

# XAI meets LLMs: A Survey of the Relation between Explainable AI and Large Language Models

Erik Cambria

*College of Computing and Data Science  
Nanyang Technological University, Singapore  
cambria@ntu.edu.sg*

Lorenzo Malandri

*Dept. of Statistics and Quantitative Methods  
University of Milano-Bicocca, Milan, Italy  
lorenzo.malandri@unimib.it*

Fabio Mercorio

*Dept. of Statistics and Quantitative Methods  
University of Milano-Bicocca, Milan, Italy  
fabio.mercorio@unimib.it*

Navid Nobani

*Dept. of Statistics and Quantitative Methods  
University of Milano-Bicocca, Milan, Italy  
navid.nobani@unimib.it*

Andrea Seveso

*Dept. of Statistics and Quantitative Methods  
University of Milano-Bicocca, Milan, Italy  
andrea.seveso@unimib.it*

**Abstract**—In this survey, we address the key challenges in Large Language Models (LLM) research, focusing on the importance of interpretability. Driven by increasing interest from AI and business sectors, we highlight the need for transparency in LLMs. We examine the dual paths in current LLM research and eXplainable Artificial Intelligence (XAI): enhancing performance through XAI and the emerging focus on model interpretability. Our paper advocates for a balanced approach that values interpretability equally with functional advancements. Recognizing the rapid development in LLM research, our survey includes both peer-reviewed and preprint (arXiv) papers, offering a comprehensive overview of XAI’s role in LLM research. We conclude by urging the research

community to advance both LLM and XAI fields together.

**Index Terms**—Explainable Artificial Intelligence, Interpretable Machine Learning, Large Language Models, Natural Language Processing

## I. INTRODUCTION

**T**HE emergence of LLMs has significantly impacted Artificial Intelligence (AI), given their excellence in several Natural Language Processing (NLP) applications. Their versatility reduces the need for handcrafted features, enabling applications

across various domains. Their heightened creativity in content generation and contextual understanding contributes to advancements in creative writing and conversational AI. Additionally, extensive pre-training on large amounts of data enables LLMs to exhibit strong generalisation capacities without further domain-specific data from the user [Zhao et al. \[2023a\]](#), [Amin et al. \[2023\]](#). For those reasons, LLMs are swiftly becoming mainstream tools, deeply integrated into many industry sectors, such as medicine (see, e.g., [Thirunavukarasu et al. \[2023\]](#)) and finance (see, e.g., [Wu et al. \[2023a\]](#)), to name a few.

However, their emergence also raises ethical concerns, necessitating ongoing efforts to address issues related to bias, misinformation, and responsible AI deployment. LLMs are a notoriously complex “black-box” system. Their inner workings are opaque, and their intricate complexity makes their interpretation challenging [Kaadoud et al. \[2021\]](#), [Cambria et al. \[2023a\]](#). Such opaqueness can lead to the production of inappropriate content or misleading outputs [Weidinger et al. \[2021\]](#). Finally, lacking visibility on their training data can further hinder trust and accountability in critical applications [Liu \[2023\]](#).

In this context, XAI is a crucial bridge between complex LLM-based systems and human understanding of their behaviour. Developing XAI frameworks for LLMs is essential for building user trust, ensuring accountability and fostering a responsible and ethical use of those models.

In this article, we review and categorise current XAI for LLMs in a structured manner. Emphasising the importance of clear and truthful explanations, as suggested by [Sevastjanova and El-Assady \[2022\]](#), this survey aims to guide future research towards enhancing LLMs’ explainability and trustworthiness in practical applications.

#### A. Contribution

The contribution of our work is threefold:

- 1) We introduce a novel categorisation framework for assessing the body of research concerning the explainability of LLMs. The framework provides

a clear and organised overview of the state of the art.

- 2) We conduct a comprehensive survey of peer-reviewed and preprint papers based on ArXiv and DBLP databases, going beyond using common research tools.
- 3) We critically assess current practices, identifying research gaps and issues and articulating potential future research trajectories.

#### B. Research questions

In this survey, we explore the coexistence of XAI methods with LLMs and how these two fields are merged. Specifically, our investigation revolves around these key questions:

- Q1 How are XAI techniques currently being integrated with LLMs?
- Q2 What are the emerging trends in converging LLMs with XAI methodologies?
- Q3 What are the gaps in the current related literature, and what areas require further research?

## II. THE NEED FOR EXPLANATIONS IN LLMs

In XAI field, the intersection with LLMs presents unique challenges and opportunities. This survey paper aims to dissect these challenges, extending the dialogue beyond the conventional understanding of XAI’s objective, which is to illuminate the inner mechanisms of opaque models for various stakeholders while avoiding the introduction of new uncertainties (See e.g., [Cambria et al. \[2023b\]](#), [Burkart and Huber \[2021\]](#)).

Despite their advancements, LLMs struggle with complexity and opacity, raising design, deployment and interpretation issues. Inspired by [Weidinger et al. \[2021\]](#), this paper categorises LLM challenges into user-visible and invisible ones.

