

# Better Call the Plumber: Orchestrating Dynamic Information Extraction Pipelines

Mohamad Yaser Jaradeh<sup>1\*</sup>[0000-0001-8777-2780], Kuldeep Singh<sup>2</sup>[0000-0002-5054-9881], Markus Stocker<sup>3</sup>[0000-0001-5492-3212], Andreas Both<sup>4</sup>[0000-0002-9177-5463], and Sören Auer<sup>3</sup>[0000-0002-0698-2864]

<sup>1</sup> L3S Research Center, Leibniz University Hannover, Hanover, Germany  
jaradeh@l3s.de

<sup>2</sup> Zerotha-Research & Cerence GmbH, Aachen, Germany  
kuldeep.singh1@cerence.com

<sup>3</sup> TIB Leibniz Information Centre for Science and Technology, Hanover, Germany  
{markus.stocker, auer}@tib.eu

<sup>4</sup> Anhalt University of Applied Sciences, Bernburg, Germany  
andreas.both@hs-anhalt.de

**Abstract.** We propose PLUMBER, the first framework that brings together the research community’s disjoint information extraction (IE) efforts. The PLUMBER architecture comprises 33 reusable components for various Knowledge Graphs (KG) information extraction subtasks, such as coreference resolution, entity linking, and relation extraction. Using these components, PLUMBER dynamically generates suitable information extraction pipelines and offers overall 264 distinct pipelines. We study the optimization problem of choosing suitable pipelines based on input sentences. To do so, we train a transformer-based classification model that extracts contextual embeddings from the input and finds an appropriate pipeline. We study the efficacy of PLUMBER for extracting the KG triples using standard datasets over two KGs: DBpedia, and Open Research Knowledge Graph (ORKG). Our results demonstrate the effectiveness of PLUMBER in dynamically generating KG information extraction pipelines, outperforming all baselines agnostics of the underlying KG. Furthermore, we provide an analysis of collective failure cases, study the similarities and synergies among integrated components, and discuss their limitations.

**Keywords:** Information Extraction · NLP Pipelines · Software Reusability · Semantic Search · Semantic Web

## 1 Introduction and Motivation

In last one decade, publicly available KGs (DBpedia [2] and Wikidata [42] ) have become rich sources of structured content used in various applications, including Question Answering (QA), relation extraction, and dialog systems [39,

---

\* Corresponding Author

4]. The research community developed numerous approaches to extract triple statements [44], keywords/topics [9], tables [45, 23, 22], or entities [35, 36] from unstructured text to complement KGs. Despite extensive research, public KGs are not exhaustive and require continuous effort to align newly emerging unstructured information to the concepts of the KGs.

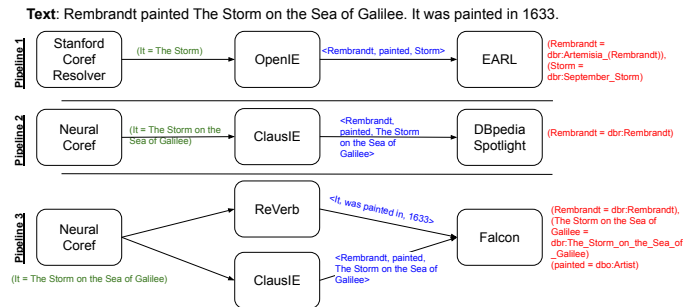
**Research Problem:** This work was motivated by an observation of recent approaches [35, 45, 14] that automatically align unstructured text to structured data on the Web. Such approaches are not viable in practice for extracting and structuring information because they only address very specific subtasks of the overall KG information extraction problem. If we consider the exemplary sentence *Rembrandt painted The Storm on the Sea of Galilee. It was painted in 1633.* (cf. Figure 1). To extract statements aligned with the DBpedia KG from the given sentences, a system must first recognize the entities and relation surface forms in the first sentence. The second sentence requires an additional step of the coreference resolution, where *It* must be mapped to the correct entity surface form (namely, *The Storm on the Sea of Galilee*). The last step requires the mapping of entity and relation surface forms to the respective DBpedia entities and predicates. There has been extensive research in aligning concepts in unstructured text to KG, including entity linking [14, 17], relation linking [36, 38, 4], and triple classification [13]. However, these efforts are disjoint, and little has been done to align unstructured text to the complete KG triples (i.e., represented as subject, predicate, object) [25]. Furthermore, many entity and relation linking tools have been reused in pipelines of QA systems [39, 26]. The literature suggests that once different approaches put forward by the research community are combined, the resulting pipeline-oriented integrated systems can outperform monolithic end-to-end systems [27]. The motivation of our work is also shared with this similar integrative effort in the software architecture community [19]. For the KG information extraction task, however, to the best of our knowledge, approaches aiming at dynamically integrating and orchestrating various existing components do not exist.

**Objective and Contributions:** Based on these observations, we build a framework that enables the integration of previously disjoint efforts on the KG-IE task under a single umbrella. We present the PLUMBER framework (cf. Figure 2) for creating Information Extraction pipelines. PLUMBER integrates 33 reusable components released by the research community for the subtasks entity linking (EL), relation linking (RL), text triple extraction (TE) (subject, predicate, object), and coreference resolution (CR). Overall, there are 264 different composable KG information extraction pipelines (generated by the possible combination of the available 33 components, i.e., for DBpedia 3 CRs, 8 TEs, 10 EL/RLs gives  $3 \cdot 8 \cdot 10 = 240$ , and  $4 \cdot 3 \cdot 2 = 24$  for the ORKG. Hence,  $240 + 24 = 264$  pipelines). PLUMBER implements a transformer-based classification algorithm that intelligently chooses a suitable pipeline based on the unstructured input text.

