

Adapter-Based Approaches to Knowledge-Enhanced Language Models: A Survey

Alexander Fichtl¹^a, Juraj Vladika²^b and Georg Groh¹

¹*Social Computing Research Group, Technical University of Munich, Boltzmannstraße 3, 85748, Garching, Germany*

²*SEBIS, Technical University of Munich, Boltzmannstraße 3, 85748, Garching, Germany*

{alexander.fichtl, juraj.vladika, georg.groh}@tum.de

Keywords: Natural Language Processing, Knowledge Engineering, Knowledge Enhancement, Knowledge Graphs, Language Models, Adapters.

Abstract: Knowledge-enhanced language models (KELMs) have emerged as promising tools to bridge the gap between large-scale language models and domain-specific knowledge. KELMs can achieve higher factual accuracy and mitigate hallucinations by leveraging knowledge graphs (KGs). They are frequently combined with adapter modules to reduce the computational load and risk of catastrophic forgetting. In this paper, we conduct a systematic literature review (SLR) on adapter-based approaches to KELMs. We provide a structured overview of existing methodologies in the field through quantitative and qualitative analysis and explore the strengths and potential shortcomings of individual approaches. We show that general knowledge and domain-specific approaches have been frequently explored along with various adapter architectures and downstream tasks. We particularly focused on the popular biomedical domain, where we provided an insightful performance comparison of existing KELMs. We outline the main trends and propose promising future directions.

1 INTRODUCTION


The field of natural language processing (NLP) has, in recent years, been dominated by the rise of large language models (LLMs). These models are usually pre-trained on large amounts of unstructured textual data, which enables them to solve complex reasoning tasks and generate new text. Still, LLMs can lack awareness of structured knowledge hierarchies, such as relations between concepts and reasoning capabilities in knowledge-intensive tasks (Rosset et al., 2020; Hu et al., 2023). These drawbacks can lead to inaccurate predictions in downstream tasks and so-called "hallucinations" (Huang et al., 2023) within text generation, making LLMs less reliable in practice, an especially precarious issue in high-risk domains like healthcare or law.


A potential solution to counteract hallucinations and improve the reliability of LLMs is knowledge enhancement: By leveraging expert knowledge from sources such as manually curated knowledge graphs (KGs), structured knowledge can be injected into LLMs. Such knowledge-enhanced language mod-

els (KELMs) are a promising approach for higher structured knowledge awareness, better factual accuracy, and less hallucinations (Colon-Hernandez et al., 2021; Wei et al., 2021; Hu et al., 2023). KGs are a vital part of *knowledge engineering*, a discipline that can be leveraged to make LLMs use advanced logic and formal expressions (Allen et al., 2023).

Unfortunately, knowledge enhancement in the form of supervised fine-tuning (SFT) can be highly computationally expensive, especially for LLMs with billions of parameters. A promising research avenue to overcome this limitation is using lightweight and efficient adapter modules (Houlsby et al., 2019; Pfeiffer et al., 2020a). These modules can enhance the task performance of LLMs and, at the same time, be a very computationally efficient solution. Despite the rising popularity of this approach and to the best of our knowledge, a comprehensive overview of adapter-based KELMs is still missing in the NLP research landscape.

To bridge this research gap, we conduct a systematic literature review (SLR) on adapter-based knowledge enhancement of LLMs. Our contributions are (1) a novel review of adapter-based knowledge enhancement, (2) a quantitative and qualitative analysis of different methods in the field, and (3) a detailed cat-

^a <https://orcid.org/0009-0001-9875-0197>

^b <https://orcid.org/0000-0002-4941-9166>

egorization of literature and identification of the most promising trends. This work is a reworked and updated version of an SLR conducted as part of a master's thesis of one of the authors (Fichtl, 2024).

2 BACKGROUND AND RELATED WORK

This section gives an overview of related work and existing surveys on knowledge enhancement. Knowledge graphs are the most common external knowledge source, so we start with their overview.

2.1 Knowledge Graphs

Knowledge Graphs (KGs): are a structured representation of world knowledge and have seen a rising prominence in NLP research over the past decade (Schneider et al., 2022). Hogan et al. (2020) define a KG as "a graph of data intended to accumulate and convey knowledge of the real world, whose nodes represent entities of interest and whose edges represent relations between these entities". Similarly, Ji et al. (2020) published a comprehensive survey on KGs and, following existing literature, defined the concept of a KG as " $\mathcal{G} = \{\mathcal{E}, \mathcal{R}, \mathcal{F}\}$, where \mathcal{E}, \mathcal{R} and \mathcal{F} are sets of entities, relations and facts, respectively; a fact is denoted as a triple $(h, r, t) \in \mathcal{F}$ ". h, r, t denote the head, relation, and tail of a triplet. An **Ontology** is a formal representation of concepts and semantic relations between them – it provides a "schema" that a KG has to adhere to, making it possible to do reasoning and derive rules from KGs (Khadir et al., 2021).

Depending on the source and purpose of a KG, entities and relations can take on various shapes. For example, in the biomedical knowledge graph UMLS (Bodenreider, 2004), a relation can take the shape of a single word like "inhibits", a short phrase like "relates to", or a compound term including, for example, chemical or medical categories such as "[protein] relates to [disease]" or "[substance] induces [physiology]". A textual connection is vital because it links the graph structure with natural language, simplifying the integration of information from KGs into language models and the associated learning processes. Other than UMLS, other examples of popular KGs are DBpedia (Auer et al., 2007) and ConceptNet (Speer et al., 2017).

2.2 Approaches to Knowledge Enhancement

At the time of writing, some reviews had already been published that gave an overview of KELMs and classified different approaches. Colon-Hernandez et al. (2021) review the existing literature and split the approaches to integrate structure knowledge with LMs into three categories: (1) input-centered strategies, centering around altering the structure of the input or selected data, which is fed into the base LLM; (2) architecture-focused approaches, which involve either adding additional layers that integrate knowledge with the contextual representations or modifying existing layers to alter parts like attention mechanisms; (3) output-focused approaches, which work by changing either the output structure or the losses used in the base model. Our study focuses on category (2) by examining the adapter-based mechanisms for injecting information into the model (see example in Figure 1), which were shown to be the most promising by the authors.

The second survey by Wei et al. (2021) reviews a large number of studies on KELMs and classifies them using three taxonomies: (1) knowledge sources, (2) knowledge granularity, and (3) application areas. Within (1), the knowledge sources include linguistic, encyclopedic, common-sense, and domain-specific knowledge. The second taxonomy (2) acknowledges the common approach of using KGs as a source of knowledge. Levels of granularity mentioned are text-based knowledge, entity knowledge, relation triples, and KG sub-graphs. Lastly, with the third taxonomy (3), the authors discuss how knowledge enhancement can improve natural language generation and understanding. They also review popular benchmarks that can be used for task evaluation of KELMs (Wei et al., 2021).

