain				Out-of-domain			
	NER	Rel	Rel+	RE	NER	Rel	Rel+	RE
<i>Supervised Baselines</i>								
PL-Marker ¹ (Ye et al., 2022)	87.40	82.55	74.52	87.88	73.87	70.27	52.27	78.85
PL-Marker ² (Ye et al., 2022)	86.33	80.45	72.48	33.69	51.93	23.51	16.52	24.79
HGERE ¹ (Yan et al., 2023)	87.63	82.10	73.93	-	74.58	70.49	52.41	-
HGERE ² (Yan et al., 2023)	86.95	80.93	73.08	-	72.76	69.48	51.48	-
<i>Zero-shot LLM</i>								
GPT-5 (Pipeline)	57.14	50.47	41.14	69.86	52.08	51.46	37.60	68.24
GPT-5 (Joint)	22.90	22.23	17.54	-	21.94	21.23	14.78	-
Claude-4.5-Opus (Pipeline)	69.34	40.56	31.90	48.01	63.81	34.37	24.90	43.33
Claude-4.5-Opus (Joint)	39.37	33.66	26.11	-	65.74	31.31	24.79	-
Llama-4 (Pipeline)	41.08	1.88	1.44	1.69	42.34	1.56	1.21	0.74
Llama-4 (Joint)	1.73	0.75	0.00	-	0.92	0.38	0.19	-
Qwen-max (Pipeline)	67.12	20.53	17.02	27.38	60.23	8.11	13.05	25.05
Qwen-max (Joint)	65.73	24.20	19.10	-	61.34	20.02	14.05	-
<i>Few-shot LLM</i>								
GPT-5 (Pipeline)	62.74	66.83	59.73	79.30	52.80	54.00	41.42	68.46
GPT-5 (Joint)	27.71	25.03	22.14	-	33.89	32.18	26.08	-
Claude-4.5-Opus (Pipeline)	85.18	67.87	63.47	77.20	73.23	53.38	41.48	64.88
Claude-4.5-Opus (Joint)	83.41	65.88	60.27	-	73.11	52.57	41.33	-
Llama-4 (Pipeline)	49.05	18.73	16.75	21.03	41.23	11.71	9.73	16.24
Llama-4 (Joint)	13.51	7.73	7.39	-	18.26	10.15	7.69	-
Qwen-max (Pipeline)	83.28	64.35	59.87	73.17	71.81	46.23	36.45	56.74
Qwen-max (Joint)	75.68	55.49	51.84	-	62.16	37.25	29.27	-
<i>LoRA LLM</i>								
Llama-3-8B (Pipeline)	86.33	75.28	68.13	81.67	75.70	62.97	49.44	73.24
Llama-3-8B (Joint)	84.71	74.86	66.58	-	74.86	62.68	46.72	-
Qwen-3-8B (Pipeline)	84.44	69.95	62.68	77.06	74.90	61.63	46.81	68.59
Qwen-3-8B (Joint)	82.92	70.70	62.96	-	73.86	59.99	44.54	-

Table 2: **Test F1 scores of different baselines on our proposed dataset.** “Joint” denotes joint IE, while “Pipeline” refers to performing NER and RE separately. “Rel” and “Rel+” indicate relation extraction from original text under boundary and strict evaluation, respectively, and “RE” denotes relation extraction with gold entities, applicable only to the pipeline methods. The number of in-context examples was tuned on the validation set, as shown in Fig. 3(A). ¹ Denotes the use of *scibert* (Beltagy et al., 2019) as the encoder; ² denotes *biobert* (Lee et al., 2020).

at the sentence level as follows. Let $\mathbb{E}$ denote a set of entity types. Given a sentence $s_i = \{w_1, w_2, \dots, w_k\}$, the NER task identifies a set of entity mentions $\{e_1, e_2, \dots, e_m\}$. Each entity mentioned $e_j = \{w_l, \dots, w_r\}$ corresponds to a contiguous span of tokens and is assigned an entity type $t_j \in \mathbb{E}$. Let $\mathbb{R}$ denotes the set of relation types. The RE task predicts a relation label $r_{jk} \in \mathbb{R} \cup \{\text{NULL}\}$ for each ordered entity pair (e_j, e_k) occurring within sentence s_i . The special label NULL indicates the absence of a semantic relation. In our RE task, we require LLMs to explicitly generate NULL to mitigate their tendency to misclassify non-relational pairs into predefined categories (Zhang et al., 2024). However, to ensure rigorous evaluation, these NULL predictions are strictly excluded from the final metric calculations.

The evaluation tasks include NER, RE from original text, and RE with gold standard entities. For NER, we conduct span-level evaluation, requiring both correct boundaries and entity types. For RE from original text, we report two metrics following prior work (Ye et al., 2022; Yan et al., 2023): (1) Boundary evaluation (Rel), which requires correct prediction of subject and object boundaries and their relation, and (2) Strict evaluation (Rel+), which additionally requires correct entity types.

3.3 Information Extraction Results

Tab. 2 reports the experimental results on both the ID and OOD test sets. Fig. 3(A) illustrates the impact of varying the number of demonstrations in the few-shot setting on RE from the original text performance (Rel+). Overall, introducing a

small number of demonstrations yields substantial performance gains. Most LLMs reach their peak performance with approximately 5–15 demonstrations, after which additional examples provide diminishing or even negative returns. These findings suggest that overly large demonstration sets may introduce noise and reduce the effectiveness of In-Context Learning (ICL). Detailed results for precision and recall are provided in the Appx. A.5.

Among supervised baselines using *scibert*, PL-Marker achieves the best ID performance (82.55 Rel, 74.52 Rel+, 87.88 RE). In contrast, HGERE demonstrates stronger OOD robustness (74.58 NER, 70.49 Rel, 52.41 Rel+). For both models, ID-to-OOD performance drops more for NER (~ 13) than for RE (10–12), indicating unseen entities are harder to recognize. Notably, substituting *scibert* with *biobert* causes a catastrophic OOD collapse in PL-Marker (e.g., Rel drops to 23.51), whereas HGERE remains robust (69.48 Rel). This suggests that HGERE’s OOD advantage is an intrinsic architectural benefit rather than solely a byproduct of the pretrained encoder.

In the zero-shot setting, LLMs exhibit substantial performance variability. GPT-5 achieves the most balanced pipeline performance on ID data (57.14 NER, 41.14 Rel+), while Claude-4.5-Opus achieves strong NER performance but weaker RE results. Llama-4 performs poorly across most RE-related metrics, and Qwen-max achieves reasonable NER performance but limited RE capability. Across all models, pipeline extraction outperforms joint extraction, highlighting the benefit of decomposing NER and RE for LLMs. Compared to supervised baselines, LLMs show more consistent performance between ID and OOD, likely due to their large-scale pretraining.

Few-shot learning yields dramatic improvements across all evaluated models. Claude-4.5-Opus with pipeline extraction achieves the largest gains, reaching 85.18 (NER) and 63.47 (Rel+) on ID data. However, improvements from ICL are generally larger on ID than OOD, as demonstrations are more similar to ID samples.

LoRA-tuned smaller LLMs demonstrate that parameter-efficient fine-tuning can rival or even surpass LLMs with ICL. Llama-3-8B with pipeline extraction achieves 86.33 (NER) and 68.13 (Rel+) on ID, approaching supervised performance while exhibiting strong generalization. Although pipeline extraction remains slightly superior to joint extraction after fine-tuning, the gap




Figure 3: **Effects of settings on IE methods.** (A) Impact of the number of the retrievals on the validation set for Rel+ F1 score. (B) Performance trends of PL-Marker under varying training-protocol proportions on the ID test set.

becomes much smaller.

Fig. 3(B) shows the performance trends of PL-Marker on NER, Rel, Rel+, and RE, with the different training scales. As the dataset size increases, RE performance improves more slowly and gradually saturates, while NER continues to show moderate gains. Rel and Rel+ also consistently improve with more training data. In particular, Rel+ increases from 47.09 (0.1 of the training set) to 59.79 (0.2), 69.24 (0.5), 73.92 (0.9), and 74.52 (1.0), with diminishing gains as the data scale grows.

4 Biomedical Experiment Understanding Validation

As a typical natural language understanding task, biomedical experiment Question Answering (QA) has become an active research area in recent years (Liu et al., 2026). To demonstrate the utility of BioPIE in enhancing experimental understanding, we implement a KG-augmented QA system and evaluate its performance on questions specifically designed to test HID and MSR capabilities.

4.1 QA System Design

The proposed QA system follows current basic paradigms for graph-based reasoning in LLMs (Sun et al., 2024). We adopt a synergistic architecture that jointly leverages unstructured text and structured graph knowledge to enhance response fidelity. To ensure computational efficiency and leverage the high density of specialized terminology in biomedical protocols, we first implement a filter. This module evaluates candidate sentence-graph pairs (s_i, g_i) for a query q using a hy-




Figure 4: Pipeline of the proposed QA system.

brid relevance score:

$$R(q, s_i, g_i) = R_t(q, s_i) \cdot \log(1 + R_g(q, g_i)), \quad (1)$$

where $R_t(q, s_i)$ denotes the semantic similarity between the query and the sentence, and $R_g(q, g_i) = \sum_{v \in \mathcal{V}_i} \mathbb{I}[v \subseteq q]$ quantifies the structural alignment by measuring the entity overlap between the query and the knowledge graph g_i ($\mathcal{V}_i$ denotes the entity set of g_i) (Kruit et al., 2024).