We perform an exhaustive evaluation of PLUMBER on the two large-scale KGs DBpedia, and Open Research Knowledge Graph (ORKG) [24] to investigate the efficacy of PLUMBER in creating KG triples from unstructured text. We

demonstrate that independent of the underlying KG; PLUMBER can find and assemble different extraction components to produce better suited KG triple extraction pipelines, significantly outperforming existing baselines. In summary, we provide the following novel contributions: i) The PLUMBER framework is the first of its kind for dynamically assembling and evaluating information extraction pipelines based on sequence classification techniques and for a given input text. PLUMBER is easily extensible and configurable, thus enabling the rapid creation and adjustment of new information extraction components and pipelines. Researchers can also use the framework for running IE components independently for specific subtasks such as triple extraction and entity linking. ii) A collection of 33 reusable IE components that can be combined to create 264 distinct IE pipelines. iii) The exhaustive evaluation and our detailed ablation study of the integrated components and composed pipelines on various input text will guide future research for collaborative KG information extraction.

We motivate our work with a running example; the sentence *Rembrandt painted The Storm on the Sea of Galilee. It was painted in 1633.* Multiple steps are required to extract these formally represented statements from the given text. First, the pronoun *it* in the second sentence should be replaced by *The Storm on the Sea of Galilee* using a coreference resolver. Next, a triple extractor should extract the correct text triples from the natural language text, i.e., `<Rembrandt, painted, The Storm on the Sea of Galilee>`, and `<The Storm on the Sea of Galilee, painted in, 1633>`. In the next step, the entity and relation linking component aligns the entity and relation surface forms extracted in the previous step to the DBpedia entities: `dbr:Rembrandt` for *Rembrandt van Rijn*, and `dbr:The_Storm_on_the_Sea_of_Galilee` for *The Storm on the Sea of Galilee*, and for relations: `dbo:Artist` for *painted*, and `dbp:year` for *painted in*. Figure 1 illustrates our running example and shows three PLUMBER IE pipelines with different results. In Pipeline 1, the coreference resolver is unable to map the pronoun *it* to the respective entity in the previous sentence. Moreover, the triple extractor generates incomplete triples, which also hinders the task of



**Fig. 1.** Three example information extraction pipelines showing different results for the same text snippet. Each pipeline consists of coreference resolution, triple extractors, and entity/relation linking components.

the entity and relation linker in the last step. Pipeline 2 uses a different set of components, and its output differs from the first pipeline. Here, the coreference resolution component is able to correctly co-relate the pronoun *it* to *The Storm on the Sea of Galilee*, and extract the text triple correctly. However, the overall result is only partially correct because the second triple is not extracted. Also, the linking component is not able to spot the second entity. Pipeline 3 correctly extracts both triples. This pipeline employs the same component as the second pipeline for coreference resolution but also includes an additional information extraction component (i.e., ReVerb [15]) and a joint entity and relation linking component, namely Falcon [35]. With this combination of components, the text triple extractors were able to compensate for the loss of information in the second pipeline by adding one more component. Using the extracted text triples, the last component of the pipeline, a joint entity and relation linking tool, can map both triple components correctly to the corresponding KG entities.

The remainder of this article is organized as follows. Related work is reviewed in Section 2. Section 3 presents PLUMBER, which is extensively evaluated in Section 4. Section 5 discusses the results, and Section 6 concludes and outlines directions for future research and work.

## 2 Related Work

In the last decade, many open source tools have been released by the research community to tackle IE tasks for KGs. These IE components are not only used for end-to-end KG triple extraction but also for various other tasks, such as:

**Text Triple Extraction:** The task of open information extraction is a well studied researched task in the NLP community [1]. It relies on NER (Named Entity Recognition) and RE (Relation Extraction). SalIE [33] uses MinIE [21] in combination with PageRank and clustering to find facts in the input text. Furthermore, OpenIE [1] leverages linguistic structures to extract self-contained clauses from the text. A comprehensive survey by Niklaus et al. [32] provides detailed about such techniques.

**Entity and Relation Linking:** Entity and relation linking is a widely studied researched topic in the NLP, Web, and Information Retrieval research communities [3, 4, 11]. Often, entity and relation linking is performed independently. DBpedia Spotlight [10] is one of the first approaches for entity recognition and disambiguation over DBpedia. TagMe [17] links entities to DBpedia using in-link matching to disambiguate candidates entities. Others tools such as RelMatch [38] do not perform entity linking and only focus on linking the relation in the text to the corresponding KG relation. Recon [4] uses graph neural networks to map relations between the entities with the assumption that entities are already linked in the text. EARL [14] is a joint linking tool over DBpedia and models the task as a generalized traveling salesperson problem. Sakor et al. [35] proposed Falcon, a linguistic rules based tool for joint entity and relation linking over DBpedia.

**Coreference Resolution:** This task is used in conjunction with other tasks in NLP pipelines to disambiguate text and resolve syntactic complexities. The

Stanford Coreference Resolver [34] uses a multi pass sieve of deterministic coreference models. Clark and Manning [8] use reinforcement learning to fine-tune a neural mention-ranking model for coreference resolution. And more recently [37].

