future internet

MDPI

Article

Leveraging Medical Knowledge Graphs and Large Language Models for Enhanced Mental Disorder Information Extraction

Chaelim Park, Hayoung Lee † and Ok-ran Jeong *

School of Computing, Gachon University, 1342 Sujeong-gu, Seongnam-si 13120, Republic of Korea; rkwhr928@gachon.ac.kr (C.P.); hhzet11@gachon.ac.kr or underzero11@etri.re.kr (H.L.)
* Correspondence: orjeong@gachon.ac.kr
† Current address: Electronics and Telecommunications Research Institute, Daejeon 34129, Republic of Korea.

Abstract: The accurate diagnosis and effective treatment of mental health disorders such as depression remain challenging owing to the complex underlying causes and varied symptomatology. Traditional information extraction methods struggle to adapt to evolving diagnostic criteria such as the Diagnostic and Statistical Manual of Mental Disorders fifth edition (DSM-5) and to contextualize rich patient data effectively. This study proposes a novel approach for enhancing information extraction from mental health data by integrating medical knowledge graphs and large language models (LLMs). Our method leverages the structured organization of knowledge graphs specifically designed for the rich domain of mental health, combined with the powerful predictive capabilities and zero-shot learning abilities of LLMs. This research enhances the quality of knowledge graphs through entity linking and demonstrates superiority over traditional information extraction techniques, making a significant contribution to the field of mental health. It enables a more fine-grained analysis of the data and the development of new applications. Our approach redefines the manner in which mental health data are extracted and utilized. By integrating these insights with existing healthcare applications, the groundwork is laid for the development of real-time patient monitoring systems. The performance evaluation of this knowledge graph highlights its effectiveness and reliability, indicating significant advancements in automating medical data processing and depression management.

Keywords: depression; knowledge graph; zero-shot information extraction; large language models; DSM-5; mental health; entity linking

check for updates

Citation: Park, C.; Lee, H.; Jeong, O.-r. Leveraging Medical Knowledge Graphs and Large Language Models for Enhanced Mental Disorder Information Extraction. *Future Internet* **2024**, *16*, 260. https://doi.org/10.3390/fi16080260

Academic Editors: Yanhao Wang, Michael Sheng and Chengcheng Yang

Received: 2 June 2024
Revised: 5 July 2024
Accepted: 11 July 2024
Published: 24 July 2024

Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

Depression is a major public health concern worldwide. The "World Mental Health Report" (2020) (World mental health report: Transforming mental health for all; https://www.who.int/publications/i/item/9789240049338 (accessed on 10 July 2024)) of the World Health Organization (WHO) revealed a significant rise in mental health conditions: depression cases reached 264 million (a 28% increase from 2019) and anxiety disorders rose to 374 million (26% increase). This surge highlights the growing societal and economic impacts of mental health, which was previously considered a private matter. In collaboration with the International Labor Organization, the WHO estimates that mental health problems incur annual economic losses of approximately $1 trillion (WHO and ILO call for new measures to tackle mental health issues at work; https://www.who.int/news/item/28-09-2022-who-and-ilo-call-for-new-measures-to-tackle-mental-health-issues-at-work (accessed on 10 July 2024)). These figures highlight the crucial role of mental health beyond social welfare.

Given the complex and multifaceted nature of depression, successful treatment and management require personalized approaches. Systematic data analysis and comprehension are vital for medical professionals and researchers. However, current biomedical databases often rely on manual information extraction from medical literature, leading to

time-consuming and inefficient processes [1]. To overcome the limitations of previous methods, large language models (LLMs), which have shown excellent performance, have been actively utilized in many areas, such as information extraction [2] and prompt tuning [3]. In particular, recent studies have focused on zero-shot information extraction (ZeroIE) [4] and relation triple extraction (ZeroRTE) [5,6], which enable pretrained models to extract information from new data.

ZeroIE and ZeroRTE play pivotal roles in depression research owing to several key advantages. First, they reduce the data-labeling costs. Depression-related data are vast and manual labeling is expensive and inefficient. ZeroIE leverages pretrained models to extract valuable information from unlabeled data, thereby saving time and resources [7]. Second, ZeroIE and ZeroRTE are versatile and scalable. Depression is a complex illness characterized by a wide range of symptoms and evolving diagnostic criteria. ZeroIE models can extract information from diverse domains and contexts, enabling their application to new symptoms or diagnostic criteria without requiring model retraining. Third, ZeroIE and ZeroRTE facilitate rapid information extraction. Timely access to the latest information is critical for depression research and treatment. ZeroIE models can process and extract information in real time, thereby empowering healthcare professionals and researchers to respond promptly. Finally, ZeroIE and ZeroRTE support personalized treatment. Personalized approaches are essential for effective depression management. ZeroIE enables the rapid extraction and analysis of patient-specific information, which can facilitate the development of personalized treatment plans that optimize outcomes [8].

This study proposes a novel method for enhancing information extraction from mental disorder data by integrating medical knowledge graphs and LLMs. Our approach leverages the structured organization of knowledge graphs tailored to the rich mental health domain. We combine these knowledge graphs with the robust predictive and zero-shot learning capabilities of LLMs to extract and organize depression-related information efficiently. Specifically, we develop a knowledge graph that comprehensively maps the symptoms and diagnostic criteria associated with depression based on the Diagnostic and Statistical Manual of Mental Disorders fifth edition (DSM-5) document, which is a medical resource that describes the criteria for designating anxiety disorders for major depressive disorder and bipolar depression [9].

Furthermore, rigorous performance evaluation experiments were conducted to validate the effectiveness and reliability of the knowledge graph. This research significantly contributes to advancing the understanding of depression management and treatment, while also redefining the analysis and use of data in the mental health field. Our approach outperforms traditional information extraction methods through advanced entity linking techniques and contributes to deeper understanding and management of complex mental health issues such as depression.

The key contributions of this research are as follows:

- Automated and enhanced information extraction accuracy: By leveraging LLMs for zero-shot learning, we automate the traditional manual information extraction process, thereby enabling faster and more accurate data handling.
- Improved data accuracy through entity linking: By employing entity linking techniques, we ensure accurate connections between textual information and entities within the knowledge graph, leading to increased data consistency and reliability.
- Expanded practical applications in medicine: The developed knowledge graph has broader applicability beyond depression, potentially encompassing other mental disorders.

This research paves the way for integration with various medical applications including real-time patient monitoring systems. The remainder of this paper is organized as follows: Section 2 reviews related studies and discusses recent advancements and trends in medical knowledge graphs and zero-shot information extraction techniques. Section 3 details the proposed methodology, with a focus on constructing medical knowledge graphs and using LLMs for entity linking. Section 4 presents the experimental results and demonstrates the effectiveness of the proposed method. Section 5 evaluates the validity of the

methodology based on the findings and discusses its potential applications in the medical field. Finally, Section 6 concludes the paper and proposes directions for future research.