Building upon this filtered context, the planner module decomposes complex natural language queries into structured execution paths. Specifically, it models the generation of an abstract relation sequence $z = \{r_1, \dots, r_l\}$ via an LLM as $P_\theta(z | q, \hat{s}, \hat{G})$, where $\hat{s}$ and $\hat{G}$ denote the sets of retrieved sentences and graphs, respectively. Guided by these plans, the retriever serves as a grounding engine, instantiating the abstract paths within the knowledge graph to identify specific trajectories $w_z = e_q \xrightarrow{r_1} e_1 \dots \xrightarrow{r_l} e_a$ that link query entities to the potential answers.

In the final stage, the reasoner synthesizes the retrieved textual evidence $\hat{s}$ with the instantiated graph descriptions $G_z = \{w_z\}$ to formulate the terminal response Y . The generation process, parameterized by θ , is optimized by maximizing the conditional probability:

$$p_\theta(Y | q, \hat{s}, G_z) = \prod_{t=1}^{|Y|} P_\theta(y_t | y_{<t}, [q, \hat{s}, G_z]). \quad (2)$$

By integrating structured logical constraints with rich linguistic context, the system ensures that the output is consistent with both textual and relational data, as illustrated in Fig. 4.

4.2 QA Baselines

To demonstrate the effectiveness of the proposed dataset and QA system, we compare our approach

	Test	HID	MSR
LLM only	14.74	16.09	15.45
LLM LoRA	12.44	11.49	9.76
BM25	63.72	65.52	52.03
LaBSE	55.60	63.22	53.66
Emb-3-large	61.52	66.09	56.10
Emb-v4	59.14	67.24	54.47
GRAG	8.83	10.34	7.32
GRAG LoRA	22.77	17.82	16.26
ToG	69.90	75.29	69.92
Ours w/o Sentence	59.84	64.94	65.04
Ours w/o Graph	55.69	66.67	60.16
Ours w/o Planner	72.29	73.56	62.60
Ours w SciERC	62.40	70.11	59.35
Ours w ChemProt	64.70	72.99	64.23
Ours	72.99	83.33	74.80

Table 3: Performance comparison across different QA systems.

against a broad range of commonly used retrieval-based QA systems. Specifically, our experiments cover text-based QA systems equipped with different retrievers, including BM25 (Robertson and Zaragoza, 2009), LaBSE (Feng et al., 2022), OpenAI’s *text-embedding-3-large* (Emb-3-large), and Qwen’s *embedding-v4* (Emb-v4). We also include GRAG (Hu et al., 2025), which relies solely on knowledge-graph-based retrieval without using raw textual corpora, for comparison, and ToG (Sun et al., 2024), which allows LLMs to iteratively explore and prune reasoning chains.

Furthermore, to investigate the impact of KGs on biomedical experiment QA, we compare graphs constructed from IE models trained on different datasets, including SciERC (Luan et al., 2018), which focuses on scientific IE, and ChemProt (Kringelum et al., 2016), which targets chemical reaction IE. In addition, we include two LLM-only baselines that do not leverage any externally retrieved knowledge for LLM: a frozen LLM, and an LLM fine-tuned using LoRA (Hu et al., 2022).

4.3 QA Evaluation Details

The QA dataset is divided into training, validation, and test sets with sizes of 1983, 159, and 1133, respectively. To further analyze model performance under challenging conditions, we construct MSR and HID question sets from the test set. The detailed QA dataset information can be found in Appx. B.1. In the evaluation phase, adhering to the standard paradigm of Knowledge Graph Question Answering (KGQA), relevant contexts

are retrieved from the comprehensive pool of all protocols (Luo et al., 2024; Sun et al., 2024).

We adopt accuracy as the evaluation metric for all experiments, *i.e.*, whether the answer appears in the model output. This is a widely recognized standard (Krithara et al., 2023). Following the established evaluation paradigm in generative QA benchmarks (Wang et al., 2023), accuracy is initially defined via substring matching to prevent severe false negative penalties caused by verbose or explanatory LLM outputs (Adlakha et al., 2024). Furthermore, we tune the number of in-context examples on the validation set. As an additional metric, we compute the retrieval hit rate on the validation set.

During the experiments, the random seed is set to zero. The sentence-level relevance function $R_t(q, s_i)$ is BM25 (Robertson and Zaragoza, 2009) in implementation. Experiments are conducted on Llama-3-8B. The prompts used for the planner, the retriever, and the reasoner are adapted from RoG (Luo et al., 2024). Specifically, we use HGERE (Yan et al., 2023) as the IE method, including NER and RE. Furthermore, we conduct ablation studies comparing textual inputs and knowledge-graph-based inputs, while keeping the retrieval strategy fixed to the proposed pipeline.

4.4 QA Evaluation Results

Tab. 3 presents results on the test set. Fig. 5 shows the effect of varying the number of retrieved texts for accuracy and hit rate on the validation set. Across all settings, the proposed QA system achieves the best overall performance. Fig. 6 illustrates two example outputs from our experimental QA system; additional examples are provided in the Appx. B.3. Detailed evaluation results, including the impact of different hops and various IE methods, are provided in Appx. B.4.

Ablation studies further confirm the complementary role of graph-based inputs. Removing graph-based knowledge (“Ours w/o Graph”) leads to a noticeable performance degradation compared to the full model. Our findings demonstrate that combining sentences with structured KGs can effectively address the challenges of HID and MSR to achieve optimal performance. To isolate specific component contributions, we hold the KG fixed and observe that advanced graph reasoners like ToG (69.90%) outperform standard retrievers like BM25 (63.72%); further, our full reasoning planner (72.99%) outperforms the planner abla-




Figure 5: Effect of the number of retrievals on QA systems performance on the validation set.

tion (72.29%, with a severe MSR drop to 62.60%).

Using graphs constructed from SciERC and ChemProt also results in reduced performance, only marginally outperforming the text-retrieval baseline. Specifically, under the identical QA pipeline, the BioPIE schema (72.99% on Test) outperforms SciERC (62.40%) and ChemProt (64.70%). This indicates that fine-grained knowledge representations tailored to biomedical experimental protocols are critical for effective biomedical experiment QA. As our evaluation protocol targets on unseen text, a Gold KG is inherently unavailable, making predicted KGs the realistic standard for these comparisons.

5 Discussion

The core strength of BioPIE lies in its procedure-centric, rather than concept-centric, design philosophy. In BioPIE, experimental operations are treated as the fundamental units, and the structured dependencies among actions, objects, parameters, and procedural steps are explicitly captured. This design enables complex biomedical experiment understanding: our system achieves 83.33% accuracy on the HID questions and 74.80% on the MSR questions, outperforming all baselines. These results demonstrate that BioPIE can effectively process content with HID and support MSR, highlighting its potential to serve as a foundation for experiment-level understanding. BioPIE is also constructed at a practically reasonable scale. The diminishing improvements in Fig. 3(B) indicate decreasing marginal returns from additional training data, with RE in particular exhibiting clear saturation behavior, while NER shows only slow improvement with markedly diminishing returns.

While closed-source LLMs might have encountered the raw texts during pretraining, this poten-




Figure 6: QA system showcases.

tial exposure does not compromise our evaluation fidelity. Our benchmarks are explicitly designed to assess operational capabilities rather than textual memorization: first, the IE task evaluates how effectively a model adapts to our newly defined, schema-constrained templates for generating fine-grained procedural KGs; second, the QA system strictly measures performance in formulating responses strictly grounded within structured evidence and local constraints.

Beyond standard QA systems, BioPIE can serve as a structured framework to analyze and decompose complex human instructions in the biomedical domain. By formalizing natural language protocols into structured representations, the dataset helps analyze intricate experimental logic and facilitates intent understanding.

As a structured human-robot interface, BioPIE can assist in bridging the gap between human-readable protocols and robotic execution scripts. It facilitates basic workflow verification, such as local parameter consistency and constraint checking. Furthermore, the dataset offers a reusable semantic reference that can support modular protocol composition and conditional adaptation. Consequently, BioPIE provides a valuable decision-making reference for LLM-based planners (*e.g.*, ReAct (Yao et al., 2023)), contributing to more reliable integration between AI assistants and robotic execution systems in automated laboratories (see Appx. B.5 for detailed discussion).

6 Conclusion

In this work, we investigate the problem of biomedical experiment QA from the perspective

of structured procedural understanding. We introduce BioPIE, a new IE dataset designed to capture fine-grained experimental entities, actions, and procedural relations while maintaining sufficient breadth to generalize across biomedical research. A comprehensive benchmark on BioPIE indicates existing supervised models and LLMs face challenges in protocol-centric IE, particularly with OOD protocols. A QA system is proposed to evaluate the QA performance enhancement with BioPIE. Both the QA evaluation results and ablation studies highlight the crucial role of BioPIE in complex reasoning, including HID and MSR, of biomedical experiment protocols.

Limitations

Despite our efforts, constructing a gold-standard dataset for IE over biomedical experimental protocols remains challenging. One limitation arises from our use of the LLMs for protocol text normalization. While normalization improves consistency, it may introduce misalignments in step references, *e.g.*, references to the product of an earlier step may be shifted to a later step after normalization. Such errors can affect fine-grained step-level grounding and temporal dependency annotation.