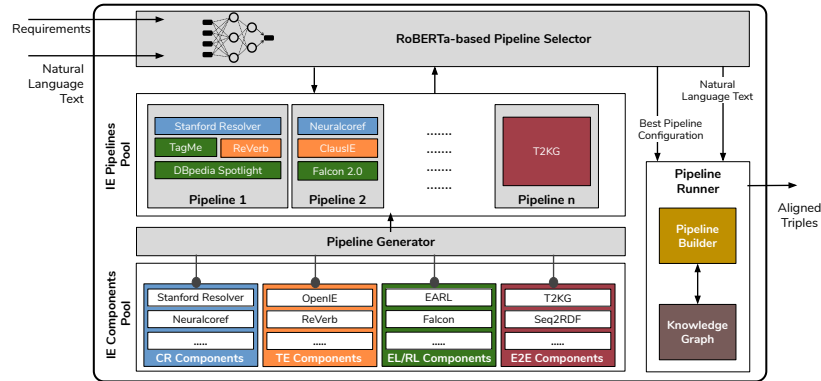
**Frameworks and Dynamic Pipelines:** There have been few attempts in various domains aiming to consolidate the disjoint efforts of the research community under a single umbrella for solving a particular task. The Gerbil platform [41] provides an easy-to-use web-based platform for the agile comparison of entity linking tools using multiple datasets and uniform measuring approaches. OK-BQA [26] is a community effort for the development of multilingual open knowledge base and QA systems. Frankenstein integrates 24 QA components to build QA systems collaboratively on-top of the Qanary integration framework [6]. Other ETL pipelines system exists such as Apache NiFi. Semantic Web Pipes [31] and LarKC [16] are other prominent examples.

**End-to-End Extraction Systems:** More recently, end-to-end systems are gaining more attention due to the boom of deep learning techniques. Such systems draw on the strengths of deep models and transformers [29]. Kertkeidkachorn and Ichise [25] present an end-to-end system to extract triples and link them to DBpedia. Other attempts such as KG-Bert [44] leverage deep transformers [29] for the triple classification task, given the entity and relation descriptions of a triple. KG-Bert does not attempt end-to-end alignment of KG triples from a given input text. Liu et al. [28] design an encoder-decoder framework with an attention mechanism to extract and align triples to a KG.

### 3 Dynamic Information Extraction Pipelining Framework

PLUMBER has a modular design (see Figure 2) where each component is integrated as a microservice. To ensure a consistent data exchange between components, the framework maps the output of each component to a homogeneous data representation using the Qanary [6] methodology. PLUMBER follows three design principles of i) *Isolation*, ii) *Reusability*, and iii) *Extensibility* inspired by [39, 41].

**Dynamic pipeline selection:** PLUMBER uses a RoBERTa [29] based classifier that given a text and a set of requirements, PLUMBER predicts a good pipeline to extract KG triples. Rather than handcrafting features to train models on, we let the RoBERTa model acts as intermediary that classifies the contextual embeddings extracted from the input text into a class which represents one of the possible pipelines. Regarding RoBERTa’s training, we run each input sequence on all possible pipelines and compute the evaluation metrics F1-score (i.e., estimated performance). RoBERTa is fed with the sentence and the sentence-level performance with the best value among all pipelines as the target class. Hence, in practice, the user points PLUMBER to a piece of text and internally it uses RoBERTa to classify the text to a class (i.e., the pipeline) to execute against the input text. We choose a transformer-based architecture due to its ability to encode the contextual knowledge from the input text, providing more accurate classification.



**Fig. 2.** Overview of PLUMBER’s architecture highlighting the components for pipeline generation, selection, and execution. PLUMBER receives an input sentence and requirement (underlying KG) from the user. The framework intelligently selects a suitable pipeline based on the contextual features captured from the input sentence.

**Architecture:** PLUMBER includes the following modules: **i) IE Components Pool:** All information extraction components that are integrated within the framework are parts of the pool. The components are divided based on their respective tasks, i.e., coreference resolution, text triple extraction, as well as entity and relation linking. These components have different input requirements and output formats; thus, PLUMBER provides standard interfaces to facilitate the interaction between pipeline components. **ii) Pipeline Generator:** This module creates possible pipelines depending on the requirements of the components (i.e., the underlying KG). Users can manually select the underlying KG and, using the metadata associated with each component, PLUMBER aggregates the components for the concerned KG. **iii) IE Pipelines Pool:** PLUMBER stores the configurations of the possible pipelines in the pool of pipelines for faster retrieval and easier interaction with other modules. **iv) Pipeline Selector:** Based on the requirements (i.e., underlying KG) and the input text, a RoBERTa based model extracts contextual embeddings from the text and classifies the input into one of the possible classes. Each class corresponds to one pipeline configuration that is held in the IE pipelines pool. **v) Pipeline Runner:** Given the input text, and the generated pipeline configuration, the module executes the pipeline and produce the final KG triples.

## 4 Evaluation

In this section, we detail the empirical evaluation of the framework in comparison to baselines on different datasets and knowledge graphs. As such, we study the following research question: *How does the dynamic selection of pipelines based on the input text affect the end-to-end information extraction task?*

## 4.1 Experimental Setup

**Knowledge Graphs** To study the effectiveness of PLUMBER in building dynamic KG information extraction pipelines, we use the following KGs during our evaluation:

**DBpedia** [2] is containing information extracted automatically from Wikipedia info boxes. DBpedia consists of approximately 11.5B triples [35].