*a) Visible User Challenges: Directly perceivable challenges for users without specialised tools.*

*b) Trust and Transparency:* Trust issues arise in crucial domains, e.g., healthcare [Mercorio et al. \[2020\]](#), [Gozzi et al. \[2022\]](#), [Alimonda et al. \[2022\]](#) or finance [Xing et al. \[2020\]](#), [Castellonovo et al. \[2023\]](#), [Yeo et al. \[2023\]](#), due to the opacity of black-box models, including LLMs.

XAI must offer transparent, ethically aligned explanations for wider acceptance, especially under stringent regulations that mandate explainability (e.g., EU’s GDPR [Novelli et al. \[2024\]](#)). This impacts regulatory compliance and public credibility, with examples in European skill intelligence projects requiring XAI for decision explanations [Malandri et al. \[2022a, 2024, 2022b,c\]](#).

*c) Misuse and Critical Thinking Impacts:* LLMs’ versatility risks misuse, such as content creation for harmful purposes and evading moderation [Shen et al. \[2023\]](#). Over-reliance on LLMs may also erode critical thinking and independent analysis, as seen in educational contexts (see, e.g. [Abd-Alrazaq et al. \[2023\]](#)).

*d) Invisible User Challenges: Challenges requiring deeper model understanding.*

*e) Ethical and Privacy Concerns:* Ethical dilemmas from LLM use, such as fairness and hate speech issues, and privacy risks like sensitive data exposure, require proactive measures and ethical guidelines [Weidinger et al. \[2021\]](#), [Yan et al. \[2023\]](#), [Salimi and Saheb \[2023\]](#).

*f) Inaccuracies and Hallucinations:* LLMs can generate false information, posing risks in various sectors like education, journalism, and healthcare. Addressing these issues involves improving LLM accuracy, educating users, and developing fact-checking systems [Rawte et al. \[2023\]](#), [Azaria and Mitchell \[2023\]](#).

### III. METHODOLOGY

Systematic Mapping Studies (SMSs) are comprehensive surveys that categorise and summarise a range of published works in a specific research area, identifying literature gaps, trends, and future research needs. They are especially useful in large or under-explored fields where a detailed Systematic Literature Review (SLR) may not be feasible.

SMS and SLR follow a three-phase method (planning, conducting, reporting) but differ in their approach, as SMSs address broader questions, cover a wider range of publications with a less detailed review, and aim to provide an overview of the research field. In contrast,

SLRs focus on specific questions, thoroughly review fewer publications, and strive for precise, evidence-based outcomes [Barn et al. \[2017\]](#).

Following [Martínez-Gárate et al. \[2023\]](#), we designed our SMS for XAI and LLMs, including peer-reviewed and preprint papers. The latter choice is because we believe in rapidly evolving fields like computer science, including preprints offering access to the latest research, essential for a comprehensive review [Oikonomidi et al. \[2020\]](#).

We followed these steps to structure our SMS: Section [I-B](#) proposes and defines the research questions, Section [III-A](#) describes how the paper retrieval has been performed; Section [III-B](#) describes the paper selection process based on the defined criteria; Section [III-C](#) explains who we dealt with false positive results and finally in Section [IV](#) we describe the obtained results.

#### A. Paper retrieval

*a) Overview:* Instead of utilising common scientific search engines such as Google Scholar, we employed a custom search methodology described in the following part. By scrutinising the titles and abstracts of the obtained papers, we conducted targeted searches using a predefined set of keywords pertinent to LLMs and XAI. This manual and deliberate search strategy was chosen to minimise the risk of overlooking relevant studies that automated search algorithms might miss and ensure our SMS dataset’s accuracy and relevance. Through this rigorous process, we constructed a well-defined corpus of literature poised for in-depth analysis and review. [Figure 1](#) provides an overview of this process.

*b) Peer-reviewed papers:* We initiated this step by identifying top-tier Q1 journals within the “Artificial Intelligence” category of 2022 (last year available at the start of the study), providing us with 58 journals from which to draw relevant publications.

Subsequently, we utilised the XML dump<sup>1</sup> from dblp computer science bibliography to get

<sup>1</sup><https://dblp.org/xml/dblp.xml.gz>

the titles of all papers published in the identified Q1 journals, except ten journals not covered by dblp. Once we gathered these paper titles, we proceeded to find their abstract. To do so, we initially used the last available citation network of AMiner<sup>2</sup> but given that this dump lacks the majority of 2023 publications, we leveraged Scopus API, a detailed database of scientific abstracts and citations, to retrieve the missing abstracts corresponding to the amassed titles.

c) *Pre-print papers*: We scraped all computer science papers presented in the Arxiv database from 2010 until October 2023, resulting in 548,711 papers. Consequently, we used the Arxiv API to get the abstracts of these papers.

### B. Paper selection

We employed a comprehensive set of keywords to filter the collected papers for relevance to LLMs and XAI. The search terms were carefully chosen to encompass the various terminologies and phrases commonly associated with each field.<sup>3</sup>

In our search, we applied a logical OR operator within the members of each list to capture any of the terms within a single category, and an AND operator was used between the two lists to ensure that only papers containing terms from both categories were retrieved for our analysis.