Adapters are part of a broader paradigm of *modular deep learning*, described in detail by Pfeiffer et al. (2024). The two field studies by Colon-Hernandez et al. (2021) and Wei et al. (2021) on the classification of KELM approaches were a valuable starting point for exploring KELMs. Although they address some adapter-based studies like K-Adapter (Wang et al., 2020), most other adapter-based KELMs are missing. This lack of coverage led us to conduct a novel systematic literature search focusing specifically on the adapter-based KELMs, considering their rising popularity and importance.

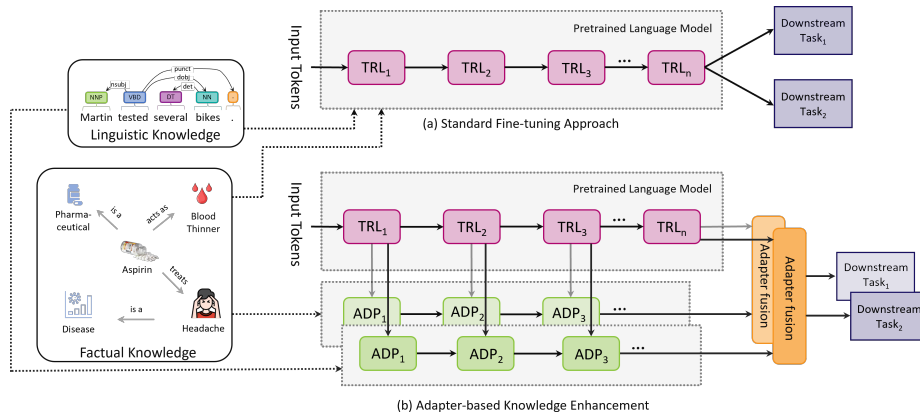


Figure 1: Illustration of a standard fine-tuning versus a knowledge enhancement process. In the example, knowledge from a KG is injected into the model via adapters.

3 ADAPTERS

This section will give an overview of LLM adapters and their applications to establish a conceptual understanding of adapter-based enhancement.

3.1 Overview

Broadly speaking, adapters are small bottleneck feed-forward layers inserted within each layer of an LLM (Houlsby et al., 2019). The small number of additional parameters allows for injecting new data or knowledge without fine-tuning the whole model. This feat is usually accomplished by freezing the layers of the large base model while only updating the adapter weights. Due to the lightweight nature of adapters, this approach leads to short training times with relatively low computing resource requirements. Adapters were mainly used for quick and cheap downstream-task fine-tuning but are now increasingly used for knowledge enhancement. Because it is possible to train adapters individually, they can also be used for multi-task training by specializing one adapter for each task or multi-domain knowledge injection by specializing adapters to different domains (Pfeiffer et al., 2020a).

Leveraging adapters in LLMs also has positive "side effects": Adapters can avoid catastrophic forgetting (the issue when an LLM suddenly deteriorates in performance after fine-tuning) by introducing new task-specific parameters (Houlsby et al., 2019; Pfeiffer et al., 2020a) and, in transfer learning, adapters have even been shown to improve stability and adversarial robustness for downstream tasks (Han et al., 2021). Additionally, adapters have been shown to operate on top of the base model's frozen representation space while largely preserving its structure rather than

on an isolated subspace (Alabi et al., 2024).

3.2 Adapter Types

Houlsby Adapter. The Houlsby Adapter (Houlsby et al., 2019) was the first adapter to be used for transfer learning in NLP. The idea was based on adapter modules initially introduced by Rebuffi et al. (2017) in the computer vision domain. The two main principles stayed the same: Adapters require a relatively small number of parameters compared to the base model and a near-identity initialization. These principles ensure that the total model size grows relatively slowly when more transfer tasks are added, while a near-identity initialization is required for stable training of the adapted model (Houlsby et al., 2019). The optimal architecture of the Houlsby Adapter was determined by meticulous experimenting and tuning; the result can be seen in figure 2. In a classical transformer structure (Vaswani et al., 2017), the adapter module is added once after the multi-headed attention and once after the two feed-forward layers. The modules project the d -dimensional layer features of the base model into a smaller dimension, m , then apply a non-linearity (like ReLU) and project back to d dimensions. The configuration also hosts a skip-connection, and the output of each sub-layer is forwarded to a layer normalization (Ba et al., 2016). Including biases, $2md + d + m$ parameters are added per layer, accounting for only 0.5 to 8 percent of the parameters of the original BERT model used by the authors when setting $m \ll d$.

Bapna and Firat Adapter. In contrast to the Houlsby Adapter, Bapna and Firat (2019) only introduce one adapter module in each transformer layer:

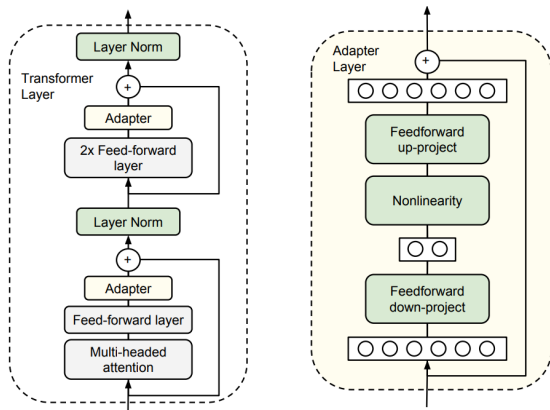


Figure 2: Location of the adapter module in a transformer layer (left) and architecture of the Housby Adapter (right). All green layers are trained on fine-tuning data, including the adapter itself, the layer normalization parameters, and the final classification layer (not shown). Image with permission from Housby et al. (2019).

they keep the adapters after the multi-headed attention (so-called "top" adapters) while dropping the adapters after the feed-forward layers (so-called "bottom" adapters) of the transformer (refer to Figure 2 for better understanding of the component positions). Moreover, while Housby et al. (2019) re-train layer normalization parameters for every domain, Bapna and Firat (2019) "simplify this formulation by leaving the parameters frozen, and introducing new layer normalization parameters for every task, essentially mimicking the structure of the transformer feed-forward layer".

