

## Highlights

### **Knowledge Prompting: How Knowledge Engineers Use Generative AI**

Elisavet Koutsiana, Johanna Walker, Michelle Nwachukwu, Bohui Zhang, Albert Meroño-Peñuela, Elena Simperl

- Large Language Models (LLMs) can enhance efficiency in knowledge engineering tasks like knowledge graph construction, but their effective use requires new skills, such as advanced prompting and responsible AI integration.
- The introduction of tools such as KG Cards can address ethical and transparency challenges in LLM-assisted knowledge engineering by guiding practitioners towards safer and more responsible practices.
- Generative AI offers promising potential for assisting less experienced users in KE tasks, but more research is needed to mitigate biases and ensure reliability throughout the KG development lifecycle.

# Knowledge Prompting: How Knowledge Engineers Use Generative AI

Elisavet Koutsiana<sup>a,\*</sup>, Johanna Walker<sup>a,\*</sup>, Michelle Nwachukwu<sup>a</sup>, Bohui Zhang<sup>a</sup>,  
Albert Meroño-Peñuela<sup>a</sup> and Elena Simperl<sup>a,\*</sup>

<sup>a</sup>Department of Informatics, King's College London, London, WC2B 4BG, United Kingdom

## ARTICLE INFO

### Keywords:

Knowledge Graph  
Knowledge Engineering  
Large Language Models  
Hackathon  
Ethnographic Study  
Interviews  
Knowledge Engineers Skills  
Bias

## ABSTRACT

Despite many advances in knowledge engineering (KE), challenges remain in areas such as engineering knowledge graphs (KGs) at scale, automating tasks, and keeping pace with evolving domain knowledge. KE has used NLP demonstrating notable advantages in knowledge-intensive tasks, but the most effective use of generative AI to support knowledge engineers across the KE activities is still in its infancy. To explore how generative AI may enhance KE and change existing KE practices, we conducted a multi-method study during a KE hackathon. We investigated participants' views on the use of generative AI, the challenges they face, the skills they may need to integrate generative AI into their practices, and how they use generative AI responsibly. We found participants felt LLMs could indeed contribute to improving efficiency when engineering KGs, but presented increased challenges around the already complex issues of evaluating KE task success. We discovered prompting to be a useful but undervalued skill for knowledge engineers working with LLMs, and note that NLP skills may become more relevant across more roles in KE workflows. Integrating generative AI into KE tasks needs to be done with awareness of potential risks and harms. Given the limited ethical training most knowledge engineers receive, solutions such as our proposed 'KG Cards' based on Data Cards could be a useful guide for KG construction. Our findings can support designers of KE AI copilots, KE researchers, and practitioners using advanced AI to develop trustworthy applications, propose new methodologies for KE and operate new technologies responsibly.

;

## 1. Introduction

Knowledge engineering (KE) is the process of capturing, representing and maintaining knowledge in a machine-readable way. It involves the development of knowledge-based systems (KBSs), and tools and principles to operationalise these [163, 32]. Since Google launched its knowledge graph (KG) over a decade ago, KGs (especially large ones with millions of entities and billions of facts in triples) have become the predominant way to represent, store, and utilise knowledge, offering several advantages across various domains. Recently they have been used to enhance pre-trained language models (PLMs) [100] and retrieval-augmented generation (RAG) systems [90] and offer explanations in artificial intelligence (AI) applications [167]. The techniques used for engineering these KGs have benefited from development in Natural Language Processing (NLP) techniques across tasks such as knowledge extraction, completion, and inconsistency detection [149]. More recently, the most advanced branch of NLP, large language models (LLMs), such as GPT models [24, 121] and the LLaMA series [168, 169], have demonstrated notable potential for

knowledge-intensive tasks, with applications for KG construction through methods such as knowledge probing [132, 6, 193]. The integration of KGs and LLMs has become a growing trend that fosters a mutually beneficial relationship: LLMs enhance KG construction and maintenance, while KGs are used to train, prompt, augment, and evaluate LLMs [123].


Although research efforts to unify LLM and KG have gained popularity, debates and concerns continue. Critical perspectives around misinformation, ethical issues, privacy and security, bias, transparency, accountability and many others have emerged at a broader level and apply to LLM-based systems [124, 128]. Regardless of improvements in automated processes, it may be that KE will always require some form of human-in-the-loop participation to build trust with end-users and stakeholders [58].

Despite the increasing adoption of LLMs in KE, there is a notable gap in user studies focusing on the user-centric aspects of LLM technology in KE. To explore the opportunities and transformations in KE practices enabled by generative AI, it is essential to analyse how researchers and knowledge engineers interact with this emerging technology in standard KE tasks. To this end, the idea of a hackathon focused on LLMs in KE originated during the Dagstuhl Seminar on "Knowledge Graphs and their role in the Knowledge Engineering of the 21st Century" held in September 2022 [58].<sup>1</sup> Building on this idea, we organised a four-day hackathon, inviting researchers and practitioners from the AI and KE communities to explore generative AI's potential in automating KE tasks.<sup>2</sup>

\*Corresponding author

 elisavet.koutsiana@kcl.ac.uk (E. Koutsiana);

johanna.walker@kcl.ac.uk (J. Walker); michelle.nwachukwu@kcl.ac.uk (M. Nwachukwu); bohui.zhang@kcl.ac.uk (B. Zhang); albert.meron@kcl.ac.uk (A. Meroño-Peñuela); elena.simperl@kcl.ac.uk (E. Simperl)

 <https://www.albertmeronyo.org> (A. Meroño-Peñuela);

<https://elenasimperl.eu> (E. Simperl)

ORCID(s): 0000-0001-6544-0435 (E. Koutsiana); 0000-0002-5498-8670

(J. Walker); 0009-0007-7546-4888 (M. Nwachukwu); 0000-0001-5430-1624

(B. Zhang); 0000-0003-4646-5842 (A. Meroño-Peñuela);

0000-0003-1722-947X (E. Simperl)

<sup>1</sup><https://www.dagstuhl.de/22372>

<sup>2</sup><https://king-s-knowledge-graph-lab.github.io/knowledge-prompting-hackathon/>

The user study presented in this work leverages the hackathon as a test bed. Our investigation included an ethnographic study conducted during the event, complemented by post-event semi-structured interviews and a review of the hackathon's documentary outputs. During the four-day workshop, we observed knowledge engineers working on KE research tasks using generative AI models. KE is a broad field involving numerous tasks, tools and practices. In this study, we aim address a set of curated topics such as knowledge graphs and their construction with human-in-the-loop. Specifically, we examined the obstacles that participants encountered, the skills they needed, and the safety testing practices they used. Our research questions are:

1. What are the main challenges experienced by knowledge engineers when using generative AI for KE tasks?
2. How do knowledge engineers evaluate generative AI output for their practices?
3. What skills do knowledge engineers need to incorporate generative AI into their practice?
4. How aware are knowledge engineers of using generative AI responsibly?
5. What factors may affect knowledge engineers' trust and uptake of generative AI technology?

We contribute one of the first user studies on the interaction of knowledge engineers with generative AI. Additionally, we identify key areas of strength and challenge for KE in the era of generative AI. Finally, we propose methods to improve transparency and explainability, including Data Cards for KGs and Model Cards for KG embedding models. It is worth noting that this study was conducted in 2023, and to this day, there is a rapid evolution in the capabilities of generative AI. However, most of the issues raised by participants in the hackathon have not been satisfactorily resolved in the intervening period, and persist as issues that need addressing. Our findings give direction to designers of KE AI copilots to create responsible applications, guide KE researchers and practitioners in developing new methodologies, and advise practitioners wishing to use advanced AI technologies responsibly.

## 2. Background and Related Works

### 2.1. Knowledge Engineering

Knowledge engineering, the branch of AI concerned with building and managing knowledge-based systems [150, 164], has changed dramatically with the latest innovations in machine learning (ML), NLP, and computer vision. The process of constructing a KG usually involves acquiring knowledge, processing it, and deploying the knowledge graph [44, 166]. And yet, as the most recent advances in NLP (especially LLMs) and generative AI demonstrate, questions of how to capture and encode domain knowledge in computational representation remain as challenging as ever [147].

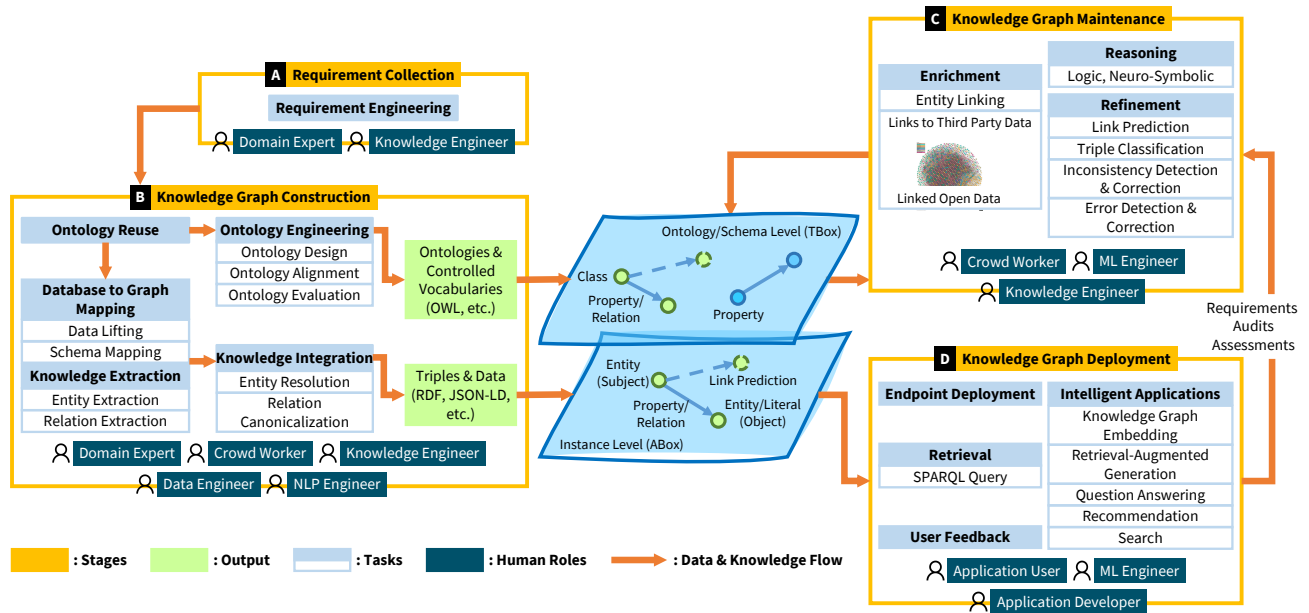
Figure 1 shows that the current KG lifecycle consists of four stages. This constitutes a mix of automated and

manual capabilities and contributions from several stakeholder groups: knowledge engineers, ML specialists, subject domain experts, online volunteers, and crowdsourcing services, as well as developers of applications using KGs. As the figure suggests, **KGs are interacting with AI capabilities in complex ways**. Human-in-the-loop tasks in the KG lifecycle increasingly use ML models with varying levels of interpretability.

On the left side of the figure, at stage **A**, which is an entry point and essential step of the KG lifecycle, knowledge engineers (e.g., ontologists) and KG stakeholders (e.g., domain experts) first determine the scope of work and the success criteria [79]. With the support of description logics, knowledge formalised within a well-structured and expressive TBox that contains terminological axioms can define the relevant vocabulary for a specific application domain. This process, known as ontology engineering (OE), is primarily manual or semi-automatic [12]. Ontology requirement engineering, which serves as the starting point for OE, typically involves writing user stories and deriving competency questions (CQs) from these stories [78].

KG construction generally involves extracting knowledge from multiple heterogeneous data sources, ranging from structured and semi-structured to unstructured data, and integrating it into KGs [166]. At stage **B**, knowledge engineers and other specialists may reuse standard or novel ontologies to build KGs through data lifting and knowledge extraction. Data from multiple sources is transformed into KGs using ML techniques for tasks such as named entity recognition [181], relation extraction [96] and entity reconciliation [152]. Core KE tasks at this stage include knowledge acquisition and representation. Technologies and end-user tools supporting knowledge acquisition have advanced significantly to meet the scale requirements of modern KGs and lead a growing trend toward adoption of end-to-end approaches [149, 182], which we discuss further in section 2.2. However, the most effective methods for knowledge representation still rely on human oversight at various levels [156, 158]. Increasingly, human input is focused on augmenting or validating algorithmic suggestions [166].

KGs can also be created on a larger scale through human collaboration, utilizing crowdsourcing and collaborative-editing platforms [68]. Crowd workers and editors have important roles throughout the KG lifecycle, especially in knowledge representation and updates, where annotation tasks such as quizzes and voting are often designed to harness their background knowledge [2, 143, 84]. Although KGs developed using these methods may face quality issues, including errors [133, 154], disagreements [85], and bias [68], similar transparency challenges exist in crowdsourcing for supervised ML. These challenges stem from the transparency limitations of the digital services commonly used for such tasks, such as Prolific and Amazon's Mechanical Turk, which are black-box, proprietary platforms, making results difficult to replicate or reproduce [139]. Furthermore, educating crowd workers to perform annotation tasks effectively is a non-trivial task [143].



**Figure 1:** The KG lifecycle today, illustrating the stages from requirement collection to deployment. Each phase — requirement elicitation, KG construction, maintenance, and deployment — involves collaboration across roles like domain experts, knowledge engineers, and ML engineers to build, enrich, and deploy KGs for intelligent applications.

On the right-hand side of the figure, KG maintenance (stage C) is prompted by source updates from stage B, and requirements, audits, and assessments from stage D. To further increase their completeness, correctness, and utility, KGs are refined by completion tasks such as link prediction [145] and error detection and correction tasks, etc [28]. Quality assessment of KGs is essential following construction and enrichment stages to ensure they are fit for specific purposes. Key dimensions for evaluation include accuracy, coverage, consistency, completeness and understandability [68]. Accuracy ensures that data represents real-world entities [188], while completeness and coverage assess how well the KG represents all domain-relevant information [31]. Validity assessment is supported by constraints expressed in shape languages, including SHACL and ShEx [51]. By measuring whether the graph conforms to the constraints, knowledge engineers can prevent issues where essential parts of the graph are missing or contradictory. Additional dimensions like coherency, conciseness, and understandability support effective usage by minimizing redundancy and enhancing interpretability, all contributing to the graph's overall utility for applications [188].

At stage D, KGs support a variety of use cases. They are utilised for querying and reasoning in applications such as search [177], question answering [27, 62], and retrieval-augmented generation [90, 50]. Information can be derived from KGs through deductive (e.g., logical rules) or inductive methods (e.g., continuous graph embeddings) [68]. To ensure trustworthiness and legal compliance, both approaches must be transparent and accountable to users [146, 16].

The roles and required skills in ontology and knowledge engineering have evolved significantly. Beyond technical expertise, knowledge engineers are expected to possess

strong communication skills, conceptual thinking, and organizational abilities [103, 112]. In a collaborative ontology engineering, each team member can play multiple roles, depending on the types of contributions and the technology used [157]. To actively involve domain experts in ontology development, technologies such as controlled natural language, semantic wikis, intelligent user interfaces and social computing are proposed [33]. However, the rise of generative AI is reshaping these practices. Collaboration between humans and AI agents is redefining ontology and knowledge engineering, streamlining workflows, and altering the responsibilities of knowledge engineers [8].

## 2.2. Automation in Knowledge Engineering

Automated methods have long played critical roles in KE, with NLP supporting tasks ranging from knowledge acquisition to completion. The emergence of LLMs has further propelled automated methods into the mainstream, leading to substantial advancements [185]. However, this progress also introduces new challenges in human-AI collaboration and quality assurance.

Once the data source is determined, the construction of KGs typically begins with information extraction, which involves identifying key terms — entities and relations. The task of identifying entities is referred to as *named entity recognition* (NER). A named entity is a word or phrase that is consistently standard for a real-world thing, e.g., an organization, person, location, etc. The task can be either formulated as a sequence labelling or a span classification task. Current automated NER methods are deep learning-based and consist of input representation at different levels (character, word, hybrid), context encoders that adopt LSTM



and Transformers, and tag decoders using models such as conditional random field (CRF) [91].