**Open Research Knowledge Graph** [24] (ORKG) collects structured scholarly knowledge published in research articles, using crowd sourcing and automated techniques. In total, ORKG consists of approximately 984K triples.

**Datasets** Throughout our evaluation, we employed a set of existing and newly created datasets for structured triple extraction and alignment to knowledge graphs: the WebNLG [20] dataset for DBpedia, and COV-triples for ORKG.

**WebNLG** is the Web Natural Language Generation Challenge. The challenge introduced the task of aligning unstructured text to DBpedia. In total, the dataset contains 46K triples with 9K triples in the testing and 37K in the training set.

**COV-triples** is a handcrafted dataset that focuses on COVID-19 related scholarly articles. The COV-triples dataset consists of 21 abstracts from peer-reviewed articles and aligns the natural language text to the corresponding KG triples into the ORKG. Three Semantic Web researchers verified annotation quality, and triples approved by all three researchers are part of the dataset. The dataset contains only 75 triples. Hence, we use the WebNLG dataset for training, and 75 triples are used as a test set.

**Components and Implementation** The PLUMBER framework integrates 33 components, the components span different IE tasks from Triple Extraction, Entity and Relation Linking, and Coreference Resolution. Most of the components used are open-sourced and they have been evaluated and used by the community in their respective publications. PLUMBER’s code and all related resources are publicly available online at <https://github.com/YaserJaradeh/ThePlumber>.

**Baselines** We include the following baselines:

**T2KG** [25] is an end-to-end static system aligns a given natural language text to DBpedia KG triples.

**Frankenstein** [39] dynamically composes Question Answering pipelines over the DBpedia KG. It employs logistic regression based classifiers for each component for predicting the accuracy and greedily composes a dynamic pipeline of the best components per task. We adapted Frankenstein for the KG information extraction over DBpedia.

## 4.2 Experiments

The section summarizes a variety of experiments to compare the PLUMBER framework against other baselines. Note, that evaluating the performance of individual components or their combination is out of this evaluation’s scope, since

they were already used, benchmarked, and evaluated in the respective publications. We report values of the standard metrics Precision (P), Recall (R), and F1 score (F1). In all experiments, end-to-end components (e.g., T2KG) are not part of PLUMBER.

**Performance of Static Pipelines** In this experiment, we report results of the static pipelines, i.e., no dynamic selection of a pipeline based on the input text is considered. We ran all 264 pipelines and Table 2 (T2KG & Static noted rows) reports the performance of the best PLUMBER pipeline against the baselines. PLUMBER static pipeline for DBpedia comprises of NeuralCoref [8] for coreference resolution, OpenIE [1] for text triple extraction, TagMe [17] for EL, and Falcon [35] for RL tasks. Also, in case of Frankenstein, we choose its best performing static pipeline. Results illustrated in the Table 2 confirm that the static pipeline composed by the components integrated in PLUMBER outperforms all baselines on DBpedia. We observe that the performance of pipeline approaches is better than an end-to-end monolithic information extraction approaches. Although the PLUMBER pipeline outperforms the baselines, the overall performance is relatively low. All our components have been trained on distinct corpora in their respective publications and our aim was to put them together to understand their collective strengths and weaknesses. Note, Frankenstein addresses the QA pipeline problem and not all components are comparable and can be applied in the context of information extraction. Thus, we integrated NeuralCoref coreference resolution component and OpenIE triple extraction component used in PLUMBER static pipeline into Frankenstein for providing the same experimental settings.

**Static Pipeline for Scholarly KG** In order to assess how PLUMBER performs on domain-specific use cases, we evaluate the static pipelines’ performance on a scholarly knowledge graph. We use the COV-triples dataset for ORKG. To the best of our knowledge, no baseline exists on information extractions of research contribution descriptions over ORKG. Hence, we execute all static pipelines in PLUMBER tailored to ORKG to select the best one as shown in Table 2 (COV-triples rows). PLUMBER pipelines over ORKG extract statements determining the reproductive number estimates for the COVID-19 infectious disease from scientific articles as shown below.

```
@prefix orkg: <http://orkg.org/orkg/resource/>.
@prefix orkgp: <http://orkg.org/orkg/property/>.

orkg:R48100    orkgp:P16022    "2.68" .
```

In this example, *orkg:R48100* refers to the city of Wuhan in China in the ORKG and *orkgp:P16022* is the property “has R0 estimate (average)”. The number “2.68” is the reproductive number estimate. Although COV-triples is a small and manually annotated dataset, we believe that it sheds some light on how PLUMBER will perform on different domains and datasets. Furthermore, it is the first step in creating such a scholarly dataset for IE tasks.



**Table 1.** 10-fold CV of pipeline selection classifiers wrt. Precision, Recall, and F1 score.

Pipeline Selection Approach	Dataset	Knowledge Graph	Classification		
			P	R	F1
Frankenstein [39]	WebNLG	DBpedia	0.732	0.751	0.741
	COV-triples	ORKG	0.832	0.858	0.845
PLUMBER	WebNLG	DBpedia	<b>0.877</b>	<b>0.900</b>	<b>0.888</b>
	COV-triples	ORKG	<b>0.901</b>	<b>0.917</b>	<b>0.909</b>

**Comparison of the Classification Approaches for Dynamic Pipeline Selection** In this experiment, we study the effect of the transformer-based pipeline selection approach implemented in PLUMBER against the pipeline selection approach of Frankenstein. For a comparable experimental setting, we re-use Frankenstein’s classification approach in PLUMBER, keeping the underlying components precisely the same. We perform a 10-fold cross-validation for the classification performance of the employed approach. Table 1 indicates that the PLUMBER pipeline selection significantly outperforms baselines across the board.