Pfeiffer Adapter and AdapterFusion. The approaches of Bapna and Firat (2019); Housby et al. (2019) did not allow information sharing between tasks. Pfeiffer et al. (2020a) introduce Adapter Fusion, a two-stage algorithm that addresses the sharing of information encapsulated in adapters trained on different tasks. In the first stage, they train the adapters in single-task or multi-task setups for N tasks similar to the Housby Adapter, but only keeping the top adapters, similar to the Bapna and Firat Adapter. As a second step, they combine the set of N adapters with AdapterFusion: They fix the parameters Θ and all adapters Φ , and finally introduce parameters Ψ that learn to combine the N task adapters for the given target task (Pfeiffer et al., 2020a): $\Psi_m \leftarrow \underset{\Psi}{\operatorname{argmin}} L_m(D_m; \Theta, \Phi_1, \dots, \Phi_N, \Psi)$

Here, Ψ_m are the learned AdapterFusion parameters for task m . In the process, the training dataset of m is used twice: once for training the adapters Φ_m and again for training Fusion parameters Ψ_m , which learn to compose the information stored in

the N task adapters (Pfeiffer et al., 2020a). With their approach of separating knowledge extraction and knowledge composition, they further improve the ability of adapters to avoid catastrophic forgetting and interference between tasks and training instabilities. The authors also find that using only a single adapter after the feed-forward layer performs on par with the Housby adapter while requiring only half of the newly introduced adapters (Pfeiffer et al., 2020a). This fact makes the Pfeiffer adapter an attractive choice for many applications, further proven by its popularity among the papers in our review.

K-Adapter. Wang et al. (2020) follow a substantially different approach where the adapters work as "outside plug-ins". In their work, an adapter model consists of K adapter layers (hence the name) that contain N transformer layers and two projection layers. Similar to the approaches above, a skip connection is added but instead applied across the two projection layers. The adapter layers are plugged in among varying transformer layers of the pre-trained model. The authors explain that they concatenate the output hidden feature of the transformer layer in the pre-trained model and the output feature of the former adapter layer as the input feature of the current adapter layer.

Adapter architectures for knowledge enhancement exist that differ from the four adapter types mentioned here. For example, the "Parallel Adapter" (He et al., 2021a) or the adapter architecture by Stickland and Murray (2019)). However, as the upcoming comprehensive literature survey will show, these architectures are either unique to specific papers or have not found broader applications in the field of KELMs. Another popular type of efficient adaptation we'd like to mention for completeness is low-rank adaptation or LoRA (Hu et al., 2022) and its quantized version QLoRA (Dettmers et al., 2023). These approaches do not add new layers but rather enforce a low-rank constraint on the weight updates of the base model's layers. This methodology enables efficient fine-tuning of LLMs and also allows for domain adaption or knowledge enhancement with KGs (Tian et al., 2024).

4 METHODOLOGY

This chapter details the methodology we employed for the systematic literature review. We largely followed the procedure of Kitchenham et al. (2009) for systematic literature reviews in software engineering.

The search strategy for the systematic literature review of this study included literature that fulfilled the following inclusion criteria:

- Peer-reviewed articles from ACM¹, ACL², and IEEE Xplore³
- Article abstracts that match the search string ("*adapter*" OR "*adapter-based*") AND ("*language model*" OR "*nlp*" OR "*natural language processing*") AND ("*injection*" OR "*knowledge*")
- Articles published after February 2, 2019 (publication of the Houlsby Adapter, the first LLM adapter)
- Articles that address the topic of adapter-based knowledge-enhanced language models

We also included three articles not found in the databases because they were fundamental works on the topic of the SLR and frequently referenced. The SLR was concluded in January 2024 and represents the state of research literature up to this point. Additional details on the methodology can be found in Appendix 7.1

5 RESULTS

This section will present the results of the SLR on adapter-based knowledge enhancement.

5.1 Overview

Table 1: Quantitative overview of the literature sources and the selection process.

Source	Initial	Abstract	Full Text
IEEE	28	6	6
ACM	10	6	5
ACL	36	16	13
Others	2	2	2
Total	76	30	26

Table 1 shows the source distribution for all included papers. Fifty-nine papers were found by applying the search string as a query on the ACL, ACM, and IEEE search engines. Due to their importance for the field, we included three additional papers from other sources. These papers were found through online search and paper references during the general research process. In summary, after the abstract screening, 31 articles met all inclusion criteria (and no exclusion criteria). After the full paper screening, 26

¹<https://dl.acm.org/>

²<https://aclanthology.org/>

³<https://ieeexplore.ieee.org/Xplore/home.jsp>

papers formed the final paper pool of the survey. Table 2 gives an overview of all papers included in the survey. It includes information on the adapter type used in the paper, the domain and scope of the paper, and the downstream NLP tasks for which it was developed.

5.2 Data Analysis

We will now give a quantitative analysis showcasing and interpreting quantitative distributions, followed by significant qualitative insights from the papers.

5.2.1 Quantitative Analysis

Yearly Distribution. There has been a significant increase in publications on adapter-based approaches to knowledge-enhanced language models in recent years (Fig. 3). While only two papers were published in 2020, eleven new papers were published in 2023. This trend suggests growing interest and research activity in the domain.




Figure 3: Yearly distribution of publications.

Adapter Type Distribution. Next, we evaluate the popularity and variety of adapter types used across the papers (Fig. 4). The “Pfeiffer” and “Houlsby” adapter types stand out as the most common, which suggests that the closely related underlying architecture is the most popular methodology in the field. This popularity is likely not only an achievement of the adapter’s performance but also due to the well-established Adapter-Hub platform (Pfeiffer et al., 2020b), which, although offering other options, uses adapters with the Pfeiffer configuration by default. This finding showcases a need and trend to build custom adapters well-suited to individual tasks. In the upcoming years, we will likely see many novel adapter architectures. The “K-Adapter” and “Bapna and Firat” adapters are the less frequently mentioned architectures, suggesting that these approaches are less well-established.