After determining entities, the next step is to extract relations that link two entities, which is called *relation extraction* (RE). The task can also be formulated as a span classification task. Given the corpus and a pair of entity strings, the task aims to extract a relation between the entities. Depending on the context given, RE models can be performed at sentence-level, dialogue-level, and document-level. Distant supervision approaches are dominant in current RE tasks, which combine the advantages of prior semi-supervised and unsupervised approaches and leverage a vast amount of semi-structured data stored in a KG [107, 159].

The knowledge extracted in previous steps may contain duplicates, requiring further cleaning and integration. The next step, *entity linking* (also known as entity alignment, resolution, or matching), is a crucial step of data and knowledge integration. Given two data sources  $A$  and  $B$ , each containing entities with a shared set of attributes, entity linking aims to identify pairs  $(e_1, e_2)$ , where  $e_1 \in A$  and  $e_2 \in B$ , such that both refer to the same real-world object. Due to the potential quadratic complexity of matching ( $|A| \times |B|$ ), a pre-processing step known as blocking [125] is often used to reduce the complexity. According to Shen et al. [153], an entity linking system typically consists of three modules: candidate entity generation, candidate entity ranking, and unlinkable mention prediction. Furthermore, PLMs have demonstrated state-of-the-art performance on large-scale, real-world entity linking tasks [92, 162].

The third category is knowledge completion. Most KGs are inherently incomplete, necessitating techniques such as *link prediction* (also referred to as KG completion or reasoning) to infer missing information. Specifically, given a query  $q$  with either missing tail entities  $(h, r, ?)$  or missing relations  $(h, ?, t)$ , the task is to complete the query by predicting the missing element. Link prediction models are broadly categorised into two main types: rule-based models, for instance, AnyBURL [104], and embedding-based models. Embedding-based models, which represent graph elements in continuous vector spaces, are further divided into conventional approaches (e.g., TransE [20], RESCAL [117], NTN [160], ConvE [35]) and graph neural network (GNN)-based approaches (e.g., R-GCN [148], KBGAT [116]) [145, 187].

The ability of LLMs to perform knowledge-intensive tasks has been extensively studied, including KG construction and completion [198, 93, 129]. One of the foundational studies, LAMA, examined KG completion by probing language models to extract facts through cloze-style (missing word completion) prompts [132]. Follow-up work after LAMA has shown improvements in performance [140, 198]. TKGCon leverages LLMs to generate high-quality KGs from theme-specific corpora [37]. Veseli et al. [172] conducted a systematic analysis of the capabilities of language models for automated KG completion. Their findings suggest that LLMs can predict facts with high precision for certain relations in Wikidata, though this is not yet

universally applicable. Concurrently, prompt engineering techniques have emerged as a focal point in efforts to elicit knowledge from language models effectively [6], further highlighting the evolving role of LLMs in KG-related tasks.

### 2.3. Responsible AI

Since AI systems are frequently applied in high-stakes contexts such as the medical domain, finance, government, and education, trust concerns have been raised about the potential harm and risks of opaque AI systems to users and stakeholders, especially in deep learning models that are notoriously difficult to interpret [3]. Integrating generative AI into the KG lifecycle may exacerbate existing issues of transparency, explainability, and bias throughout the development and use chain [86]. This underscores the need for a thorough understanding of how to use such technology responsibly. Within the scope of this study, the principles of responsible AI most relevant to KE practitioners include transparency, explainability, fairness, and privacy.

*Transparency* refers to the degree to which information about an AI system and its outputs is accessible to individuals interacting with the system — regardless of whether they are aware of doing so. It seeks to answer the question of “what happened” in the system. Effective transparency ensures that appropriate levels of information are available at different stages of the AI lifecycle, tailored to the roles and expertise of the stakeholders engaging with the system [4].

Despite its importance, ensuring transparency in the use of generative AI presents significant challenges. Studies indicate that LLMs lack sufficient accuracy and reliability to operate without human oversight, particularly regarding the lack of provenance and propensity to hallucinate [73]. The scale, probabilistic, and inherent black-box nature of LLMs obscure their internal logic, making their behaviour unpredictable and limitations difficult for users to anticipate [61, 196]. Additionally, the opacity regarding which data some of the LLMs are trained on not only reduces explainability but means accountability is difficult to establish and increases the risk of the inclusion of personal data and even the reproduction of such data in outputs [115]. Given that the training data is obtained from large-scale web crawls [168, 169], LLMs often reproduce and amplify the unfairness and biases inherent in these sources [110, 199].

To address these challenges, several methods have been proposed that emphasise comprehensive reporting of model information, including model details, metrics, evaluation and training data, etc. Mitchell et al. proposed Model Cards, a framework for transparent model reporting that includes metadata on model details, intended use, factors, demographics, performance, evaluation data, training data, supplementary analysis and ethical considerations [109]. It can be tailored depending on the context and the needs of stakeholders, with additional sections such as interpretability methods. Similarly, Pushkarna et al. introduced Data Cards for the transparent and human-centred documentation of ML datasets [138]. A Data Card is a structured collection of summaries of dataset metadata, such as attributes, intended use

cases, provenance, along with explanations and rationales. With their high degree of interpretive flexibility, Data Cards can act as boundary objects [161], enabling stakeholders to collaborate by analysing data using shared dimensions and a common vocabulary. Specifically for NLP systems, [15] suggested Data Statements to provide context about datasets, specifically to make NLP users of the dataset aware of issues related to exclusion and biases in the data, or that might emerge with the use of the data.

Achieving transparency in ML models can also be achieved through *explainability*. eXplainable AI (XAI) encompasses multiple notions, including interpretability, responsibility, and accountability [3]. According to DARPA's XAI program<sup>3</sup>, XAI systems are designed to explain their rationale, characterize their strengths and weaknesses, and help human users to understand “how” they will behave in the future.

Several algorithms and models have been predominantly used for XAI. Bach et al. proposed layer-wise relevance propagation (LRP), which decomposes non-linear classifiers like neural networks into several layers of computation and computes relevance scores at each layer that can be used for image classification [13]. LIME, proposed by Ribeiro et al., learns an interpretable model locally around the prediction [144]. Lundberg and Lee proposed a unified framework that uses SHAP values to measure the feature importance [102]. Certain explainers target specific types of deep learning models. Ying et al. developed GNNExplainer [186], which identifies subgraphs and relevant node features through counterfactual-based importance measurement, applicable across GNN architectures and graph-based ML tasks. With the rise of attention mechanisms [171], attention values are increasingly used for explanations. However, debates persist on their adequacy, with critiques arguing that attention is not inherently explanatory [71] and that more effective alternatives exist [14].

In the context of this study, Danilevsky et al. surveyed state-of-the-art XAI models in NLP [30]. Tiddi and Schlobach's systematic literature review [167] focused on the integration of KGs into explainable machine learning, where KGs are used as domain knowledge for explanations. Zhang et al. conducted the first examination of XAI methods for KG construction [192]. Beyond the technical scope, Miller et al. explore XAI from a sociotechnical perspective, drawing insights from philosophy, cognitive science, and social psychology [106].

*Fairness* in AI is a key focus in this study, addressing concerns of equality and equity by tackling issues such as harmful bias and unfair outcomes that discriminate against certain demographic groups [4, 120]. Friedman and Nissenbaum defined bias to refer to computer systems that systematically and unfairly discriminate against certain individuals or groups of individuals in favour of others [46]. Bias exists in various forms, including gender, representativeness, and selectivity, and may be encoded into KGs through

automated processes. Given that LLMs can generate stereotypical bias, their use in automated KE tasks deserves special attention [113].

For example, bias arises in information extraction (IE), NER, and semantic role labelling (SRL). In IE, spurious correlations between entities and classes can skew results, as models often rely on biased statistical dependencies due to unbalanced data distributions [114, 53]. NER is biased in terms of name regularity and demographics, influenced by biased dictionary definitions or neural network-based extraction methods that exclude protective variables [53, 194, 108, 126]. SRL, which analyses predicate-argument structures for input sentences, also exhibits demographic biases, such as gender and race bias [197], and commonsense biases, failing to account for basic commonsense knowledge [89]. However, mitigating bias is challenging, as humans are naturally prone to various biases [77, 86]. Technical approaches to addressing bias include causal inference, which modifies target entities [189, 175, 176], and counterfactual IE (CFIE) [114], which incorporates structural information to capture non-local interactions critical for IE tasks [195, 74].

Beyond the above mentioned principles, Responsible AI encompasses several other key themes [105]. *Safety* in AI focuses on preventing harm, which includes various themes such as robustness against adversarial attacks, malicious use, reliability and reproducibility, etc [76]. AI safety can describe harms in multiple ways, such as the existential safety risk of the fast development of AI without regulation [25], and the safety risks of current AI systems, which may undermine trust in their deployment [155]. This study specifically focuses on the latter. *Privacy* ensures data protection throughout the lifecycle of AI systems, both in terms of securing personal information and safeguarding individuals from unwanted exposure [40].

### 3. Methodology

During August 2023 we organised a four-day research hackathon for knowledge engineers and AI researchers to investigate KE with prompt engineering.<sup>4</sup> The hackathon format offered a tool to explore emerging practices as KE processes and approaches are disrupted by generative AI. The project topics (see Table 1) were defined in a community process with 30 experts [58]. The hackathon hosted 39 participants from 15 different institutes and various environments from academia and industry, with 31 PhD students, 2 postdocs, 1 lecturers, 2 professors, and 3 industry members. The hackathon attendees were selected based on their background and experience with KE. They were from European research labs that have a strong profile in either publishing in KE venues, or in offering well-used KE industry products. During the hackathon, participants are divided into 7 groups of 5 to 7 members, based on their prior experience in KE. Each group explored one of the topics listed in Table 1. While participants were provided with high-level project

<sup>3</sup><https://www.darpa.mil/program/explainable-artificial-intelligence>

<sup>4</sup><https://king-s-knowledge-graph-lab.github.io/knowledge-prompting-hackathon/>

**Table 1**

Overview of hackathon topics with associated tools and methodologies. Detailed descriptions and references for the tools and methodologies listed here can be found in Appendix C and Table 7.

Hackathon Topics	KE task	KE Tools & Methodologies	GenAI Tools & Methodologies
Determine if generative AI can extract knowledge structures, including inference rules, to go with facts for KG construction	Use prompting to fill in triples, for example Coldplay - BandHasMember - ?	Triples investigation	ChatGPT, Few-shot prompt
Create a framework providing tools for collaborative human-AI ontology engineering and requirements elicitation	If there is a user story, use LLMs to create competency questions and then build an ontology	Competency Questions, eXtreme Design methodology	ChatGPT, Multiple prompting techniques
Determine how KE tasks can be supported by generative AI	Use ChatGPT to re-create the wine ontology	NeOn methodology, Competency Questions, HerMiT Reasoner, OOPS!	PaLM, Llama, ChatGPT, Few-shot prompt
Determine if generative AI perform reasoning tasks completely in natural language	Develop an agent (i.e. LLM) to collaboratively create a comprehensive knowledge graph with a domain expert. This involves crafting both the graph's structure and content interactively, with the agent taking the lead	Triples investigation	ChatGPT, Multiple prompting techniques
Determine if generative AI can perform ontology alignment (i.e. identify and match entities between ontologies)	Given a request to an LLM, explore whether two items in Wikidata should be aligned	OAEI	ChatGPT, Zero-shot prompt
Determine if generative AI can be used towards the construction of multimodal KGs	Given a painting, explore whether LLMs can fill in a KG with information, such as painting information, painter's image, painter's information, art historian	Investigate triples	mPLUG-Owl, InstructBLIP, Text only prompt (no text + image)
Determine if we can perform ontology refinement (i.e. techniques for KG completion and correction) using generative AI	Use OntoClean within an iterative process to improve an ontology and to better represent knowledge	OntoClean	ChatGPT, Llama, Claud, Few-shot prompt

**Table 2**

Number of documents and descriptions for each category of documents in the dataset.

Selected Method	Collected Documents
Ethnographic study	2 sets of observation notes from two of the authors
Hackathon reports	7 reports from the Hackathon groups
Hackathon presentations	7 presentations and slides from the Hackathon groups
Interviews	14 semi-structured interviews with Hackathon participants

descriptions, they had the flexibility to select the tools, methodologies, domains, and LLM techniques they wanted to focus on.

### 3.1. Data Collection and Preparation

Data was collected using three approaches to ensure comprehensive coverage. A detailed list of datasets is provided in Table 2.

*1. Ethnographic observation* Two researchers performed an ethnographic observation study, using the ‘observer as

participant’ technique for ethnographic research [94]. The observer as participant role moves the observer into the study environment and closer to the activity of interest. The participants are aware of the observer, but the observer is not engaging with the participants and does not ask or interact with anyone. The researchers observed the various groups and kept detailed notes of the 7 groups’ interactions and decisions.

*2. Documents collected from the hackathon* Research output in the form of reports and documentation from the

**Table 3**

Themes identified through the inductive analysis.

Themes	Description
Background experience	The participants' experience prior to the hackathon
Evaluation	Evaluation issues and suggestions for generative AI outputs
Skills	The skills and qualities participants have or need to interact with generative AI
Challenges	Challenges that participants face while using generative AI for KE tasks
Bias	Participants opinions related to AI safety testing
Generative AI interaction opinions	Participants opinion about the use of generative AI for KE tasks

hackathon was collected. Hackathon groups produced reports regarding their experience answering specific questions and slide presentations with their strategies and results. This resulted in 7 presentations and reports.

These two data collection activities were conducted contemporaneously with the Hackathon.

**3. Semi-structured interviews** After the hackathon, we contacted 14 participants who had volunteered to participate in follow-up interviews to talk about their opinions and experiences interacting with generative AI for KE tasks. Interviewees had backgrounds in KG construction, KG explainability, ontologies, and machine learning, as well as biology and bioinformatics, data management, FAIR data [38], and digital humanities. We conducted these interviews virtually between 17 August 2023 and 6 September 2023 using Microsoft Teams.<sup>5</sup> The interview guide can be found in Appendix A.

We used Otter.ai<sup>6</sup> to transcribe the interview recordings. The transcribed data sources were then uploaded into NVIVO<sup>7</sup>.

### 3.2. Data Analysis

To address our research questions, we conducted thematic analysis across all the data. We first used an inductive method to search for key themes from the interview topic guide (see Appendix A). Table 3 shows the main themes identified. We then used a deductive method. We read the data corpus for grounded themes that emerged in the data. The final codebook is shown in Appendix B.

This analysis was performed by two researchers, who first read the interview topic guide and together identified the key themes. Then they read the data corpus for the ground themes. Following this, the two researchers, advised by another senior researcher, discussed and confirmed the codebook presented in Appendix B.

### 3.3. Ethics

The ethnographic study was approved by Authors' institution's Ethical Advisory Committee via the Full Application Form. Informed electronic (using Microsoft Forms<sup>8</sup>

and storing in CSV) consent was given by the participants. Thirty out of the 39 hackathon participants consented to the ethnographic study. Participants who did not consent were not observed. No personal information was collected in the analysis. The documentary data were collected as an output of the hackathon.

The interview study was approved by Authors' institution's Ethical Advisory Committee via the Minimal Risk Procedure. Informed verbal (audio recordings and chat transcripts) consent was given by the participants. No personal information was used in the analysis.

## 4. Results

### 4.1. RQ1: What are the main challenges experienced by knowledge engineers when using generative AI for KE tasks?

Testing the use of generative AI for KE tasks was challenging for all participants. Our observations during the hackathon revealed that the main concerns for participants were identifying appropriate datasets to use, prompting LLMs efficiently, and evaluating the LLM outputs. We asked interviewees which of these they considered important.

**Dataset.** Interviewees (1), (14), (12) emphasised the importance of datasets as a starting point, and groups used prompting as a means to create a dataset (Observer 1). A dataset was used by groups, for example, to design ontologies based on a domain, to extend existing KGs, and to test reasoning tasks and ontology alignment techniques. Without a dataset, they felt they could not continue with prompting and evaluation (14). "There are many codes online...to interact with language models that wasn't any problem at all. And the evaluation given that you know, you possess a ground truth,...was not too difficult. So, the most challenging case work for us was finding a dataset" (8). Finding a dataset was challenging for specific group tasks mainly because of the time pressure (14) and fast planning (1) needed to find and set domain knowledge for the task. On the other hand, an interviewee noted that they felt LLMs would be supportive for establishing datasets. "The dataset is not what concerns me because if anything, I'm pretty impressed with what LLMs potentially could do in terms of extracting, like I was saying, concepts, relationships, maybe even constraints or parameters or universal restrictions or existential restrictions, from unstructured data" (13).