### C. Dealing with false positives

Upon completion of the initial retrieval phase, we identified a total of 1,030 manuscripts. Since

<sup>2</sup>[https://originalfileserver.aminer.cn/misc/dblp\\_v14.tar.gz](https://originalfileserver.aminer.cn/misc/dblp_v14.tar.gz)

<sup>3</sup>The keywords for XAI included: ['xai', 'explain', 'explanation', 'interpret', 'black box', 'black-box', 'blackbox', 'transparent model understanding', 'feature importance', 'accountable ai', 'ethical ai', 'trustworthy ai', 'fairness', 'ai justification', 'causal inference', 'ai audit']

While for LLMs, the keywords are: ['llm', 'large language model', 'gpt-3', 'gpt-2', 'gpt3', 'gpt2', 'bert', 'language model pre-training', 'fine-tuning language models', 'generative pre-trained transformer', 'llama', 'bard', 'roberta', 'T5', 'xl-net', 'megatron', 'electra', 'deberta', 'ernie', 'albert', 'bart', 'blenderbot', 'open pre-trained transformer', 'mt-nlg', 'turing-nlg', 'pegasus', 'gpt-3.5', 'gpt-4', 'gpt3.5', 'gpt4', 'cohere', 'claude', 'jurassic-1', 'openllama', 'falcon', 'dolly', 'mpt', 'guanaco', 'bloom', 'alpaca', 'openchatkit', 'gpt4all', 'flan-t5', 'orca']

some research keywords possess a broad meaning, for instance the words 'explain' and 'interpret' can be used in contexts different from the one of XAI, we retrieved few false positive papers, i.e., papers not dealing with both XAI and LLMs. We excluded the false positives—publications that address only XAI or LLMs independently or none of them. To do so, we manually analysed the title and abstract of each paper. This meticulous vetting process resulted in 233 papers relevant to XAI and LLMs.

Given that including all these papers in our survey was not feasible, we have selected the most relevant ones, based on their average number of citations per year. The whole research process resulted in 35 articles selected.

## IV. RETRIEVAL RESULTS

We divide papers into two macro-categories of *Application papers*, i.e., papers that somehow generated explanations, either towards explainability or to use them as a feature for another task, and *Discussion papers*, i.e., papers that do not engage with explanation generation but address an issue or research gap regarding the explainable LLM models.

### A. Application Papers

The first macro-category includes papers using LLMs in a methodology, tool, or task. Based on how LLMs are used, we further divide this category into two sub-categories as follows: "*To explain*", i.e., papers which try to explain how LLMs work and provide an insight into the opaque nature of these models. The second sub-category of papers called "*As feature*", uses the explanations and features generated by LLMs to improve the results of various tasks. The following parts discuss these sub-categories:

1) *To Explain*: Most papers, i.e., 17 out of 35, fit into this sub-category, with most addressing the need for more interpretable and transparent LLMs.

For instance, Vig [2019] introduces a visualisation tool for understanding the attention mechanism in Transformer models like BERT and

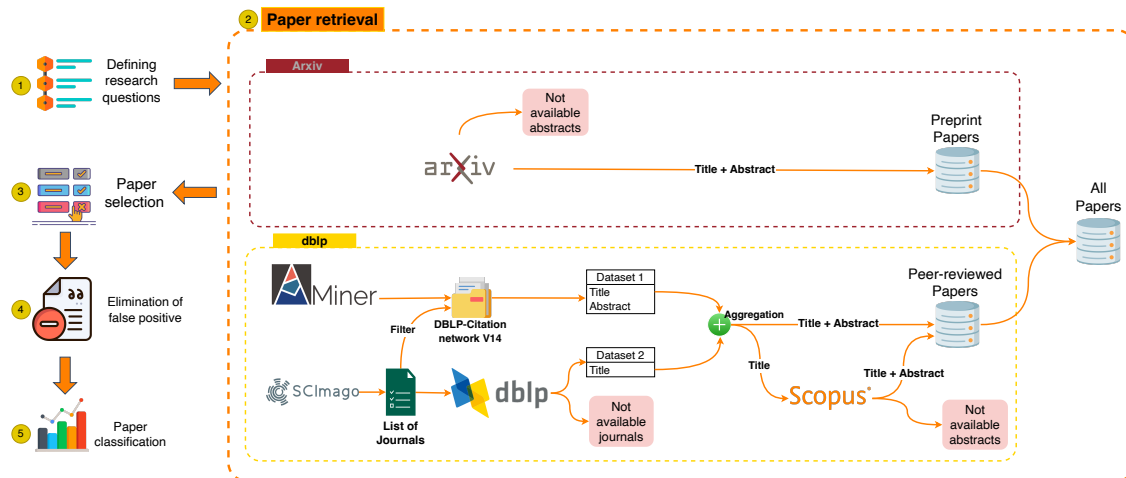


Fig. 1: The process used for getting the papers related to our keywords, including the definition of research questions, paper retrieval, paper selection, elimination of false positives and classifying papers in the pre-defined categories.

GPT-2. Their proposed tool provides insights at multiple scales, from individual neurons to whole model layers, helping to detect model bias, locate relevant attention heads, and link neurons to model behaviour.

Swamy et al. [2021] presents a methodology for interpreting the knowledge acquisition and linguistic skills of BERT-based language models by extracting knowledge graphs from these models at different stages of their training. Knowledge graphs are often used for explainable extrapolation reasoning Lin et al. [2023].

Wu et al. [2021] propose Polyjuice, a general-purpose counterfactual generator. This tool generates diverse, realistic counterfactuals by fine-tuning GPT-2 on multiple datasets, allowing for controlled perturbations regarding type and location.