**Performance Comparison for KG Information Extraction Task** Our third experiment focuses on comparing the performance of PLUMBER against previous baselines for an end-to-end information extraction task. The results in Table 2 illustrate that the dynamic pipelines built using PLUMBER for KG information extraction outperform the best static pipelines of PLUMBER as well as the dynamically selected pipelines by Frankenstein (rows noted with dynamic). The end-to-end baselines, such as Kertkeidka-chorn and Ichise [25]. We also observe that in cross-domain experiments for COV-triples datasets, dynamically selected pipelines perform better than the static pipeline. In the cross-domain experiment, the static and dynamic PLUMBER pipelines are relatively better performing than the other two KGs. Unlike components for DBpedia, components integrated into PLUMBER for ORKG are customized for KG triple extraction. We conclude that when components are integrated into a framework such as PLUMBER aiming for the KG information extraction task, it is crucial to select the pipeline based on the input text dynamically. The superior performance of PLUMBER shows that the dynamic pipeline selection has a positive impact agnostic of the underlying KG and dataset. This also answers our overall research question.

### 4.3 Ablation Studies

PLUMBER and baselines render relatively low performance on all the employed datasets. Hence, in the ablation studies our aim is to provide a holistic picture of underlying errors, collective success, and failures of the integrated components.

In the first study, we calculate the proportion of errors in PLUMBER. The modular architecture of the proposed framework allows us to benchmark each component independently. We consider the erroneous cases of PLUMBER on the

test set of the WebNLG dataset. We calculate the performance (F1 score) of the PLUMBER dynamic pipeline (cf. Table 2) at each step in the pipeline. The results show that the coreference resolution components caused 21.54% of the errors, 33.71% are caused by text triple extractors, 18.17% by the entity linking components, and 26.58% are caused by the relation linking components.

We conclude that the text triple extractor components contribute to the largest chunk of the errors over DBpedia. One possible reason for their limited performance is that open-domain information extracting components were not initially released for the KG information extraction task. Also, these components do not incorporate any schema or prior knowledge to guide the extraction. We observe that the errors mainly occur when the sentence is complex (with more than one entity and predicate), or relations are not explicitly mentioned in the sentence. We further analyze the text triple extractor errors. The error analysis at the level of the triple subject, predicate, and object showed that most errors are in predicates (40.17%) followed by objects (35.98%) and subjects (23.85%).

**Further Analysis** Aiming to understand why IE pipelines perform with low accuracy, we conduct a more in-depth analysis per IE task. In the first analysis, we evaluated each component independently on the WebNLG dataset. Researchers [12, 40] proposed several criterion for micro-benchmarking tools/components for KG tasks (entity linking, relation linking, etc.) based on the linguistic features of a sentence. We motivate our analysis based on the following:

I) *Text Triple Extraction*: We consider the number of words ( $wc$ ) in the input sentence (a sentence is termed by “simple” with average word length of 7.41 [39]. Sentences with higher number of words than seven are complex sentences). Furthermore, having a comma in a sentence (sub-clause) to separate clauses is another factor. Atomic sentences (e.g., “*cats have tails*”) are a type of sentence that also affects triples extractors’ behavior. Moreover, nominal relation as in “*Durin, son of Thorin*” is another impacting factor on the performance. Uppercase and lowercase mentions of the words (i.e., correct capitalization of the first character and not the entire word) in a sentence are standard errors for entity linking components. We consider this as a micro-benchmarking criteria.

**Table 2.** Overall performance comparison of static and dynamic pipelines for the KG information extraction task.

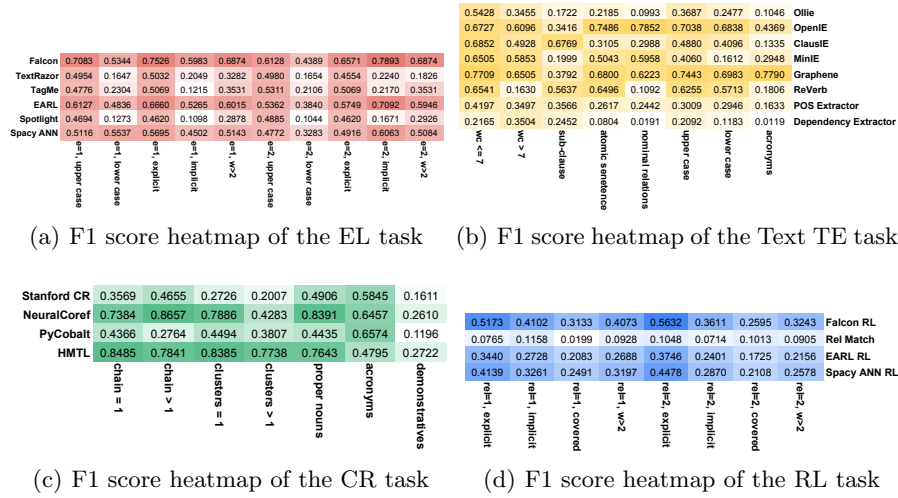
System	Dataset	Knowledge Graph	Performance		
			P	R	F1
T2KG [25]	WebNLG	DBpedia	0.133	0.140	0.135
Frankenstein (Static) [39]	WebNLG	DBpedia	0.177	0.189	0.181
PLUMBER (Static)	WebNLG	DBpedia	0.210	0.225	0.215
	COV-triples	ORKG	0.403	0.423	0.413
Frankenstein (Dynamic) [39]	WebNLG	DBpedia	0.199	0.208	0.203
	COV-triples	ORKG	0.403	0.424	0.413
PLUMBER (Dynamic)	WebNLG	DBpedia	<b>0.287</b>	<b>0.307</b>	<b>0.297</b>
	COV-triples	ORKG	<b>0.411</b>	<b>0.437</b>	<b>0.424</b>