Table 2: Overview of the results for the literature survey, including all papers and their references. The task and source acronyms are explained in the appendix. The dotted lines separate the database sources: First come the IEEE papers, then ACM, ACL, and finally, the papers from other sources. For the definition of all task acronyms, see Appendix 7.1.

paper & nickname	adapter type	scope	main source	task
K-MBAN (Zou et al., 2022)	K-Adapter	open	T-REx (Wiki)	RC
/ (Moon et al., 2021)	Houlsby	open	WMT20	MT
CSBERT (Yu and Yang, 2023)	Unique	open	diverse	SL
/ (Qian et al., 2022)	Unique	open	AESRC2020	SR
/ (Li et al., 2023)	Houlsby	closed (multiple)	diverse	SF
CPK (Liu et al., 2023)	K-Adapter	closed (biomed)	Wikipedia	RC, ET, QA
CKGA (Lu et al., 2023)	Unique	open	DBpedia	SC
/ (Nguyen-The et al., 2023)	Pfeiffer	open	diverse	SA
KEBLM (Lai et al., 2023)	Pfeiffer	closed (biomed)	UMLS	QA, NLI, EL
/ (Guo and Guo, 2022)	Unique	open	Ch. Lexicon	NER
/ (Tiwari et al., 2023)	Unique	closed (biomed)	Vis-MDD	TS
AdapterSoup (Chronopoulou et al., 2023)	Bapna and Firat	closed (multiple)	diverse	LM
/ (Wold, 2022)	Houlsby	open	ConceptNet	LAMA
/ (Chronopoulou et al., 2022)	Unique	closed (multiple)	diverse	LM
DS-TOD (Hung et al., 2022)	Pfeiffer	closed (multiple)	CCNet	TOD
/ (Emelin et al., 2022)	Houlsby	closed (multiple)	MultiWOZ	TOD
KnowExpert (Xu et al., 2022)	Bapna and Firat	open	WoW	KGD
mDAPT (Kær Jørgensen et al., 2021)	Pfeiffer	closed (multiple)	WMT20	NER, STC
DAKI (Lu et al., 2021)	K-Adapter	closed (biomed)	UMLS	NLI
/ (Majewska et al., 2021)	Pfeiffer	open	VerbNet	EE
/ (Lauscher et al., 2020)	Houlsby	open	ConceptNet	GLUE
TADA (Hung et al., 2023)	Unique	open	CCNet	TOD, NER, NLI
LeakDistill (Vasylenko et al., 2023)	StructAdapt	open	AMR graph	SMATCH
MixDA (Diao et al., 2023)	Houlsby, Pfeiffer	closed (multiple)	diverse	GLUE, TXM
MoP (Meng et al., 2021)	Pfeiffer	closed (biomed)	UMLS	BLURB
K-Adapter (Wang et al., 2020)	K-Adapter	open	T-REx (Wiki)	RCL, ET, QA




Figure 4: Distribution of adapter types used in the papers.

Domain Analysis. Third, we analyze the distribution of papers across the domain scope and coverage to understand domain-specific preferences in the literature. The open-domain scope is the most popular, with many papers exploring adapter-based approaches within the open domain. The popularity is likely caused by the interest in creating LLMs with a common-sense or world knowledge. Furthermore, the single- and multi-domain approaches are split evenly within the closed-domain papers. Finally, only six papers focus on the biomedical domain, but relative to other domains, the biomedical field is by far the most prominent of all domain-specific approaches. The popularity likely comes

down to the availability of large biomedical KGs, and medicine historically being one of the most active research fields in general science (Cimini et al., 2014).

Task and Source Distribution. A highly diverse range of tasks and sources is being explored throughout the papers, which signifies the versatility and potential of adapter-based approaches across different NLP tasks and domains. Tasks such as Reading Comprehension (RC), Named Entity Recognition (NER), and Question Answering (QA) appear to be popular areas of focus in the literature. This could be because these tasks are the most demanding regarding structural knowledge requirements. The knowledge source distributions show only minimal overlap.

5.2.2 Qualitative Analysis

This section of the analysis highlights recurring themes and individual insights from the papers. Fully summarizing all articles was outside the scope of this survey. However, we still provide an overview of the most common patterns.

General Knowledge. The quantitative analysis showed that open-domain approaches are more pop-

Table 3: performance reports for tasks with highest overlap in the biomedical domain. The metric for HoC is Micro F1; for NCBI, it is F1, while for the other three, it is accuracy. The best results for every task are in bold. "†" denotes that improvements are observed compared to the base model. "†" denotes a statistically significant better result over the base model (T-test, $p < 0.05$), but not all papers report their scores. The baseline performance of the models is taken from the original papers if given. Otherwise, the scores are taken from the MoP results.

↓ model—dataset →	HoC	PubMedQA	BioASQ7b	MedNLI	NCBI
SciBERT-base	80.52 \pm 0.60	57.38 \pm 4.22	75.93 \pm 4.20	81.19 \pm 0.54	88.57
+ MoP	81.79 \pm 0.66 † ↑	54.66 \pm 3.10	78.50 \pm 4.06 † ↑	81.20 \pm 0.37 † ↑	/
+ KEBLM	/	59.00 \pm †	/	82.14 \pm †	93.50 †
+ DAKI	/	/	/	/	/
+ CPK	/	/	/	/	/
BioBERT-base	81.41 \pm 0.59	60.24 \pm 2.32	77.50 \pm 2.92	82.42 \pm 0.59	88.30
+ MoP	82.53 \pm 1.08 † ↑	61.04 \pm 4.81 † ↑	80.79 \pm 4.40 † ↑	82.93 \pm 0.55 † ↑	/
+ KEBLM	/	68.00 † ↑	/	84.24 † ↑	93.20†
+ DAKI	/	/	/	83.41 † ↑	89.01†
+ CPK	/	/	/	81.65	88.42†
PubMedBERT-base	82.25 \pm 0.46	55.84 \pm 1.78	87.71 \pm 4.25	84.18 \pm 0.19	87.82
+ MoP	83.26 \pm 0.32 † ↑	62.84 \pm 2.71 † ↑	90.64 \pm 2.43 † ↑	84.70 \pm 0.19 † ↑	/

ular than their close-domain counterparts. Subsequently, there is also a large variety in the used frameworks, knowledge sources, and overall goals. Two commonly used KGs for general knowledge are ConceptNet (Speer et al., 2017) for common-sense and DBpedia (Auer et al., 2007) for encyclopedic world knowledge. Two example works that use these KGs are Wold (2022) and the CKGA ("knowledge graph-based adapter") by Lu et al. (2023). Wold (2022) train adapter modules on sub-graphs of ConceptNet to inject factual knowledge into LLMs. They evaluate their framework on the Concept-Net Split of the LAMA Probe (Petroni et al., 2019) and see increasing performance while only adding 2.1% of new parameters to the original models. CKGA (Lu et al., 2023) tackles aspect-level sentiment classification by leveraging knowledge from DBpedia. They link aspects to DBpedia and extract an aspect-related sub-graph. Then, a PLM and the KG embedding are utilized to encode the common-sense of entities, where the corresponding knowledge is extracted with graph convolutional networks (Lu et al., 2023).