<sup>5</sup><https://www.microsoft.com/en-gb/microsoft-teams/group-chat-software>

<sup>6</sup><https://otter.ai/>

<sup>7</sup><https://help-nv.qsrinternational.com/20/win/Content/about-nvivo/about-nvivo.htm>

<sup>8</sup><https://forms.office.com/Pages/DesignPageV2.aspx>



**Prompting.** An interviewee believed prompting is challenging if knowledge engineers' lack NLP training and experience, "I think prompting templates are going to be, generally more difficult for Semantic Web people, because we are not necessarily natural language processing people" (1). This was also emphasised in the skills that a knowledge engineer needs to work efficiently with LLMs. Furthermore, many participants noted the time-consuming phase of iteratively testing prompts to receive desirable outcomes (3), (Observer 1). Some suggested the usefulness of possible "templates for prompting" [6, 193] (Observer 1), (Observer 2), (3), (13). In addition, "once we settled on a specific prompt, we said okay, it seems like it works good enough, the most tricky part of it was to make the language model consistent in its responses" (3). This was particularly problematic when the process was automated with thousands of prompting iterations. Furthermore, syntactical errors also break the scripts and prevent automation (3). In contrast, an interviewee pointed out that prompting engineering is "less scientific" and models "with billions of parameters are not controllable,..., resulting in an output that you cannot reproduce one hour later, two days later. So, this might be a problem" (9).

**Evaluation.** Participants believed that evaluation is challenging mostly due to the lack of benchmarks for specific KE activities, which depends on manual intervention, "so there are papers doing auto prompting...I'm a little biased towards evaluation simply because in my tasks evaluation pretty much boils down to a manual evaluation, there's no other way to do it...it's not that we have solved the dataset problem, but we can synthesise some datasets, or we can you know, create them artificially somehow. Evaluation still remains a challenge because it has to be done manually" (5). In the case of using a specific controlled dataset for KG construction, the accuracy and coverage of the KG can be automatically evaluated (2), (10). However, in tasks such as designing competency questions or an ontology, there is not always a gold standard to compare with, and manual work from an expert is needed to evaluate the output. Nevertheless, exploring the use of holistic LLM benchmarks such as HELM [93] was suggested (5). Recent progress in developing benchmarks for LLMs in KE reflects the community's strong emphasis on establishing standardised methods for their evaluation [131, 75].

#### 4.2. RQ2. How do knowledge engineers evaluate generative AI output for their practices?

We further asked interviewees whether current evaluation techniques for KE tasks could be used to evaluate generative AI outputs. Most respondents believed that evaluation mainly depends on the task, but highlighted that current evaluation metrics are not sufficient for many KE activities. Following this, we asked them what a new benchmark or metric would look like.

**Current evaluation techniques.** Some interviewees consider that for specific tasks, like using a controlled dataset (13),(14) and designing ontologies with simple taxonomies (12), evaluation metrics like F1 score, precision and recall

(8), and comparison with gold standards are sufficient metrics for evaluation. However, others felt that regarding KE task evaluation, "it's not [only] a problem that concerns only interaction with LLMs, it's a broader problem that needs to be solved" (9). Many KE tasks, like ontology semantic and reasoning errors, require human evaluation (1), (5) because "there is no automatic way [to assess]...how well the ontology represents the domain knowledge" (11), and this is important because "if it's mistaking...in relations which creates the hierarchy...[that] is much more important than making a mistake on some data." (2)

In addition, others highlighted the need to evaluate the additional knowledge, specifically the extra knowledge produced by LLMs (10) and the knowledge is missing from KGs (4),(6),(14). An interviewee (10) highlighted the extra knowledge produced by LLMs based on their training, for example, when we asked the LLMs to create triples from an expression like "X was born in Y", we expected "X-is a-person", "Y-is a-country", and "X-birthplace-Y, but we also may get "Y-part of-European Union". These outputs include knowledge that may be needed to complete a KG, but it requires evaluation. In addition, a way to evaluate the knowledge missed in the LLMs' outputs is needed. An interviewee pointed out that, "by using these classical evaluation metrics...[we check] what is inside the [input] text and [we miss] what is not in the text" (4).

Finally, interviewees emphasised the difficulty in processing LLMs output mainly because LLMs return "strings" (10) of text and they lacked a way to evaluate the "generative text" (7). It is common in KE tasks to use a specific format like JSON or the expression of an ontology. Even in the cases where participants asked the LLM to give output in these particular formats, they faced difficulties because, "the generated string contains inconsistencies, you really need to parse it and get what you want from the string itself, and then you need to convert it into a list, whatever you need" (10). Recent advances in LLMs' structured generation capabilities have sought to address this challenge, as such requirements have emerged across domains. Approaches such as token-level constrained decoding [83, 178, 39] and prompt-level output structuring [98] can enable LLMs to produce outputs in specific, predefined formats. Complicating this is the fact that LLM outputs are not all the same type (4), for example, a KG includes dates, values, images, URLs etc. It is not possible to have a united way to evaluate the different types. Furthermore, another factor in output evaluation is the similarity of the results (4). Current evaluation techniques primarily rely on exact matching rather than similarity, making them less effective at handling variations and leaving LLMs prone to inconsistency.

**New evaluation techniques.** Most of the interviewees noted that current evaluation techniques are not sufficient and suggested possible alternatives. Some focused on the ontology design (2), (5), suggesting the development of a set of ontologies to be used as gold standards (observer 1). Others suggested the creation of toolkits (1) for LLMs to

review ontology errors similar to existing ontology toolkits, like OntoClean [60] and Oops [136]. However, they acknowledged that this cannot be applied to large-scale KG like Wikidata, and manual work will still be required. Humans could be more efficient in evaluation if they have the information which prompt is related to the results (Observer 1). Recent advances in KG evaluation demonstrate that integrating LLMs with human-in-the-loop approaches can be effective [170].

Another suggestion was to test the consistency of data and semantics in a real scenario, “for a company using this data and querying this data, and the simplicity of querying this data, and also the performances of these KGs in queries per second. And also the errors and the consistency of the data that is created,...making sure that the semantics of the data stored in the KG is consistent and reliable” (9).

In addition, some interviewees were more specific, suggesting the use of techniques used in other fields, like fact-checking (10), adversarial algorithms (13), and self-play from reinforcement learning (12). An interviewee suggested using fact-checking techniques to establish that the extra knowledge received is correct and can be used for the KG construction (10). Another interviewee was fascinated by the concept of adversarial algorithms, a technique used to test machine learning by misguiding a model with malicious input. They suggested using this technique to test LLM results (13). In addition, another suggested the use of self-play technique, a reinforcement learning technique for improving agents’ performance by playing “against themselves”. They suggested to “let the model come up with, a whole host of answers, then compare them between each other [and]...get a more refined answer”, but this will require “more time and computing” (12).

Two interviewees (3), (6) commented that we should see generative AI as assistants making the processes in KE smoother and faster because, “we know that as humans, we are lazy to write down some stuff like the competency question, so, if [you have a suggestion of] competency question you decide [faster], actually, which one you want to use” (6) and the same can apply every step of the way. They suggested that evaluation should focus on human satisfaction, “so the language model should probably optimise not really the performances, but the satisfaction of the user at least dealing with that” (3).

One interviewee suggested that a more “open-minded” approach is needed in order to advance the field of evaluation. If metrics are, “novel, it takes longer to review. If it’s a straightforward paper with a slight modification in the method and a new row, it’s easier to evaluate the merit of the work. That’s not necessarily the way to go. So there’s a lot of work to be done in evaluation” (5).

### 4.3. RQ3. What skills does a knowledge engineer need to incorporate generative AI into their practice?

**Skills possessed and of use.** As it is all collaborative practices, communication within teams was pointed out as

an important skill in multi-disciplinary groups such as those at the hackathon (4), (5). To keep the project moving, the ability to explain and listen to ideas for an efficient workflow are required (4). One interviewee talked about the risk of getting lost in the implementation details, stating, “I kind of knew how to steer people in a direction that would eventually be useful for the final goals” (5).

Having skill in building ontologies was found to be useful in adding context to the prompts (10), (6). For example, using the knowledge structure of a dataset, such as Wikidata (6), and information on how to form triples (10): “adding...context...as part of the prompt actually helped with the result” (6). In contrast, one interviewee found that the added context did not make a difference so despite having this skill, they did not consider it important (3).

Building ontologies was also found to be useful in defining the overall goals for the task (5), (9): “I could contextualise the tasks defined by the hackathon. And also, from a more practical point of view, I knew before starting coding, how the final output would have looked like” (9).

Having prompting skills and understanding how prompts can be composed and revised to achieve a desired outcome was found to be important (10), (12), (8): “I think I have a good understanding of the prompt components. So, in a prompt, first you should assign a role, a persona to the language models, and then short task description is really useful. And these models shine when you give a few examples in context learning” (10). In turn, these participants were able to streamline their approach and prompt more efficiently due to their understanding of the different components of prompts. Even just limited experience was found to be useful for forming prompts with one interviewee describing their skill as “not scientific” but found this helped in “understanding how these models respond to small changes” (12). Interviewees acknowledged that rules are needed to prompt efficiently. There is a sweet spot in how much information is given to LLMs in order to get a desirable output, with too much information being as problematic as too little.

Coding was the main technical skill that many interviewees mentioned as being important (2), (3), (8), (9). It was found to be useful in interacting with the LLM through the respective API (Application Programming Interface) and not the interface, which was important in querying large batches of prompts. Along with this, knowledge in using repositories of models like Hugging Face<sup>9</sup> and PyTorch<sup>10</sup> was also noted as useful by an interviewee (2). A useful skill was in version control systems like Git<sup>11</sup> to be able to clone, test and debug projects efficiently (2). Some interviewees also mentioned how their coding skill helped them with datasets, such as, scraping websites for datasets (3) and “[building] a dataset” (4).

Developing a scientific framework was found from one interviewee to help define the task’s workflow, describing

<sup>9</sup><https://huggingface.co/>

<sup>10</sup><https://pytorch.org/>

<sup>11</sup><https://git-scm.com/>

it as “defining objective, defining hypotheses, defining an evaluation strategy, designing so an experiment” (7).

**Skills gap.** Many of the interviewees who lacked the skill in ontology design and engineering stated its importance (2), (12), (13) specifically for the evaluation of the outputs of the LLM (1), (7). They felt that without this, it is difficult to know what a desired output would look like.

For those who did not have prior prompting experience, there were two main distinctions on how they viewed this skill. The first way was that some interviewees did not see prompting as a skill and felt that it was too simple of a task with multiple participants, restricting it to a trial-and-error exercise, not acknowledging the possibility to go beyond this. They deemed prompting as requiring no expertise where KE and machine learning were categorised as more skillful (7), (13). One interviewee who described themselves as being familiar with LLMs did not deem prompting as a specialist skill in the KE process as they believed any developer could do it (11).

On the other hand some interviewees were interested in gaining more experience to go beyond trial-and-error (4), (3), (1), (14). One interviewee acknowledged that prompting can be an intricate process. Some interviewees felt that time for experimentation with generative AI is important to build an understanding of prompting and this familiarity with LLMs can improve performance (5). This can also allow for quicker execution of prompting to streamline the workflow, one interviewee noting, “I’m pretty sure that a lot of tricks or tips that I don’t know about” (6).

Various technical skills were mentioned by several interviewees all of which were specific to running LLMs more efficiently. For example, one interviewee mentioned learning to use computer clusters to fine-tune LLMs (10), which needs experience and very high computational requirements. Another interviewee mentioned using LLaMA and Hugging Face for the use of APIs with different models as an important skill (12). There was also an interviewee who felt they needed an enhancement of their coding skills and hardware skills, such as using GPU (9).

#### 4.4. RQ3. How aware are knowledge engineers of using generative AI responsibly?

Our question on responsible AI was initially framed in terms of “safety”, offering the respondents the opportunity to engage with any aspect of AI safety they felt relevant. However, although two respondents raised data security (1) and insufficient data (13) as safety concerns, the majority were not clear on what concerns they might be considering. In these cases, we suggested they focus on bias as most people are familiar with this concept. However, this familiarity may also breed contempt. The authors of this paper observed, “when asked about harms, they said it wasn’t discussed – assumes humans are biased anyway so there will be bias anyway, to me sounds like implying that there is no need to consider bias” (Observer 2). One respondent noted that they felt it would be “interesting and helpful” to apply safety testing to LLMs and so far had read papers on it (9).

When asked to consider bias, interviewees made comments such as that they had only fleetingly interacted with the issues. “I’m familiar with the terms, but I haven’t worked on it directly” (10). They described themselves as have experienced limited exposure (2), possessing a lack of knowledge (12) and interpreting safety as a risk to researchers rather than other publics (14). One respondent had previously worked as an engineer on a bias detection algorithm but felt this was purely a “theoretical thing”.

One reason for this lack of engagement appeared to be the perception of safety was seen as a siloed activity rather than integrated with engineering: “that’s not my field” (4). This was not necessarily due to a lack of interest, but perhaps also to a lack of opportunity: that participants had been interested in safety work but “haven’t been able to get to it” (5).

The many variants of bias mean that this is a number of problems not one (6). There was some awareness that understanding the type of bias is key to addressing it (4). In terms of specific types of bias, a few interviewees had awareness of what these might be, such as provenance (13), selectivity (5), and gender (6). These had also been brought up with one group in the hackathon, with the observer noting, “organiser [told the group] – make sure you don’t bias against gender, race, geographical concepts – if competency questions are biased then ontology is biased” (Observer 2). It is difficult to tell, therefore, whether this knowledge was something the interviewees had applied to previous work, or gained in the relatively recent past of the hackathon.

Correspondingly, there was variation in awareness of bias mitigation: “[I have] no idea how many safety checking solutions have already been provided” (1); “I have some doubts on how these testings are carried out” (9). One interviewee felt that bias mitigation was more important in commercial products than in the lab (12). There was awareness that addressing bias may change relationships in a KG (4).

Particular issues related to LLMs regarding provenance and source were noted, for instance, “there are so many confounding issues in the training of an LLM that even if you address bias in a prompt you can’t see the resources it is using to make a decision” (6). This respondent suggested that there could be ways to mitigate this by mapping the results of repeated prompting. One noted that harm can arise from use, even in a well-tested model. “So even though [LLMs] pass [bias] tests, I think there’s still risk if you prompt them in a hateful or discriminative way” (10).

One interviewee suggested that, given the difficulties of mitigating bias, we should at a minimum “put an array [of] practices in place to monitor what’s coming out [and] how unbiased, how ethical, it is” (5). The idea of ‘co-creating’ bias mitigation with users was suggested, with one interviewee suggesting not only reviewing the model for bias before use but also asking users for feedback on perceived bias as they used it (7) and another suggesting that, “either we fix the data, or maybe we fix the tools somehow that [users] can detect this bias and root it out during their working” (5).

Respondents concerns about difficulties with addressing bias included a lack of knowledge of LLM training data and processes (13), a fear that tests for bias would be ineffective (9), and concern that tests alone are insufficient if appropriate action is not subsequently taken (10). There was also concern about a lack of consensus around safety issues (13) and a feeling that other problems in the model (e.g. accuracy) may be more pressing (2).

Some interviewees felt there might be a tension between attempting to reduce bias and model efficacy (2), (12) "There are evidences that if you debias a language model, it usually downgrades in performance" (8). Whereas most respondents felt bias was found in data, a few perceived bias might also lie in both the LLM (7) and KGs, specifically in the semantics (9). One respondent felt that safety was not a concern in KGs (14). Ever since we conducted the interviews, work has surveyed biases and fairness in LLMs [48], specifically building taxonomies of papers that propose metrics, datasets, and methods for LLM-based bias evaluation; but without considering LLMs at the task of KG construction. Since 2024, the Semantic Web community [130] has studied LLMs for KG construction tasks, e.g. ontology alignment [9, 54], class membership [7], and conversational ontology engineering [191]; although these works recognise the importance of bias and fairness in using LLMs for KG, none of them address them directly.

#### 4.5. RQ5. What factors may affect knowledge engineers' trust and uptake of the generative AI technology?

Interviewees have varied opinions on the use of generative AI in their practices. Some find their use promising, and others are more sceptical.