2. Related Works

2.1. Medical Knowledge Graphs

Medical knowledge graphs have garnered significant attention owing to their excellent performance in intelligent healthcare applications. As diverse medical departments proliferate within hospitals, numerous medical knowledge graphs are being constructed for various diseases. For instance, ref. [10] built a graph encompassing over one million clinical concepts from 20 million clinical notes, covering drugs, diseases, and procedures, while ref. [11] constructed a graph detailing 156 diseases and 491 symptoms based on emergency room visits. In addition, ref. [12] developed an EMR-based medical knowledge network (EMKN) with 67,333 nodes and 154,462 edges, focusing on symptom-based diagnostic models to explore the application and performance of knowledge graphs.

Various datasets have led to diverse methods of constructing medical knowledge graphs. The authors of [11] utilized logistic regression, naïve Bayes classifiers, and Bayesian networks with noisy OR gates to automatically generate knowledge graphs using maximum likelihood estimation. In addition, a combination of bootstrapping and support vector machines was applied [13] to extract relationships among entities to build an obstetrics- and gynecology-related knowledge graph. Recent studies have increasingly focused on automated entity and relation extraction using deep-learning techniques [14,15].

Existing medical knowledge graphs have provided insights into individual disease domains and expanded their application areas for diagnosis and recommendation. However, constructing medical knowledge graphs requires expert review and labeling, which, while detailed, demands considerable time and resources. This requirement makes it challenging to extend knowledge graphs to new or different diseases. This is particularly a problem in the mental health sector, where research progress has been slow due to the complexity and sensitivity of the data. Our study aims to overcome these limitations by automating the extraction of entities and relationships using high-performance LLMs.

2.2. Zero-Shot Information Extraction

As AI models are trained and produce output based on datasets, their performance depends on the structure of the model as well as the quality of the data. Expert-reviewed labeling of large amounts of data is an essential process; however, it is labor intensive and time consuming. Therefore, to reduce the time and labor spent on datasets, substantial work has been conducted on extracting relations [16] and arguments [17] from a few resources based on zero-/few-shot techniques [18].

Particularly in medicine, where clinical notes, medical reports, and patient information often lack the necessary annotations, the cost of using AI models is prohibitive, thereby delaying their application and limiting their adaptability to other medical fields. Recent advancements include the application of zero-/few-shot information extraction techniques in the medical field, simplifying medical reports using ChatGPT [19] and extracting information from radiology reports [20].

Our study builds on these methodologies, targeting depression-related information extraction with a focus on the accurate linking of entities and relationships using detailed annotation guidelines and zero-shot learning.

2.3. Entity Linking

Entity linking plays a crucial role in preprocessing and preparing model inputs by identifying key entities in the text data and linking them to appropriate identifiers in knowledge bases [21]. This process is vital for converting text mentions into structured data, reducing redundancy, and enhancing data consistency. Recent trends have highlighted the use of LLMs such as GPT-3 for performing entity linking [22,23], utilizing their extensive

general knowledge and language comprehension capabilities to identify and link entities accurately, even in complex contexts.

Our research employs the gpt-3.5-turbo-instruct model [24] for effective entity identification and linking during the preprocessing stage, which significantly enhances the accuracy and relevance of the extracted information and contributes to the efficiency and precision of the information extraction and knowledge graph construction processes.

This approach offers a high level of automation and scalability, especially when working with limited data, and provides a foundational basis for exploring the applicability of LLMs to extend the scope of entity linking across various data types.

3. Proposed Method

This section describes each component of the proposed framework and the data processing in detail. This process includes the schema definition and annotation guideline setting, information extraction and entity linking, guideline-based model usage, triple extraction, final output, and knowledge graph structure.

We propose a novel zero-shot approach that integrates medical knowledge graphs and LLMs to enhance information extraction from mental health disorder data. The data processing pipeline of the proposed framework is illustrated in Figure 1. To improve the extraction efficiency, we use LLMs to extract the desired information from a document corpus and identify and link relevant entities based on schema definitions and annotation guidelines. For zero-shot extraction, we provide high-quality gold annotations to a guideline-based model using the results of information extraction and entity linking, and convert them into a final knowledge graph through triple extraction. Therefore, the proposed framework focuses on effectively automating and optimizing complex data processing tasks.

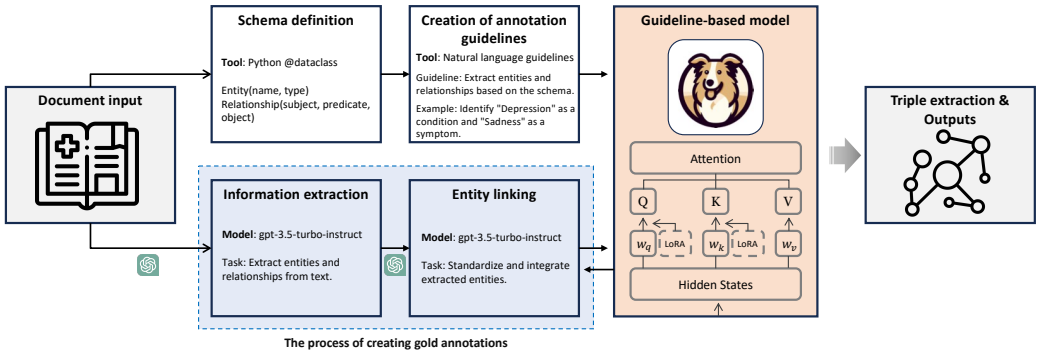


Figure 1. Guideline-Based Model Pipeline for Zero-Shot Information Extraction and Entity Linking.

Methodology and Techniques

Our zero-shot approach includes the following key steps:

Schema definition involves defining entity and relationship types to ensure structured and consistent data processing. We use Python @dataclass to define these entities and relationships, which clarifies the structure and maintains consistency in data processing. The creation of annotation guidelines involves writing natural language guidelines that the model follows during information extraction. These guidelines ensure consistency and accuracy in the extracted data.

In the information extraction phase, we use the gpt-3.5-turbo-instruct model [24] to extract entities and relationships from the input documents. This model excels at understanding the context and identifying the appropriate entities and relationships. The extracted entities are then standardized and integrated during the entity linking phase to maintain consistency. This process resolves issues of varied expressions for the same concept and converts the data into a unified format.

The guideline-based model application involves processing the standardized data according to the established guidelines using the GoLLIE model [25]. The GoLLIE model, trained with guideline-based learning, ensures consistent data processing and enhances

data accuracy and consistency by leveraging high-quality gold annotations. This model improves the reliability of the analysis results by ensuring data consistency and accuracy.