Wang et al. [2022] investigates the mechanistic interpretability of GPT-2 small, particularly its ability to identify indirect objects in sentences. The study involves circuit analysis and reverse engineering of the model’s computational graph, identifying specific attention heads and their roles in this task.

Menon and Vondrick [2022] introduce a novel

approach for visual classification using descriptions generated by LLMs. This method, which they term “classification by description,” involves using LLMs like GPT-3 to generate descriptive features of visual categories. These features are then used to classify images more accurately while providing more transparent results than traditional methods that rely solely on category names.

Gao et al. [2023a] examines ChatGPT’s capabilities in causal reasoning using tasks like Event Causality Identification (ECI), Causal Discovery (CD), and Causal Explanation Generation (CEG). The authors claim that while ChatGPT is effective as a causal explainer, it struggles with causal reasoning and often exhibits causal hallucinations. The study also investigates the impact of In-Context Learning (ICL) and Chain-of-Thought (CoT) techniques, concluding that ChatGPT’s causal reasoning ability is highly sensitive to the structure and wording of prompts.

Pan et al. [2023] is a framework that aims to enhance LLMs with explicit, structured knowledge from KGs, addressing issues like hallucinations and lack of interpretability. The paper

Paper and Tool	Star	Fork	Update	Target	Agnostic	Goal				
Vig [2019] BertViz	6.1k	734	08/23	Transformers	✓	C	E	IMP	INT	R
Swamy et al. [2021] Experiments	19	2	05/22	BERT-based LM	✗	C	E	IMP	INT	R
Wu et al. [2021] Polyjuice	90	16	08/22	-	✓	C	E	IMP	INT	R
Wang et al. [2022] TransformerLens	48	161	01/23	GPT2-small	✗	C	E	IMP	INT	R
Menon and Vondrick [2022] -	-	-	-	Vision-LM	✓	C	E	IMP	INT	R
Gao et al. [2023a] Experiments	17	0	10/23	ChatGPT	✗	C	E	IMP	INT	R
Pan et al. [2023] -	-	-	-	LLMs	✓	C	E	IMP	INT	R
Conmy et al. [2023] ACDC	105	23	11/23	Transformers	✓	C	E	IMP	INT	R
He et al. [2022] RR	38	2	02/23	LLMs	✓	C	E	IMP	INT	R
Yoran et al. [2023] MCR	71	9	01/24	LLMs	✓	C	E	IMP	INT	R
Sarti et al. [2023] Inseq	250	26	01/24	SeqGen models	✓	C	E	IMP	INT	R
Wu et al. [2023b] Boundless DAS	0	17	01/24	LLMs	✓	C	E	IMP	INT	R
Li et al. [2023] XICL	1	3	11/23	LLMs	✓	C	E	IMP	INT	R
Chen et al. [2023] LMExplainer	-	-	-	LLMs	✓	C	E	IMP	INT	R
Gao et al. [2023b] Chat-REC	-	-	-	Rec. systems	✗	C	E	IMP	INT	R
Zhang et al. [2022] DSRLM	9	1	07/23	LLMs	✓	C	E	IMP	INT	R
Singh et al. [2023] SASC	61	14	01/24	LLMs	✓	C	E	IMP	INT	R
Li et al. [2022] -	-	-	-	LLMs	✓	C	E	IMP	INT	R
Ye and Durrett [2022] TextualExplnContext	11	2	02/23	LLMs	✓	C	E	IMP	INT	R
Turpin et al. [2023] Experiments	25	9	03/23	LLMs	✓	C	E	IMP	INT	R
Kang et al. [2023] AutoSD	-	-	-	Debugging models	✗	C	E	IMP	INT	R
Krishna et al. [2023] AMPLIFY	-	-	-	LLMs	✓	C	E	IMP	INT	R
Yang et al. [2023] Labo	51	4	12/23	CBM	✗	C	E	IMP	INT	R
Bitton-Guetta et al. [2023] WHOOPS!	-	-	-	LLMs	✓	C	E	IMP	INT	R
Shi et al. [2023] Chatgraph	2	0	07/23	LLMs	✓	C	E	IMP	INT	R

TABLE I: Synthesis of recent application papers, summarising engagement indicators as of January 2024, update timelines, model specificity, and the overarching aims of each study. In the first section of the table, *To Explain* papers are listed, and *As Feature* works in the second. Stars, forks, and last updates are not reported (-) for papers lacking associated repositories. Target is the specific focus of the study, such as a particular type of language model. Agnostic indicates whether the study is model-agnostic or not. The goal represents the primary objective of each study: comparison of models (C), explanation (E), improvement (IMP), interpretability (INT), and reasoning (R).

outlines three main approaches: KG-enhanced LLMs, LLM-augmented KGs, and synergised LLMs with KGs. This unification improves the performance and explainability of AI systems in various applications.

Conmy et al. [2023] focuses on automating a part of the mechanistic interpretability workflow in neural networks. Using algorithms like Automatic Circuit Discovery (ACDC), the authors automate the identification of sub-graphs in neural models that correspond to specific behaviours or functionalities.