⚠ Warning. Reproducing the biomedical experiments described in BioPIE **must only be carried out under the direct supervision of qualified domain experts**, as many procedures **involve significant safety hazards** and may **pose serious risks to personnel, equipment, and the environment** if performed improperly. The biomedical protocols provided are strictly for reference purposes only and are not intended to serve as standalone or executable experimental instructions. This is consistent with their presentation in the original publications.

Ethical Statement

The original natural language descriptions are sourced from three websites, including Nature¹, Cell², and JoVE³. We further performed data cleaning and annotation on these descriptions. We carefully ensured that all experimental protocols incorporated into our corpus strictly comply with open-access policies and are distributed under Creative Commons licenses. This guarantees full adherence to copyright and intellectual property regulations, without any infringement or

unauthorized use of protected materials.

Reproducibility

Both the BioPIE and QA datasets are available at <https://sites.google.com/view/biopie>.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 52475001). The authors would like to thank Linkerbot Co., Ltd. for providing the dexterous robotic hand used in this study. The authors also thank Yu-Zhe Shi for assistance with early-stage data collection and Jiawen Liu for helpful discussions related to the figures.

References

- Vaibhav Adlakha, Parishad BehnamGhader, Xing Han Lu, Nicholas Meade, and Siva Reddy. 2024. [Evaluating correctness and faithfulness of instruction-following models for question answering](#). *Transactions of the Association for Computational Linguistics*, 12:681–699.
- Vahan Arsenyan, Spartak Bughdaryan, Fadi Shaya, Kent Wilson Small, and Davit Shahnazaryan. 2024. [Large language models for biomedical knowledge graph construction: Information extraction from EMR notes](#). In *Proceedings of the 23rd Workshop on Biomedical Natural Language Processing*, pages 295–317, Bangkok, Thailand. Association for Computational Linguistics.
- Anita Bandrowski, Ryan Brinkman, Mathias Brochhausen, Matthew H. Brush, Bill Bug, Marcus C. Chibucos, Kevin Clancy, Mélanie Courtot, Dirk Derom, Michel Dumontier, Liju Fan, Jennifer Foster, Gilberto Frago, Frank Gibson, Alejandra Gonzalez-Beltran, Melissa A. Haendel, Yongqun He, Mervi Heiskanen, Tina Hernandez-Boussard, and 25 others. 2016. [The ontology for biomedical investigations](#). *PLOS ONE*, 11(4):1–19.
- Iz Beltagy, Kyle Lo, and Arman Cohan. 2019. [SciBERT: A pretrained language model for scientific text](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3615–3620, Hong Kong, China. Association for Computational Linguistics.
- Benjamin Burger, Phillip M. Maffettone, Vladimir V. Gusev, Catherine M. Aitchison, Yang Bai, Xiaoyan Wang, Xiaobo Li, Ben M. Alston, Buyi Li, Rob Clowes, Nicola Rankin, Brandon Harris, Reiner Sebastian Sprick, and Andrew I. Cooper. 2020. [A mobile robotic chemist](#). *Nature*, 583(7815):237–241.
- Yupeng Chang, Xu Wang, Jindong Wang, Yuan Wu, Linyi Yang, Kaijie Zhu, Hao Chen, Xiaoyuan Yi, Cunxiang Wang, Yidong Wang, Wei Ye, Yue Zhang, Yi Chang, Philip S. Yu, Qiang Yang, and Xing Xie. 2024. [A survey on evaluation of large language models](#). *ACM Trans. Intell. Syst. Technol.*, 15(3).
- Mark Davies and Joseph L Fleiss. 1982. [Measuring agreement for multinomial data](#). *Biometrics*, pages 1047–1051.
- George Doddington, Alexis Mitchell, Mark Przybocki, Lance Ramshaw, Stephanie Strassel, and Ralph Weischedel. 2004. [The automatic content extraction \(ACE\) program – tasks, data, and evaluation](#). In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC’04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Jingyuan Ma, Rui Li, Heming Xia, Jingjing Xu, Zhiyong Wu, Baobao Chang, Xu Sun, Lei Li, and Zhifang Sui. 2024. [A survey on in-context learning](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 1107–1128, Miami, Florida, USA. Association for Computational Linguistics.
- Tianqing Fang, Zeming Chen, Yangqiu Song, and Antoine Bosselut. 2024. [Complex reasoning over logical queries on commonsense knowledge graphs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11365–11384, Bangkok, Thailand. Association for Computational Linguistics.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2022. [Language-agnostic BERT sentence embedding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 878–891, Dublin, Ireland. Association for Computational Linguistics.
- Giacomo Frisoni, Gianluca Moro, and Lorenzo Balzani. 2022. [Text-to-text extraction and verbalization of biomedical event graphs](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 2692–2710, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- María Herrero-Zazo, Isabel Segura-Bedmar, Paloma Martínez, and Thierry Declercq. 2013. [The ddi corpus: An annotated corpus with pharmacological substances and drug–drug interactions](#). *Journal of Biomedical Informatics*, 46(5):914–920.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and

¹<https://www.nature.com/>

²<https://star-protocols.cell.com/>

³<https://www.jove.com/>

- Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Yuntong Hu, Zhihan Lei, Zheng Zhang, Bo Pan, Chen Ling, and Liang Zhao. 2025. [GRAG: Graph retrieval-augmented generation](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 4145–4157, Albuquerque, New Mexico. Association for Computational Linguistics.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. [PubMedQA: A dataset for biomedical research question answering](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Martin Krallinger, Obdulia Rabal, Antonio Miranda-Escalada, and Alfonso Valencia. 2021. [Drugprot corpus: biocreative vii track 1-text mining drug and chemical-protein interactions](#). *Zenodo*, Jun, 29.
- Jens Kringelum, Sonny Kim Kjaerulff, Søren Brunak, Ole Lund, Tudor I. Oprea, and Olivier Taboureau. 2016. [Chemprot-3.0: a global chemical biology diseases mapping](#). *Database*, 2016:bav123.
- Anastasia Krithara, Anastasios Nentidis, Konstantinos Bougiatiotis, and Georgios Paliouras. 2023. [BioASQ-QA: A manually curated corpus for biomedical question answering](#). *Scientific Data*, 10(1):170.
- Benno Kruit, Yiming Xu, and Jan-Christoph Kalo. 2024. [Retrieval-based question answering with passage expansion using a knowledge graph](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 14063–14072, Torino, Italia. ELRA and ICCL.
- Po-Ting Lai, Elisabeth Coudert, Lucila Aimò, Kristian Axelsen, Lionel Breuza, Edouard de Castro, Marc Feuermann, Anne Morgat, Lucille Pourcel, Ivo Pedruzzi, Sylvain Poux, Nicole Redaschi, Catherine Rivoire, Anastasia Sveshnikova, Chih-Hsuan Wei, Robert Leaman, Ling Luo, Zhiyong Lu, and Alan Bridge. 2024. [EnzChemRED, a rich enzyme chemistry relation extraction dataset](#). *Scientific Data*, 11(1):982.
- Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang. 2020. [Biobert: a pre-trained biomedical language representation model for biomedical text mining](#). *Bioinformatics*, 36(4):1234–1240.
- Guowei Li, Luciana P. F. Abbade, Ikunna Nwosu, Yanling Jin, Alvin Leenus, Muhammad Maaz, Mei Wang, Meha Bhatt, Laura Zielinski, Nitika Sanger, Bianca Bantoto, Candice Luo, Ieta Shams, Hamnah Shahid, Yaping Chang, Guangwen Sun, Lawrence Mbuagbaw, Zainab Samaan, Mitchell A. H. Levine, and 2 others. 2018. [A systematic review of comparisons between protocols or registrations and full reports in primary biomedical research](#). *BMC Medical Research Methodology*, 18(1):9.
- Yuyang Liu, Liuzhenghao Lv, Xiancheng Zhang, Jingya Wang, Li Yuan, and Yonghong Tian. 2026. [Bioprobench: A corpus and benchmark for biological protocol reasoning in autonomous science](#). In *Forty-third International Conference on Machine Learning*.
- Yi Luan, Luheng He, Mari Ostendorf, and Hannaneh Hajishirzi. 2018. [Multi-task identification of entities, relations, and coreference for scientific knowledge graph construction](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3219–3232, Brussels, Belgium. Association for Computational Linguistics.
- Ling Luo, Po-Ting Lai, Chih-Hsuan Wei, Cecilia N Arighi, and Zhiyong Lu. 2022. [BioRED: a rich biomedical relation extraction dataset](#). *Briefings in Bioinformatics*, 23(5):bbac282.
- Linhao Luo, Yuan-Fang Li, Reza Haffari, and Shirui Pan. 2024. [Reasoning on graphs: Faithful and interpretable large language model reasoning](#). In *International Conference on Learning Representations*, volume 2024, pages 14400–14423.
- James Malone, Ele Holloway, Tomasz Adamusiak, Misha Kapushesky, Jie Zheng, Nikolay Kolesnikov, Anna Zhukova, Alvis Brazma, and Helen Parkinson. 2010. [Modeling sample variables with an experimental factor ontology](#). *Bioinformatics*, 26(8):1112–1118.
- S. Hessam M. Mehr, Matthew Craven, Artem I. Leonov, Graham Keenan, and Leroy Cronin. 2020. [A universal system for digitization and automatic execution of the chemical synthesis literature](#). *Science*, 370(6512):101–108.
- Zara Nasar, Syed Waqar Jaffry, and Muhammad Kamran Malik. 2021. [Named entity recognition and relation extraction: State-of-the-art](#). *ACM Comput. Surv.*, 54(1).
- Humza Naveed, Asad Ullah Khan, Shi Qiu, Muhammad Saqib, Saeed Anwar, Muhammad Usman, Naveed Akhtar, Nick Barnes, and Ajmal Mian. 2025. [A comprehensive overview of large language models](#). *ACM Transactions on Intelligent Systems and Technology*, 16(5):1–72.
- Le Peng, Gaoxiang Luo, Sicheng Zhou, Jiandong Chen, Ziyue Xu, Ju Sun, and Rui Zhang. 2024. [An in-depth evaluation of federated learning on biomedical natural language processing for information extraction](#). *NPJ Digital Medicine*, 7(1):127.
- Nadeesha Perera, Matthias Dehmer, and Frank Emmert-Streib. 2020. [Named entity recognition and](#)