In the triple extraction and final output generation phase, we utilize the results from the GoLLIE model to generate triples using the standardized entities and relationships. These triples are then converted into a knowledge graph using visualization tools and databases. This final knowledge graph visually represents complex data relationships, aiding researchers and healthcare professionals in quickly accessing the required information.

By integrating these steps, our approach enables precise and automated processing of unstructured text data, particularly in the mental health domain, maximizing the utility of the zero-shot approach.

3.1. Schema Definition and Annotation Guidelines

We propose a zero-shot-based information extraction method. We provide the model with the required schema and annotation guidelines required to extract the correct information at the document level without any pretrained relevant information [26]. Thus, we aim to achieve successful zero-shot extraction by providing guidelines based on the performance results of the LLM model.

The schema template we utilized is designed to be universally applicable for triple extraction in zero-shot scenarios. This template ensures flexibility and robustness across various contexts. However, it is important to note that different scenarios might benefit from using tailored templates, depending on the specific requirements and characteristics of the data involved.

Schema Definition: The schema definition is an essential initial step in setting the standards for the information extraction process. It clearly defines the types of entities and relationships to be extracted from the documents, thereby enabling structured and consistent data processing and establishing the semantic structure of the data through entities and relationships [27]. By defining the schema, we ensure that the extracted data are relevant and accurately categorized, which is essential for maintaining data integrity and reliability. The schema acts as a blueprint that guides the model in identifying and classifying entities and relationships within the text, thereby enhancing the overall accuracy of the information extraction process.

Entity Definition: Each entity represents a specific concept or object, and entity types and attributes are structured using Python's @dataclass. These entities function as key components in the data extraction and analysis.

```
@dataclass
class Entity:
"""
Represents an entity in a document, which could be a person, location, organization, or~any other named object.
The 'name' attribute holds the canonical name of the entity as identified in the document~text.

Attributes:
name (str): The canonical name of the entity.
"""
name: str
```

Relationship Definition: Relationships specify the interactions between entities and provide connections among them. The relationship types and examples are defined precisely to ensure consistency in the annotation guidelines.

```
@dataclass
class Triple:
"""
Represents a triple consisting of two entities and a relation in a semantic graph derived from textual data.
```

```
A triple is typically used in knowledge representation and information
    extraction tasks to model factual~statements.

Attributes:
Subject (Entity): The subject entity of the triple,
usually the one performing or causing an action.
Predicate (str): The relationship that connects
the subject and object entities.
Object (Entity): The object entity of the triple,
usually the one experiencing or affected by the action.
"""

Subject: Entity
Predicate: str
Object: Entity
```

Annotation Guidelines and Model Utilization: The annotation guidelines are based on the schema definition to ensure that the data are processed consistently and accurately within the documents. To enable the LLM to perform a specific task, we provide natural language descriptions of that task, known as prompts, along with the input data. This allows the model to recognize the in-context information within the guidelines and generate the appropriate output based on the format [28,29].

The use of annotation guidelines as prompts within the schema framework is essential for several reasons. Firstly, prompts help the model understand the context and generate appropriate outputs by leveraging natural language descriptions. This ensures consistency in the data annotation process and enhances the accuracy of data analysis. Secondly, the GoLLIE model, which we employed in our research, has demonstrated exceptional performance as a guideline-based model. The GoLLIE model outperforms other zero-shot state-of-the-art methods, largely due to its use of guidelines.

By providing detailed annotation guidelines, we ensure that the GoLLIE model can recognize and process information accurately and consistently. Our research provides an annotation guideline for entity definition as follows: "Represents an entity in a document, which could be a person, location, organization, or any other named object. The 'name' attribute holds the canonical name of the entity as identified in the document text". This approach enhances data consistency and reliability, leading to more precise and reliable data analysis. In summary, integrating schema definitions and annotation guidelines as prompts is essential for leveraging the full capabilities of the GoLLIE model, ensuring superior performance in zero-shot information extraction tasks.

3.2. Information Extraction and Entity Linking

Information Extraction: This stage also utilizes the gpt-3.5-turbo-instruct model [24], leveraging its NLP capabilities to identify key information and entities within the documents. The model excels at recognizing meaningful patterns and structures in text, particularly within specialized domains such as mental health text, ensuring high accuracy in information extraction [30]. This is achieved as follows:

- Context-aware recognition: The model analyzes the surrounding text to determine the precise scope and types of entities. For instance, it can distinguish whether "depression" refers to a clinical condition or a general emotional state [31].
- Entity classification: The extracted information is categorized based on predefined classes (e.g., symptoms, diagnoses, and treatments) aligned with the established schema.

Entity Linking: Entity linking is crucial to guarantee the quality and accuracy of depression-related data annotations. This study employs the NLP capabilities of the gpt-3.5-turbo-instruct model to perform entity-linking tasks effectively. The model identifies relevant entities within the data and accurately maps their relationships. These high-quality annotations, termed "gold annotations", are then fed into the guideline-based model described in the following section [32]. This step is instrumental in enabling more refined

data analysis and information extraction, which are essential for the efficient processing of complex mental health data.

The entity linking process consists of the following key steps:

1. Symptom Integration and Standardization: Similar or redundant symptoms are integrated and standardized to ensure consistency. This step addresses the issue of varied expressions for the same concept, converting data into a unified format.
2. Triple Construction: New triples are constructed using the standardized symptoms and durations. This process generates new triples based on integrated symptoms, maintaining data consistency.
3. Entity Linking: The constructed triples are linked to standardized terms, ensuring data consistency and reliability. Entity linking connects the extracted information with standardized terms, enhancing data consistency and improving interoperability among data collected from various sources.

To illustrate the entity linking process, we present a case study on the diagnosis of Major Depressive Disorder (MDD).

Case Study: Entity Linking for Major Depressive Disorder

This case study demonstrates the entity linking process applied to the diagnosis of Major Depressive Disorder (MDD). The raw text used is as follows:

> "A diagnosis of Major Depressive Disorder requires that the patient experiences profound sadness or a loss of interest or pleasure most of the time for at least two weeks".

First, the gpt-3.5-turbo-instruct model is used to extract information from the raw text. The extracted triples are as follows:

- ("Major Depressive Disorder", "requires", "patient experiences profound sadness")
- ("Major Depressive Disorder", "requires", "patient experiences a loss of interest")
- ("Major Depressive Disorder", "requires", "patient experiences a loss of pleasure")
- ("Major Depressive Disorder", "requires", "symptoms last at least two weeks")

Next, similar or redundant symptoms are integrated and standardized, and the triples are reconstructed as follows:

- "profound sadness" is standardized to "Profound Sadness"
- "loss of interest" and "loss of pleasure" are integrated into "Loss of Interest or Pleasure"
- "symptoms last at least two weeks" is standardized to "At Least Two Weeks"

Finally, new triples are constructed using the standardized terms and linked to standardized terms. The final triples are as follows:

- ("Major Depressive Disorder", "manifests as", "Profound Sadness")
- ("Major Depressive Disorder", "manifests as", "Loss of Interest or Pleasure")
- ("Major Depressive Disorder", "lasts", "At Least Two Weeks")

This entity linking process ensures the standardization and consistency of information, enhancing the reliability of the data and enabling more accurate data analysis. The detailed process can be visually confirmed in Figure 2.