**Promising.** Interviewees believe that generative AI cannot replace humans (3), but they can support KE activities. Using LLMs we can reduce time spent on tasks (11), (hackathon report), improve communication between experts in different disciplines (hackathon report), perform tasks even without extended experience (12), for example, in ontology design, and incorporate NLP pipelines in our practices without the need to develop them from the beginning (7) [64, 26]. However, "we need to use them carefully by providing human oversight, [or] human in the loop, or our pipelines need to add options [where] human can intervene... to avoid passing those errors produced by LLMs to the downstream tasks" (1).

In this vein, interviewees pointed out that LLMs could improve KE practices (9),(7), (13). However, their use is task-specific and interviewees raised concerns that effective prompting requires experience and there are reproducibility concerns (hackathon report). As well as assisting with processes, there was a suggestion for the incorporation of LLMs to improve existing KE tools (9) or to create new and advanced tools to support KE (13).

**Sceptic.** Some interviewees were found to be more sceptical about the use of generative AI for KE tasks. They felt that LLMs without using any search or retrieval-augmented

generation are not up-to-date with the latest knowledge and that hallucination is very common (1). Moreover, LLMs respond with a prediction of what they think is the answer without giving information about how "confident" they are in the answer (13). Recent work has sought to address this through approaches such as confidence elicitation [180, 52], self-assessment [142], chain-of-verification [36], and repeated sampling [23]. In reverse, trusted KGs already support AI with trusted knowledge and can also support LLMs (1).

Other interviewees highlighted that LLMs cannot perform all KE activities (11), and the lack of evaluation techniques (6) lead some of the participants to find no value in their use. An interviewee expressed, "Yes, it can help you to quickly get some data just for testing purposes, for example. But for production, I don't think that I will use it...since I never know which part of it is actually correct" (6). During the hackathon, participants felt that LLMs can perform some tasks or simple examples, but in complex scenarios, the output cannot be controlled (i.e. evaluation). An interviewee expressed that integrating KGs and LLMs was not required as LLMs would entirely replace KGs.

## 5. Discussion

In this section, we consider some of the key issues that arose from our results. We examine in more depth issues related to skills, in particular prompting and bias detection and mitigation, and conclude by emphasising the importance of transparency and explainability in KGs. Finally, we introduce the concept of "KG Cards" and we suggest the use of the existing Model Cards for the KG embedding models.

### 5.1. Issues in Using Generative AI for Knowledge Engineering

A fundamental challenge in using generative AI for KE stems from the ambiguity and variability of its outputs. This is most evident in the inherent irreproducibility of results from single, directly prompted LLMs, where generated relations or term similarities are based on statistical probability rather than consistent fact. Such inconsistency undermines scientific rigour and complicates error tracing, making control and refinement difficult. (Since we conducted the interviews, this has been at least partly mitigated with retrieval augmented generation (RAG) [90], which enables LLMs to retrieve relevant documents and integrate their information to user prompts before providing query responses; however, issues of bias and fairness around the integration of specific external knowledge sources and the models themselves persist [80, 179].) These concerns align with prior findings that knowledge extraction remains challenging despite advances in NLP [8, 119]. Compared to pre-LLM human-in-the-loop KE tools [49], LLM-assisted approaches offer more end-to-end workflows and streamlined human-machine interaction, but also obscure the origin of errors. However, recent advances—such as tool calling [141], retrieval-augmented generation [90, 50], confidence elicitation [180, 52], and



self-assessment [142]—have significantly mitigated these issues. These developments could lead to more reliable, and reproducible LLM-assisted KE workflows, thus change the attitude of knowledge engineers towards LLM-assisted systems.

Another notable application of LLMs in knowledge engineering is their potential to generate structured datasets from various unstructured data sources, particularly for data integration purposes. Research has shown that, when given a prompt containing the term “dataset”, in the absence of a dataset, ChatGPT may attempt to “curate” a dataset like an approximation from unstructured data [173]. Therefore, this application seems highly plausible for KE and, particularly, knowledge extraction in a structured format which is familiar to existing practices. Following the usual actions of cleaning and transforming, the dataset can be used for KG construction. Recent advances in the structured generation capabilities of LLMs have further strengthened their potential for such applications [83, 178, 39, 98]. Yet, the accuracy of LLM results remains an issue.

Following the knowledge extraction and integration, another particular area of concern was evaluation. Attempting to evaluate LLM support for KE tasks would simply exacerbate existing evaluation challenges, which were many. These include evaluating *accuracy* (i.e., syntactic and semantic) and *coherency* (i.e. consistency and validity), emphasising that we may assess accuracy automatically, but most of the time, coherency requires manual interventions. Furthermore, concerns about *coverage* (i.e. completeness and representativeness) and *succinctness* (i.e. conciseness and understandability) are also relevant to input as well as output, and should be included as a consideration. This is a similar problem to that experienced in crowd-sourced KGs such as Wikidata [1], where variations in the input (such as background and skill of contributors) also demands evaluation as much as the output. Recently, the community has devoted increased attention to evaluation, emphasising both the development of robust evaluation benchmarks and the adoption of human-centered principles, strategies, and paradigms [190, 101, 97, 65, 63]. Recent research has explored the use of LLMs as a judge to evaluate the results of competency questions, but suggests these are still at the stage of being useful as supporting human inputs, rather than as trustworthy standalone tools [82]. Regarding the evaluation of generative AI outputs, one interesting suggestion was made related to using adversarial algorithms. “Would something or a product like that stand the scrutiny of an adversarial network, who is trying to maybe test or break the new information derived by the existing system. So can there be a counter point counter punch that tries to break it, and if it cannot be broken or falsified?” An adversarial attack on BARD [57] asking to repeat one iteratively one word eventually got the model to reveal gigabits of its training data, demonstrating the usefulness of such processes in testing systems [115].

Another proposed aspect of evaluation was human satisfaction. Other research has shown that the natural language

element of LLMs, often read by humans as possessed of emotion, can make people feel that the process of using the model was successful even if the end result was not [173]. This could be potentially problematic for KE tasks, which requires logic over feelings. A comparable, well-researched situation is that of emotion in financial markets, which can move prices up or down based on public mood rather than underlying economic fundamentals.

## 5.2. Key Skills for Effective Prompting

Prompting was a fascinating area of discussion. Although many interviewees spoke of how they found previous skills useful in developing effective prompts (and equally, that it was very possible to use ineffective prompts by accident), some interviewees also felt that prompting skills were not important. This conflict in opinions may be explained by the different tasks and the experience levels of interviewees in KE and in research.

However, ultimately, prompting was utilised for all the different KE tasks. This flags a new era and requirement for a new skill set in the field of KE. This suggests that the creation of templates (e.g., [6, 99, 193]) for the KE tasks could benefit the field towards KG construction and fill in gaps in prompting skill level. This has parallels in the introduction of crowd-sourcing to KG construction around a decade ago. Using crowdsourcing means for large-scale KG construction was a successful innovation supporting many intelligent applications today [2, 143, 84]. Furthermore, leveraging (semi-)automatic prompting methods [183] and incorporating domain-specific prompting strategies [67] can help address skill gaps and improve performance. Very recent research suggests that although reasoning has received a lot of attention as a prompting strategy, for some major LLMs this does not provide substantial improvement over simple instruction and task demonstration for knowledge extraction from text [135]. This implies that in this area at least, proficiency in prompting may easily be acquired.

## 5.3. Key Skills to Mitigate Bias

Knowledge engineers and domain experts involved in generative AI-supported KE projects are key stakeholders in promoting responsible AI practices. Ensuring that AI systems are developed, assessed, and deployed in a safe, trustworthy, and ethical manner is critical for fostering appropriate trust among these professionals. There was a distinct lack of engagement with the processes of bias mitigation. This tended to arise from limited exposure to techniques or theories for identifying bias in the model. While this was sometimes due to a lack of opportunity for knowledge engineers to acquire capabilities, in other cases it was due to a lack of a sense of responsibility for bias mitigation. It is difficult to argue that the identification and eradication of bias should be left to a specific set of experts, rather than integrating core safety facets into the training of those responsible for building AI. Given the variety of safety and ethical aspects and potential harms in AI, it may be that space for “ethical deliberation” in the KE process is more

valuable than simply providing, for instance, a list of biases for engineers to look out for [56].

While capacity building (training) and bringing awareness to the importance of taking responsibility for ownership of AI safety issues appear to be two separate issues, they are contingent upon each other and should be treated as two sides of the same coin. Proper exposure to bias awareness and mitigation techniques should include training on when and where to review aspects of safety – thus integrating the idea of responsibility for safety with knowledge of the actions to take.

a

Our results illustrated a data-centric approach to bias in KE, where our participants felt that the source of bias was more likely to be the data than the graph model itself, and relatively few recognised that the model may be a source of bias. Including generative AI in the construction of KGs will introduce a further source of bias [69]. The integration of generative AI into the KE processes may additionally create further potential for harm, such as when, intentionally or unintentionally, the model can be manipulated to overcome privacy (as shown in the adversarial attack example above [115]). This suggests an integrated review of the possible sources of harms introduced to KE when generative AI models are utilised would be valuable.

A number of the issues raised by participants in the hackathon - particularly those around evaluation and ethics – speak to deeper, more sustained challenges within the discipline. Our findings show that many knowledge engineers feel ethical issues are the responsibility of 'someone else', and that bias mainly resides in data and not in any part of model construction. This finding is similar to other studies with AI engineers or practitioners, who feel either that the responsibility for mitigating harm does not lie with them, or that as individuals they cannot do much about it [87]. A great deal of recent work on LLMs for KE emphasises that research into bias in this area is of great importance [48, 130, 9, 54, 191]. However, no study has directly addressed bias and fairness in LLM-based KG construction and there are no concrete proposals on how to address it. This suggests that work needs to be done to embed a greater understanding of where bias lies, and how KE can address it. LLMs are not neutral and have bias up and down the chain that must be understood and appraised [80, 179]. Benchmarks for evaluating LLM-generated ontologies have been proposed [134], but are limited to assessing the presence/absence of classes and properties and their level of abstraction, not their possible issues around bias and fairness. A starting point would be measuring bias and fairness in LLM-based KG outputs, for which inspiration could be taken from concepts such as Vendi (diversity evaluation) scores and the multi-metric HELM approach [47, 93].

#### 5.4. KG Cards to Enhance Transparency

As described in Section 2.3, prior research aimed at improving transparency and explainability has introduced frameworks such as Model Cards [109], Data Cards [138],

and Data Statements [15]. Implementing our review of harms we point out the need for similar practices for KGs. The aforementioned Cards can not fully serve the KG's needs because they do not consider the unique design aspects of KGs. These Cards fail to describe the ontology and the KG data, such as relations and schemas. Building on these studies, we propose the concept of Knowledge Graph Cards (**KG Cards**) as a documentation framework to promote transparency and explainability, ensuring the creation and use of trustworthy KGs. Drawing inspiration from the flexibility and extensibility of Data Cards, and considering the foundational representation format of most KGs, KG Cards are designed to enhance transparency and trustworthiness by providing comprehensive and structured documentation.

As shown in Figure 2. It contains 6 main themes:

- **Knowledge.** The Knowledge section aims to provide an overview of the domain, the languages represented, and the range of knowledge included. It may also document ontology development artifacts, such as competency questions, which help define the knowledge graph's (KG) scope and support ontology testing.
- **Provenance.** The Provenance section is intended to capture detailed provenance information about the KG. Various platforms and frameworks, such as the W3C Provenance Ontology (PROV-O) [88], have been proposed for managing provenance in KGs [81]. Additional provenance details can be documented here, including information about crowdsourcing (e.g., platform used, demographics of crowdworkers, and domain experts), references to datasets (and Data Cards, where applicable) used in KG creation, and other contextual information.
- **Graph Construction.** The Graph Construction section outlines the details of the data preprocessing methods and the approaches used during the construction and curation process. This includes techniques for knowledge extraction, integration, and the use of KG embedding models.
- **Quality.** A crucial aspect of KG transparency and explainability comprises the quality assessment. Considering the KG quality assessment dimensions described in Section 2.1 and quality dimensions for linked data [188], we introduce the Quality section. This includes descriptions about accuracy, coverage, coherency, and succinctness. Many studies have focused on completeness [66, 127, 151, 70], which is also a vital dimension for explainability. Syntactic and semantic accuracy describes how errors are handled, while timeliness states the pace of updates. Coverage looks at the network details, stating statistics which can offer an understanding of the KG coverage and potentially can support the next card about safety tests. Moreover, coherency expresses contradictions in the data and succinctness describes the human-readable

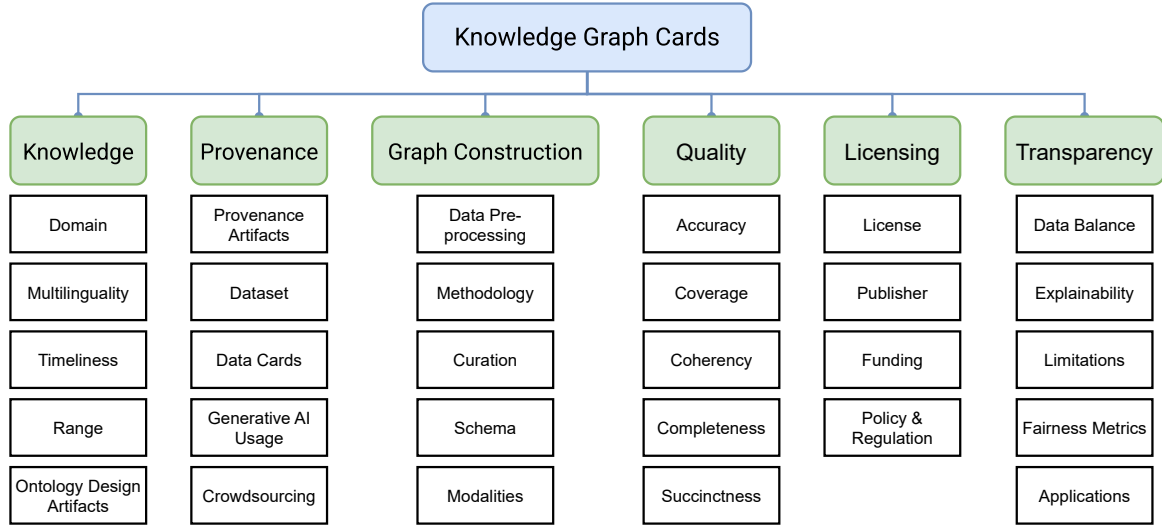


Figure 2: Six suggested sections of KG Cards with related details.

aspects, the URIs details, SPARQL query services, and community contact options.

- **Licensing.** The Licensing section provides details about the licenses, publishers, funding sources, and associated policies.
- **Transparency.** While previous documentation methods provide detailed transparency principles, they often lack emphasis on fairness aspects. “Fairness” as a concept is vague, culturally dependent and difficult to operationalise [111, 38]. For KG Cards, we suggest the Transparency Card with a set of transparency and fairness checks related to data, the intended and unintended application of the KG, and explainability. With this, we aim to provide explanations and support for bias and ethics.

An example of the use of KG cards is presented in Appendix D, Table 8. The presented KG Cards represent a preliminary framework for enhancing transparency in KGs. Like previous model and dataset documentation methods, the proposed KG Cards are adaptable, allowing more sections and properties to be included based on the specific needs of the stakeholders in various contexts. Further research could offer valuable insights through extended analysis, helping to validate usability and identify potential modifications or enhancements to the structure.

The use of well-structured documentation for transparency and explainability has inspired initiatives like the Hugging Face platform’s adoption of Dataset Cards and the Croissant metadata format for its shared datasets [5].<sup>12</sup> Table 4 shows the alignment between Hugging Face’s Dataset Cards, Croissant metadata format, and our suggested KG Cards. Most cards match different needs and directions depending on the product (i.e., dataset or KG); however, we

suggest incorporating state-of-the-art metrics and emphasising the importance of ethical issues.

### 5.5. Model Cards to Reduce Harms

Besides the Data Cards framework, Hugging Face has also adapted the set of Model Cards for transparency and detailed documentation of their models.<sup>13</sup> We suggest a similar tactic for the explainability of KG embedding models in the KE field. The set of Model Cards [109] can be used to support documentation with a metadata template and enhance model transparency. Table 5 presents the list of Model Cards with descriptions and one example of the KG embedding model, TransE [20]. TransE is a KG embedding model that models relationships in the graph by interpreting them as translations applied to the low-dimensional embeddings of entities. Inspired by the TransE model, several variant models have been developed, such as TransR (rotating) [95] and TransF (folding) [43]. The TransE example in the table shows that KG embedding model studies offer detailed documentation for model training and evaluation, but the lack of demographic details and ethical considerations flag the need, similar to datasets, to adapt specialised instructions for model documentation in the field of KE.