Linguistic Knowledge. Instead of only including factual knowledge, some works inject additional linguistic knowledge into adapters (Majewska et al., 2021; Zou et al., 2022; Yu and Yang, 2023; Wang et al., 2020). While LLMs already encode a range of syntactic and semantic properties of language, Majewska et al. (2021) explain that LLMs "are still prone to fall back on superficial cues and simple heuristics to solve downstream tasks, rather than leverage deeper linguistic information". Their paper explores the interplay between verb meaning and argument structure. They use knowledge to enhance LLMs with Pfeiffer Adapters to improve English event extraction

and machine translation in other languages. Another example is the work of Zou et al. (2022) on machine reading comprehension (MRC). They proposed the K-MBAN model to integrate linguistic and factual external knowledge into LLMs through K-Adapters.

Domain-Specific Knowledge. Chronopoulou et al. (2022) propose a parameter-efficient approach to domain adaptation using adapters. They represent domains as a hierarchical tree structure where each node in the tree is associated with a set of adapter weights. Their work focused on specializing adapters in website domains like *booking.com* and *yelp.com*. In another instance, Chronopoulou et al. (2023) propose "AdapterSoup". In this framework, they also use adapters for domain-specific tasks but use "an approach that performs weight-space averaging of adapters trained on different domains". AdapterSoup can be helpful in various domain-specific approaches in low-resource settings, especially when only a small amount of data on a specific subdomain is obtainable and closely related adapters are available instead. Earlier, we saw that the biomedical domain is the most prevalent among the closed-domain approaches to adapter-based KELMs. We will briefly examine the relevant works in the following.

Biomedical Knowledge. We have found the works of DAKI (Lu et al., 2021), MoP (Meng et al., 2021), and KEBLM (Lai et al., 2023) to be the most impactful. According to the results of our literature survey, DAKI ("Diverse Adapters for Knowledge Integration") was the first work to use adapters specifically for knowledge enhancement in the biomedical domain. Lu et al. (2021) leverage data from the UMLS meta-thesaurus and UMLS Semantic Network

groups concepts, but also from Wikipedia articles for diseases as proposed by He et al. (2020). Meng et al. (2021) recognize that KGs like UMLS, which can be several gigabytes large, are very expensive to train on in their entirety. They propose to use a "Mixture of Partitions" (MoP), which splits the KG into sub-graphs and combines later with AdapterFusion (Pfeiffer et al., 2020a). Finally, the KEBLM framework's trademark is that it allows the inclusion of a variety of knowledge types from multiple sources into biomedical LLMs. In contrast to DAKI, which uses more than one source, KEBLM includes a knowledge consolidation phase after the knowledge injection, where they teach the fusion layers to effectively combine knowledge from both the original PLM and newly acquired external knowledge by using an extensive collection of unannotated texts (Lai et al., 2023). For completeness, we refer to Kær Jørgensen et al. (2021) for information on the m-DAPT framework, which addresses multi-lingual domain adaptation for biomedical LLMs and KeBioSum (Xie et al., 2022), who state their work is the first study exploring knowledge injection for biomedical extractive summarization.

Performance Insights. In the papers covered by this survey, the performance of adapter-based KELMs on downstream tasks is consistently shown to be better than that of base LMs. For example, Diao et al. (2023) show an increase of +9% on Common-senseCommon-sensettalmor-et-al-2019-commonsenseqa with their mixture-of-adapters approach, while Kær Jørgensen et al. (2021) improve financial text classification on OMP-9 (Schabus et al., 2017) by +4%. While the task variation across domains is too diverse to be shown systematically in our survey, we report in detail on performance comparison in the biomedical domain in Table 3 in the next paragraph. Another interesting insight is found by He et al. (2021b), who show that adapter-based tuning mitigates forgetting issues better than regular fine-tuning since it yields representations with less deviation from those generated by the initial pre-trained language model.

Performance Comparison (Biomedical). Table 3 gives an overview of the downstream task performance of several papers that are included in this survey. The focus lies on the biomedical domain, so the task overlap is high enough for an insightful comparison. The scores are reported for five downstream tasks, namely HoC (Baker et al., 2015), PubMedQA (Gu et al., 2020), BioASQ7b (Nentidis et al., 2020), MedNLI (Romanov and Shivade, 2018), and NCBI (Dogan et al., 2014), as well as three com-

mon biomedical language models (SciBERT (Beltagy et al., 2019), BioBERT (Lee et al., 2019), and PubMedBERT (Gu et al., 2020)). While adapter-based KELMs consistently improve performance in almost all instances, performance boosts across different tasks and models vary strongly. In this specific setting, we recommend the MoP (Meng et al., 2021) and KEBLM (Lai et al., 2023) frameworks since they show the highest performance boosts (e.g., PubMedQA (Jin et al., 2019) accuracy increase of around +7% and +8%, respectively) and overshadow the lower performing CPK and DAKI frameworks in all instances. MoP, in particular, is being continually used for biomedical knowledge enhancement, even in 2024 (Vladika. et al., 2024).

6 CURRENT AND FUTURE TRENDS

In this section, we outline the most important findings and trends and point out promising future directions:

- Adapter-based KELMs are a recent development in NLP, but interest in them is rising fast, with a linear yearly increase of published papers. We predict the growth trend to continue.
- Various adapter architectures exist and are advanced by researchers to be more efficient while preserving task performance. This advancement has temporarily peaked with the Pfeiffer adapter, the most popular type. We expect future work to focus their updates on adapter architecture by overcoming the latency of sequential data processing in adapters and enabling hardware parallelism.
- Research focuses on the open domain – injecting general world knowledge into models. Within the closed domain, the biomedical domain is the most popular, owing to the existence of large biomedical KGs. We foresee the potential to apply adapter-based KELMs to other highly structured domains, such as the legal or financial domain (documents with rigid structure).
- The largest improvements in task performance is seen in knowledge-intensive tasks like question answering and text classification, with more minor improvements for reasoning tasks like entailment recognition. Generative tasks, other than dialogue modeling, are rather unexplored. We envision a future popular use case that could use knowledge enhancement to improve the factuality and informativeness of generated text.

7 CONCLUSION

In this paper, we conducted a systematic literature review on approaches to enhancing language models with external knowledge using adapter modules. We portrayed which adapter-based approaches exist and how they compare to each other. We showed there is a steady growth of interest in this domain with each new year and highlighted the most popular adapter architectures (with "Pfeiffer" as the predominant one). We discovered a balance in popularity between open-domain approaches focusing on integrating general world knowledge into models and closed-domain approaches focusing on specialized fields, with biomedical as the most popular domain. With our review, we contribute a novel and extensive resource for this nascent yet fast-growing field, and we hope it will be a useful entry point for other researchers in the future.