3.3. Guideline-Based Model

We utilize a pretrained GoLLIE model [25] to analyze the depression data. This model is strong in zero-shot information extraction using annotation guidance and can effectively handle NLP and code-based input and output, making it well-suited for processing medical data such as depression data [33]. The model was trained on large text datasets collected from various domains including news, biomedicine, and social media.

Notably, the training process incorporated normalization techniques, such as dropout and batch normalization, enabling the model to learn data variations and prevent overfitting. This approach allowed the GoLLIE model to reflect real-world language-usage patterns effectively and adapt to a broader range of data. Utilizing a pretrained model

facilitates swift and accurate depression data analysis, bypassing the complexities of the initial model setup and lengthy training times. The model leverages the extensive language knowledge acquired during training to identify and analyze meaningful patterns and relationships within the depression data efficiently.

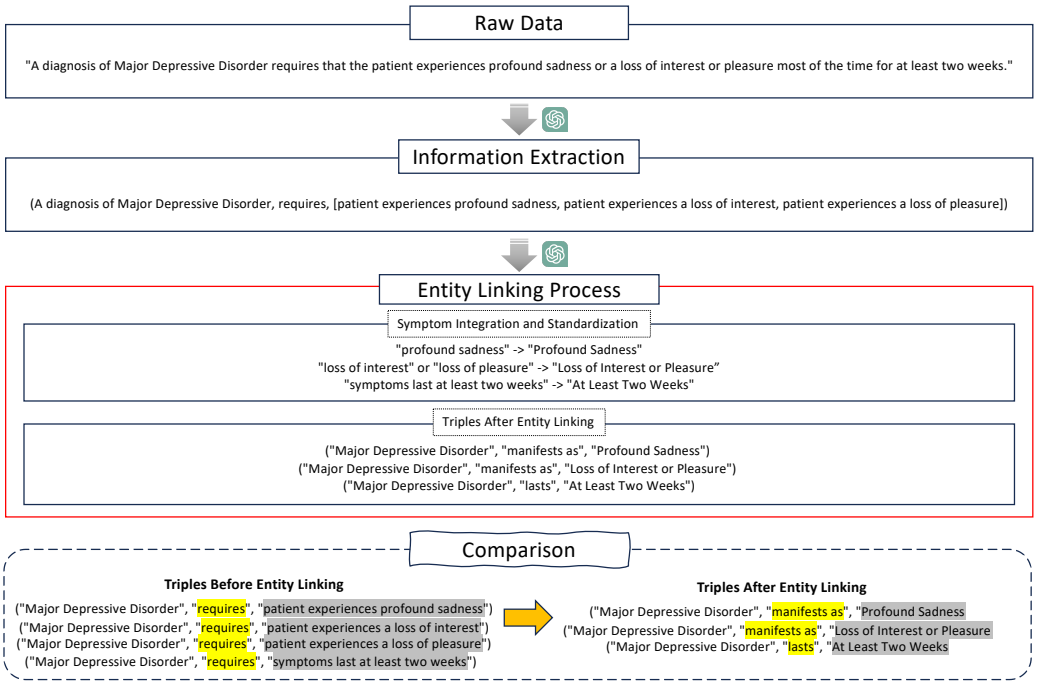


Figure 2. Impact of entity linking on information extraction for major depressive disorder.

3.4. Triple Extraction and Output

Triple Extraction: The triple extraction process utilizes the gold annotations generated by the LLM model above to convert the interrelationships between the entities extracted from the document into the triple form [34]. These triples consist of three elements: subject, predicate, and object. They semantically link each piece of information to create a machine-readable representation of the relationships. For instance, the relationship between "major depressive disorder" (subject) and "difficulty thinking" (object) can be expressed as a triple with the predicate "is characterized by".

The results of the triple extraction are stored as structured data, which can be used as the basis for knowledge graphs. The resulting knowledge graph can help researchers and healthcare professionals to access the required information rapidly by visually representing complex data relationships and making them easy to navigate. This process maximizes the value of the data and plays an important role in adding depth to medical decision-making, work, and education.

Output: The extracted triples are integrated into the knowledge graph, making the information readily usable in various applications. This knowledge graph plays a crucial role in visualizing complex relationships and enabling easy access to information [35]. Figure 3 shows a sample of the knowledge graph developed in this study. It illustrates the interrelationships and characteristics of major depressive disorder and disruptive mood dysregulation disorder, as defined in the DSM-5. This visual representation demonstrates how the proposed method for constructing a knowledge graph can be applied to capture the complexities of mental health issues.

This section details the triple extraction process and the methods for constructing and utilizing the knowledge graph. It highlights the methodological contributions of this study by demonstrating how data structuring can be applied across various fields, particularly for managing complex medical information.

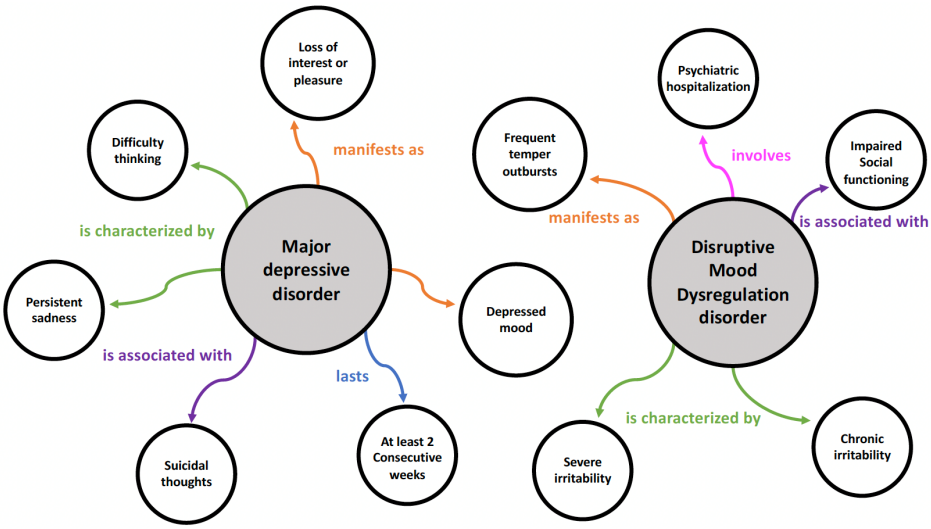


Figure 3. Illustrative representation of the constructed knowledge graph comparing major depressive disorder and disruptive mood dysregulation disorder.