II) *Coreference Resolution*: We focus on the length of the coreference chain (i.e., the number of aliases for a single mention). Additionally, the number of clusters is another criterion in the analysis. A cluster refers to the groups of mentions that require disambiguation (e.g., "mother bought a new phone, she is so happy about it" where the first cluster is *mother* → *she* and the second is *phone* → *it*). The presence of proper nouns in the sentence is studied as well as acronyms. Furthermore, the demonstrative nature of the sentence is also observed as a factor. Demonstrative sentences are the ones that contain demonstrative pronouns (this, that, etc.).

III) *Entity Linking*: The number of entities in a sentence ( $e=1,2$ ) is a crucial observation for the entity linking task. Capitalization of the surface form is another criterion for micro-benchmarking entity linking tools. An entity is termed as an explicit entity when the entity’s surface form in a sentence matches the KG label. An entity is implicit when there is a vocabulary mismatch. For example, in the sentence "The wife of Obama is Michelle Obama.", the surface form *Obama* is expected to be linked to `dbr:Barack_Obama` and considered as an implicit entity [40]. The last linguistic feature is the number of words ( $w$ ) in an entity label (e.g., *The Storm on the Sea of Galilee* has seven words).

IV) *Relation Linking*: Similar to the entity linking criteria, we focus on the number of relations in a sentence ( $rel=1,2$ ). The type of relation (i.e., explicit, or implicit) is another parameter. Covered relation (sentences without a predicate surface form) is also used as a feature for micro-benchmarking: "Which companies have launched a rocket from Cape Canaveral Air Force station?" where the `dbo:manufacturing` relation is not mentioned in the sentence. Covered relations highly depend on common sense knowledge (i.e., reasoning) and the structure of the KG [40]. Lastly, the number of words ( $w \leq N$ ) in a predicate surface form is also considered.

Figure 3 illustrates micro-benchmarking of various PLUMBER components per task. We observe that across IE tasks, the F1 score of the components varies significantly based on the sentence’s linguistic features. In fact, there exist no single component which performs equally well on all the micro-benchmarking criteria. This observation further validates our hypothesis to design PLUMBER for building dynamic information extraction pipelines based on the strengths and weaknesses of the integrated components. We also note in Figure 3 that all the CR components report limited performance for the demonstrative sentences (*demonstratives*). When there is more than one coreference cluster in a sentence, all other CR components observe a discernible drop in F1 score. The NeuralCoref [8] component performs best for *proper nouns*, whereas PyCobalt [18] performs best for the *acronyms* feature (almost being tied by NeuralCoref). In the TE task, Graphene [7] shows the most stable performance across all categories. However, the performance of all components (except Dependency Parser) drops significantly when the number of words in a sentence exceeds seven ( $wc > 7$ ). Case sensitivity also affects the performance and all components observe a noticeable drop in F1 score for lowercase entity mentions in the sentence. Similar behavior is observed for entity linking components where case sensitivity is a significant



**Fig. 3.** Comparison of F1 scores per component for different IE tasks based on the various linguistic features of an input sentence (number of entities, word count in a sentence, implicit vs. explicit relation, etc.). Darker colors indicate a higher F1 score.

cause of poor performance. When the sentence has one entity and it is implicit ( $e=1$ , implicit); all entity linking components face challenges in correctly linking the entities to the underlying KG. Relation linking components also report lower performance for implicit relations.

## 5 Discussion

Even though the dynamic pipelines of PLUMBER outperforms static pipelines, the overall performance of PLUMBER and baselines for the KG information extraction task remains low. Our detailed and exhaustive ablation studies suggest that when individual components are plugged together, their individual performance is a major error source. However, this behavior is expected, considering earlier research works in other domains also observe a similar trend. As in 2015 Gerbil framework [41] and in 2018 Frankenstein [39]. Within two years, the community has released several components dedicated to solving entity linking and relation linking [35, 14, 30], which were two loopholes identified by [39] for the QA task.

We observe that state of the art components for information extraction still have much potential to improve their performance (both in terms of runtime and F1 score). It is essential to highlight that some of the issues observed in our ablation study are very basic and repeatedly pointed out by researchers in the community. For instance, Derczynski et al. [12] in 2015, followed by Singh et al. [39] in 2018, showed that case sensitivity is a main challenge for EL tools. Our observation in Figure 3 again confirms that case sensitivity of entity surface

forms remains an open issue even for newly released components. In contrast, on specific datasets such as CoNLL-AIDA, several EL approaches reported F1 scores higher than 0.90 [43], showing that EL tools are highly customized to particular datasets. In a real-world scenario like ours, the underlying limitations of approaches are uncovered.