7.1 Limitations

Our literature search methodology follows a strict search string and exclusion criteria. Subsequently, we might have overlooked some relevant work on adapter-based KELMs. Also, some reviewed papers were not adequately analyzed in this work due to space constraints, leading to potentially missing insights and a non-complete representation of the state of research on adapter-based enhancement. Additionally, due to the variety of applications and domains, we could not give precise guidelines on what methods to use under which circumstances. Still, we aimed to report on the most common patterns and trends discovered in the literature, which can serve as a basis for future research.

REFERENCES

- Alabi, J., Mosbach, M., Eyal, M., Klakow, D., and Geva, M. (2024). The hidden space of transformer language adapters. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*, pages 6588–6607. Association for Computational Linguistics.
- Allen, B. P., Stork, L., and Groth, P. (2023). Knowledge Engineering Using Large Language Models. *Transactions on Graph Data and Knowledge*, 1(1):3:1–3:19.
- Auer, S., Bizer, C., Kobilarov, G., Lehmann, J., Cyganiak, R., and Ives, Z. (2007). Dbpedia: A nucleus for a web of open data. In *The Semantic Web*, pages 722–735, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Ba, J., Kiros, J. R., and Hinton, G. E. (2016). Layer normalization. *ArXiv*, abs/1607.06450.
- Baker, S., Silins, I., Guo, Y., Ali, I., Höggberg, J., Stenius, U., and Korhonen, A. (2015). Automatic semantic classification of scientific literature according to the hallmarks of cancer. *Bioinformatics*, 32(3):432–440.
- Bapna, A. and Firat, O. (2019). Simple, scalable adaptation for neural machine translation. In *Proceedings of the 2019 Conference on EMNLP-IJCNLP*, pages 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Beltagy, I., Lo, K., and Cohan, A. (2019). Scibert: A pre-trained language model for scientific text. In *Conference on Empirical Methods in Natural Language Processing*.
- Bodenreider, O. (2004). The unified medical language system (umls): integrating biomedical terminology. *Nucleic acids research*, 32.
- Budzianowski, P., Wen, T.-H., Tseng, B.-H., Casanueva, I., Ultes, S., Ramadan, O., and Gašić, M. (2018). MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 5016–5026, Brussels, Belgium. Association for Computational Linguistics.
- Cai, S. and Knight, K. (2013). Smatch: an evaluation metric for semantic feature structures. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 748–752, Sofia, Bulgaria. ACL.
- Chronopoulou, A., Peters, M., and Dodge, J. (2022). Efficient hierarchical domain adaptation for pretrained language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1336–1351, Seattle, United States. Association for Computational Linguistics.
- Chronopoulou, A., Peters, M., Fraser, A., and Dodge, J. (2023). AdapterSoup: Weight averaging to improve generalization of pretrained language models. In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 2054–2063, Dubrovnik, Croatia. Association for Computational Linguistics.
- Cimini, G., Gabrielli, A., and Labini, F. (2014). The scientific competitiveness of nations. *PLoS one*, 9.
- Colon-Hernandez, P., Havasi, C., Alonso, J. B., Huggins, M., and Breazeal, C. (2021). Combining pre-trained language models and structured knowledge. *ArXiv*, abs/2101.12294.
- Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. (2023). Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*.
- Diao, S., Xu, T., Xu, R., Wang, J., and Zhang, T. (2023). Mixture-of-domain-adapters: Decoupling and injecting domain knowledge to pre-trained language models' memories. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*, pages 5113–5129, Toronto, Canada. Association for Computational Linguistics.
- Dinan, E., Roller, S., Shuster, K., Fan, A., Auli, M., and Weston, J. (2018). Wizard of wikipedia: Knowledge-powered conversational agents. *ArXiv*, abs/1811.01241.
- Dogan, R. I., Leaman, R., and Lu, Z. (2014). Ncbi disease

- corpus: A resource for disease name recognition and concept normalization. *Journal of biomedical informatics*, 47:1–10.
- Elsahar, H. (2017). T-Rex : A Large Scale Alignment of Natural Language with Knowledge Base Triples [NIF SAMPLE].
- Emelin, D., Bonadiman, D., Alqahtani, S., Zhang, Y., and Mansour, S. (2022). Injecting domain knowledge in language models for task-oriented dialogue systems. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11962–11974. Association for Computational Linguistics.
- Fichtl, A. (2024). Evaluating adapter-based knowledge-enhanced language models in the biomedical domain. Master's thesis, Technical University of Munich, Munich, Germany.
- Gu, Y., Tinn, R., Cheng, H., Lucas, M. R., Usuyama, N., Liu, X., Naumann, T., Gao, J., and Poon, H. (2020). Domain-specific language model pre-training for biomedical natural language processing. *ACM Transactions on Computing for Healthcare (HEALTH)*, 3:1 – 23.
- Guo, Q. and Guo, Y. (2022). Lexicon enhanced chinese named entity recognition with pointer network. *Neural Computing and Applications*.
- Han, W., Pang, B., and Wu, Y. N. (2021). Robust transfer learning with pretrained language models through adapters. *ArXiv*, abs/2108.02340.
- He, J., Zhou, C., Ma, X., Berg-Kirkpatrick, T., and Neubig, G. (2021a). Towards a unified view of parameter-efficient transfer learning. *ArXiv*, abs/2110.04366.
- He, R., Liu, L., Ye, H., Tan, Q., Ding, B., Cheng, L., Low, J.-W., Bing, L., and Si, L. (2021b). On the effectiveness of adapter-based tuning for pretrained language model adaptation.
- He, Y., Zhu, Z., Zhang, Y., Chen, Q., and Caverlee, J. (2020). Infusing Disease Knowledge into BERT for Health Question Answering, Medical Inference and Disease Name Recognition. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4604–4614, Online. Association for Computational Linguistics.
- Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., de Melo, G., Gutiérrez, C., Kirrane, S., Gayo, J. E. L., Navigli, R., Neumaier, S., Ngomo, A.-C. N., Polleres, A., Rashid, S. M., Rula, A., Schmelzeisen, L., Sequeda, J., Staab, S., and Zimmermann, A. (2020). Knowledge graphs. *ACM Computing Surveys (CSUR)*, 54:1 – 37.
- Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., de Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. (2019). Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. (2022). LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*.
- Hu, L., Liu, Z., Zhao, Z., Hou, L., Nie, L., and Li, J. (2023). A survey of knowledge enhanced pre-trained language models.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., and Liu, T. (2023). A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, and Open Questions. *arXiv e-prints*, page arXiv:2311.05232.
- Hung, C.-C., Lange, L., and Strötgen, J. (2023). TADA: Efficient task-agnostic domain adaptation for transformers. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 487–503, Toronto, Canada. Association for Computational Linguistics.
- Hung, C.-C., Lauscher, A., Ponzetto, S., and Glavaš, G. (2022). DS-TOD: Efficient domain specialization for task-oriented dialog. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 891–904. Association for Computational Linguistics.
- Ji, S., Pan, S., Cambria, E., Marttinen, P., and Yu, P. S. (2020). A survey on knowledge graphs: Representation, acquisition, and applications. *IEEE Transactions on Neural Networks and Learning Systems*, 33:494–514.
- Jin, Q., Dhingra, B., Liu, Z., Cohen, W., and Lu, X. (2019). PubMedQA: A dataset for biomedical research question answering. In *Proceedings of the 2019 Conference on EMNLP-IJCNLP*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Kær Jørgensen, R., Hartmann, M., Dai, X., and Elliott, D. (2021). mDAPT: Multilingual domain adaptive pretraining in a single model. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3404–3418. Association for Computational Linguistics.
- Khadir, A. C., Aliane, H., and Guessoum, A. (2021). Ontology learning: Grand tour and challenges. *Computer Science Review*, 39:100339.
- Kitchenham, B., Pearl Brereton, O., Budgen, D., Turner, M., Bailey, J., and Linkman, S. (2009). Systematic literature reviews in software engineering – a systematic literature review. *Information and Software Technology*, 51(1):7–15. Special Section - Most Cited Articles in 2002 and Regular Research Papers.
- Lai, T. M., Zhai, C., and Ji, H. (2023). Keblm: Knowledge-enhanced biomedical language models. *Journal of Biomedical Informatics*, 143:104392.
- Lauscher, A., Majewska, O., Ribeiro, L. F. R., Gurevych, I., Rozanov, N., and Glavaš, G. (2020). Common sense or world knowledge? investigating adapter-based knowledge injection into pretrained transformers. In *Proceedings of Deep Learning Inside Out (DeeLIO)*, pages 43–49. Association for Computational Linguistics.
- Lee, J., Yoon, W., Kim, S., Kim, D., Kim, S., So, C. H., and Kang, J. (2019). BioBERT: a pre-trained biomedical language representation model for biomedical text mining. *Bioinformatics*, 36(4):1234–1240.
- Li, B., Hwang, D., Huo, Z., Bai, J., Prakash, G., Sainath, T. N., Chai Sim, K., Zhang, Y., Han, W., Strohmaier, T., and Beaufays, F. (2023). Efficient domain adaptation for speech foundation models. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5.