3.5. Structure of the Knowledge Graph

As a result of the zero-shot information extraction process proposed in this study, a medical knowledge graph was constructed by systematically categorizing complex depression symptoms and relationships. This graph encompasses 381 nodes and 505 relationships and provides a detailed representation of the clinical characteristics of depression. The knowledge graph comprises two primary node types: subject and object.

- Subject nodes: These nodes represent specific medical conditions or pathological states. They categorize various forms of depression according to the DSM-5 criteria, including major depressive disorder, persistent depressive disorder, premenstrual dysphoric disorder, and substance/medication-induced depressive disorder. Each subject node is labeled with the name of the disorder and a unique identifier.
- Object nodes: These nodes depict the symptoms or characteristics that a subject node may exhibit. Examples include emotional or behavioral responses such as depressed mood, fatigue, and loss of interest or pleasure. Each object node includes the name of the symptom and a unique identifier.
- Interaction between nodes: Within the knowledge graph, each subject node is linked to one or more object nodes that represent the manifestation of disease traits or symptoms. For instance, the "manifests as" relationship indicates how major depressive disorder might manifest as irritability, while the "is characterized by" relationship suggests that it may be characterized by a depressed mood.

This structured organization and defined relationships allow for an in-depth analysis and understanding of the various symptoms of depression and their interactions. Researchers, physicians, and treatment specialists can use this information to gain a deeper understanding of the causes, manifestations, and characteristics of the diseases. This information will enable the formulation of more effective diagnostic and treatment plans. Table 1 displays a subset of the data used in the knowledge graph, encompassing 488 unique relationships.

Table 1. Overview of relationships in depressive disorders.

Number	Subject	Object	Relation
1	Major depressive disorder	Irritability	manifests as
2	Major depressive disorder	Depressed mood	manifests as
3	Major depressive disorder	Loss of interest or pleasure	manifests as
4	Major depressive disorder	Changes in sleep patterns	manifests as

Table 1. *Cont.*

Number	Subject	Object	Relation
5	Major depressive disorder	Decreased energy levels	manifests as
...	...	...	...
483	Unspecified depressive disorder	Appetite change	lasts
484	Unspecified depressive disorder	Weight change	lasts
485	Unspecified depressive disorder	Sexual interest or desire	lasts
486	Unspecified depressive disorder	Sleep disturbance	includes
487	Unspecified depressive disorder	Psychomotor changes	includes

4. Experiments

This section details the evaluations to assess the performance of the proposed approach. Two primary experiments were performed.

- Zero-shot information extraction on healthcare datasets: This experiment evaluated the effectiveness of zero-shot information extraction on various biomedical datasets.
- Zero-shot relationship extraction at the document level: This experiment focused on extracting relationships from unstructured data at the document level using a zero-shot approach.

4.1. Datasets and Settings

We employed several biomedical datasets to evaluate the effectiveness of zero-shot information extraction from healthcare data. In addition, the DocRED [36] and Re-DocRED datasets [37] were used for document-level information extraction, with a specific focus on extracting and systematizing depression-related information through complex relationship extraction.

4.1.1. BioCreative Datasets

- BC5-Chemical and BC5-Disease [38]: Derived from the BioCreative V chemical-disease relation corpus, these datasets focus on exploring interactions between drugs and diseases. Each dataset includes 1500 PubMed abstracts (evenly split) for training, development, and testing. We used a preprocessed version by Crichton et al., focusing on the named entity recognition (NER) of chemicals and diseases without relationship labels.
- NCBI Disease [39]: Provided by the National Center for Biotechnology Information (NCBI), this dataset includes 793 PubMed abstracts with 6,892 disease mentions linked to 790 unique disease entities. We used a preprocessed version by Crichton et al. for training, development, and testing splits.
- BC2GM [40]: Originating from the BioCreative II gene mention corpus, this dataset consists of sentences from PubMed abstracts with manually tagged genes and alternative gene entities. Our study focused on gene entity annotations using a version of the dataset separated for development by Crichton et al.
- JNLPBA [41]: Designed for applications in molecular biology, this dataset focuses on NER for entity types such as proteins, DNA, RNA, cell lines, and cell types. We focused on entity mention detection without differentiating between entity types, using the same splits as those of Crichton et al.

4.1.2. Document-Level Relationship Extraction Datasets

- DocRED [36]: This dataset was constructed using Wikipedia and Wikidata for document-level relationship extraction. It contains 9228 documents and 57,263 relationship triples, covering 96 predefined relationship types. DocRED was used to evaluate the ability to extract relationships from complex texts spanning multiple sentences.
- Re-DocRED [37]: Re-DocRED, which is an expanded version of DocRED, includes additional positive cases (11,854 documents and 70,608 relationship triples) and in-

corporates relationship types and scenarios that are not addressed by DocRED. This dataset is useful for research aimed at identifying diverse and in-depth relationship patterns within documents.

Table 2 summarizes the statistics for all of the datasets. This structured approach allowed a comprehensive evaluation of the proposed zero-shot information extraction methodology, particularly in the context of mental health data. This ensured a robust assessment and valuable insights into its applicability.

Table 2. Summary of statistics for all datasets.

Category	Dataset Name	Source	Document Count	Primary Entity Type	Entity Count
Biomedical Dataset	BC5-Chemical [38]	PubMed Abstracts	1500	Chemicals (Drugs)	N/A
	BC5-Disease [38]	PubMed Abstracts	1500	Diseases	N/A
	NCBI-Disease [39]	PubMed Abstracts	793	Diseases	6,892
	BC2GM [40]	PubMed Abstracts	N/A	Genes and Alternative Gene Products	N/A
	JNLPBA [41]	PubMed Abstracts	N/A	Proteins, DNA, RNA, Cell Lines, Cell Types	N/A
Document-level Information Extraction	DocRED [36]	Wikipedia	9228	Relationships	57,263 Triples
	Re-DocRED [37]	Wikipedia	11,854	Relationships	70,608 Triples

4.2. Experimental Results

This section presents the findings of the experiments evaluating our zero-shot information extraction approach. The F1 score is calculated as follows and is expressed as a percentage by multiplying the final value by 100:

$$\text{F1} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \times 100\%$$

This approach makes the interpretation of the results clearer and facilitates easier comparison of small differences.

4.2.1. Zero-Shot Information Extraction on Healthcare Datasets

This section presents the evaluation results of our proposed zero-shot information extraction method compared to traditional supervised learning models and other zero-shot models. The performance comparison is detailed in the table.

Our study represents significant advancements in zero-shot information extraction. Unlike traditional methods, our proposed method demonstrates the ability to make accurate predictions without training data. The proposed method achieves significant performance improvements on the NCBI Disease [39], BC5-Disease, and BC5-Chemical [38] datasets, surpassing existing zero-shot models such as GPT-3 In-context learning [1], GPT-3.5-Turbo, and Flan-T5-XXL [42]. These results suggest that our research has the potential to significantly expand the boundaries of zero-shot learning in the medical domain.