He et al. [2022] presents a novel post-processing approach for LLMs that leverages

external knowledge to enhance the faithfulness of explanations and improve overall performance. This approach, called Rethinking with Retrieval, uses CoT prompting to generate reasoning paths refined with relevant external knowledge. The authors claim that their method significantly improves the performance of LLMs on complex reasoning tasks by producing more accurate and reliable explanations.

Multi-Chain Reasoning (MCR) introduced by Yoran et al. [2023] improves question-answering in LLMs by prompting them to meta-reason over multiple reasoning chains. This approach helps select relevant facts, mix informa-

tion from different chains, and generate better explanations for the answers. The paper demonstrates MCR’s superior performance over previous methods, especially in multi-hop question-answering.

Inseq [Sarti et al. \[2023\]](#) is a Python library that facilitates interpretability analyses of sequence generation models. The toolkit focuses on extracting model internals and feature importance scores, particularly for transformer architectures. It centralises access to various feature attribution methods, intuitively representable with visualisations such as heatmaps [Aminimehr et al. \[2023\]](#), promoting fair and reproducible evaluations of sequence generation models.

Boundless Distributed Alignment Search (Boundless DAS) introduced by [Wu et al. \[2023b\]](#) is a method for identifying interpretable causal structures in LLMs. In their paper, the authors demonstrate that the Alpaca model, a 7B parameter LLM, solves numerical reasoning problems by implementing simple algorithms with interpretable boolean variables.

[Li et al. \[2023\]](#) investigate how various demonstrations influence ICL in LLMs by exploring the impact of contrastive input-label demonstration pairs, including label flipping, input perturbation, and adding complementary explanations. The study employs saliency maps to qualitatively and quantitatively analyse how these demonstrations affect the predictions of LLMs.

LMExplainer [Chen et al. \[2023\]](#) is a method for interpreting the decision-making processes of LMs. This approach combines a knowledge graph and a graph attention neural network to explain the reasoning behind an LM’s predictions.

[Gao et al. \[2023b\]](#) propose a novel recommendation system framework, Chat-REC, which integrates LLMs for generating more interactive and explainable recommendations. The system converts user-profiles and interaction histories into prompts for LLMs, enhancing the recommendation process with the ICL capabilities of LLMs.

DSR-LM proposed by [Zhang et al. \[2022\]](#) is a framework combining differentiable symbolic

reasoning with pre-trained language models. The authors claim their framework improves logical reasoning in language models through a symbolic module that performs deductive reasoning, enhancing accuracy on deductive reasoning tasks.

2) *As Feature*: Papers in this sub-category do not directly aim to provide more transparent models or explain LLM-based models. Instead, they use LLMs to generate reasoning and descriptions, which are used as input to a secondary task.

For instance, [Li et al. \[2022\]](#) explore how LLMs’ explanations can enhance the reasoning capabilities of smaller language models (SLMs). They introduce a multi-task learning framework where SLMs are trained with explanations from LLMs, leading to improved performance in reasoning tasks.

[Ye and Durrett \[2022\]](#) evaluates the reliability of explanations generated by LLMs in few-shot learning scenarios. The authors claim that LLM explanations often do not significantly improve learning performance and can be factually unreliable by highlighting the potential misalignment between LLM reasoning and factual correctness in their explanations.

[Turpin et al. \[2023\]](#) investigates the reliability of CoT reasoning. The authors claim that while CoT can improve task performance, it can also systematically misrepresent the true reason behind a model’s prediction. They demonstrate this through experiments showing how biasing features in model inputs, such as reordering multiple-choice options, can heavily influence CoT explanations without being acknowledged in the explanation itself.

[Kang et al. \[2023\]](#) introduce an approach for automating the debugging process called Automated Scientific Debugging (AutoSD). This approach leverages LLMs to generate hypotheses about bugs in code and uses debuggers to interact with the buggy code. This approach leads to automated conclusions and patch generation and provides clear explanations for the debugging decisions, potentially leading to more efficient and accurate decisions by developers.

[Krishna et al. \[2023\]](#) present a framework

called Amplifying Model Performance by Leveraging In-Context Learning with Post Hoc Explanations (AMPLIFY), aiming to improve the performance of LLMs on complex reasoning and language understanding tasks by automating the generation of rationales. It leverages post hoc explanation methods, which output attribution scores indicating the influence of each input feature on model predictions, to construct natural language rationales. These rationales provide corrective signals to LLMs.

Yang et al. [2023] introduces Language Guided Bottlenecks (LaBo), a method for constructing high-performance Concept Bottleneck Models (CBMs) without manual specification of concepts. LaBo leverages GPT-3 to generate factual sentences about categories, forming candidate concepts for CBMs. These concepts are then aligned with images using CLIP Radford et al. [2021] to form a bottleneck layer. The method efficiently searches for bottlenecks using a submodular utility, focusing on discriminative and diverse information. The authors claim their method outperforms black box linear probes in few-shot classification tasks across 11 diverse datasets, showing comparable or better performance with more data.

Bitton-Guetta et al. [2023] introduces WHOOPS!, a new dataset and benchmark designed to test AI models' visual commonsense reasoning abilities. The dataset comprises images intentionally defying commonsense, created using image generation tools like Midjourney. The paper assesses AI models on tasks such as image captioning, cross-modal matching, visual question answering, and the challenging task of explanation generation, where models must identify and explain the unusualness of an image. Results show that even advanced models like GPT3 and BLIP2 struggle with these tasks, highlighting a gap in AI's visual commonsense reasoning compared to human performance.