## 6. Conclusion

Generative AI heralds the emergence of a new era for many fields. It is inevitable that they will change KE in many ways. In what direction this will lead us is not clear. The hype around conversational generative AI is such that may one day render KGs obsolete. Currently the volatility and unreliability of generative AI means that KGs will continue to be a trusted source, although it is impossible to tell if this will still be true if LLMs become more reliable in the future. Today, however, the added efficiency in terms of addressing volume based tasks is of great potential for use

<sup>12</sup><https://huggingface.co/docs/hub/datasets-cards>

<sup>13</sup><https://huggingface.co/docs/hub/model-cards>

**Table 4**

Alignment between KG Cards, Hugging Face dataset cards, and Croissant metadata format.

KG Cards	Hugging Face Dataset Cards	Croissant Metadata Format
Knowledge	Dataset Description	The Dataset Metadata Layer
Provenance	Dataset Creation	The Resources Layer
Construction	Dataset Creation	The Structure Layer
Quality	Dataset Structure	The Dataset Metadata Layer
Licensing	Additional Information	The Dataset Metadata Layer
Transparency	Considerations of Using the Data	The Dataset Metadata Layer (RAI properties), The Semantic Layer

**Table 5**

Model Cards by Michelle et al. [109], with descriptions and the alignment for one example of TransE model.

Model Cards	Description	Example Model (TransE)
Model details	Basic information about the model	Bordes et al. [21] 2013 Translation-based model {bordes, nusunier, agarciad}@utcm.fr {jweston, oksana}@google.com
Intended use	Use cases that were envisioned during development	Link prediction (details in [21] section 4.3)
Factors	Factors could include demographic or phenotypic groups, environmental conditions, technical attributes, or others	-
Metrics	Metrics should be chosen to reflect potential real-world impacts of the model	Head and Tail example (details in [21] Table 5)
Evaluation data	Details on the dataset(s) used for the quantitative analyses in the card	WordNet [42] and Freebase [18] (details in [21] section 4.1)
Training data	May not be possible to provide in practice. When possible, this section should mirror Evaluation Data. If such detail is not possible, minimal allowable information should be provided here, such as details of the distribution over various factors in the training datasets	Trained on a large-scale split of Freebase containing 1M entities, 25k relationships and more than 17M training samples (details in [21] section 4.1)
Quantitative analysis	Unitary and intersectional results	Mean and Hits metrics, compared with RESCAL [118], SE [22], SME[19], LFM [72] (details in [21] section 4.2)
Ethical considerations	Demonstrate the ethical considerations that went into model development, surfacing ethical challenges and solutions to stakeholders	-
Caveats and recommendations	List additional concerns that were not covered in the previous sections	-

alongside more traditional methods. In terms of adoption of new ways of creating KGs this recent development is somewhat analogous to the introduction of crowd sourcing to KE; an innovation that has subsequently proved of immense value. However, it is crucial that those working in the field are appropriately skilled and supported to work with generative AI in an effective, productive and safe fashion. As well as training, this implies the need for documentation to accompany LLM-assisted KGs such as prompt templates and KG cards.

## 6.1. Limitations

Participants were selected from European research labs that have a strong profile in either publishing in KE venues, or in offering well-used KE industry products. Given the number of participants not all KE tasks were investigated, the experiments were relatively small scale and groups worked for only three days on the tasks. Furthermore, most interviewees were different-level PhD students, missing the opinion of researchers with more than 10 years of experience in the field. It is worth mentioning as a limitation that there is a gap between the completion of data analysis and the publication of this study, which may influence the contemporaneity of the findings. However, the core insights remain



robust and continue to provide meaningful contributions to how knowledge engineers use generative AI.

## 6.2. Contributions and Future Work

This is one of the first attempts to understand the interaction between knowledge engineers and generative AI. We have identified and presented key areas of strength and challenge for KE working with LLMs, and we have built on the suggestion of implementing Data Cards for KGs and propose extending these to more comprehensively address safety within KGs developed with generative AI.

Future work could explore how generative AI can solve other KE tasks and assist less experienced users in KE tasks, building on its established applications in other domains such as programming, report generation, and content creation. Additionally, further efforts are needed to deepen our understanding of the role generative AI plays in specific KE tasks and throughout the KG lifecycle. A particular focus could be placed on conducting comprehensive studies to identify practical strategies for detecting and mitigating biases in LLM-assisted KGs across the entire development pipeline. Key to this is understanding the scale of whether KE's lack knowledge about safety, or whether the real barrier is obstacles to the implementation of this knowledge.

## 6.3. Acknowledgments

This paper is produced as part of the MuseIT project which has been co-funded by the EU under the Grant Agreement number 101061441. The authors would like to thank all the participants in the 2023 Knowledge Prompting Hackathon at King's College London.

## References

- [1] Abián, D., Guerra, F., Martínez-Romanos, J., Trillo-Lado, R., 2017. Wikidata and dbpedia: a comparative study, in: Semantic Keyword-based Search on Structured Data Sources, Springer. pp. 142–154.
- [2] Acosta, M., Zaveri, A., Simperl, E., Kontokostas, D., Auer, S., Lehmann, J., 2013. Crowdsourcing linked data quality assessment, in: Alani, H., Kagal, L., Fokoue, A., Groth, P., Biemann, C., Parreira, J.X., Aroyo, L., Noy, N., Welty, C., Janowicz, K. (Eds.), The Semantic Web – ISWC 2013, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 260–276.
- [3] Adadi, A., Berrada, M., 2018. Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). IEEE Access 6, 52138–52160. doi:10.1109/ACCESS.2018.2870052.
- [4] AI, N., 2023. Artificial intelligence risk management framework (ai rmf 1.0).
- [5] Akhtar, M., Benjelloun, O., Conforti, C., Gijssbers, P., Giner-Miguel, J., Jain, N., Kuchnik, M., Lhoest, Q., Marcenac, P., Maskey, M., Mattson, P., Oala, L., Ruysen, P., Shinde, R., Simperl, E., Thomas, G., Tykhonov, S., Vanschoren, J., van der Velde, J., Vogler, S., Wu, C.J., 2024. Croissant: A Metadata Format for ML-Ready Datasets, Association for Computing Machinery, New York, NY, USA. p. 1–6. URL: <https://doi.org/10.1145/3650203.3663326>.
- [6] Alivanistos, D., Santamaría, S.B., Cochez, M., Kalo, J.C., van Krieken, E., Thanapalasingam, T., 2023. Prompting as Probing: Using Language Models for Knowledge Base Construction. URL: <https://arxiv.org/abs/2208.11057>, arXiv:2208.11057.
- [7] Allen, B.P., Groth, P.T., 2024. Evaluating class membership relations in knowledge graphs using large language models. URL: <https://arxiv.org/abs/2404.17000>, arXiv:2404.17000.
- [8] Allen, B.P., Stork, L., Groth, P., 2023. Knowledge engineering using large language models. arXiv preprint arXiv:2310.00637.
- [9] Amini, R., Norouzi, S.S., Hitzler, P., Amini, R., 2024. Towards complex ontology alignment using large language models. URL: <https://arxiv.org/abs/2404.10329>, arXiv:2404.10329.
- [10] Anil, R., Dai, A.M., Firat, O., Johnson, M., Lepikhin, D., Passos, A., Shakeri, S., Taropa, E., Bailey, P., Chen, Z., et al., 2023. Palm 2 technical report. arXiv preprint arXiv:2305.10403.
- [11] Anthropic, 2024. Introducing claude. <https://www.anthropic.com/news/introducing-claude>. Accessed: April 9, 2024.
- [12] Baader, F., Horrocks, I., Lutz, C., Sattler, U., 2017. Introduction to description logic. Cambridge University Press.
- [13] Bach, S., Binder, A., Montavon, G., Klauschen, F., Müller, K.R., Samek, W., 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. PLOS ONE 10, 1–46. URL: <https://doi.org/10.1371/journal.pone.0130140>, doi:10.1371/journal.pone.0130140.
- [14] Bastings, J., Filippova, K., 2020. The elephant in the interpretability room: Why use attention as explanation when we have saliency methods?, in: Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP, Association for Computational Linguistics, Online. pp. 149–155. URL: <https://aclanthology.org/2020.blackboxnlp-1.14>, doi:10.18653/v1/2020.blackboxnlp-1.14.
- [15] Bender, E.M., Friedman, B., 2018. Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. Transactions of the Association for Computational Linguistics 6, 587–604.
- [16] Bianchi, F., Rossiello, G., Costabello, L., Palmonari, M., Minervini, P., 2020. Knowledge Graph Embeddings and Explainable AI. CoRR abs/2004.14843. URL: <https://arxiv.org/abs/2004.14843>, arXiv:2004.14843.
- [17] Blomqvist, E., Hammar, K., Presutti, V., 2016. Engineering ontologies with patterns-the extreme design methodology. Ontology Engineering with Ontology Design Patterns 25, 23–50.
- [18] Bollacker, K., Evans, C., Paritosh, P., Sturge, T., Taylor, J., 2008. Freebase: a collaboratively created graph database for structuring human knowledge, in: Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data, Association for Computing Machinery, New York, NY, USA. p. 1247–1250. URL: <https://doi.org/10.1145/1376616.1376746>, doi:10.1145/1376616.1376746.
- [19] Bordes, A., Glorot, X., Weston, J., Bengio, Y., 2014. A semantic matching energy function for learning with multi-relational data: Application to word-sense disambiguation. Machine learning 94, 233–259.
- [20] Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O., 2013a. Translating embeddings for modeling multi-relational data, in: Burges, C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2013/file/1cecc7a77928ca8133fa24680a88d2f9-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2013/file/1cecc7a77928ca8133fa24680a88d2f9-Paper.pdf).
- [21] Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., Yakhnenko, O., 2013b. Translating embeddings for modeling multi-relational data. Advances in neural information processing systems 26.
- [22] Bordes, A., Weston, J., Collobert, R., Bengio, Y., 2011. Learning structured embeddings of knowledge bases, in: Proceedings of the AAAI conference on artificial intelligence, pp. 301–306.
- [23] Brown, B., Juravsky, J., Ehrlich, R., Clark, R., Le, Q.V., Ré, C., Mirhoseini, A., 2024. Large Language Monkeys: Scaling Inference Compute with Repeated Sampling. URL: <https://arxiv.org/abs/2407.21787>, arXiv:2407.21787.
- [24] Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish,

- S., Radford, A., Sutskever, I., Amodei, D., 2020. Language models are few-shot learners. CoRR abs/2005.14165. URL: <https://arxiv.org/abs/2005.14165>, arXiv:2005.14165.
- [25] Bucknall, B.S., Dori-Hacohen, S., 2022. Current and near-term ai as a potential existential risk factor, in: Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society, pp. 119–129.
- [26] Caufield, J.H., Hegde, H., Emonet, V., Harris, N.L., Joachimiak, M.P., Matentzoglou, N., Kim, H., Moxon, S., Reese, J.T., Haendel, M.A., Robinson, P.N., Mungall, C.J., 2024. Structured prompt interrogation and recursive extraction of semantics (spires): a method for populating knowledge bases using zero-shot learning. *Bioinformatics* 40, btac104. URL: <https://doi.org/10.1093/bioinformatics/btac104>, doi:10.1093/bioinformatics/btac104, arXiv:<https://academic.oup.com/bioinformatics/article-pdf/40/3/btac104/6441111>
- [27] Chowdhery, A., Narang, S., Devlin, J., Bosma, M., Mishra, G., Roberts, A., Barham, P., Chung, H.W., Sutton, C., Gehrmann, S., et al., 2023. Palm: Scaling language modeling with pathways. *Journal of Machine Learning Research* 24, 1–113.
- [28] Cimiano, P., Paulheim, H., 2017. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semant. Web* 8, 489–508. URL: <https://doi.org/10.3233/SW-160218>, doi:10.3233/SW-160218.
- [29] Dai, W., Li, J., Li, D., Tiong, A.M.H., Zhao, J., Wang, W., Li, B., Fung, P.N., Hoi, S., 2024. Instructblip: Towards general-purpose vision-language models with instruction tuning. *Advances in Neural Information Processing Systems* 36.
- [30] Danilevsky, M., Qian, K., Aharonov, R., Katsis, Y., Kawas, B., Sen, P., 2020. A Survey of the State of Explainable AI for Natural Language Processing. CoRR abs/2010.00711. URL: <https://arxiv.org/abs/2010.00711>, arXiv:2010.00711.
- [31] Darari, F., Nutt, W., Pirrò, G., Razniewski, S., 2018. Completeness management for rdf data sources. *ACM Transactions on the Web (TWEB)* 12, 1–53.
- [32] Debenham, J., 2012. Knowledge Engineering: Unifying Knowledge Base and Database Design. Springer Science & Business Media.
- [33] Denaux, R., Dolbear, C., Hart, G., Dimitrova, V., Cohn, A.G., 2011. Supporting domain experts to construct conceptual ontologies: A holistic approach. *Journal of Web semantics* 9, 113–127.
- [34] Dessì, D., Osborne, F., Reforgiato Recupero, D., Buscaldi, D., Motta, E., 2022. Cs-kg: A large-scale knowledge graph of research entities and claims in computer science, in: International Semantic Web Conference, Springer. pp. 678–696.
- [35] Dettmers, T., Minervini, P., Stenetorp, P., Riedel, S., 2018. Convolutional 2d knowledge graph embeddings. *Proceedings of the AAAI Conference on Artificial Intelligence* 32. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/11573>, doi:10.1609/aaai.v32i1.11573.
- [36] Dhuliawala, S., Komeili, M., Xu, J., Raileanu, R., Li, X., Celikyilmaz, A., Weston, J., 2024. Chain-of-Verification Reduces Hallucination in Large Language Models, in: Ku, L.W., Martins, A., Srikumar, V. (Eds.), Findings of the Association for Computational Linguistics: ACL 2024, Association for Computational Linguistics, Bangkok, Thailand. pp. 3563–3578. URL: <https://aclanthology.org/2024.findings-acl.212/>, doi:10.18653/v1/2024.findings-acl.212.
- [37] Ding, L., Zhou, S., Xiao, J., Han, J., 2024. Automated Construction of Theme-specific Knowledge Graphs. arXiv preprint arXiv:2404.19146.
- [38] Dodds, L., 2020. Fair, fairer, fairest? URL: <https://blog.ldodds.com/2020/07/30/fair-fairer-fairest/>.
- [39] Dong, Y., Ruan, C.F., Cai, Y., Lai, R., Xu, Z., Zhao, Y., Chen, T., 2024. XGrammar: Flexible and Efficient Structured Generation Engine for Large Language Models. CoRR abs/2411.15100. URL: <https://doi.org/10.48550/arXiv.2411.15100>, doi:10.48550/ARXIV.2411.15100, arXiv:2411.15100.
- [40] Elliott, D., Soifer, E., 2022. Ai technologies, privacy, and security. *Frontiers in Artificial Intelligence* 5, 826737.
- [41] Euzenat, J., Meilicke, C., Stuckenschmidt, H., Shvaiko, P., Trojahn, C., 2011. Ontology alignment evaluation initiative: six years of experience. *Journal on data semantics XV*, 158–192.
- [42] Fellbaum, C., 2010. Wordnet, in: Theory and applications of ontology: computer applications. Springer, pp. 231–243.
- [43] Feng, J., Huang, M., Wang, M., Zhou, M., Hao, Y., Zhu, X., 2016. Knowledge Graph Embedding by Flexible Translation, in: Fifteenth International Conference on the Principles of Knowledge Representation and Reasoning.
- [44] Fensel, D., Simsek, U., Angele, K., Huaman, E., Kärle, E., Panasiuk, O., Toma, I., Umbrich, J., Wahler, A., 2020. Knowledge graphs. Springer.
- [45] Fleiss, J.L., Nee, J.C., Landis, J.R., 1979. Large sample variance of kappa in the case of different sets of raters. *Psychological bulletin* 86, 974.
- [46] Friedman, D., Nissenbaum, H., 1996. Bias in computer systems. *ACM Trans. Inf. Syst.* 14, 330–347. URL: <https://doi.org/10.1145/230538.230561>, doi:10.1145/230538.230561.
- [47] Friedman, D., Dieng, A.B., 2023. The vendi score: A diversity evaluation metric for machine learning. URL: <https://arxiv.org/abs/2210.02410>, arXiv:2210.02410.
- [48] Gallegos, I.O., Rossi, R.A., Barrow, J., Tanjim, M.M., Kim, S., Derroncourt, F., Yu, T., Zhang, R., Ahmed, N.K., 2024. Bias and fairness in large language models: A survey. *Computational Linguistics* 50, 1097–1179. URL: <https://aclanthology.org/2024.cl-3.8/>, doi:10.1162/coli\_a\_00524.
- [49] Gangemi, A., 2013. A comparison of knowledge extraction tools for the semantic web, in: Cimiano, P., Corcho, O., Presutti, V., Hollink, L., Rudolph, S. (Eds.), The Semantic Web: Semantics and Big Data, Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 351–366.
- [50] Gao, Y., Xiong, Y., Gao, X., Jia, K., Pan, J., Bi, Y., Dai, Y., Sun, J., Guo, Q., Wang, M., Wang, H., 2023. Retrieval-augmented generation for large language models: A survey. CoRR abs/2312.10997. URL: <https://doi.org/10.48550/arXiv.2312.10997>, doi:10.48550/ARXIV.2312.10997, arXiv:2312.10997.
- [51] Gayo, J.E.L., Prud’Hommeaux, E., Boneva, I., Kontokostas, D., 2017. Validating RDF data. Morgan & Claypool Publishers.
- [52] Geng, J., Cai, F., Wang, Y., Koepl, H., Nakov, P., Gurevych, I., 2024. A Survey of Confidence Estimation and Calibration in Large Language Models. URL: <https://arxiv.org/abs/2311.08298>, arXiv:2311.08298.
- [53] Ghaddar, A., Langlais, P., Rashid, A., Rezagholizadeh, M., 2021. Context-aware adversarial training for name regularity bias in named entity recognition. *Transactions of the Association for Computational Linguistics* 9, 586–604.
- [54] Giglou, H.B., D’Souza, J., Engel, F., Auer, S., 2024. Lms4om: Matching ontologies with large language models. URL: <https://arxiv.org/abs/2404.10317>, arXiv:2404.10317.
- [55] Glimm, B., Horrocks, I., Motik, B., Stoilos, G., Wang, Z., 2014. Hermit: an owl 2 reasoner. *Journal of automated reasoning* 53, 245–269.
- [56] Gogoll, J., Zuber, N., Kacianka, S., Greger, T., Pretschner, A., Nida-Rümelin, J., 2021. Ethics in the software development process: from codes of conduct to ethical deliberation. *Philosophy & Technology*, 1–24.
- [57] Google, 2024. Bard main page. URL: <https://bard.google.com/chat>, accessed on January, 2024.
- [58] Groth, P., Simperl, E., van Erp, M., Vrandečić, D., 2023. Knowledge graphs and their role in the knowledge engineering of the 21st century (dagstuhl seminar 22372), in: Dagstuhl Reports, Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [59] Guarino, N., Welty, C., 2002. Evaluating ontological decisions with ontoclean. *Communications of the ACM* 45, 61–65.
- [60] Guarino, N., Welty, C.A., 2009. An overview of ontoclean. *Handbook on ontologies*, 201–220.
- [61] Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., Pedreschi, D., 2018. A Survey of Methods for Explaining Black Box Models. *ACM Comput. Surv.* 51. URL: <https://doi.org/10.1145/3236009>, doi:10.1145/3236009.