- Liu, C., Zhang, S., Li, C., and Zhao, H. (2023). Cpk-adapter: Infusing medical knowledge into k-adapter with continuous prompt. In *2023 8th International Conference on Intelligent Computing and Signal Processing (ICSP)*, pages 1017–1023, Los Alamitos, CA, USA. IEEE Computer Society.
- Lu, G., Yu, H., Yan, Z., and Xue, Y. (2023). Commonsense knowledge graph-based adapter for aspect-level sentiment classification. *Neurocomputing*, 534:67–76.
- Lu, Q., Dou, D., and Nguyen, T. H. (2021). Parameter-efficient domain knowledge integration from multiple sources for biomedical pre-trained language models. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3855–3865. Association for Computational Linguistics.
- Majewska, O., Vulić, I., Glavaš, G., Ponti, E. M., and Korhonen, A. (2021). Verb knowledge injection for multilingual event processing. In *Proceedings of the 59th Annual Meeting of the ACL-IJCNLP*, pages 6952–6969. ACL.
- Meng, Z., Liu, F., Clark, T. H., Shareghi, E., and Collier, N. (2021). Mixture-of-partitions: Infusing large biomedical knowledge graphs into bert. *ArXiv*, abs/2109.04810.
- Moon, H., Park, C., Eo, S., Seo, J., and Lim, H. (2021). An empirical study on automatic post editing for neural machine translation. *IEEE Access*, 9:123754–123763.
- Nentidis, A., Bougiatiotis, K., Krithara, A., and Paliouras, G. (2020). Results of the seventh edition of the bioasq challenge. In *Machine Learning and Knowledge Discovery in Databases*, pages 553–568, Cham. Springer International Publishing.
- Nguyen-The, M., Lamghari, S., Bilodeau, G.-A., and Rockemann, J. (2023). Leveraging sentiment analysis knowledge to solve emotion detection tasks. In *Pattern Recognition, Computer Vision, and Image Processing. ICPR 2022 International Workshops and Challenges*, pages 405–416. Springer Nature Switzerland.
- Petroni, F., Rocktäschel, T., Lewis, P., Bakhtin, A., Wu, Y., Miller, A. H., and Riedel, S. (2019). Language models as knowledge bases? *ArXiv*, abs/1909.01066.
- Pfeiffer, J., Kamath, A., Rücklé, A., Cho, K., and Gurevych, I. (2020a). Adapterfusion: Non-destructive task composition for transfer learning. *ArXiv*, abs/2005.00247.
- Pfeiffer, J., Rücklé, A., Poth, C., Kamath, A., Vulić, I., Ruder, S., Cho, K., and Gurevych, I. (2020b). Adapterhub: A framework for adapting transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54.
- Pfeiffer, J., Ruder, S., Vulić, I., and Ponti, E. M. (2024). Modular deep learning. *Transactions of Machine Learning Research*.
- Qian, Y., Gong, X., and Huang, H. (2022). Layer-wise fast adaptation for end-to-end multi-accent speech recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 30:2842–2853.
- Rebuffi, S.-A., Bilen, H., and Vedaldi, A. (2017). Learning multiple visual domains with residual adapters. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 506–516.
- Romanov, A. and Shivade, C. (2018). Lessons from natural language inference in the clinical domain. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1586–1596, Brussels, Belgium. ACL.
- Rosset, C., Xiong, C., Phan, M., Song, X., Bennett, P., and Tiwary, S. (2020). Knowledge-Aware Language Model Pretraining. *arXiv e-prints*, page arXiv:2007.00655.
- Schabus, D., Skowron, M., and Trapp, M. (2017). One million posts: A data set of german online discussions. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, page 1241–1244, New York, NY, USA. Association for Computing Machinery.
- Schneider, P., Schopf, T., Vladika, J., Galkin, M., Simperl, E., and Matthes, F. (2022). A decade of knowledge graphs in natural language processing: A survey. In *Proceedings of the 2nd AACL-IJCNLP*, pages 601–614. Association for Computational Linguistics.
- Schuler, K. K. (2006). *VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon*. PhD thesis, University of Pennsylvania.
- Shi, X., Yu, F., Lu, Y., Liang, Y., Feng, Q., Wang, D., Qian, Y., and Xie, L. (2021). The accented english speech recognition challenge 2020: Open datasets, tracks, baselines, results and methods. *ICASSP 2021 - IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6918–6922.
- Speer, R., Chin, J., and Havasi, C. (2017). Conceptnet 5.5: An open multilingual graph of general knowledge. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*. AAAI Press.
- Stickland, A. C. and Murray, I. (2019). BERT and PALs: Projected attention layers for efficient adaptation in multi-task learning. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5986–5995. PMLR.
- Tian, S., Luo, Y., Xu, T., Yuan, C., Jiang, H., Wei, C., and Wang, X. (2024). KG-adapter: Enabling knowledge graph integration in large language models through parameter-efficient fine-tuning. In *Findings of the Association for Computational Linguistics ACL 2024*, pages 3813–3828, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Tiwari, A., Manthena, M., Saha, S., Bhattacharyya, P., Dhar, M., and Tiwari, S. (2022). Dr. can see: Towards a multi-modal disease diagnosis virtual assistant. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, page 1935–1944, New York, NY, USA. Association for Computing Machinery.
- Tiwari, A., Saha, A., Saha, S., Bhattacharyya, P., and Dhar, M. (2023). Experience and evidence are the eyes of an excellent summarizer! towards knowledge infused multi-modal clinical conversation summarization. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, page 2452–2461. Association for Computing Machinery.