- relation detection for biomedical information extraction. *Frontiers in cell and developmental biology*, 8:673.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: BM25 and beyond](#). *Foundations and Trends in Information Retrieval*, 4(1-2):1–174.
- Simon Rohrbach, Mindaugas Šiaučiulis, Greig Chisholm, Petrisor-Alin Pirvan, Michael Saleeb, S. Hessam M. Mehr, Ekaterina Trushina, Artem I. Leonov, Graham Keenan, Aamir Khan, Alexander Hammer, and Leroy Cronin. 2022. [Digitization and validation of a chemical synthesis literature database in the ChemPU](#). *Science*, 377(6602):172–180.
- Yu-Zhe Shi, Haofei Hou, Zhangqian Bi, Fanxu Meng, Xiang Wei, Lecheng Ruan, and Qining Wang. 2024. [AutoDSL: Automated domain-specific language design for structural representation of procedures with constraints](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12177–12214, Bangkok, Thailand. Association for Computational Linguistics.
- Yu-Zhe Shi, Mingchen Liu, Fanxu Meng, Qiao Xu, Zhangqian Bi, Kun He, Lecheng Ruan, and Qining Wang. 2025. [Hierarchically encapsulated representation for protocol design in self-driving labs](#). In *The Thirteenth International Conference on Learning Representations*.
- Yu-Zhe Shi, Qiao Xu, Yanjia Li, Mingchen Liu, Huamin Qu, Lecheng Ruan, and Qining Wang. 2026. [Automated constraint specification for job scheduling by regulating generative model with domain-specific representation](#). *IEEE Transactions on Automation Science and Engineering*, 23:3683–3698.
- Larisa N Soldatova, Daniel Nadis, Ross D King, Piyali S Basu, Emma Haddi, Véronique Baumlé, Nigel J Saunders, Wolfgang Marwan, and Brian B Rudkin. 2014. [EXACT2: the semantics of biomedical protocols](#). *BMC Bioinformatics*, 15(Suppl 14):S5.
- Sebastian Steiner, Jakob Wolf, Stefan Glatzel, Anna Andreou, Jarosław M. Granda, Graham Keenan, Trevor Hinkley, Gerardo Aragon-Camarasa, Philip J. Kitson, Davide Angelone, and Leroy Cronin. 2019. [Organic synthesis in a modular robotic system driven by a chemical programming language](#). *Science*, 363(6423):eaav2211.
- Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun’ichi Tsujii. 2012. [brat: a web-based tool for NLP-assisted text annotation](#). In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 102–107, Avignon, France. Association for Computational Linguistics.
- Jiashuo Sun, Chengjin Xu, Lumingyuan Tang, Saizhuo Wang, Chen Lin, Yeyun Gong, Lionel Ni, Heung-Yeung Shum, and Jian Guo. 2024. [Think-on-graph: Deep and responsible reasoning of large language model on knowledge graph](#). In *International Conference on Learning Representations*, volume 2024, pages 3868–3898.
- Nathan J. Szymanski, Bernardus Rendy, Yuxing Fei, Rishi E. Kumar, Tanjin He, David Milsted, Matthew J. McDermott, Max Gallant, Ekin Dogus Cubuk, Amil Merchant, Haegyeom Kim, Anubhav Jain, Christopher J. Bartel, Kristin Persson, Yan Zeng, and Gerbrand Ceder. 2023. [An autonomous laboratory for the accelerated synthesis of inorganic materials](#). *Nature*, 624(7990):86–91.
- Cunxiang Wang, Sirui Cheng, Qipeng Guo, Yuanhao Yue, Bowen Ding, Zhikun Xu, Yidong Wang, Xiangkun Hu, Zheng Zhang, and Yue Zhang. 2023. [Evaluating open-QA evaluation](#). In *Advances in Neural Information Processing Systems*, volume 36, pages 77013–77042. Curran Associates, Inc.
- Xiang Wei, Xingyu Cui, Ning Cheng, Xiaobin Wang, Xin Zhang, Shen Huang, Pengjun Xie, Jinan Xu, Yufeng Chen, Meishan Zhang, Yong Jiang, and Wenjuan Han. 2023. [ChatIE: Zero-shot information extraction via chatting with ChatGPT](#). *arXiv preprint arXiv:2302.10205*.
- Zhaohui Yan, Songlin Yang, Wei Liu, and Kewei Tu. 2023. [Joint entity and relation extraction with span pruning and hypergraph neural networks](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7512–7526, Singapore. Association for Computational Linguistics.
- Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. [ReAct: Synergizing reasoning and acting in language models](#). In *The Eleventh International Conference on Learning Representations*.
- Deming Ye, Yankai Lin, Peng Li, and Maosong Sun. 2022. [Packed levitated marker for entity and relation extraction](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4904–4917, Dublin, Ireland. Association for Computational Linguistics.
- Dongxu Zhang, Sunil Mohan, Michaela Torkar, and Andrew McCallum. 2022. [A distant supervision corpus for extracting biomedical relationships between chemicals, diseases and genes](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 1073–1082, Marseille, France. European Language Resources Association.
- Qi Zhang, Zhijia Chen, Huitong Pan, Cornelia Caragea, Longin Jan Latecki, and Eduard Dragut. 2024. [SciER: An entity and relation extraction dataset for datasets, methods, and tasks in scientific documents](#).

In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 13083–13100, Miami, Florida, USA. Association for Computational Linguistics.

Xiaoyan Zhao, Yang Deng, Min Yang, Lingzhi Wang, Rui Zhang, Hong Cheng, Wai Lam, Ying Shen, and Ruifeng Xu. 2024. [A comprehensive survey on relation extraction: Recent advances and new frontiers](#). *ACM Comput. Surv.*, 56(11).

A Details of BioPIE

A.1 Development of the Annotation Scheme

A.1.1 Scope of the Proposed Dataset

Specifically, BioPIE is bound to foundational laboratory explorations, rather than clinical cohort studies or medical interventions targeting human patients, with a particular emphasis on mainstream experimental setups and methodologies established between 2020 and 2025. The dataset is tailored for wet-lab workflows that involve physical interactions among reagents and materials, excluding bioinformatics data analyses (dry-lab). Accordingly, the experimental subjects represented in our dataset encompass cells, tissues, animal models, plants, and microorganisms—spanning both in vivo and in vitro paradigms for living systems—while excluding human subject groups. By concentrating our source corpus on operation-intensive branches such as genetic manipulation, molecular interactions, physiological signal detection, and regenerative medicine, BioPIE models fine-grained laboratory procedures; however, while it aims to achieve broad applicability, generalizability across all peripheral biomedical disciplines cannot be guaranteed.

A.1.2 Design Process

Our scheme was developed through an iterative, corpus-driven process. We began with pilot annotations on a representative subset of experimental protocols, then refined entity and relation types through multiple rounds of adjudication between the annotators and domain experts. Each revision was motivated by observed annotation conflicts or coverage gaps rather than arbitrary design choices. Some entity or relation types were merged or split based on inter-annotator disagreement analysis, which is standard practice in corpus annotation methodology.

A.1.3 Comparison with Existing Biomedical Ontologies

The scheme was initially designed with reference to established formal ontologies for experimental procedures, namely EXACT2 (Soldatova et al., 2014) and OBI (Bandrowski et al., 2016). However, this alignment remains descriptive and ontology-inspired rather than operationally ontology-grounded. While EXACT2 and OBI provide machine-readable semantic representations of experimental procedures, our scheme is a text-span annotation scheme intended for information extraction from free-text protocols.

The correspondences are as follows:

- EXACT2’s action participants (reagents, materials, instruments) correspond to our entity types `chemical`, `biomaterial`, `protein`, `nucleic acid`, `cell`, *etc.*
- EXACT2’s quantitative parameter slots map directly to our `temperature`, `time`, `volume`, `concentration`, `speed`, `force`, `mass`, and `length` entity types.
- OBI’s material entity hierarchy (`organism`, `anatomical part`, `cell line`) informs our `organ`, `animal`, `plant`, and `cell` types.
- At the relation level, EXACT2’s `has_object`, `has_instrument`, `has_reagent`, and `has_parameter` slots correspond to our `is_object_of`, `use_device`, `use_reagent`, and `have_parameter` relations, respectively. Action sequencing in EXACT2 maps to our `next_step` relation.