- [62] Guo, J., Li, J., Li, D., Tiong, A.M.H., Li, B., Tao, D., Hoi, S.C.H., 2022. From images to textual prompts: Zero-shot vqa with frozen large language models. URL: <https://arxiv.org/abs/2212.10846>, doi:10.48550/ARXIV.2212.10846.
- [63] Hannah, G., de Berardinis, J., Payne, T.R., Tamma, V., Mitchell, A., Piercy, E., Johnson, E., Ng, A., Rostron, H., Konev, B., 2025. Large Language Models as Knowledge Evaluation Agents, in: 2nd Workshop on Evaluation of Language Models in Knowledge Engineering (ELMKE), co-located with ESWC-25, to appear. URL: <https://ceur-ws.org/Vol-3977/elmke-4.pdf>.
- [64] He, Y., Chen, J., Dong, H., Horrocks, I., Allocca, C., Kim, T., Sapkota, B., 2024. DeepOnto: A Python package for ontology engineering with deep learning. Semantic Web 15, 1991–2004.
- [65] Heim, D., Meyer, L.P., Schröder, M., Frey, J., Dengel, A., 2025. How do Scaling Laws Apply to Knowledge Graph Engineering Tasks? The Impact of Model Size on Large Language Model Performance, in: 2nd Workshop on Evaluation of Language Models in Knowledge Engineering (ELMKE), co-located with ESWC-25, to appear. URL: <https://ceur-ws.org/Vol-3977/elmke-3.pdf>.
- [66] Heist, N., Hertling, S., Ringler, D., Paulheim, H., 2020. Knowledge graphs on the web-an overview. Knowledge Graphs for eXplainable Artificial Intelligence , 3–22.
- [67] Heston, T.F., Khun, C., 2023. Prompt engineering in medical education. International Medical Education 2, 198–205. URL: <https://www.mdpi.com/2813-141X/2/3/19>, doi:10.3390/ime2030019.
- [68] Hogan, A., Blomqvist, E., Cochez, M., d'Amato, C., Melo, G.d., Gutierrez, C., Kirrane, S., Gayo, J.E.L., Navigli, R., Neumaier, S., et al., 2022. Knowledge graphs .
- [69] Huang, D., Bu, Q., Zhang, J., Xie, X., Chen, J., Cui, H., 2023. Bias assessment and mitigation in llm-based code generation. arXiv preprint arXiv:2309.14345 .
- [70] Issa, S., Adekunle, O., Hamdi, F., Cherfi, S.S.S., Dumontier, M., Zaveri, A., 2021. Knowledge graph completeness: A systematic literature review. IEEE Access 9, 31322–31339.
- [71] Jain, S., Wallace, B.C., 2019. Attention is not explanation. CoRR abs/1902.10186. URL: <http://arxiv.org/abs/1902.10186>, arXiv:1902.10186.
- [72] Jenatton, R., Roux, N., Bordes, A., Obozinski, G.R., 2012. A latent factor model for highly multi-relational data. Advances in neural information processing systems 25.
- [73] Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P., 2023. Survey of hallucination in natural language generation. ACM Computing Surveys 55, 1–38.
- [74] Jie, Z., Lu, W., 2019. Dependency-guided lstm-crf for named entity recognition, in: Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 3862–3872.
- [75] Kalo, J.C., Nguyen, T.P., Razniewski, S., Zhang, B., 2024. Preface: LM-KBC Challenge 2024, in: 2nd Workshop on Knowledge Base Construction from Pre-Trained Language Models, CEUR. ws.
- [76] Kazim, E., Koshiyama, A.S., 2021. A high-level overview of ai ethics. Patterns 2.
- [77] Keidar, D., Zhong, M., Zhang, C., Shrestha, Y.R., Paudel, B., 2021. Towards automatic bias detection in knowledge graphs. arXiv preprint arXiv:2109.10697 .
- [78] Kendall, E.F., McGuinness, D.L., 2019a. Ontology engineering. Morgan & Claypool Publishers.
- [79] Kendall, E.F., McGuinness, D.L., 2019b. Requirements and Use Cases. Springer International Publishing, Cham. pp. 25–44. URL: [https://doi.org/10.1007/978-3-031-79486-5\\_3](https://doi.org/10.1007/978-3-031-79486-5_3), doi:10.1007/978-3-031-79486-5\_3.
- [80] Kim, T., Springer, J.M., Raghunathan, A., Sap, M., 2025. Mitigating bias in RAG: Controlling the embedder, in: Che, W., Nabende, J., Shutova, E., Pilehvar, M.T. (Eds.), Findings of the Association for Computational Linguistics: ACL 2025, Association for Computational Linguistics, Vienna, Austria. pp. 18999–19024. URL: <https://aclanthology.org/2025.findings-acl.974/>, doi:10.18653/v1/2025.findings-acl.974.
- [81] Kleinstauber, E., Al Mustafa, T., Zander, F., König-Ries, B., Babalou, S., 2024. Managing provenance data in knowledge graph management platforms. Datenbank-Spektrum 24, 43–52.
- [82] Kommineni, V.K., König-Ries, B., Samuel, S., 2024. From human experts to machines: An llm supported approach to ontology and knowledge graph construction. URL: <https://arxiv.org/abs/2403.08345>, arXiv:2403.08345.
- [83] Koo, T., Liu, F., He, L., 2024. Automata-based constraints for language model decoding, in: First Conference on Language Modeling. URL: <https://openreview.net/forum?id=B0BdbImyzY>.
- [84] Kou, Z., Zhang, Y., Zhang, D., Wang, D., 2022. CrowdGraph: A Crowdsourcing Multi-modal Knowledge Graph Approach to Explainable Fauxtography Detection. Proc. ACM Hum.-Comput. Interact. 6. URL: <https://doi.org/10.1145/3555178>, doi:10.1145/3555178.
- [85] Koutsiana, E., Yadav, T., Jain, N., Meroño-Peñuela, A., Simperl, E., 2023. Agreeing and disagreeing in collaborative knowledge graph construction: An analysis of wikidata. CoRR abs/2306.11766. URL: <https://doi.org/10.48550/arXiv.2306.11766>, doi:10.48550/ARXIV.2306.11766, arXiv:2306.11766.
- [86] Kraft, A., Usbeck, R., 2022. The Lifecycle of "Facts": A Survey of Social Bias in Knowledge Graphs. arXiv preprint arXiv:2210.03353 .
- [87] Lancaster, C.M., Schulenberg, K., Flathmann, C., McNeese, N.J., Freeman, G., 2024. "it's everybody's role to speak up... but not everyone will": Understanding ai professionals' perceptions of accountability for ai bias mitigation. ACM J. Responsib. Comput. 1. URL: <https://doi.org/10.1145/3632121>, doi:10.1145/3632121.
- [88] Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J., 2013. Prov-o: The prov ontology: W3c recommendation 30 april 2013.
- [89] Lent, H., Søgaaard, A., 2021. Common sense bias in semantic role labeling, in: Proceedings of the seventh workshop on noisy user-generated text (w-nut 2021), pp. 114–119.
- [90] Lewis, P.S.H., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., Kiela, D., 2020. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. CoRR abs/2005.11401. URL: <https://arxiv.org/abs/2005.11401>, arXiv:2005.11401.
- [91] Li, J., Sun, A., Han, J., Li, C., 2018. A survey on deep learning for named entity recognition. CoRR abs/1812.09449. URL: <http://arxiv.org/abs/1812.09449>, arXiv:1812.09449.
- [92] Li, Y., Li, J., Suhara, Y., Doan, A., Tan, W.C., 2020. Deep entity matching with pre-trained language models. arXiv preprint arXiv:2004.00584 .
- [93] Liang, P., Bommasani, R., Lee, T., Tsipras, D., Soylu, D., Yasunaga, M., Zhang, Y., Narayanan, D., Wu, Y., Kumar, A., et al., 2022. Holistic evaluation of language models. arXiv preprint arXiv:2211.09110 .
- [94] emerald publishing limited, 2023. Ethnography techniques. URL: <https://www.emeraldgrouppublishing.com/how-to/observation/use-ethnographic-methods-participant-observation>, . accessed on January, 2024.
- [95] Lin, Y., Liu, Z., Sun, M., Liu, Y., Zhu, X., 2015. Learning Entity and Relation Embeddings for Knowledge Graph Completion. Proceedings of the AAAI Conference on Artificial Intelligence 29. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/9491>, doi:10.1609/aaai.v29i1.9491.
- [96] Lin, Y., Shen, S., Liu, Z., Luan, H., Sun, M., 2016. Neural relation extraction with selective attention over instances, in: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), Association for Computational Linguistics, Berlin, Germany. pp. 2124–2133. URL: <https://aclanthology.org/P16-1200>, doi:10.18653/v1/P16-1200.
- [97] Lippolis, A.S., Saeedizade, M.J., Keskiarkka, R., Gangemi, A., Blomqvist, E., Nuzzolese, A.G., 2025. Assessing the Capability of Large Language Models for Domain-Specific Ontology Generation, in: 2nd Workshop on Evaluation of Language Models in Knowledge