Additionally, compared to traditional supervised learning models like PubMedBERT [43], the proposed method maintains competitive performance with significantly fewer data points. PubMedBERT is a widely used pre-trained model in biomedical text analysis and is directly relevant to our research objectives. By selecting such domain-specific models, we demonstrate the advantages and practicality of our zero-shot approach.

The proposed method achieves an F1 score of 85.4% on the NCBI Disease dataset, closely matching PubMedBERT's 87.8%. On the BC5-Disease dataset, the proposed method scores 87.3%, outperforming PubMedBERT (85.6%). In the BC5-Chemical dataset, the proposed method scores 88.5%, significantly higher than GPT-3 (43.6%), GPT-3.5-Turbo (66.5%), and Flan-T5-XXL (67.3%).

Particularly, on the BC2GM dataset, the proposed method achieves an F1 score of 67.2%, surpassing GPT-3 (41.1%), GPT-3.5-Turbo (47.7%), and Flan-T5-XXL (42.4%). Lastly, on the JNLPBA dataset, the proposed method scores 49.7%, demonstrating higher performance compared to other zero-shot models.

These results clearly indicate that our proposed zero-shot method can achieve results comparable to state-of-the-art supervised learning models while significantly outperforming other zero-shot models. The robust performance across diverse datasets illustrates the model's adaptability and effectiveness, supporting its applicability in various real-world scenarios.

Our study provides compelling evidence of the robustness of the proposed method. The results shown in the table highlight the competitive performance of our zero-shot approach compared to traditional supervised learning models and other zero-shot models across different medical datasets. Specifically, the zero-shot learning model demonstrates F1 scores comparable to those of supervised models trained on specific datasets. This underscores the model's ability to generalize and perform well even without fine-tuning on the target datasets. The strong performance under the constraints of zero-shot learning further proves the strength and reliability of our approach in real-world applications.

4.2.2. Zero-Shot Relationship Extraction at the Document Level

Our investigation of document-level relationship extraction demonstrates the potential of zero-shot approaches in advancing this field. The relevant results are presented in Table 3. Implementing entity linking significantly improved the performance compared to not using it. For example, on the DocRED dataset [36], the F1 score increased from 7.803 (without entity linking) to 9.844 (with entity linking). Similar improvements were observed for Re-DocRED [37] (7.527 to 9.150).

Table 3. Performance Comparison of Zero-Shot Information Extraction on Healthcare Datasets.

	Task-Specific SOTA	**Zero-Shot Evaluation**			
	PubMedBERT	**GPT3**	**GPT-3.5-Turbo**	**Flan-T5-XXL**	**Proposed Method**
NCBI	87.8	51.4	47.5	51.8	85.4
BC5-disease	85.6	73.0	67.2	54.7	87.3
BC5-chem	93.3	43.6	66.5	67.3	88.5
BC2GM	84.5	41.1	47.7	42.4	67.2
JNLPBA	79.1	48.3	42.0	38.9	49.7

These results highlight the importance of entity linking in enhancing the overall annotation accuracy, particularly in zero-shot learning scenarios, where model adaptability is crucial Table 4. Our model not only competed well with established models such as LLaMA2-7B, LLaMA2-13B [44], and Flan-T5-XXL [42], but also exhibited consistent competitive strength across the datasets. Notably, the inclusion of annotation and entity-linking processes significantly improved the performance of our model. This approach enables more accurate capturing of the contextual nuances and complex interactions between entities within documents.

Table 4. Performance comparison of document-level zero-shot relationship extraction methods.

	Existing Large Language Model			Proposed Method	
	LLaMA2-7B	**Flan-T5-XXL**	**LLaMA2-13B**	**Without Entity Linking**	**With Entity Linking**
DocRED	1.2	4.4	4.0	7.8	**9.8**
RE-DocRED	1.9	4.3	3.5	7.5	**9.2**

Annotations help the model to develop a deeper understanding of specific cases within datasets, foster better connections between entities, and enhance its grasp of their interrelationships within documents. This method is particularly effective for zero-shot learning. Moreover, the entity-linking process strengthens the semantic connections and significantly improves the accuracy of the extracted information. This method is especially crucial for complex datasets containing various entities and relationships. By incorporating these additional steps, our proposed model achieved higher F1 scores than existing models. This highlights the importance of additional information in extracting complex relationships at the document level, thereby validating the suitability of our model for constructing depression-related knowledge graphs.

The capability of the model to extract and link entities and their relationships within complex medical data accurately is essential for systematically mapping the interactions among various symptoms, diagnoses, and treatment methods associated with depression. Our approach integrates this information, providing medical professionals and researchers with rich resources to improve the understanding and management of depression. Therefore, the technological innovations and performance of our model are expected to contribute significantly to in-depth research on data-driven treatment strategies for depression.

5. Discussion

This study has investigated the effectiveness of entity linking combined with zero-shot information extraction using LLMs in the medical data domain. The findings highlight the critical role of entity linking in this context while also revealing some methodological limitations that offer opportunities for future research and applications in healthcare.

- Performance evaluation and interpretation
  The results confirm that incorporating entity linking with LLMs significantly enhances zero-shot information extraction for medical data. The performance of our model surpassed that of traditional zero-shot approaches, demonstrating that these techniques can compete with conventional supervised learning methods. This is particularly valuable in healthcare, where data labeling is expensive and data privacy is paramount.
- Importance of entity linking
  Entity linking plays a vital role in ensuring data consistency and boosting the model performance. In this study, this went beyond simple identification tasks. By significantly improving the overall data accuracy, entity linking underscores its importance in maintaining the integrity and usefulness of medical information systems.
- Methodological limitations and future directions
  This study used a limited set of datasets, which potentially affected the generalizability of the findings. Future studies should address this issue by exploring a broader range of medical datasets and incorporating a wider variety of entity types. This will help to validate and extend the applicability of the proposed method.
- Potential applications in healthcare
  The constructed medical knowledge graph serves as a critical tool for the systematic analysis of complex medical data and disease states. It has the potential to be integrated into real-time patient management systems to improve both diagnosis and ongoing patient care. Furthermore, the model of synergized LLMs and knowledge graphs suggests the potential benefits of integrating LLMs with knowledge graphs [24].

This approach can leverage the NLP capabilities of LLMs to interpret complex medical data and provide more accurate disease diagnosis and treatment predictions. This holds promise for more precise analysis of the diverse manifestations of depression and the development of effective personalized treatment plans.