We additionally introduce entity types (`software`, `data`, `position`, `blend`) and relation types (`during`, `for_each`, `based_on`, `in_condition_of`, `is_goal_of`, *etc.*) not prominently covered by EXACT2 or OBI, motivated by patterns frequently observed in our corpus. The scheme also underwent data-driven iterative refinement during annotation, some types were merged or split based on inter-annotator disagreement analysis, which is standard practice in corpus annotation.

A.1.4 Recovering Investigation-level Information from Fine-grained Procedural Annotations

BioPIE is designed as a procedure-centric information-extraction scheme rather than a

Design / Focus	Original Protocol Evidence	Annotation Evidence (Examples)	Investigation-level Interpretation
Dose-control-readout design	- “Apply physical dose levels of 1 Gy and 0.5 Gy, resulting in expected carbon ion beam fluences of $8.43 \times 10^6 \text{ cm}^{-2}$ and $4.16 \times 10^6 \text{ cm}^{-2}$, respectively.” - “Place potential stopping material (e.g., PMMA)...” - “Shield control well positions with additional stopping material” - “Ensure identical positions are used for correlation...” - “Analyze RIF for single cells over time” - “Quantify RIF number, size, and signal intensity.”	["physical dose levels", "have_property", "1 Gy"] ["Shield", "apply_material", "additional stopping material"] ["RIF number", "is_object_of", "Quantify"]	Supports the reconstruction of two irradiation conditions, shielded control positions, dose-related variables, matched measurement settings, and single-cell RIF readouts associated with irradiation conditions.
Donor-to-recipient provenance	- “Isolate spleens from CTL-donor mice” - “Generate a single cell suspension by extruding the spleen...” - “Inject 200 μL (2.5×10^6 cells) intravenously into mice...” - “Inject antigen with or without adjuvant intravenously into mice 16 hours after adoptive T cell transfer” - “Isolate spleens from injected mice after 10 hours.” - “Determine the number of labeled cells...”	["Isolate", "from", "CTL-donor mice"] ["adoptive T cell transfer", "next_step", "16 hours"] ["number of labeled cells", "is_object_of", "Determine"]	Maps the provenance chain (CTL-donor → spleen → adoptive transfer → recipient spleen), compared treatment conditions, sampling times, and quantitative readouts.
Genetic perturbations & mechanisms	- “Electrophoresis lab-purified recombinant proteins MITOL (WT and CD) in SDS-PAGE” - “Determine the exact concentrations of MITOL WT and MITOL CD...” - “Add 1 μg Flag-PolγA, 1 μg Myc-MITOL WT, 1 μg MITOL CD, and 100 ng His-Ubiquitin plasmids...” - “Analyze ubiquitylation of substrate protein by Western blot...” - “Determine that equal amounts of PolγA... were used”	["recombinant proteins MITOL", "contain", "WT"] ["Analyze", "use_method", "Western blot"] ["exact concentrations...", "is_object_of", "Determine"]	Preserves manipulated MITOL catalytic states, alternative Pol γ A substrate variants, calibrated protein inputs, an equal-input check, and the ubiquitylation readout.
Animal treatment & tissue lineage	- “Adjust the airflow to deliver a mixture of 2.3% sevoflurane...” - “Continuously deliver sevoflurane, medical air, and oxygen for 2 hours.” - “Quickly extract the brain and keep it in ice-cold carbonated aCSF” - “Cut coronal slices...” - “Carefully remove hippocampal slices from the coronal slices...” - “Load 10 μL of each protein fraction in the same gel for comparison by Western blotting”	["mixture", "contain", "sevoflurane"] ["hippocampal slices", "is_object_of", "remove"] ["comparison", "use_method", "Western blotting"]	Connects animal-level anesthetic exposure to the lineage (treated mouse → brain → coronal slices → hippocampal slices) and matched protein fraction comparisons.
Comparative methods & statistics	- “Blend cells into the bioink using the passive mixing unit technique 1, 2, or 3 times...” - “Blend cells into a separate bioink using manual mechanical mixing...” - “Isolate and cryopreserve primary human nasal chondrocytes (hNC) from patients” - “Calculate cell viability based on the ratio of live cells...” - “Analyze the data via one-way ANOVA followed by Tukey’s multiple comparisons test”	["Blend", "repetitions", "1"] ["primary human nasal chondrocytes", "from", "patients"] ["Analyze", "use_method", "one-way ANOVA"]	Exposes comparison groups by mixing method, patient-derived cell provenance, quantitative definition for cell viability, inferential analysis, and longitudinal endpoints.

Table A1: Original protocol sentences and selected relation annotations supporting the corresponding investigation-level interpretations.

complete investigation ontology. Nevertheless, this does not mean that investigation-level information is absent from the annotations. When such information is explicitly stated in a protocol, it can often be compositionally recovered from the serialized local steps, because BioPIE jointly preserves actions, biological and non-biological participants, quantitative parameters, typed semantic relations, and procedural order. In particular, alternative materials, dosages, variants, and processing conditions provide evidence for experimental designs and treatment/control arms; *from*, *to*, and *step-sequencing* relations preserve biosample trajectories; measurement actions and data-bearing entities identify outcome variables and analysis procedures; and relations such as *based_on*, *is_goal_of*, *during*, and *next_step* connect interventions, observations, and downstream readouts.

Tab. A1 quotes the original protocol sentences in the dataset and show selected relation annotations that support the corresponding investigation-level interpretations.

These examples clarify the intended scope of BioPIE. The scheme directly represents protocol actions, participants, instruments, quantities, conditions, local semantic relations, and procedural dependencies. From these annotations, treatment and control conditions, biosample trajectories, manipulated variables, measurement variables, time points, comparison structures, and intervention–readout evidence chains can be reconstructed when they are explicitly described in the source text. However, BioPIE does not normalize these elements into dedicated global objects such as *ExperimentalDesign*, *ControlArm*, *Outcome*, or *Mechanism*, nor does it aim to encode complete hypotheses, cohort-level metadata, randomization or blinding procedures, raw measurement matrices, statistical results, causal conclusions, or formal biological-mechanism graphs unless these are directly stated as procedural content.

Accordingly, BioPIE should be understood as a procedure-grounded intermediate IE layer that is complementary to OBI and EXACT2. Its dense text-span annotations provide grounded entities and relations that can support downstream aggregation into investigation-level representations, whereas formal ontologies provide the global conceptual structures that BioPIE intentionally does not duplicate. We therefore use “biomedical exper-

iment understanding” to refer specifically to understanding how biomedical experiments are operationally executed, conditioned, sampled, and measured, rather than to a complete formal representation of every dimension of a biomedical investigation.

A.1.5 Complementarity with the Experimental Factor Ontology

The Experimental Factor Ontology (EFO) provides standardized concepts for experimental variables, such as biological materials, chemical treatments, assays, and measured phenotypes (Mallone et al., 2010). BioPIE is complementary to EFO: EFO supports concept normalization, while BioPIE captures how these variables are used in protocol steps through actions, parameters, semantic relations, and procedural order.

- **Paired treatment and vehicle control.**

Original protocol text: “Topically apply PMA (100 μ M, 20 μ L) to the left ear of each mouse.” “Topically apply vehicle (ethanol, 20 μ L) to the right ear of each mouse.” “Measure the photon fluxes in each ROI.”

Complementarity: EFO-compatible concepts can normalize PMA, ethanol, ear anatomy, and photon flux. BioPIE additionally links each material to its application action, dose, volume, anatomical site, and downstream measurement. These relations support reconstruction of the within-animal treatment–control design.

- **Cell type and electrophysiological measurement conditions.**

Original protocol text: “Load 1 mL of Fluo-4-AM working solution (2 μ M) to human DA neurons.” “Hold the cell at -60 mV in voltage-clamp configuration.” “Record Ca^{2+} currents by stepping from -60 mV to $+50$ mV for 250 ms in 10 mV increments.”

Complementarity: EFO and related ontologies can normalize the cell type, chemical probe, assay, and measured variable. BioPIE preserves the loading concentration and volume, target cells, holding potential, voltage range, pulse duration, increment, and Ca^{2+} -current readout.

- **Stage-specific differentiation factors.**

Source	Train	Dev	ID Test	OOD Test
Cell	115	16	32	0
Nature	82	9	20	17
JoVE	43	11	16	28

Table A2: **Distribution and source journals of the BioPIE protocols.**

	Train	Val	ID	OOD
Genetic Manipulation	60	14	28	0
Molecular Interactions	73	4	13	0
Physiological Signal Detection	55	8	15	0
Regenerative Medicine	52	10	12	0
Plant-based Biopharmaceuticals	0	0	0	28
Functional Imaging	0	0	0	10
Epidemiological Analysis	0	0	0	7

Table A3: **Domain distribution and dataset splits in BioPIE protocols.** The dataset includes four primary domains for training and ID testing, and three specialized domains reserved exclusively for OOD evaluation to test model robustness across diverse biomedical fields.

Original protocol text: “When human iPS cell colonies reach 80% confluence, DE induction is performed.” “Mix RPMI medium/Glutamax with Activin A (100 ng/mL) and Wnt 3a (25 ng/mL).” “Mix DMEM/F12 basal medium with Noggin (200 ng/mL) and SB-431542 (10 μ M).” “Mix DMEM/F12 basal medium with FGF2 (250 ng/mL), FGF7 (100 ng/mL), and FGF10 (100 ng/mL).”