- Engineering (ELMKE), co-located with ESWC-25, to appear. URL: <https://ceur-ws.org/Vol-3977/elmke-2.pdf>.
- [98] Liu, J., Contributors, 2024. Instructor: A library for structured outputs from large language models. URL: <https://github.com/instructor-ai/instructor>.
- [99] Liu, P., Yuan, W., Fu, J., Jiang, Z., Hayashi, H., Neubig, G., 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys* 55, 1–35.
- [100] Liu, W., Zhou, P., Zhao, Z., Wang, Z., Ju, Q., Deng, H., Wang, P., 2020. K-BERT: Enabling Language Representation with Knowledge Graph, in: *Proceedings of AAAI 2020*.
- [101] Llugiqi, M., Ekaputra, F.J., Sabou, M., 2025. From Experts to LLMs: Evaluating the Quality of Automatically Generated Ontologies, in: *2nd Workshop on Evaluation of Language Models in Knowledge Engineering (ELMKE)*, co-located with ESWC-25, to appear. URL: <https://ceur-ws.org/Vol-3977/elmke-1.pdf>.
- [102] Lundberg, S.M., Lee, S.I., 2017. A unified approach to interpreting model predictions, in: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc., pp. 4765–4774. URL: <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>.
- [103] McGraw, K.L., Harbison-Briggs, K., 1989. *Knowledge acquisition: Principles and guidelines*. Prentice-Hall, Inc.
- [104] Meilicke, C., Chekol, M.W., Fink, M., Stuckenschmidt, H., 2020. Reinforced anytime bottom up rule learning for knowledge graph completion. *CoRR abs/2004.04412*. URL: <https://arxiv.org/abs/2004.04412>, arXiv:2004.04412.
- [105] Mikalef, P., Conboy, K., Lundström, J.E., Popovič, A., 2022. Thinking responsibly about responsible AI and ‘the dark side’ of AI.
- [106] Miller, T., 2019. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence* 267, 1–38. URL: <https://www.sciencedirect.com/science/article/pii/S0004370218305988>, doi:<https://doi.org/10.1016/j.artint.2018.07.007>.
- [107] Mintz, M., Bills, S., Snow, R., Jurafsky, D., 2009. Distant supervision for relation extraction without labeled data, in: *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, Association for Computational Linguistics, Suntec, Singapore. pp. 1003–1011. URL: <https://aclanthology.org/P09-1113>.
- [108] Mishra, S., He, S., Belli, L., 2020. Assessing demographic bias in named entity recognition. *arXiv preprint arXiv:2008.03415*.
- [109] Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I.D., Gebru, T., 2019. Model cards for model reporting, in: *Proceedings of the conference on fairness, accountability, and transparency*, pp. 220–229.
- [110] Morales, S., Clarisó, R., Cabot, J., 2024. A DSL for Testing LLMs for Fairness and Bias, in: *Proceedings of the ACM/IEEE 27th International Conference on Model Driven Engineering Languages and Systems*, pp. 203–213.
- [111] Munn, L., 2023. The uselessness of ai ethics. *AI and Ethics* 3, 869–877.
- [112] Mykityn Jr, P.P., Mykityn, K., Raja, M., 1994. Knowledge acquisition skills and traits: a self-assessment of knowledge engineers. *Information & Management* 26, 95–104.
- [113] Nadeem, M., Bethke, A., Reddy, S., 2021. Stereoset: Measuring stereotypical bias in pretrained language models, in: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 5356–5371.
- [114] Nan, G., Zeng, J., Qiao, R., Guo, Z., Lu, W., 2021. Uncovering main causalities for long-tailed information extraction, in: *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 9683–9695.
- [115] Nasr, M., Carlini, N., Hayase, J., Jagielski, M., Cooper, A.F., Ippolito, D., Choquette-Choo, C.A., Wallace, E., Tramèr, F., Lee, K., 2023. Scalable extraction of training data from (production) language models. *arXiv preprint arXiv:2311.17035*.
- [116] Nathani, D., Chauhan, J., Sharma, C., Kaul, M., 2019. Learning attention-based embeddings for relation prediction in knowledge graphs. *CoRR abs/1906.01195*. URL: <http://arxiv.org/abs/1906.01195>, arXiv:1906.01195.
- [117] Nickel, M., Tresp, V., Kriegel, H.P., 2011a. A three-way model for collective learning on multi-relational data, in: *Proceedings of the 28th International Conference on International Conference on Machine Learning*, Omnipress, Madison, WI, USA. p. 809–816.
- [118] Nickel, M., Tresp, V., Kriegel, H.P., et al., 2011b. A three-way model for collective learning on multi-relational data., in: *Icml*, pp. 3104482–3104584.
- [119] Noy, N., Gao, Y., Jain, A., Narayanan, A., Patterson, A., Taylor, J., 2019. Industry-scale knowledge graphs: Lessons and challenges: Five diverse technology companies show how it’s done. *Queue* 17, 48–75.
- [120] Ntoutsis, E., Fafalios, P., Gadiraju, U., Iosifidis, V., Nejd, W., Vidal, M.E., Ruggieri, S., Turini, F., Papadopoulos, S., Krasanakis, E., et al., 2020. Bias in data-driven artificial intelligence systems—an introductory survey. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 10, e1356.
- [121] OpenAI, 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- [122] OpenAI, 2024. Chatgpt. URL: <https://openai.com/chatgpt>, accessed: April 9, 2024.
- [123] Pan, J.Z., Razniewski, S., Kalo, J.C., Singhania, S., Chen, J., Dietze, S., Jabeen, H., Omeliyanenko, J., Zhang, W., Lissandrini, M., Biswas, R., de Melo, G., Bonifati, A., Vakaj, E., Dragoni, M., Graux, D., 2023a. Large Language Models and Knowledge Graphs: Opportunities and Challenges. *Transactions on Graph Data and Knowledge* 1, 2:1–2:38. URL: <https://drops.dagstuhl.de/entities/document/10.4230/TGDK.1.1.2>, doi:10.4230/TGDK.1.1.2.
- [124] Pan, Y., Pan, L., Chen, W., Nakov, P., Kan, M.Y., Wang, W.Y., 2023b. On the risk of misinformation pollution with large language models. *arXiv preprint arXiv:2305.13661*.
- [125] Papadakis, G., Skoutas, D., Thanos, E., Palpanas, T., 2020. Blocking and filtering techniques for entity resolution: A survey. *ACM Comput. Surv.* 53. URL: <https://doi.org/10.1145/3377455>, doi:10.1145/3377455.
- [126] Paparidis, E., Kotis, K., 2021. Towards engineering fair ontologies: Unbiasing a surveillance ontology, in: *2021 IEEE International Conference on Progress in Informatics and Computing (PIC)*, IEEE. pp. 226–231.
- [127] Paulheim, H., 2017. Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web* 8, 489–508.
- [128] Peng, B., Chen, K., Li, M., Feng, P., Bi, Z., Liu, J., Niu, Q., 2024. Securing large language models: Addressing bias, misinformation, and prompt attacks. *arXiv preprint arXiv:2409.08087*.
- [129] Peng, H., Wang, X., Hu, S., Jin, H., Hou, L., Li, J., Liu, Z., Liu, Q., 2022. COPEN: Probing conceptual knowledge in pre-trained language models, in: Goldberg, Y., Kozareva, Z., Zhang, Y. (Eds.), *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Abu Dhabi, United Arab Emirates. pp. 5015–5035. URL: <https://aclanthology.org/2022.emnlp-main.335>, doi:10.18653/v1/2022.emnlp-main.335.
- [130] Peñuela, A.M., Dimou, A., Troncy, R., Hartig, O., Acosta, M., Alam, M., Paulheim, H., Lisena, P., 2024. The Semantic Web: 21st International Conference, ESWC 2024, Hersonissos, Crete, Greece, May 26–30, 2024, *Proceedings, Part I*. volume 14664. Springer Nature.
- [131] Petroni, F., Piktus, A., Fan, A., Lewis, P., Yazdani, M., De Cao, N., Thorne, J., Jernite, Y., Karpukhin, V., Maillard, J., Plachouras, V., Rocktäschel, T., Riedel, S., 2021. KILT: a Benchmark for Knowledge Intensive Language Tasks, in: Toutanova, K., Rumshisky,



- A., Zettlemoyer, L., Hakkani-Tur, D., Beltagy, I., Bethard, S., Cotterell, R., Chakraborty, T., Zhou, Y. (Eds.), Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online. pp. 2523–2544. URL: <https://aclanthology.org/2021.naacl-main.200/>, doi:10.18653/v1/2021.naacl-main.200.
- [132] Petroni, F., Rocktäschel, T., Lewis, P.S.H., Bakhtin, A., Wu, Y., Miller, A.H., Riedel, S., 2019. Language models as knowledge bases? CoRR abs/1909.01066. URL: <http://arxiv.org/abs/1909.01066>, arXiv:1909.01066.
- [133] Piscopo, A., Simperl, E., 2019. What we talk about when we talk about wikidata quality: a literature survey, in: Proceedings of the 15th International Symposium on Open Collaboration, Association for Computing Machinery, New York, NY, USA. URL: <https://doi.org/10.1145/3306446.3340822>, doi:10.1145/3306446.3340822.
- [134] Plu, J., Moreno Escobar, O., Trouilleux, E., Gapin, A., Troncy, R., 2024. A comprehensive benchmark for evaluating llm-generated ontologies, in: CEUR (Ed.), ISWC 2024, 23rd International Semantic Web Conference, 11-15 November 2024, Baltimore, USA, Baltimore. CEUR.
- [135] Polat, F., Tiddi, I., Groth, P., 2025. Testing prompt engineering methods for knowledge extraction from text. Semantic Web 16, SW-243719. URL: <https://journals.sagepub.com/doi/abs/10.3233/SW-243719>, doi:10.3233/SW-243719, arXiv:<https://journals.sagepub.com/doi/pdf/10.3233/SW-243719>.
- [136] Poveda-Villalón, M., Gómez-Pérez, A., Suárez-Figueroa, M.C., 2014. Oops!(ontology pitfall scanner!): An on-line tool for ontology evaluation. International Journal on Semantic Web and Information Systems (IJSWIS) 10, 7–34.
- [137] Poveda-Villalón, M., Suárez-Figueroa, M.C., Gómez-Pérez, A., 2012. Validating ontologies with oops!, in: Knowledge Engineering and Knowledge Management: 18th International Conference, EKAW 2012, Galway City, Ireland, October 8-12, 2012. Proceedings 18, Springer. pp. 267–281.
- [138] Pushkarna, M., Zaldivar, A., Kjartansson, O., 2022. Data cards: Purposeful and transparent dataset documentation for responsible ai, in: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, pp. 1776–1826.
- [139] Qarout, R., Checco, A., Demartini, G., Bontcheva, K., 2019. Platform-related factors in repeatability and reproducibility of crowdsourcing tasks, in: Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, pp. 135–143.
- [140] Qin, G., Eisner, J., 2021. Learning How to Ask: Querying LMs with Mixtures of Soft Prompts, in: Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Association for Computational Linguistics, Online. pp. 5203–5212. URL: <https://aclanthology.org/2021.naacl-main.410>, doi:10.18653/v1/2021.naacl-main.410.
- [141] Qin, Y., Liang, S., Ye, Y., Zhu, K., Yan, L., Lu, Y., Lin, Y., Cong, X., Tang, X., Qian, B., Zhao, S., Hong, L., Tian, R., Xie, R., Zhou, J., Gerstein, M., Li, D., Liu, Z., Sun, M., 2023. ToolLLM: Facilitating Large Language Models to Master 16000+ Real-world APIs. URL: <https://arxiv.org/abs/2307.16789>, arXiv:2307.16789.
- [142] Ren, J., Zhao, Y., Vu, T., Liu, P.J., Lakshminarayanan, B., 2023. Self-Evaluation Improves Selective Generation in Large Language Models. URL: <https://arxiv.org/abs/2312.09300>, arXiv:2312.09300.
- [143] Revenko, A., Sabou, M., Ahmeti, A., Schauer, M., 2018. Crowdsourced knowledge graph extension: A belief revision based approach., in: HCOMP (WIP&Demo).
- [144] Ribeiro, M.T., Singh, S., Guestrin, C., 2016. ‘Why should I trust you?’: Explaining the predictions of any classifier. CoRR abs/1602.04938. URL: <http://arxiv.org/abs/1602.04938>, arXiv:1602.04938.
- [145] Rossi, A., Firmani, D., Matinata, A., Merialdo, P., Barbosa, D., 2020. Knowledge graph embedding for link prediction: A comparative analysis. CoRR abs/2002.00819. URL: <https://arxiv.org/abs/2002.00819>, arXiv:2002.00819.
- [146] Rossi, A., Firmani, D., Merialdo, P., Teofili, T., 2022. Explaining link prediction systems based on knowledge graph embeddings, in: Proceedings of the 2022 International Conference on Management of Data, Association for Computing Machinery, New York, NY, USA. p. 2062–2075. URL: <https://doi.org/10.1145/3514221.3517887>, doi:10.1145/3514221.3517887.
- [147] Sarker, M.K., Zhou, L., Eberhart, A., Hitzler, P., 2021. Neuro-symbolic artificial intelligence. AI Communications 34, 197–209.
- [148] Schlichtkrull, M., Kipf, T.N., Bloem, P., van den Berg, R., Titov, I., Welling, M., 2018. Modeling relational data with graph convolutional networks, in: Gangemi, A., Navigli, R., Vidal, M.E., Hitzler, P., Troncy, R., Hollink, L., Tordai, A., Alam, M. (Eds.), The Semantic Web, Springer International Publishing, Cham. pp. 593–607.
- [149] Schneider, P., Schopf, T., Vladika, J., Galkin, M., Simperl, E., Matthes, F., 2022. A decade of knowledge graphs in natural language processing: A survey, in: Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), Association for Computational Linguistics, Online only. pp. 601–614. URL: <https://aclanthology.org/2022.aacl-main.46>.
- [150] Schreiber, G., 2000. Knowledge Engineering and Management: The CommonKADS Methodology. A Bradford book, MIT Press. URL: [https://books.google.co.uk/books?id=H1XOW\\_1fsIEC](https://books.google.co.uk/books?id=H1XOW_1fsIEC).
- [151] Seo, S., Cheon, H., Kim, H., Hyun, D., 2022. Structural quality metrics to evaluate knowledge graph quality.
- [152] Sevgili, Ö., Shelmanov, A., Arkhipov, M.Y., Panchenko, A., Biemann, C., 2020. Neural Entity Linking: A Survey of Models based on Deep Learning. CoRR abs/2006.00575. URL: <https://arxiv.org/abs/2006.00575>, arXiv:2006.00575.
- [153] Shen, W., Wang, J., Han, J., 2015. Entity linking with a knowledge base: Issues, techniques, and solutions. IEEE Transactions on Knowledge and Data Engineering 27, 443–460. doi:10.1109/TKDE.2014.2327028.
- [154] Shenoy, K., Ilievski, F., Garijo, D., Schwabe, D., Szekely, P.A., 2021. A study of the quality of wikidata. CoRR abs/2107.00156. URL: <https://arxiv.org/abs/2107.00156>, arXiv:2107.00156.
- [155] Shneiderman, B., 2020. Human-centered artificial intelligence: Reliable, safe & trustworthy. International Journal of Human–Computer Interaction 36, 495–504.
- [156] Simperl, E., Luczak-Rösch, M., 2014. Collaborative ontology engineering: a survey. The Knowledge Engineering Review 29, 101–131. doi:10.1017/S0269888913000192.
- [157] Simperl, E., Luczak-Rösch, M., 2014. Collaborative ontology engineering: a survey. The Knowledge Engineering Review 29, 101–131.
- [158] Simsek, U., Kärle, E., Angele, K., Huaman, E., Opdenplatz, J., Sommer, D., Umbrich, J., Fensel, D., 2022. A knowledge graph perspective on knowledge engineering. SN Comput. Sci. 4. URL: <https://doi.org/10.1007/s42979-022-01429-x>, doi:10.1007/s42979-022-01429-x.
- [159] Smirnova, A., Cudré-Mauroux, P., 2018. Relation extraction using distant supervision: A survey. ACM Comput. Surv. 51. URL: <https://doi.org/10.1145/3241741>, doi:10.1145/3241741.
- [160] Socher, R., Chen, D., Manning, C.D., Ng, A., 2013. Reasoning with neural tensor networks for knowledge base completion, in: Burges, C., Bottou, L., Welling, M., Ghahramani, Z., Weinberger, K. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2013/file/b337e84de8752b27eda3a12363109e80-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2013/file/b337e84de8752b27eda3a12363109e80-Paper.pdf).
- [161] Star, S.L., Griesemer, J.R., 1989. Institutional Ecology, ‘Translations’ and Boundary Objects: Amateurs and Professionals in Berkeley’s Museum of Vertebrate Zoology, 1907-39. Social Studies of Science 19, 387–420. URL: <https://doi.org/10.1177/030631289019003001>, doi:10.1177/030631289019003001, arXiv:<https://doi.org/10.1177/030631289019003001>.