A deeper understanding of depression and other complex health conditions can be gained by advancing the methodologies outlined in this study. This will provide richer resources for healthcare professionals and researchers, ultimately improving diagnostic and treatment strategies. The synergy between LLMs and knowledge graphs not only fosters richer data interaction, but also lays the groundwork for transformative changes in medical research and practice, paving the way for innovative healthcare solutions and improved patient outcomes.

6. Conclusions

This study has proposed a novel approach that integrates medical knowledge graphs with LLMs to enhance information extraction for mental health disorders. Our methodology effectively addresses complex mental health conditions, such as depression, by leveraging the structural advantages of knowledge graphs and the robust predictive capabilities of LLMs. This approach significantly improves the accuracy and consistency of information extraction, particularly using entity linking and zero-shot information extraction techniques.

We acknowledge the limitations of this study, particularly the limited dataset. Future research plans include validating the versatility of our methodology across diverse medical datasets and expanding the types of entities involved. These efforts are crucial to strengthen the validity of our approach further and explore its practical applicability in the healthcare sector.

In conclusion, this study demonstrates significant advancements in the field of mental health through the use of medical knowledge graphs and LLMs for information extraction. It provides a powerful tool that can contribute to the diagnosis and treatment of various diseases, marking a notable step forward in integrating advanced AI technologies into healthcare.

Author Contributions: Conceptualization, C.P., H.L. and O.-r.J.; methodology, C.P.; software, C.P.; validation, C.P.; formal analysis, C.P., H.L. and O.-r.J.; investigation, C.P.; resources, C.P.; data curation, C.P.; writing—original draft preparation, C.P.; writing—review and editing, C.P., H.L. and O.-r.J.; visualization, C.P.; supervision, O.-r.J.; project administration, C.P., H.L. and O.-r.J. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available in this article.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Gutierrez, B.J.; McNeal, N.; Washington, C.; Chen, Y.; Li, L.; Sun, H.; Su, Y. Thinking about gpt-3 in-context learning for biomedical ie? think again. *arXiv* **2022**, arXiv:2203.08410.
2. Wang, Y.; Zhao, Y.; Petzold, L. Are large language models ready for healthcare? A comparative study on clinical language understanding. In *Machine Learning for Healthcare Conference*; PMLR: New York, NY, USA, 2023; pp. 804–823.
3. Li, Q.; Wang, Y.; You, T.; Lu, Y. BioKnowPrompt: Incorporating imprecise knowledge into prompt-tuning verbalizer with biomedical text for relation extraction. *Inf. Sci.* **2022**, *617*, 346–358. [CrossRef]
4. Kartchner, D.; Ramalingam, S.; Al-Hussaini, I.; Kronick, O.; Mitchell, C. Zero-Shot Information Extraction for Clinical Meta-Analysis using Large Language Models. In Proceedings of the 22nd Workshop on Biomedical Natural Language Processing and BioNLP Shared Tasks, Toronto, ON, Canada, 13 July 2023; pp. 396–405.
5. Chia, Y.K.; Bing, L.; Poria, S.; Si, L. RelationPrompt: Leveraging prompts to generate synthetic data for zero-shot relation triplet extraction. *arXiv* **2022**, arXiv:2203.09101.
6. Wang, C.; Liu, X.; Chen, Z.; Hong, H.; Tang, J.; Song, D. Zero-shot information extraction as a unified text-to-triple translation. *arXiv* **2021**, arXiv:2109.11171.