Complementarity: EFO-compatible terms can normalize the cell type, differentiation process, growth factors, and inhibitors. BioPIE additionally preserves the induction condition, exact concentrations, factor combinations, and their positions in successive differentiation stages.

These examples show that EFO and BioPIE operate at complementary levels. EFO provides standardized identifiers and concept hierarchies, whereas BioPIE provides text-grounded procedural context, including actions, targets, parameters, and step order.

A.2 Data Source

To construct BioPIE, we curated a corpus of biological protocols from three primary high-quality sources: Cell, Nature, and JoVE, as shown in Tab. A2. We targeted protocols published between 2020 and 2025. Our sampling strategy was designed to cover a wide range of biological sub-domains, as shown in Tab. A3.

A.3 The Impact of Normalization on Protocols

A.3.1 Qualitative Analyses

The raw protocol text sourced from Nature Protocols is often written in a mixed style—containing compound sentences, parenthetical reagent details, conditional instructions, passive constructions, and non-actionable commentary. The normalization step, powered by Qwen-max, decomposes and rewrites this text into discrete, imperative, subject-free sentences, each representing a single atomic action. This standardized format is essential for downstream annotation tasks such as action-type labeling and argument extraction.

Example 1: Confocal Microscopy

Raw text:

"Acquire Z-sections using a custom-made spinning disk confocal microscope as described above."

Normalized output:

"Acquire Z-section images using a custom-made spinning disk confocal microscope."

Here the normalization removes the vague cross-reference ("as described above") while preserving the core imperative action, making the step self-contained and annotation-ready.

Example 2: Haplotype-Fusion PCR

Raw text:

"Transfer completed PCR reactions to a 1.5 ml tube and add 400 μ l hexane. Vortex thoroughly and centrifuge at full speed for 3 minutes. Remove and discard hexane layer (upper layer). Repeat with 400 μ l fresh hexane."

Normalized output (4 discrete steps):

"Transfer completed PCR reactions to a 1.5 ml tube and add 400 μ l hexane."
 "Vortex thoroughly and centrifuge at full speed for 3 minutes."
 "Remove and discard the upper hexane layer."
 "Repeat with 400 μ l fresh hexane."

This example illustrates the most common normalization operation: splitting compound sentences joined by conjunctions into the separate atomic steps. The parenthetical clarification "(upper layer)" is also integrated naturally into the imperative sentence rather than left as a parenthetical.

Example 3: Gap1 Integrative Vector

Raw text:

"Amplify a pUC fragment containing the bacterial origin of replication and the kanamycin resistance marker by PCR from PCR-BluntII-TOPO (Invitrogen, Carlsbad, CA) and ligate with a PvuII fragment of pDRf1 containing the fl origin of replication, the PMA1 promoter and ADH3 terminator producing pDL001."

Normalized output (2 discrete steps):

"Amplify a pUC fragment containing the bacterial origin of replication and the kanamycin resistance marker by PCR from PCR-BluntII-TOPO."

"Ligate the amplified pUC fragment with a PvuII fragment of pDRf1 to produce pDL001."

Here, the normalization does two things simultaneously: it separates two sequential actions (amplify, then ligate) that were conjoined in one sentence, and it removes the vendor information ("Invitrogen, Carlsbad, CA") which is not part of the procedural action itself. The second normalized step also adds a clarifying reference ("the amplified pUC fragment") to preserve logical coherence between steps.

Example 4: Simplified Behavioural Tests

Raw text:

"Begin the first trial by introducing a stimulus female into the home cage of a male mouse for a 1-min interaction. Record investigation time of the female by the male mouse . At the end of the 1-min trial, remove the stimulus animal and return to an individual holding cage . Score behaviour during experiments and record behaviour on a videotape . Scoring should be done by trained raters (inter-rater reliability >90%)."

Normalized output:

"Introduce a stimulus female into the home cage of a male mouse for a 1-minute interaction."

"Remove the stimulus animal after the 1-minute trial and return the male to an individual holding cage."

"Score behavior during the experiment and record on videotape."

This example demonstrates that normalization also filters out non-procedural meta-commentary

	Action Verbs	Descriptive Verbs
Original	135	52
Normalized	125	11
Retention Rate	92.59%	21.15%

Table A4: **Quantitative comparison of verb distribution between the original and normalized protocols.**

	Test	HID	MSR
Original	60.90	68.26	53.66
Normalized	61.52	66.09	56.10

Table A5: **QA performance evaluation for text normalization.**

("Scoring should be done by trained raters with inter-rater reliability >90%") that describes quality criteria rather than an executable action, keeping only genuinely actionable steps.

A.3.2 Quantitative Analyses

To verify that the transformation of raw protocols into normalized, stepwise imperative sentences preserves the integrity of experimental information, we evaluate the normalization process from two perspectives: content fidelity and semantic consistency.

Quantitative Impact on Information Content. We randomly sampled 102 pairs (original vs. normalized) and manually annotated the occurrence of two functional verb types: (1) action verbs, which denote executable experimental operations (*e.g.*, centrifuge, incubate, add); (2) descriptive verbs, which primarily explain experimental principles, background, or non-procedural context (*e.g.*, is, are, recommend).

The results (Table A4) indicate that the normalized text retains 92.59% of the action verbs, ensuring that the core procedural logic remains intact. Manual inspection reveals that this marginal reduction primarily stems from the consolidation of redundant operations and the exclusion of verbs (*e.g.*, ensure, check) that do not alter the core experimental execution. Conversely, there is a significant reduction in descriptive verbs, suggesting that the normalization process primarily impacts descriptive content. By filtering out redundancy, the process effectively distills the text while maintaining the essential operational density required for precise reasoning. Our analysis also shows a 98.08% retention rate for numerical parameters and reagents, suggesting that safety-critical facts are preserved.

Impact on Semantic Representation. To fur-

ther assess whether normalization preserves the core operational logic, we conducted a comparative analysis using our QA dataset. We employed a text embedding model OpenAI’s *text-embedding-3-large* as a retriever to construct a vector database from both the raw and the normalized corpora.

Experimental results demonstrate that the QA system leveraging normalized text maintains and exceeds the performance of the raw text baseline in several cases (Tab. A5). This leads to two conclusions:

The imperative structure used in normalization maintains a high degree of procedural fidelity to the original descriptive prose, allowing embedding models to capture core experimental facts accurately.

By eliminating linguistic noise and standardizing sentence structures, normalization enhances the signal-to-noise ratio. This assists the retriever in identifying precise experimental parameters (*e.g.*, specific temperatures or durations), thereby providing cleaner evidence for HID and MSR reasoning. The slight decline on the HID set is not statistically significant ($p \geq 0.0625 > 0.05$), indicating minor variance rather than a meaningful loss of semantic information.

A.4 Annotation Guideline, Data Scheme Definition, and IE Prompt

This section provides the annotation guidelines for the proposed dataset, covering the data scheme definition and the procedures used for consistent annotation.

The prompt for LLM-based joint extraction is the guideline shown below, while the prompt for pipeline extraction is obtained by splitting the following prompt.

You are given a piece of text describing biomedical experiments or laboratory workflows.

Your task is to identify all factual entities and all relationships between these entities.

The possible entity types are listed below.

- verb: Actions performed in a procedure . (*e.g.*, Fix, Osmicate, Dehydrate)
- part: Specific sections of an object. (*e.g.*, upper surface of the specimen , plunger, plunger of the bioink syringe)

- container: Objects used to hold substances. (*e.g.*, original culture plate, cartridge, well plate)
- force: Physical force or weight applied. (*e.g.*, 500 g, 17,000 × g, 226 × g)
- device: Tools used in experiments. (*e.g.*, fume hood, aluminum stub, underlying aluminum stub)
- method: Techniques for conducting experiments. (*e.g.*, simultaneously, direct, trypan blue exclusion method)
- chemical: Substances used in a process excluding proteins and polymers (*e.g.*, TAG, Karnovsky, aqueous osmium tetroxide)
- concentration: Ratio of a substance in a solution. (*e.g.*, 1% (wt/vol), 50%, 70%)
- consumable: Materials used up in experiments. (*e.g.*, sticky sellotape tabs, copper tape, silver paint)
- state: Condition of a material or system. (*e.g.*, continuous contact, recorded, sterile)
- volume: Measurement of liquid quantity . (*e.g.*, volumes, 12 mL, 1 mL)
- temperature: Heat level in a process. (*e.g.*, room temperature, 4°C)
- time: Duration of an action or waiting . (*e.g.*, 2 hours, overnight)
- process: Series of actions in a procedure. (*e.g.*, air dry, cross-linking, additional blends)
- times: Number of repetitions. (*e.g.*, three times, two, 1)
- cell: Basic biological unit in living organisms. (*e.g.*, cell monolayers, samples, sample)
- nucleic acid: DNA or RNA sequences used in biological experiments. (*e.g.* ., genomic DNA, T7-RT primer, first-strand cDNA)
- biomaterial: Biological substances in use. (*e.g.*, bioink)
- software: Programs for analysis or instrument control. (*e.g.*, SmartSEM software, Nikon Imaging Software)
- number: Countable values in a process. (*e.g.*, two, total number of cells)
- energy: Measure of work or electrical energy. (*e.g.*, 3–5 KV, 400 mJ)
- speed: Rate of motion or process. (*e.g.* ., controlled rate, 20 rpm)