## B. Discussion Papers

Unlike the Application papers, this category includes papers that target the argument of XAI through LLMs and vice versa but do not necessarily provide any specific methodology, framework or application. This category, in turn, is divided into two subcategories of **Issues**, or works which mention a *concern* and **Benchmark and Metrics**, which mainly focus on evaluation and assessment of XAI methods in LLM field.

1) *Issues*: Bowman [2023] critically examines LLMs, highlighting their unpredictability and the emergent nature of their capabilities with scaling. They underscore the challenges in steering and interpreting LLMs and the necessity for a nuanced understanding of their limitations and potential.

Liu et al. [2023] offers a survey and set of guidelines for assessing the alignment of LLMs with human values and intentions. They categorise and detail aspects of LLM trustworthiness, including reliability, safety, fairness, resistance to misuse, explainability, adherence to social norms, and robustness.

Liao and Vaughan [2023] emphasise the need for transparency in LLMs from a human-centred perspective. The authors discuss the unique challenges of achieving transparency with LLMs, differentiating them from smaller, more specialised models. The paper proposes a roadmap for research, emphasising the importance of understanding and addressing the transparency needs of diverse stakeholders in the LLM ecosystem. It advocates for developing and designing transparency approaches that consider these stakeholder needs, the novel applications of LLMs, and their various usage patterns and associated challenges.

Lastly, Xie et al. [2023] highlights the limitations of ChatGPT in explainability and stability in the context of financial market analysis through a zero-shot analysis. The authors suggest the need for more specialised training or fine-tuning.

2) *Benchmark and Metrics*: Lu et al. [2022] introduce SCIENCEQA, a new dataset for multimodal science question answering. This dataset



includes around 21k questions with diverse science topics and annotations, featuring lectures and explanations to aid in understanding the reasoning process. The authors demonstrate how language models, particularly LLMs, can be trained to generate these lectures and explanations as part of a CoT process, enhancing their reasoning capabilities. The study shows that CoT improves question-answering performance and provides insights into the potential of LLMs to mimic human-like multi-step reasoning in complex, multimodal domains.

[Golovneva et al. \[2022\]](#) introduce ROSCOE, a set of metrics designed to evaluate the step-by-step reasoning of language models, especially in scenarios without a golden reference. This work includes a taxonomy of reasoning errors and a comprehensive evaluation of ROSCOE against baseline metrics across various reasoning tasks. The authors demonstrate ROSCOE’s effectiveness in assessing semantic consistency, logicity, informativeness, fluency, and factuality in model-generated rationales.

[Zhao et al. \[2023b\]](#) presents a comprehensive survey on explainability techniques for LLMs, focusing on Transformer-based models. It categorises these techniques based on traditional fine-tuning and prompting paradigms, detailing methods for generating local and global explanations. The paper addresses the challenges and potential directions for future research in explainability, highlighting LLMs’ unique complexities and capabilities compared to conventional deep-learning models. Nevertheless, the survey mainly focuses on XAI in general and has minimal coverage of the relationship between XAI and LLMs.

## V. DISCUSSION

Our analysis indicates that a limited number of the reviewed publications directly tackle the challenges highlighted in Section II. For example, the work by [Liu et al. \[2023\]](#) focuses on trust-related concerns in LLMs, whereas [Gao et al. \[2023a\]](#) investigates the issue of misinformation propagation by LLMs. This scant attention to the

identified problems suggests an imperative for substantial engagement from the XAI community to confront these issues adequately.

*a) Open-Source Engagement:* Our survey study shows that more studies are moving beyond the traditional approach of merely describing methodologies in text. Instead, they release them as tangible tools or open-source code, frequently hosted on platforms such as GitHub. This evolution is a commendable step toward enhancing transparency and reproducibility in computer science research. The trend suggests a growing inclination among authors to release their code and publicly publish their tools, a notable change from a few years ago. However, we should also mention the inconsistency in the level of community engagement with these repositories. While some repositories attract substantial interest, fostering further development and improvement, others remain underutilised. This disparity in engagement raises important questions about the factors influencing community interaction with these resources.

*b) Target:* Predominantly, most works have directed their attention towards LLMs rather than concentrating on more specialised or narrower subjects within AI-based systems. This broad approach contrasts the relatively few studies that focus specifically on Transformers or are confined to examining particular categories of systems, such as recommendation systems. This overarching focus on LLMs represents a positive and impactful trend within the AI community. Given the rapid development and increasing prominence of LLM systems in academic and practical applications, this broader focus is timely and crucial for driving our understanding and capabilities in this domain forward. It ensures that research keeps pace with the advancements in the field, fostering a comprehensive and forward-looking approach essential for AI technologies’ continued growth and evolution.