- Vaswani, A., Shazeer, N. M., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. In *NIPS*.
- Vasylenko, P., Huguet Cabot, P. L., Martínez Lorenzo, A. C., and Navigli, R. (2023). Incorporating graph information in transformer-based AMR parsing. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1995–2011, Toronto, Canada. Association for Computational Linguistics.
- Vladika, J., Fichtl, A., and Matthes, F. (2024). Diversifying knowledge enhancement of biomedical language models using adapter modules and knowledge graphs. In *Proceedings of the 16th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART*, pages 376–387. INSTICC, SciTePress.
- Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., and Bowman, S. R. (2019). Glue: A multi-task benchmark and analysis platform for natural language understanding.
- Wang, R., Tang, D., Duan, N., Wei, Z., Huang, X., Ji, J., Cao, G., Jiang, D., and Zhou, M. (2020). K-adapter: Infusing knowledge into pre-trained models with adapters.
- Wei, X., Wang, S., Zhang, D., Bhatia, P., and Arnold, A. O. (2021). Knowledge enhanced pretrained language models: A comprehensive survey. *ArXiv*, abs/2110.08455.
- Wenzek, G., Lachaux, M.-A., Conneau, A., Chaudhary, V., Guzmán, F., Joulin, A., and Grave, E. (2020). CCNet: Extracting high quality monolingual datasets from web crawl data. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4003–4012.
- Wold, S. (2022). The effectiveness of masked language modeling and adapters for factual knowledge injection. In *Proceedings of TextGraphs-16*, pages 54–59, Gyeongju, Republic of Korea. Association for Computational Linguistics.
- Xie, Q., Bishop, J. A., Tiwari, P., and Ananiadou, S. (2022). Pre-trained language models with domain knowledge for biomedical extractive summarization. *Knowledge-Based Systems*, 252:109460.
- Xu, Y., Ishii, E., Cahyawijaya, S., Liu, Z., Winata, G. I., Madotto, A., Su, D., and Fung, P. (2022). Retrieval-free knowledge-grounded dialogue response generation with adapters. In *Proceedings of the Second Dial-Doc Workshop on Document-grounded Dialogue and Conversational Question Answering*, pages 93–107. Association for Computational Linguistics.
- Yu, S. and Yang, Y. (2023). A new feature fusion method based on pre-training model for sequence labeling. In *2023 6th International Conference on Data Storage and Data Engineering (DSDE)*, pages 26–31.
- Zou, D., Zhang, X., Song, X., Yu, Y., Yang, Y., and Xi, K. (2022). Multiway bidirectional attention and external knowledge for multiple-choice reading comprehension. In *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 694–699.

APPENDIX

Supplementary Survey Data

Methodology

Articles on the following topics were excluded:

- Articles published before February 2, 2019
- Duplicate versions of the same article (when multiple versions of an article were found in different journals, only the most recent version was included)
- Articles where Adapters were used for NLP, but for use-cases other than knowledge-enhancement (such as few-shot learning or model debiasing)
- Articles written in a language other than English

Extracted data:

- Publication source, date, and full reference
- Main topic area
- Facts of interest such as adapter architecture, domain, and downstream tasks within the papers
- A short summary of the study, including the main research questions and the answers

Acronyms

- AESRC2020: Accented English Speech Recognition Challenge 2020 (Shi et al., 2021)
- BLURB: Biomedical Language Understanding and Reasoning Benchmark (Gu et al., 2020)
- CCNet: Common Crawl Net (Wenzek et al., 2020)
- EE: Event Extraction
- EL: Entity Linking
- ES: Extractive Summarization
- ET: Entity Typing
- GLUE: General Language Understanding Evaluation (Wang et al., 2019)
- IE: Information Extraction
- KGD: Knowledge-grounded Dialogue
- LAMA: Concept-Net Split of LAMA Probe (Petroni et al., 2019)
- LM: Language Modeling
- MT: Machine Translation
- MultiWOZ: Multi-Domain Wizard-of-Oz dataset (Budzianowski et al., 2018)
- NER: Named Entity Recognition
- NLI: Natural Language Inference
- OOD: Out-of-domain Detection
- QA: Question Answering
- RC: Reading Comprehension
- RE: Relation Extraction
- RCL: Relation Classification
- SA: Sentiment Analysis
- SC: Sentiment Classification
- SF: Speech Foundation
- SL: Sequence Labelling
- SMATCH: Semantic Match Score (Cai and Knight, 2013)
- SR: Speech Recognition
- STC: Sentence Classification
- TC: Text Classification
- TOD: Task-Oriented dialogue

- T-REx: A Large Scale Alignment of Natural Language with Knowledge Base Triples (Elsahar, 2017)
- VerbNet: A Broad-Coverage, Comprehensive Verb Lexicon (Schuler, 2006)
- Vis-MDD: Visual Medical Disease Diagnosis (Tiwari et al., 2022)
- WMT20: Workshop on Machine Translation 2020
- WoW: Wizard-of-Wikipedia (Dinan et al., 2018)