	In-domain				Out-of-domain			
	NER	Rel	Rel+	RE	NER	Rel	Rel+	RE
<i>Supervised Baselines</i>								
PL-Marker ¹ (Ye et al., 2022)	87.99	83.99	75.82	88.72	77.24	75.83	56.40	80.81
PL-Marker ² (Ye et al., 2022)	86.87	81.97	73.84	80.14	88.24	70.85	49.78	71.79
HGERE ¹ (Yan et al., 2023)	87.01	82.92	74.68	-	74.85	72.59	53.97	-
HGERE ² (Yan et al., 2023)	86.93	81.79	73.86	-	73.76	72.36	53.62	-
<i>Zero-shot LLM</i>								
GPT-5 (Pipeline)	47.61	47.83	39.00	64.81	41.96	47.98	35.06	65.14
GPT-5 (Joint)	45.89	61.74	48.70	-	50.28	61.09	42.53	-
Claude-4.5-Opus (Pipeline)	71.55	41.87	32.93	49.73	66.28	35.31	25.58	45.45
Claude-4.5-Opus (Joint)	71.31	35.11	27.23	-	67.76	32.55	25.77	-
Llama-4 (Pipeline)	31.86	7.14	5.46	16.67	32.91	8.82	6.86	13.33
Llama-4 (Joint)	58.62	23.08	0.00	-	54.55	28.57	14.29	-
Qwen-max (Pipeline)	70.28	22.82	18.93	29.69	64.77	20.90	15.05	27.61
Qwen-max (Joint)	69.34	26.39	20.83	-	66.64	22.29	15.64	-
<i>Few-shot LLM</i>								
GPT-5 (Pipeline)	51.12	62.89	56.20	75.80	42.33	51.15	39.23	62.92
GPT-5 (Joint)	54.20	65.50	57.95	-	49.55	63.06	51.11	-
Claude-4.5-Opus (Pipeline)	85.62	68.67	63.29	77.25	75.82	53.88	41.86	65.35
Claude-4.5-Opus (Joint)	83.89	66.01	60.38	-	72.28	53.11	41.75	-
Llama-4 (Pipeline)	35.81	18.82	16.83	21.12	30.10	10.69	8.88	14.81
Llama-4 (Joint)	53.57	32.09	30.70	-	57.09	26.40	20.00	-
Qwen-max (Pipeline)	85.14	65.67	61.11	73.52	74.94	48.63	38.34	57.86
Qwen-max (Joint)	84.68	61.80	57.73	-	72.85	43.72	34.36	-
<i>LoRA LLM</i>								
Llama-3-8B (Pipeline)	86.54	75.45	68.29	82.52	77.38	64.44	50.60	73.70
Llama-3-8B (Joint)	84.73	75.52	67.16	-	76.80	64.58	48.13	-
Qwen-3-8B (Pipeline)	84.70	70.65	63.31	78.14	77.23	64.07	48.67	69.15
Qwen-3-8B (Joint)	83.38	72.20	64.30	-	76.42	62.97	46.76	-

Table A6: **Test precision scores of different baselines on our proposed dataset.** ¹ Denotes the use of *scibert* (Beltagy et al., 2019) as the encoder; ² denotes *biobert* (Lee et al., 2020).

- mass: Quantity of matter. (e.g., final cell density, 2 μ g)
- environment: Conditions affecting an experiment. (e.g., dust-free environment, standard conditions)
- length: Measurement of distance. (e.g., approximately 1 nm, working distance of 4 mm)
- data: Recorded experimental information. (e.g., TIFF images, digital image files)
- organ: Biological structures in research. (e.g., spleen, spleens)
- animal: Living organisms in studies. (e.g., mice, CTL-donor mice)
- protein: Functional biomolecules. (e.g., trypsin/EDTA solution, BSA/PBS solution)
- polymer: Large molecular compounds. (e.g., nanocellulose/alginate, agarose gel)
- position: Spatial location of an

- object or material. (e.g., in the printed construct, on the dispensing unit)
- size: Dimensional magnitude of an object. (e.g., approximate size of the plate, 220 $\times$ 220)
- plant: Botanical specimens or components used in experiments. (e.g., red beet, spinach)
- blend: Mixed substances. (e.g., bioink-cell mixture, blend, cell/bioink)

The possible relation types are listed below.

- is_object_of: Describes that an object is the target of an action. (e.g., cell monolayers is_object_of Fix)
- contain: Indicates that something contains another thing. (e.g., Zeiss Sigma microscope contain in-lens SE1 electron detector)
- use_method: Specifies the method used

	In-domain				Out-of-domain			
	NER	Rel	Rel+	RE	NER	Rel	Rel+	RE
<i>Supervised Baselines</i>								
PL-Marker ¹ (Ye et al., 2022)	86.82	81.16	73.27	87.08	70.78	65.48	48.71	76.98
PL-Marker ² (Ye et al., 2022)	85.80	78.99	71.16	21.32	36.80	14.09	9.90	14.99
HGERE ¹ (Yan et al., 2023)	88.25	81.29	73.20	-	74.31	68.51	50.94	-
HGERE ² (Yan et al., 2023)	86.97	80.08	72.31	-	71.79	66.82	49.51	-
<i>Zero-shot LLM</i>								
GPT-5 (Pipeline)	71.44	53.41	43.54	75.75	68.62	55.47	40.53	71.65
GPT-5 (Joint)	15.26	13.56	10.69	-	14.03	12.84	8.94	-
Claude-4.5-Opus (Pipeline)	67.27	39.34	30.94	46.40	61.53	33.49	24.26	41.39
Claude-4.5-Opus (Joint)	67.53	32.34	25.08	-	63.84	30.16	23.88	-
Llama-4 (Pipeline)	57.78	1.08	0.83	0.89	59.37	0.86	0.67	0.38
Llama-4 (Joint)	0.88	0.38	0.00	-	0.46	0.19	0.10	-
Qwen-max (Pipeline)	64.23	18.65	15.47	25.40	56.28	15.98	11.51	22.93
Qwen-max (Joint)	62.47	22.34	17.63	-	56.82	18.17	12.75	-
<i>Few-shot LLM</i>								
GPT-5 (Pipeline)	81.19	71.29	63.72	83.13	70.16	57.18	43.86	75.07
GPT-5 (Joint)	18.61	15.47	13.69	-	25.75	21.60	17.51	-
Claude-4.5-Opus (Pipeline)	84.74	69.06	63.65	77.15	71.70	52.90	41.10	64.41
Claude-4.5-Opus (Joint)	82.94	65.75	60.15	-	71.78	52.05	40.91	-
Llama-4 (Pipeline)	77.84	18.65	16.68	20.94	65.38	12.94	10.75	17.98
Llama-4 (Joint)	7.73	4.39	4.20	-	10.87	6.28	4.76	-
Qwen-max (Pipeline)	81.49	63.08	58.69	72.82	68.93	44.05	34.73	55.66
Qwen-max (Joint)	68.40	50.35	47.04	-	54.20	32.45	25.50	-
<i>LoRA LLM</i>								
Llama-3-8B (Pipeline)	86.13	75.11	67.98	80.84	74.09	61.56	48.33	72.79
Llama-3-8B (Joint)	84.69	74.22	66.00	-	73.01	60.89	45.39	-
Qwen-3-8B (Pipeline)	84.18	69.26	62.06	76.00	72.71	59.37	45.10	68.03
Qwen-3-8B (Joint)	82.47	69.26	61.68	-	71.47	57.28	42.53	-

Table A7: **Test recall scores of different baselines on our proposed dataset.** ¹ Denotes the use of *scibert* (Beltagy et al., 2019) as the encoder; ² denotes *biobert* (Lee et al., 2020).

- for an action. (e.g., Dehydrate use_method incubating) - use_device: Specifies the device or tool used for an action. (e.g., air dry use_device fume hood) - use_reagent: Specifies the reagent or chemical used in an action. (e.g., Fix use_reagent TAG) - have_property: Describes a property of an object. (e.g., aqueous osmium tetroxide have_property 1% (wt/vol)) - apply_material: Specifies a material applied during an action. (e.g., stick apply_material sticky sellotape tabs) - is_goal_of: Describes that a goal is the purpose of an action. (e.g., make is_goal_of Use) - for_each: Specifies that an action applies to each specific object. (e.g., Place for_each sample) - next_step: Indicates the next step	- after an action or process. (e.g., 50% next_step 70%) - to: Container or position to which an object or solution is transferred. (e.g., stick to aluminum stub) - or: Represents alternative options. (e.g., TAG or Karnovsky) - have_parameter: Specifies an action's or process's parameter. (e.g., Fix have_parameter room temperature) - repetitions: Indicates the number of times an action is repeated. (e.g., Blend repetitions 1) - use_software: Specifies software used. (e.g., Acquire use_software SmartSEM software) - from: Indicates the source of something. (e.g., specimens from original culture plate) - in_condition_of: Specifies the condition under which an action occurs. (e.g., Acquire
--	--

- in_condition_of 3-5 KV)
- not: Denotes negation or exclusion. (e.g., Mix not cartridge)
- during: Indicates that an event happens within the time frame of another. (e.g., Balance during choosing)
- equal: Expresses equivalence between two values or objects. (e.g., one equal syringes)
- based_on: Indicates dependence or derivation from something. (e.g., Calculate based_on total number of constructs desired)

The following rules define the annotation standards for Named-Entity Recognition (NER) and Relation Extraction (RE) in this dataset. Annotators should strictly adhere to these guidelines to ensure consistency and reproducibility.