*c) Goal:* Our analysis, as delineated in Table I, reveals a bifurcation in the objectives of the LLM studies under review. On the one hand, a subset of these works is primarily dedicated to

explaining and enhancing the interpretability of these 'black box' models. On the other hand, a larger contingent is more task-oriented, focusing on augmenting specific tasks and models, with interpretability emerging merely as a byproduct. This dichotomy in research focus underscores a pivotal trend: a pressing need to shift more attention towards demystifying the inner workings of LLMs. Rather than solely leveraging these models to boost task performance, their inherently opaque nature should not be overlooked. The pursuit of performance improvements must be balanced with efforts to unravel and clarify the underlying mechanisms of LLMs. This approach is crucial for fostering a deeper understanding of these complex systems, ensuring their application is effective and transparent. Such a balanced focus is essential for advancing the field technically and maintaining ethical and accountable AI development.

## VI. CONCLUSION

Our SMS reveals that only a handful of works are dedicated to developing explanation methods for LLM-based systems. This finding is particularly salient, considering the rapidly growing prominence of LLMs in various applications. Our study, therefore, serves a dual purpose in this context. Firstly, it acts as a navigational beacon for the XAI community, highlighting the fertile areas where efforts to create interpretable and transparent LLM-based systems can effectively address the challenges the broader AI community faces. Secondly, it is a call to action, urging researchers and practitioners to venture into this relatively underexplored domain. The need for explanation methods in LLM-based systems is not just a technical necessity but also a step towards responsible AI practice. By focusing on this area, the XAI community can contribute significantly to making AI systems more efficient, trustworthy and accountable.

Our **call for action** is as follows: Firstly, **researchers employing LLM models** must acknowledge and address the potential long-term

challenges posed by the opacity of these systems. The importance of explainability should be elevated from a mere 'nice-to-have' feature to an integral aspect of the development process. This involves a proactive approach to incorporate explainability in the design and implementation phases of LLM-based systems. Such a shift in perspective is essential to ensure that these models are effective, transparent and accountable. Secondly, we urge **researchers in the XAI field** to broaden their investigative scope. The focus should not only be on devising methodologies capable of handling the complexity of LLM-based systems but also on enhancing the presentation layer of these explanations. Currently, explanations provided are often too complex for non-technical stakeholders. Therefore, developing approaches that render these explanations more accessible and understandable to a wider audience is imperative. This dual approach will make LLMs more understandable and user-friendly and bridge the gap between technical efficiency and ethical responsibility in AI development.

## REFERENCES

- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. A survey of large language models. *arXiv:2303.18223*, 2023a.
- Mostafa Amin, Erik Cambria, and Björn Schuller. Can ChatGPT's responses boost traditional natural language processing? *IEEE Intelligent Systems*, 38(5):5–11, 2023.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. Large language models in medicine. *Nature medicine*, pages 1–11, 2023.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabrovolski, Mark Dredze, Sebastian Gehrmann, Prabhakaran Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance. *arXiv:2303.17564*, 2023a.
- Ikram Chraïbi Kaadoud, Lina Fahed, and Philippe Lenca. Explainable ai: a narrative review at

- the crossroad of knowledge discovery, knowledge representation and representation learning. In *MRC*, volume 2995, pages 28–40. ceur-ws. org, 2021.
- Erik Cambria, Rui Mao, Melvin Chen, Zhaoxia Wang, and Seng-Beng Ho. Seven pillars for the future of artificial intelligence. *IEEE Intelligent Systems*, 38(6):62–69, 2023a.
- Laura Weidinger, John Mellor, Maribeth Rauh, Conor Griffin, Jonathan Uesato, Po-Sen Huang, Myra Cheng, Mia Glaese, Borja Balle, Atoosa Kasirzadeh, et al. Ethical and social risks of harm from language models. *arXiv:2112.04359*, 2021.
- Yang Liu. The importance of human-labeled data in the era of llms. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pages 7026–7032, 2023.
- Rita Sevastjanova and Mennatallah El-Assady. Beware the rationalization trap! when language model explainability diverges from our mental models of language, 2022.
- Erik Cambria, Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, and Navid Nobani. A survey on xai and natural language explanations. *Information Processing & Management*, 60(1): 103111, 2023b.
- Nadia Burkart and Marco F Huber. A survey on the explainability of supervised machine learning. *Journal of Artificial Intelligence Research*, 70: 245–317, 2021.
- Fabio Mercorio, Mario Mezzanica, and Andrea Seveso. exdil: A tool for classifying and explaining hospital discharge letters. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 159–172. Springer, 2020.
- Noemi Gozzi, Lorenzo Malandri, Fabio Mercorio, and Alessandra Pedrocchi. Xai for myo-controlled prosthesis: Explaining emg data for hand gesture classification. *Knowledge-Based Systems*, 240:108053, 2022.
- Nicola Alimonda, Luca Guidotto, Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, and Giovanni Tosi. A survey on xai for cyber physical systems in medicine. In *2022 IEEE International Conference on Metrology for Extended Reality, Artificial Intelligence and Neural Engineering (MetroXRINE)*, pages 265–270. IEEE, 2022.
- Frank Xing, Lorenzo Malandri, Yue Zhang, and Erik Cambria. Financial sentiment analysis: an investigation into common mistakes and silver bullets. In *Proceedings of the 28th international conference on computational linguistics*, pages 978–987, 2020.
- Alessandro Castelnovo, Nicole Inverardi, Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, and Andrea Seveso. Leveraging group contrastive explanations for handling fairness. In *World Conference on Explainable Artificial Intelligence*, pages 332–345. Springer, 2023.
- Wei Jie Yeo, Wihan van der Heever, Rui Mao, Erik Cambria, Ranjan Satapathy, and Gianmarco Mengaldo. A comprehensive review on financial explainable ai. *arXiv preprint arXiv:2309.11960*, 2023.
- Claudio Novelli, Federico Casolari, Philipp Hacker, Giorgio Spedicato, and Luciano Floridi. Generative ai in eu law: Liability, privacy, intellectual property, and cybersecurity. *EU Law: Liability, Privacy, Intellectual Property, and Cybersecurity (January 14, 2024)*, 2024.
- Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, Navid Nobani, and Andrea Seveso. ContrXT: Generating contrastive explanations from any text classifier. *Inf. Fusion*, 81:103–115, 2022a. doi: 10.1016/j.inffus.2021.11.016. URL <https://doi.org/10.1016/j.inffus.2021.11.016>.
- Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, and Andrea Seveso. Model-contrastive explanations through symbolic reasoning. *Decision Support Systems*, 176:114040, 2024.
- Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, Navid Nobani, and Andrea Seveso. Contrastive explanations of text classifiers as a service. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: System Demonstrations*, pages 46–53, 2022b.
- Lorenzo Malandri, Fabio Mercorio, Mario Mezzanica, Navid Nobani, Andrea Seveso, et al. The good, the bad, and the explainer: a tool for