- [162] Steiner, A., Peeters, R., Bizer, C., 2024. Fine-tuning large language models for entity matching. arXiv preprint arXiv:2409.08185 .
- [163] Studer, R., Benjamins, V.R., Fensel, D., 1998a. Knowledge engineering: Principles and methods. Data & knowledge engineering 25, 161–197.
- [164] Studer, R., Benjamins, V.R., Fensel, D., 1998b. Knowledge engineering: Principles and methods. Data & knowledge engineering 25, 161–197.
- [165] Suárez-Figueroa, M.C., Gómez-Pérez, A., Fernández-López, M., 2011. The neon methodology for ontology engineering, in: Ontology engineering in a networked world. Springer, pp. 9–34.
- [166] Tamašauskaitė, G., Groth, P., 2023. Defining a knowledge graph development process through a systematic review. ACM Transactions on Software Engineering and Methodology 32, 1–40.
- [167] Tiddi, I., Schlobach, S., 2022. Knowledge graphs as tools for explainable machine learning: A survey. Artificial Intelligence 302, 103627. URL: <https://www.sciencedirect.com/science/article/pii/S0004370221001788>, doi:<https://doi.org/10.1016/j.artint.2021.103627>.
- [168] Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al., 2023a. Llama: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971 .
- [169] Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al., 2023b. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288 .
- [170] Tsaneva, S., Dessì, D., Osborne, F., Sabou, M., 2025. Knowledge graph validation by integrating LLMs and human-in-the-loop. Information Processing & Management 62, 104145. URL: <https://www.sciencedirect.com/science/article/pii/S030645732500086X>, doi:<https://doi.org/10.1016/j.ipm.2025.104145>.
- [171] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need. CoRR abs/1706.03762. URL: <http://arxiv.org/abs/1706.03762>, arXiv:1706.03762.
- [172] Veseli, B., Singhania, S., Razniewski, S., Weikum, G., 2023. Evaluating Language Models For Knowledge Base Completion, in: The Semantic Web: 20th International Conference, ESWC 2023, Herssonissos, Crete, Greece, May 28–June 1, 2023, Proceedings, Springer-Verlag, Berlin, Heidelberg. p. 227–243. URL: [https://doi.org/10.1007/978-3-031-33455-9\\_14](https://doi.org/10.1007/978-3-031-33455-9_14), doi:10.1007/978-3-031-33455-9\_14.
- [173] Walker, J., Koutsiana, E., Massey, J., Theurmer, G., Simperl, E., 2023. Prompting datasets: Data discovery with conversational agents. arXiv preprint arXiv:2312.09947 .
- [174] Wang, K., Shen, Z., Huang, C., Wu, C.H., Dong, Y., Kanakia, A., 2020. Microsoft Academic Graph: When experts are not enough. Quantitative Science Studies 1, 396–413. URL: [https://doi.org/10.1162/qss\\_a\\_00021](https://doi.org/10.1162/qss_a_00021), doi:10.1162/qss\_a\_00021.
- [175] Wang, Z., Culotta, A., 2020. Identifying spurious correlations for robust text classification, in: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 3431–3440.
- [176] Wang, Z., Culotta, A., 2021. Robustness to spurious correlations in text classification via automatically generated counterfactuals, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 14024–14031.
- [177] Wiegmann, M., Völske, M., Stein, B., Potthast, M., 2022. Language models as context-sensitive word search engines, in: Proceedings of the First Workshop on Intelligent and Interactive Writing Assistants (In2Writing 2022), Association for Computational Linguistics, Dublin, Ireland. pp. 39–45. URL: <https://aclanthology.org/2022.in2writing-1.5>, doi:10.18653/v1/2022.in2writing-1.5.
- [178] Willard, B.T., Louf, R., 2023. Efficient guided generation for large language models. arXiv preprint arXiv:2307.09702 .
- [179] Wu, X., Li, S., Wu, H.T., Tao, Z., Fang, Y., 2025. Does RAG introduce unfairness in LLMs? evaluating fairness in retrieval-augmented generation systems, in: Rambow, O., Wanner, L., Apidianaki, M., Al-Khalifa, H., Eugenio, B.D., Schockaert, S. (Eds.), Proceedings of the 31st International Conference on Computational Linguistics, Association for Computational Linguistics, Abu Dhabi, UAE. pp. 10021–10036. URL: <https://aclanthology.org/2025.coling-main.669/>.
- [180] Xiong, M., Hu, Z., Lu, X., Li, Y., Fu, J., He, J., Hooi, B., 2024. Can LLMs Express Their Uncertainty? An Empirical Evaluation of Confidence Elicitation in LLMs, in: The Twelfth International Conference on Learning Representations. URL: <https://openreview.net/forum?id=gjeQKfXpZ>.
- [181] Yadav, V., Bethard, S., 2019. A survey on recent advances in named entity recognition from deep learning models. CoRR abs/1910.11470. URL: <http://arxiv.org/abs/1910.11470>, arXiv:1910.11470.
- [182] Ye, H., Zhang, N., Chen, H., Chen, H., 2022. Generative knowledge graph construction: A review. CoRR abs/2210.12714. URL: <https://doi.org/10.48550/arXiv.2210.12714>, doi:10.48550/arXiv.2210.12714, arXiv:2210.12714.
- [183] Ye, Q., Axmed, M., Pryzant, R., Khani, F., 2023a. Prompt engineering a prompt engineer. arXiv preprint arXiv:2311.05661 .
- [184] Ye, Q., Xu, H., Xu, G., Ye, J., Yan, M., Zhou, Y., Wang, J., Hu, A., Shi, P., Shi, Y., et al., 2023b. mplug-owl: Modularization empowers large language models with multimodality. arXiv preprint arXiv:2304.14178 .
- [185] Yin, D., Dong, L., Cheng, H., Liu, X., Chang, K.W., Wei, F., Gao, J., 2022. A Survey of Knowledge-Intensive NLP with Pre-Trained Language Models. arXiv preprint arXiv:2202.08772 .
- [186] Ying, R., Bourgeois, D., You, J., Zitnik, M., Leskovec, J., 2019. GNN explainer: A tool for post-hoc explanation of graph neural networks. CoRR abs/1903.03894. URL: <http://arxiv.org/abs/1903.03894>, arXiv:1903.03894.
- [187] Zamini, M., Reza, H., Rabiei, M., 2022. A review of knowledge graph completion. Information 13. URL: <https://www.mdpi.com/2078-2489/13/8/396>, doi:10.3390/info13080396.
- [188] Zaveri, A., Rula, A., Maurino, A., Pietrobon, R., Lehmann, J., Auer, S., 2016. Quality Assessment for Linked Data: A Survey. Semantic Web 7, 63–93.
- [189] Zeng, X., Li, Y., Zhai, Y., Zhang, Y., 2020. Counterfactual generator: A weakly-supervised method for named entity recognition, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), pp. 7270–7280.
- [190] Zhang, B., Alharbi, R., He, Y., 2025. Evaluation of Language Models in Knowledge Engineering (ELMKE) Workshop, in: 2nd Workshop on Evaluation of Language Models in Knowledge Engineering (ELMKE), co-located with ESWC-25, to appear. URL: <https://ceur-ws.org/Vol-3977/elmke-preface.pdf>.
- [191] Zhang, B., Carriero, V.A., Schreiberhuber, K., Tsaneva, S., González, L.S., Kim, J., de Berardinis, J., 2024. Ontochat: a framework for conversational ontology engineering using language models. arXiv preprint arXiv:2403.05921 .
- [192] Zhang, B., Meroño Peñuela, A., Simperl, E., 2023a. Towards Explainable Automatic Knowledge Graph Construction with Human-in-the-Loop, in: HHAI 2023: Augmenting Human Intellect. IOS Press, pp. 274–289.
- [193] Zhang, B., Reklós, I., Jain, N., Peñuela, A.M., Simperl, E., 2023b. Using Large Language Models for Knowledge Engineering (LLMKE): A Case Study on Wikidata. arXiv preprint arXiv:2309.08491 .
- [194] Zhang, W., Lin, H., Han, X., Sun, L., 2021. De-biasing distantly supervised named entity recognition via causal intervention, in: Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pp. 4803–4813.
- [195] Zhang, Y., Qi, P., Manning, C.D., 2018. Graph convolution over pruned dependency trees improves relation extraction, in: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 2205–2215.

- [196] Zhao, H., Chen, H., Yang, F., Liu, N., Deng, H., Cai, H., Wang, S., Yin, D., Du, M., 2024. Explainability for Large Language Models: A Survey. *ACM Trans. Intell. Syst. Technol.* 15. URL: <https://doi.org/10.1145/3639372>, doi:10.1145/3639372.
- [197] Zhao, J., Wang, T., Yatskar, M., Ordonez, V., Chang, K.W., 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. *arXiv preprint arXiv:1707.09457*.
- [198] Zhong, Z., Friedman, D., Chen, D., 2021. Factual probing is [mask]: Learning vs. learning to recall. *arXiv preprint arXiv:2104.05240*.
- [199] Zhu, S., Wang, W., Liu, Y., 2024. Quite good, but not enough: Nationality bias in large language models—a case study of chatgpt. *arXiv preprint arXiv:2405.06996*.

## A. Supplementary Materials for Interview Study

*Consent* Consent text (to read to interviewees):

With this research, we aim to investigate how knowledge engineers and practitioners use generative AI to develop and maintain KGs, ontologies, and knowledge bases. We will not collect any personal data. The responses to the questions are fully anonymous. The collected data will be shared with Otter AI, an automated transcription service hosted on an AWS server in the US. Your data will be processed under the terms of UK data protection law (including the UK General Data Protection Regulation (UK GDPR) and the Data Protection Act 2018). Transcripts will be kept until the publication of the research.

If you consent to participate in this study, please say loud and clear “Yes I consent”.

*Introduction* Give a bit of an intro of what are the main topics of the questions (e.g., your experience during Hackathon, challenges, skills, automation, evaluation, responsible AI). In the end we will ask you about anything important you want to mention, and you will have time to add comments. Mention that there is no wrong answer.

### Questions

1. What is your experience with KE and building KGs and ontologies?
2. How did you find the process of interacting with LLMS to solve KE tasks?
3. We found that during the Hackathon it was challenging for the groups to (i) find a dataset for testing their pipelines, (ii) find the right prompt questions for the LLMS, and (iii) evaluate the quality of each step toward building an ontology. Which one do you think is the most important?
4. What do you think were the most important skills you have to complete the tasks during the Hackathon?
  - (a) Do you feel you had those skills, or do you need to gain experience to complete several tasks?
5. Was there any automated process in the project you worked on during the Hackathon?
  - (a) What was the purpose of this automation?
6. Currently, to evaluate ontologies and KG constructions, we use metrics like precision and recall against gold standards, or human evaluation (semantic metric). Do you think this is enough?
  - (a) What they would look like new benchmarks and metrics?
7. Have you previously used safety testing for responsible AI? Safety testing can be Bias testing against discrimination, Ethical considerations, risk assessment, test for long-term effects.
  - (a) Do you think this can be used in the LLMS KE scenarios or we need to develop specific solutions?
8. Do you think there is anything else important for us to know related to LLMS and KE?

## B. Codebook for Thematic Analysis

See Table 6.

## C. List of Tools and Methodologies

See Table 7.

## D. Example of KG Cards

See Table 8.



**Table 6**

Codebook created by applying inductive and deductive thematic analysis.

First level code	Second level code	Third level code
Background Experience	KG construction KG explainability Machine Learning Ontologies Multidisciplinary background	
Bias	Bias as risk Bias mitigation awareness Difficulty of assessing bias Limited awareness Other safety issues Removing bias Types of bias	
Challenges	Dataset Evaluation  Prompting	KE tasks still demand manual evaluation Lack in KE evaluation techniques
Evaluation	Current evaluation techniques  New evaluation techniques	Evaluate the new knowledge produced by LLMs KE evaluation techniques are not sufficient LLMs output is hard to process
LLM use opinions	Promising  Skeptic	LLMs can support humans LLMs could improve KE tools  LLMs are not trusted LLMs cannot perform all KE tasks and still need human-in-the-loop
Skills	Skills had and helped  Skills need to have	Communication Knowledge or ontology engineering LLM Technical skills  LLM Ontology design or ontology engineering Technical skills

**Table 7**

KE and generative AI tools and methodologies with description and references.

(a) Knowledge Engineering	
Tools & Methodologies	Description
eXtremeDesign methodology [17]	A framework for pattern based ontology design
NeOn methodology [165]	A scenario-based methodology that supports the collaborative and dynamic aspects of ontology development, including in distributed environments
HermiT Reasoner [55]	A reasoner for ontologies that, given an OWL file, can determine if the ontology consistent, identify subsumption relationships between classes, and more
OOPS [137]	A web-based tool, independent of the ontology development environment, used for detecting modelling errors
OAEI [41]	Aims at comparing ontology matching systems on precisely defined test cases which can be based on ontologies of different levels of complexity and use different evaluation modalities
OntoClean [59]	A methodology for validating the ontological adequacy of taxonomic relationships based on formal, domain-independent properties of classes
(b) Generative AI	
Tools & Methodologies	Description
ChatGPT [122]	A language model developed by OpenAI
PaLM [10]	A transformer-based model developed by Google
Llama [168]	A family of transformer-based autoregressive causal language models developed by Meta
mPLUG-Owl [184]	A multimodal model that integrates images, text, and other information for comprehension and to generate responses
InstructBLIP [29]	A vision-language model leveraging instruction tuning to achieve state-of-the-art performance
Claude [11]	A language model developed by Anthropic with use cases such as summarisation, Q&A, and coding

**Table 8**

KG Cards for the CS-KG [34]. The CS-KG is accessed through <https://scholkg.kmi.open.ac.uk/>. The dumps include the updated versions of the KG. The methods used for the KG's development, the data, and some statistics are presented in the CS-KG scientific publication [34]. The publication describes the first version of the KG.

**(a) Knowledge**

Card	Description
Domain	Computer Science
Multilinguality	English
Timeliness	The CS-KG includes statements and entities extracted from scientific papers in the period 2010-2021. The team plans to keep maintaining and updating CS-KG in the following several years. The team created a fully automatic pipeline that we will run every six months to produce new versions of CS-KG that will include recent papers from OpenAlex [34].
Range	3.9M entities are classified as <i>Methods</i> , 1.3M as <i>Tasks</i> , 450M as <i>Materials</i> , 215K as <i>Metrics</i> , and 4M are associated with the <i>OtherEntity</i> .
Ontology Design Artifacts	The CS-KG ontology builds on top of SKOS ( <a href="https://www.w3.org/2004/02/skos/">https://www.w3.org/2004/02/skos/</a> ) and PROV-O ( <a href="https://www.w3.org/TR/prov-o/">https://www.w3.org/TR/prov-o/</a> ). The design of the object properties started from a set of 39 high level predicates (e.g., uses, analyzes, includes) produced by the knowledge graph generation pipeline (see Section 4.2). The team then associate specific domain and range constraints to them, which are used to drive and correct the automatic extraction process [34].

**(b) Provenance**

Card	Description
Provenance Artifacts	PROV-O. Each statement in CS-KG includes: <code>provo:wasDerivedFrom</code> , which provides provenance information and lists the MAG IDs (now OpenAlex IDs) of the articles from which the statement was extracted; <code>provo:wasGeneratedBy</code> , which provides provenance and versioning information of the tools used to detect the statement [34].
Dataset & Data Cards	Scientific articles selected by considering all papers from 2010 to 2019 with at least 1 citation (as of December 2021) and all the papers in 2020-2021 period from the set of articles from MAG [174] associated with the Field of Study "Computer Scienc". Since MAG has been decommissioned in 2021, the following versions will adopt OpenAlex ( <a href="https://openalex.org/">https://openalex.org/</a> ), which offers a comparable publication coverage [34].
Generative AI Usage	No information
Crowdsourcing	No use of crowdsourcing.

**(c) Graph Construction**

Card	Description
Data preprocessing	No information
Methodology	Based on the website ( <a href="https://scholkg.kmi.open.ac.uk/">https://scholkg.kmi.open.ac.uk/</a> ), the KG was generated by applying an automatic pipeline that extracts entities and relationships using four tools: DyGIE++, Stanford CoreNLP, the CSO Classifier, and a new PoS Tagger. CS-KG publication [34] presents the full pipeline. This pipeline was evaluated on a manually crafted gold standard yielding competitive results. It then integrates and filters the resulting triples using a combination of deep learning and semantic technologies in order to produce a high quality KG. The CS-KG publication [34] states that 1200 statements were selected for evaluation. The set was manually annotated by 3 senior computer science researchers. The Fleiss' kappa agreement [45] between the annotators was 0.435, indicating a moderate agreement. The majority vote schema was employed to generate the gold standard. In order to show the advantage of the hybrid method that builds on top of multiple tools, the authors compared the CS-KG pipeline against DyGIEpp, OpenIE, PoST, and against the union of their results (DyGIEpp + OpenIE + PoST). Table 2 in [34] reports the results of the evaluation in terms of precision, recall, and f-measure. The CS-KG pipeline outperforms all the other tools yielding an overall f-measure of 0.76.
Curation	No information
Schema	The schema is described in <a href="https://scholkg.kmi.open.ac.uk/cskg/documentation.php">https://scholkg.kmi.open.ac.uk/cskg/documentation.php</a> and <a href="https://scholkg.kmi.open.ac.uk/cskg/ontology">https://scholkg.kmi.open.ac.uk/cskg/ontology</a> .
Modalities	Text

**(d) Quality**

Card	Description
Accuracy	The CS-KG publication [34] in Sections 4.2 and 4.3 describe in detail the automatic methods used to handle the syntactic and semantic errors.
Coverage	67M statements, 24M entities, 219 relations ( <a href="https://scholkg.kmi.open.ac.uk/">https://scholkg.kmi.open.ac.uk/</a> )
Coherency	The CS-KG publication [34] in Section 4.4 describes in detail the automatic methods used to handle the error triples produced. There is no mention for constrain violations or other errors.
Completeness	No information
Succinctness	(i) SPARQL endpoint support ( <a href="https://scholkg.kmi.open.ac.uk/sparql/">https://scholkg.kmi.open.ac.uk/sparql/</a> ), (ii) human-readable labels are provided, (iii) URLs, and (iv) the contact information is available at <a href="https://scholkg.kmi.open.ac.uk">https://scholkg.kmi.open.ac.uk</a> .

**(e) Licensing**

Card	Description
License	CS-KG is licensed under a Creative Commons Attribution 4.0 International License.
Publisher	Team: Danilo Dessi, Francesco Osborne, Diego Reforgiato Recupero, Davide Buscaldi, Enrico Motta. This work is a collaboration of Scholarly Knowledge Modeling, Knowledge Media Institute, The Open University, University of Cagliari, FIZ Karlsruhe, Université Sorbonne Paris Nord.
Funding	No information
Policy & Regulation	CS-KG is aligned with the initiative of the Knowledge Graph Construction W3C Community Group for producing benchmarks, resources, and tools to support the semi-automatic generation of knowledge graphs from documents.

**(f) Transparency**

<b>Card</b>	<b>Description</b>
Data Balance	Statistics for the first version of CS-KG are provided in Section 5 of the publication [34].
Explainability	No information
Limitations	The main limitation of CS-KG is that it was produced with a fully automatic methodology, so the specific statements are not revised by humans, as in manually crafted KGs. The authors also working on developing an entity linking tool for automatically mapping documents (e.g., articles, patents, educational material) to entities and statements in CS-KG. Finally, they plan to further extend the ontology and the entity typing process, in particular by providing a more granular categorization of entity types.
Fairness Metrics	No information
Applications	CS-KG can support a variety of intelligent services, such as advanced literature search, document classification, article recommendation, trend forecasting, hypothesis generation, and many others [34].