General Principles

1. All annotations should preserve the original surface form as it appears in the text, without normalization or correction.
2. When uncertainty exists, prioritize precision over recall and omit questionable annotations rather than guessing.

Named-Entity Recognition (NER)

3. For NER, annotate all entity mentions and output only entity category pairs, one per line, in the following format:

```
...
entity: category
...
```
4. The entity span must be minimal and precise. Do not include determiners or function words such as "the", "a", or "this" within the entity span.
5. When both a full name and its abbreviation appear in the text, annotate each occurrence separately as independent entities.
6. Annotate every occurrence of an entity in the text, even if the same entity appears multiple times.
7. If an entity mention is ambiguous, assign the category that is most directly supported by the local

context.

8. Overlapping or nested entity spans are permitted when they correspond to valid and distinct entity mentions.

Relation Extraction (RE)

9. For RE, annotate only explicitly stated or clearly implied relationships and output only relation triplets, one per line, in the following format:

```
...
head: head_entity tail: tail_entity
relation: relationship
...
```
10. Both the head and tail entities must be annotated entity mentions present in the text.
11. Do not infer, assume, or hallucinate relations that are not directly supported by the text.
12. If multiple relations are expressed between the same entity pair, annotate each relation separately.
13. If the same relation involves an entity that appears in multiple positions in the text (e.g., via pronouns, or alternative mentions), annotate the relation only for the most salient or primary occurrence of that entity.

A.5 Detailed IE Benchmark Results

Precision in the IE benchmark is listed in Tab. A6. Recall in the IE benchmark is listed in Tab. A7.

B QA System Evaluation

B.1 QA Dataset

We extracted 3275 sub-protocols from the complete collection of textual protocols (excluding those already structured as part of the IE dataset) and automatically constructed the corresponding QA pairs; the construction process is detailed in the Appx. B.2. The dataset is divided into training, validation, and test sets with sizes of 1983, 159, and 1133, respectively. The training set is used for model training, the validation set is utilized for model hyper-parameter selection (the number of few-shot examples), and the test set is employed to evaluate the performance of the model.

To further analyze model performance under challenging conditions, we construct two subsets

from the test set. The first subset consists of 230 HID questions, which are the questions in the test set with the top 230 highest relation counts. For these questions, the corresponding sentences from which they were generated contain an average of 10.40 relations, substantially higher than the overall average of 4.62 reported in Tab. 1. The second subset comprises 123 Multi-Step Reasoning (MSR) questions (the average number of reasoning hops is 2.22), which includes all questions in the test set requiring more than one reasoning step.

B.2 QA Dataset Construction

The BioPIE QA dataset was constructed automatically to validate the utility of the IE dataset in supporting complex reasoning tasks while maintaining evaluation fairness. The process followed three rigorous steps:

Step 1 KG-based Question Generation: We utilized the annotated KGs as logical skeletons. An LLM was used to generate natural language questions by identifying key entities and their target answers within the triplet structures.

Step 2 Reasoning Path Refinement: To ensure logical integrity, we screened the reasoning paths. If entities from the reasoning path appeared directly in the initial question, we adjusted the path or question to guarantee the correctness of the answer and the reasoning process.

Step 3 Human Verification: To ensure the quality, practical relevance, and independence of the validation, we implemented a comprehensive phase of human verification. Every single auto-generated question underwent manual screening by human annotators with expertise in biomedical fields. We filtered out any ambiguous, unreasonable, or erroneous content to ensure that the benchmark strictly aligns with the genuine needs, logic, and expectations of real-world biomedical specialists.

B.3 QA Showcases

To better illustrate the full complexity and reasoning depth supported by our approach, we provide the following representative examples drawn from our QA dataset.

Example 1: Multi-step context reasoning

Q: When performing the Longitudinal two-photon calcium imaging experiment, after transferring the mouse to an animal cage, what condition should

the mouse recover from during the subsequent steps?

A: anesthesia

This example requires the model to track the experimental subject’s state across sequential procedural steps and infer the relevant recovery condition through cross-sentence coreference resolution, rather than extracting a locally available answer.

Example 2: Parameter-dependent inference

Q: When performing the COVseq experiment, if we use a PCR thermocycler with the lid set at 80 °C, at what temperature should the incubation be performed?

A: 50 °C

This QA pair reflects a conditional dependency between device configuration and a downstream protocol parameter. The model must distinguish between two co-occurring temperature values and correctly identify which is determined by the given device setting.

Example 3: Operation-specific reasoning

Q: When performing the Fast in vitro protocol for visualization and quantitative high-throughput analysis of sprouting angiogenesis experiment, given that we are imaging beads in large numbers, what unit or entity should the scanning time of about 3 minutes be applied to?

A: bead

Answering correctly requires jointly grounding the imaging target, the quantitative scale ("large numbers"), and the time specification — making it a multi-element integration task rather than a simple fact lookup.

Example 4: Protocol optimization logic

Q: When performing the Panel Optimization for High-Dimensional Immunophenotyping experiment, if the overall pattern is incorrect and we need to adjust other laser lines, what should be reduced as part of this adjustment process?

A: respective detector arrays

This example involves procedural branching logic: the question presupposes a failure condition ("incorrect pattern") and asks the model to reason




Figure A1: **BioPIE enables knowledge integration in the lab automation.** BioPIE can be used to extract large volume of biomedical protocols into structured knowledge, which can then be used by knowledge systems.

Method	2 Hop	3 Hop	4 Hop	MSR
LLM only	18.00	5.26	0.00	15.45
LLM LoRA	12.00	0.00	0.00	9.76
BM25	56.00	36.84	25.00	52.03
LaBSE	59.00	31.58	25.00	53.66
Emb-3-large	63.00	26.32	25.00	56.10
Emb-v4	63.00	21.05	0.00	54.47
GRAG	9.00	0.00	0.00	7.32
GRAG LoRA	18.00	10.53	0.00	16.26
ToG	72.00	63.16	50.00	69.92
Ours w/o Sentence	67.00	57.89	50.00	65.04
Ours w/o Graph	67.00	57.89	50.00	60.16
Ours w/o Planner	62.00	52.63	50.00	62.60
Ours w SciERC	62.00	52.63	25.00	59.35
Ours w ChemProt	68.00	52.63	25.00	64.23
Ours	76.00	73.68	50.00	74.80

Table A8: **Performance comparison across different reasoning depths on the MSR dataset.**

about the corrective action within a multi-laser optimization workflow, requiring understanding of protocol-level decision structure.

B.4 Detailed QA Evaluation Results

Experiments are conducted on Llama-3-8B. For the IE component of our method, we employ the best-performing supervised and LLM-based extraction approaches under the strict OOD evaluation setting. Specifically, we use HGERE (Yan et al., 2023) as the supervised IE method and Llama-3-8B (Pipeline) as the LLM IE method.

We have included the performance of the LLM relative to increasing reasoning depth (see Tab. A8, KGs come from supervised IE), which demonstrates a downward trend in accuracy with increasing reasoning complexity.

Tab. A9 reports the overall performance comparison across different QA systems. Our method

	Supervised IE	LLM IE
LLM only		14.74
LLM LoRA		12.44
BM25	63.72	
LaBSE	55.60	
Emb-3-large	61.52	
Emb-v4	59.14	
GRAG	8.83	21.54
GRAG LoRA	22.77	26.83
ToG	69.90	65.31
Ours w/o Sentence	59.84	58.25
Ours w/o Graph	55.69	55.60
Ours w/o Planner	72.29	71.76
Ours w SciERC	62.40	63.55
Ours w ChemProt	64.70	62.93
Ours	72.99	71.76

Table A9: **Performance comparison across different QA systems. Bold numbers indicate the best performance among all models.**

maintains strong performance under both supervised and LLM-based IE settings. Although supervised extraction generally performs slightly better, the performance gap remains small, indicating that the proposed framework is robust to different IE strategies.

Although GRAG leverages structured KGs, its performance remains substantially lower than that of text-based Retrieval-Augmented Generation (RAG) methods. This can be attributed to the use of average pooling for aggregating node representations, which may limit the model’s ability to capture fine-grained and localized subgraph semantics. Consequently, the retrieved subgraphs often provide insufficient descriptive information, leading to consistently lower retrieval hit rates (see Fig. 5(B)). In contrast, biomedical experimental QA typically involves a large number of domain-

specific terms, in which text-based retrievers naturally achieve a higher recall and more reliable evidence retrieval.

B.5 Application Showcase

In this subsection, we demonstrate a practical application of BioPIE in the field of lab automation. As illustrated in Fig. A1, BioPIE serves as a knowledge integration engine that bridges the gap between high-level biomedical protocols and robotic execution. By extracting structured knowledge from vast volumes of scientific literature, the system can provide precise answers to critical procedural questions—such as identifying the correct reagents for cell dissociation or the appropriate dyes for viability assays. Such structured understanding enables an automated planner to orchestrate a sequence of precise laboratory actions, ranging from trypsinization to cell counting using a hemocytometer, thereby facilitating the complex bioprinting of cartilage and skin tissue analogs with minimal human intervention.