## 6 Conclusion and Future Work

In this paper, we presented the PLUMBER approach and framework for information extraction. PLUMBER effectively selects the a suitable pipeline for a given input sentence using the sentential contextual features and a state-of-the-art transformer-based classification model. PLUMBER has a service-oriented architecture which is scalable, extensible, reusable, and agnostic of the underlying KG. The core idea of PLUMBER is to combine the strengths of already existing disjoint research for KG information extraction and build a foundation for a platform to promote reusability for the construction of large-scale and semantically structured KGs. Our empirical results suggest that the performance of the individual components directly impacts the end-to-end information extraction accuracy.

This article does not focus on internal system architecture or employed algorithms in a particular IE component to analyze the failures. The focus of the ablation studies is to holistically study the collective success and failure cases for the various tasks. Our studies provide the research community with insightful results over two knowledge graphs, 33 components, 264 pipelines. Our work is a step in the larger research agenda of offering the research community an effective way for synergistically combining and orchestrating various focused IE approaches balancing their strengths and weaknesses taking different application domains into account. We plan to extend our work in the following directions: i) extending PLUMBER to other KGs such as UMLS [5] and Wikidata [42]. ii) addressing multilinguality with PLUMBER, and iii) creating high performing RL components.

**Acknowledgements.** This work was co-funded by the European Research Council for the project ScienceGRAPH (Grant agreement ID: 819536) and the TIB Leibniz Information Centre for Science and Technology.

## References

1. Angeli, G., Johnson Premkumar, M.J., Manning, C.D.: Leveraging linguistic structure for open domain information extraction. pp. 344–354. ACL (2015)
2. Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., Ives, Z.: Dbpedia: A nucleus for a web of open data. In: The Semantic Web. pp. 722–735 (2007)
3. Balog, K.: Entity linking. In: Entity-Oriented Search, pp. 147–188. Springer (2018)
4. Bastos, A., Nadgeri, A., Singh, K., Mulang, I.O., Shekarpour, S., Hoffart, J., Kaul, M.: Recon: Relation extraction using knowledge graph context in a graph neural network. In: Proceedings of The Web Conference (WWW). p. :N/A (2021)

5. Bodenreider, O.: The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research* **32**, D267–D270 (2004)
6. Both, A., Diefenbach, D., Singh, K., Shekarpour, S., Cherix, D., Lange, C.: Qanary - A methodology for vocabulary-driven open question answering systems. vol. 9678, pp. 625–641 (2016)
7. Cetto, M., Niklaus, C., Freitas, A., Handschuh, S.: Graphene: Semantically-linked propositions in open information extraction. In: *Proceedings of the 27th COLING*. pp. 2300–2311 (2018)
8. Clark, K., Manning, C.D.: Deep reinforcement learning for mention-ranking coreference models. In: *Proceedings of the 2016 EMNLP*. pp. 2256–2262 (2016)
9. Cui, W., Liu, S., Wu, Z., Wei, H.: How hierarchical topics evolve in large text corpora. *IEEE TVCG* **20**(12), 2281–2290 (2014)
10. Daiber, J., Jakob, M., Hokamp, C., Mendes, P.N.: Improving efficiency and accuracy in multilingual entity extraction. In: *Proceedings of the 9th I-Semantics* (2013)
11. Delpeuch, A.: Opentapioca: Lightweight entity linking for wikidata (2019)
12. Derczynski, L., Maynard, D., Rizzo, G., Van Erp, M., Gorrell, G., Troncy, R., Petrak, J., Bontcheva, K.: Analysis of named entity recognition and linking for tweets. *Information Processing & Management* **51**, 32–49 (2015)
13. Dong, T., Wang, Z., Li, J., Bauckhage, C., Cremers, A.B.: Triple classification using regions and fine-grained entity typing. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 33, pp. 77–85 (2019)
14. Dubey, M., Banerjee, D., Chaudhuri, D., Lehmann, J.: EARL: Joint entity and relation linking for question answering over knowledge graphs. In: *Lecture Notes in Computer Science*, pp. 108–126. Springer International Publishing (2018)
15. Fader, A., Soderland, S., Etzioni, O.: Identifying relations for open information extraction. In: *Proceedings of the 2011 EMNLP*. pp. 1535–1545 (Jul 2011)
16. Fensel, D., van Harmelen, F., Andersson, B., Brennan, P., Cunningham, H., Della Valle, E., Fischer, F., Huang, Z., Kiryakov, A., Lee, T.K., Witbrock, M., Zhong, N.: Towards larkc: A platform for web-scale reasoning. In: *IEEE ICSC*. pp. 524–529 (2008)
17. Ferragina, P., Scaiella, U.: TAGME: on-the-fly annotation of short text fragments (by wikipedia entities). pp. 1625–1628 (2010)
18. Freitas, A., Bermeitinger, B., Handschuh, S.: Lambda-3/pycobalt: Coreference resolution in python. <https://github.com/Lambda-3/PyCobalt>
19. Garcia, J., Mirakhorli, M., Xiao, L., Zhao, Y., Mujhid, I., Pham, K., Okutan, A., Malek, S., Kazman, R., Cai, Y., Medvidovic, N.: Constructing a shared infrastructure for software architecture analysis and maintenance. In: *ICSA* (2021)
20. Gardent, C., Shimorina, A., Narayan, S., Perez-Beltrachini, L.: Creating training corpora for NLG micro-planners. pp. 179–188 (2017)
21. Gashteovski, K., Gemulla, R., del Corro, L.: MinIE: Minimizing facts in open information extraction. In: *Proceedings of the 2017 EMNLP*. pp. 2630–2640 (2017)
22. Hou, Y., Jochim, C., Gleize, M., Bonin, F., Ganguly, D.: Identification of tasks, datasets, evaluation metrics, and numeric scores for scientific leaderboards construction. In: *Proceedings of the 57th ACL*. pp. 5203–5213 (2019)
23. Ibrahim, Y., Riedewald, M., Weikum, G., Zeinalipour-Yazti, D.: Bridging quantities in tables and text. In: *2019 IEEE 35th ICDE*. pp. 1010–1021 (2019)
24. Jaradeh, M.Y., Oelen, A., Farfar, K.E., Prinz, M., D’Souza, J., Kismihók, G., Stocker, M., Auer, S.: Open Research Knowledge Graph: Next Generation Infrastructure for Semantic Scholarly Knowledge. *Marina Del K-CAP* **19** (2019)