7. Li, J.; Jia, Z.; Zheng, Z. Semi-automatic data enhancement for document-level relation extraction with distant supervision from large language models. *arXiv* **2023**, arXiv:2311.07314.
8. Gyrard, A.; Boudaoud, K. Interdisciplinary iot and emotion knowledge graph-based recommendation system to boost mental health. *Appl. Sci.* **2022**, *12*, 9712. [CrossRef]
9. Svenaeus, F. Diagnosing mental disorders and saving the normal: American Psychiatric Association, 2013. Diagnostic and statistical manual of mental disorders, American Psychiatric Publishing: Washington, DC. 991 pp., ISBN: 978-0890425558. Price: $122.70. *Med. Health Care Philos.* **2014**, *17*, 241–244. [CrossRef]
10. Finlayson, S.G.; LePendu, P.; Shah, N.H. Building the graph of medicine from millions of clinical narratives. *Sci. Data* **2014**, *1*, 140032. [CrossRef] [PubMed]
11. Rotmensch, M.; Halpern, Y.; Tlimat, A.; Horng, S.; Sontag, D. Learning a health knowledge graph from electronic medical records. *Sci. Rep.* **2017**, *7*, 5994. [CrossRef] [PubMed]
12. Zhao, C.; Jiang, J.; Xu, Z.; Guan, Y. A study of EMR-based medical knowledge network and its applications. *Comput. Methods Programs Biomed.* **2017**, *143*, 13–23. [CrossRef] [PubMed]
13. Zhang, K.; Li, K.; Ma, H.; Yue, D.; Zhuang, L. Construction of MeSH-like obstetric knowledge graph. In Proceedings of the 2018 International Conference on Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC), Zhengzhou, China, 18–20 October 2018; IEEE: New York, NY, USA, 2018; pp. 160–1608.
14. He, K.; Yao, L.; Zhang, J.W.; Li, Y.; Li, C. Construction of genealogical knowledge graphs from obituaries: Multitask neural network extraction system. *J. Med. Internet Res.* **2021**, *23*, e25670. [CrossRef]
15. Sun, H.; Xiao, J.; Zhu, W.; He, Y.; Zhang, S.; Xu, X.; Hou, L.; Li, J.; Ni, Y.; Xie, G.; et al. Medical knowledge graph to enhance fraud, waste, and abuse detection on claim data: Model development and performance evaluation. *JMIR Med. Inform.* **2020**, *8*, e17653. [CrossRef] [PubMed]
16. Sainz, O.; de Lacalle, O.L.; Labaka, G.; Barrena, A.; Agirre, E. Label verbalization and entailment for effective zero-and few-shot relation extraction. *arXiv* **2021**, arXiv:2109.03659.
17. Sainz, O.; Gonzalez-Dios, I.; de Lacalle, O.L.; Min, B.; Agirre, E. Textual entailment for event argument extraction: Zero-and few-shot with multi-source learning. *arXiv* **2022**, arXiv:2205.01376.
18. Wei, X.; Cui, X.; Cheng, N.; Wang, X.; Zhang, X.; Huang, S.; Xie, P.; Xu, J.; Chen, Y.; Zhang, M.; et al. Zero-shot information extraction via chatting with chatgpt. *arXiv* **2023**, arXiv:2302.10205.
19. Jeblick, K.; Schachtner, B.; Dexl, J.; Mittermeier, A.; Stüber, A.T.; Topalis, J.; Weber, T.; Wesp, P.; Sabel, B.O.; Ricke, J.; et al. ChatGPT makes medicine easy to swallow: An exploratory case study on simplified radiology reports. *Eur. Radiol.* **2023**, *34*, 2817–2825. [CrossRef] [PubMed]
20. Hu, D.; Liu, B.; Zhu, X.; Lu, X.; Wu, N. Zero-shot information extraction from radiological reports using ChatGPT. *Int. J. Med. Inform.* **2024**, *183*, 105321. [CrossRef] [PubMed]
21. Al-Moslmi, T.; Ocaña, M.G.; Opdahl, A.L.; Veres, C. Named entity extraction for knowledge graphs: A literature overview. *IEEE Access* **2020**, *8*, 32862–32881. [CrossRef]
22. Peeters, R.; Bizer, C. Using chatgpt for entity matching. In Proceedings of the European Conference on Advances in Databases and Information Systems, Barcelona, Spain, 4–7 September 2023; Springer: Cham, Switzerland, 2023; pp. 221–230.
23. Pan, S.; Luo, L.; Wang, Y.; Chen, C.; Wang, J.; Wu, X. Unifying large language models and knowledge graphs: A roadmap. *IEEE Trans. Knowl. Data Eng.* **2024**, *36*, 3580–3599. [CrossRef]
24. Ye, J.; Chen, X.; Xu, N.; Zu, C.; Shao, Z.; Liu, S.; Cui, Y.; Zhou, Z.; Gong, C.; Shen, Y.; et al. A comprehensive capability analysis of gpt-3 and gpt-3.5 series models. *arXiv* **2023**, arXiv:2303.10420.
25. Sainz, O.; García-Ferrero, I.; Agerri, R.; de Lacalle, O.L.; Rigau, G.; Agirre, E. Gollie: Annotation guidelines improve zero-shot information-extraction. *arXiv* **2023**, arXiv:2310.03668.
26. Wang, X.; Zhou, W.; Zu, C.; Xia, H.; Chen, T.; Zhang, Y.; Zheng, R.; Ye, J.; Zhang, Q.; Gui, T.; et al. InstructUIE: Multi-task instruction tuning for unified information extraction. *arXiv* **2023**, arXiv:2304.08085.
27. Zhang, X.; Peng, B.; Li, K.; Zhou, J.; Meng, H. Sgp-tod: Building task bots effortlessly via schema-guided llm prompting. *arXiv* **2023**, arXiv:2305.09067.
28. Wei, J.; Bosma, M.; Zhao, V.Y.; Guu, K.; Yu, A.W.; Lester, B.; Du, N.; Dai, A.M.; Le, Q.V. Finetuned language models are zero-shot learners. *arXiv* **2021**, arXiv:2109.01652.
29. Zhou, W.; Zhang, S.; Gu, Y.; Chen, M.; Poon, H. Universalner: Targeted distillation from large language models for open named entity recognition. *arXiv* **2023**, arXiv:2308.03279.
30. Chen, Y.; Jiang, H.; Liu, L.; Shi, S.; Fan, C.; Yang, M.; Xu, R. An empirical study on multiple information sources for zero-shot fine-grained entity typing. In Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Virtual, 7–11 November 2021; pp. 2668–2678.
31. Carta, S.; Giuliani, A.; Piano, L.; Podda, A.S.; Pompianu, L.; Tiddia, S.G. Iterative zero-shot llm prompting for knowledge graph construction. *arXiv* **2023**, arXiv:2307.01128.
32. McCusker, J. LOKE: Linked Open Knowledge Extraction for Automated Knowledge Graph Construction. *arXiv* **2023**, arXiv:2311.09366.
33. Li, P.; Sun, T.; Tang, Q.; Yan, H.; Wu, Y.; Huang, X.; Qiu, X. Codeie: Large code generation models are better few-shot information extractors. *arXiv* **2023**, arXiv:2305.05711.

34. Papaluca, A.; Krefl, D.; Rodriguez, S.M.; Lensky, A.; Suominen, H. Zero-and Few-Shots Knowledge Graph Triplet Extraction with Large Language Models. *arXiv* **2023**, arXiv:2312.01954.
35. Wu, X.; Duan, J.; Pan, Y.; Li, M. Medical knowledge graph: Data sources, construction, reasoning, and applications. *Big Data Min. Anal.* **2023**, **6**, 201–217. [CrossRef]
36. Yao, Y.; Ye, D.; Li, P.; Han, X.; Lin, Y.; Liu, Z.; Liu, Z.; Huang, L.; Zhou, J.; Sun, M. DocRED: A large-scale document-level relation extraction dataset. *arXiv* **2019**, arXiv:1906.06127.
37. Tan, Q.; Xu, L.; Bing, L.; Ng, H.T.; Aljunied, S.M. Revisiting DocRED–Addressing the False Negative Problem in Relation Extraction. *arXiv* **2022**, arXiv:2205.12696.
38. Li, J.; Sun, Y.; Johnson, R.J.; Sciaky, D.; Wei, C.-H.; Leaman, R.; Davis, A.P.; Mattingly, C.J.; Wiegers, T.C.; Lu, Z. BioCreative V CDR task corpus: A resource for chemical disease relation extraction. *Database* **2016**, *2016*, baw068. [CrossRef] [PubMed]
39. Doğan, R.I.; Leaman, R.; Lu, Z. NCBI disease corpus: A resource for disease name recognition and concept normalization. *J. Biomed. Inform.* **2014**, *47*, 1–10. [CrossRef] [PubMed]
40. Smith, L.; Tanabe, L.K.; Ando, R.J.n.; Kuo, C.-J.; Chung, I.-F.; Hsu, C.-N.; Lin, Y.-S.; Klinger, R.; Friedrich, C.M.; Ganchev, K.; et al. Overview of BioCreative II gene mention recognition. *Genome Biol.* **2008**, *9*, S2. [CrossRef] [PubMed]
41. Collier, N.; Ohta, T.; Tsuruoka, Y.; Tateisi, Y.; Kim, J.-D. Introduction to the bio-entity recognition task at JNLPBA. In Proceedings of the International Joint Workshop on Natural Language Processing in Biomedicine and Its Applications (NLPBA/BioNLP), Geneva, Switzerland, 28–29 August 2004; pp. 73–78.
42. Chung, H.W.; Hou, L.; Longpre, S.; Zoph, B.; Tay, Y.; Fedus, W.; Li, Y.; Wang, X.; Dehghani, M.; Brahma, S.; et al. Scaling instruction-finetuned language models. *J. Mach. Learn. Res.* **2024**, *25*, 1–53.
43. Gu, Y.; Tinn, R.; Cheng, H.; Lucas, M.; Usuyama, N.; Liu, X.; Naumann, T.; Gao, J.; Poon, H. Domain-specific language model pretraining for biomedical natural language processing. *ACM Trans. Comput. Healthc. (HEALTH)* **2021**, *3*, 1–23. [CrossRef]
44. Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv* **2023**, arXiv:2307.09288.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.