25. Kertkeidkachorn, N., Ichise, R.: T2kg: An end-to-end system for creating knowledge graph from unstructured text. In: *AAAI Workshops*. vol. WS-17 (2017)
26. Kim, J.D., Unger, C., Ngomo, A.C.N., Freitas, A., Hahm, Y.g., Kim, J., Nam, S., Choi, G.H., Kim, J.u., Usbeck, R., et al.: OKBQA Framework for collaboration on developing natural language question answering systems (2017)
27. Liang, S., Stockinger, K., de Farias, T.M., Anisimova, M., Gil, M.: Querying knowledge graphs in natural language (2020)
28. Liu, Y., Zhang, T., Liang, Z., Ji, H., McGuinness, D.: Seq2rdf: An end-to-end application for deriving triples from natural language text (2018)
29. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Zettlemoyer, L., Stoyanov, V.: Roberta: A robustly optimized bert pretraining approach (2019)
30. Mihindukulasooriya, N., Rossiello, G., Kapanipathi, P., Abdelaziz, I., Ravishankar, S., Yu, M., Gliozzo, A., Roukos, S., Gray, A.: Leveraging semantic parsing for relation linking over knowledge bases. *ISWC* (2020)
31. Morbidoni, C., Polleres, A., Tummarello, G., Le-Phuoc, D.: Semantic web pipes (2007)
32. Niklaus, C., Cetto, M., Freitas, A., Handschuh, S.: A survey on open information extraction. In: *Proceedings of the 27th COLING*. pp. 3866–3878 (2018)
33. Ponza, M., Del Corro, L., Weikum, G.: Facts that matter. In: *Proceedings of the 2018 EMNLP*. pp. 1043–1048. *ACL* (2018)
34. Raghunathan, K., Lee, H., Rangarajan, S., Chambers, N., Surdeanu, M., Jurafsky, D., Manning, C.: A multi-pass sieve for coreference resolution. In: *EMNLP* (2010)
35. Sakor, A., Onando Mulang, I., Singh, K., Shekarpour, S., Esther Vidal, M., Lehmann, J., Auer, S.: Old is gold: Linguistic driven approach for entity and relation linking of short text. pp. 2336–2346. *ACL* (2019)
36. Sakor, A., Singh, K., Patel, A., Vidal, M.E.: Falcon 2.0: An entity and relation linking tool over wikidata. In: *CIKM* (2020)
37. Sanh, V., Wolf, T., Ruder, S.: A hierarchical multi-task approach for learning embeddings from semantic tasks. *Proceedings of the AAAI* **33**, 6949–6956 (2019)
38. Singh, K., Mulang, I.O., Lytra, I., Jaradeh, M.Y., Sakor, A., Vidal, M., Lange, C., Auer, S.: Capturing knowledge in semantically-typed relational patterns to enhance relation linking. In: *Proceedings of the Knowledge Capture Conference, K-CAP 2017, Austin, TX, USA, December 4-6, 2017*. pp. 31:1–31:8 (2017)
39. Singh, K., Radhakrishna, A.S., Both, A., Shekarpour, S., Lytra, I., Usbeck, R., Vyas, A., Khikmatullaev, A., Punjani, D., Lange, C., Vidal, M.E., Lehmann, J., Auer, S.: Why reinvent the wheel: Let’s build question answering systems together. p. 1247–1256. *WWW ’18* (2018)
40. Singh, K., Saleem, M., Nadgeri, A., Conrads, F., Pan, J.Z., Ngomo, A.C.N., Lehmann, J.: Qaldgen: Towards microbenchmarking of question answering systems over knowledge graphs. In: *ISWC*. pp. 277–292 (2019)
41. Usbeck, R., Röder, M., et al., N.N.: Gerbil: general entity annotator benchmarking framework. In: *Proceedings of the 24th WWW*. pp. 1133–1143 (2015)
42. Vrandečić, D., Krötzsch, M.: Wikidata: A Free Collaborative Knowledgebase. *Communications of the ACM* **57**(10), 78–85 (2014)
43. Yang, X., Gu, X., Lin, S., Tang, S., Zhuang, Y., Wu, F., Chen, Z., Hu, G., Ren, X.: Learning dynamic context augmentation for global entity linking. In: *EMNLP-IJCNLP*. pp. 271–281 (2019)
44. Yao, L., Mao, C., Luo, Y.: Kg-bert: Bert for knowledge graph completion (2019)
45. Yu, W., Li, Z., Zeng, Q., Jiang, M.: Tablepedia: Automating pdf table reading in an experimental evidence exploration and analytic system. p. 3615–3619. *WWW ’19*