- contrastive explanations of text classifiers. In *IJCAI*, pages 5936–5939. AAAI Press, 2022c.
- Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. ”do anything now”: Characterizing and evaluating in-the-wild jailbreak prompts on large language models. *arXiv:2308.03825*, 2023.
- Alaa Abd-Alrazaq, Rawan AlSaad, Dari Alhuwail, Arfan Ahmed, Padraig Mark Healy, Syed Latifi, Sarah Aziz, Rafat Damseh, Sadam Alabed Alrazak, Javaid Sheikh, et al. Large language models in medical education: Opportunities, challenges, and future directions. *JMIR Medical Education*, 9(1):e48291, 2023.
- Lixiang Yan, Lele Sha, Linxuan Zhao, Yuheng Li, Roberto Martinez-Maldonado, Guanliang Chen, Xinyu Li, Yueqiao Jin, and Dragan Gašević. Practical and ethical challenges of large language models in education: A systematic scoping review. *British Journal of Educational Technology*, 2023.
- Ali Salimi and Hady Saheb. Large language models in ophthalmology scientific writing: Ethical considerations blurred lines or not at all? *American Journal of Ophthalmology*, 2023.
- Vipula Rawte, Amit Sheth, and Amitava Das. A survey of hallucination in large foundation models, 2023.
- Amos Azaria and Tom Mitchell. The internal state of an llm knows when its lying. *arXiv:2304.13734*, 2023.
- Balbir Barn, Souvik Barat, and Tony Clark. Conducting systematic literature reviews and systematic mapping studies. In *Innovations in Software Engineering Conference*, pages 212–213, 2017.
- Ángel Antonio Martínez-Gárate, José Alfonso Aguilar-Calderón, Carolina Tripp-Barba, and Aníbal Zaldívar-Colado. Model-driven approaches for conversational agents development: A systematic mapping study. *IEEE Access*, 2023.
- Theodora Oikonomidi, Isabelle Boutron, Olivier Pierre, Guillaume Cabanac, Philippe Ravaud, and Covid-19 Nma Consortium. Changes in evidence for studies assessing interventions for covid-19 reported in preprints: meta-research study. *BMC medicine*, 18:1–10, 2020.
- Jesse Vig. A multiscale visualization of attention in the transformer model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 37–42, 2019.
- Vinitra Swamy, Angelika Romanou, and Martin Jaggi. Interpreting language models through knowledge graph extraction. In *NeurIPS*, 2021.
- T Wu, M Tulio Ribeiro, J Heer, and D Weld. Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models. In *ACL-IJCNLP*, 2021.
- Kevin Ro Wang, Alexandre Variengien, Arthur Conmy, Buck Shlegeris, and Jacob Steinhardt. Interpretability in the wild: a circuit for indirect object identification in gpt-2 small. In *NeurIPS ML Safety Workshop*, 2022.
- Sachit Menon and Carl Vondrick. Visual classification via description from large language models. In *The Eleventh International Conference on Learning Representations*, 2022.
- Jinglong Gao, Xiao Ding, Bing Qin, and Ting Liu. Is chatgpt a good causal reasoner? a comprehensive evaluation. *arXiv:2305.07375*, 2023a.
- Shirui Pan, Linhao Luo, Yufei Wang, Chen Chen, Jiapu Wang, and Xindong Wu. Unifying large language models and knowledge graphs: A roadmap. *arXiv:2306.08302*, 2023.
- Arthur Conmy, Augustine N Mavor-Parker, Aengus Lynch, Stefan Heimersheim, and Adrià Garriga-Alonso. Towards automated circuit discovery for mechanistic interpretability. *arXiv:2304.14997*, 2023.
- Hangfeng He, Hongming Zhang, and Dan Roth. Rethinking with retrieval: Faithful large language model inference. *arXiv:2301.00303*, 2022.
- Ori Yoran, Tomer Wolfson, Ben Bogin, Uri Katz, Daniel Deutch, and Jonathan Berant. Answering questions by meta-reasoning over multiple chains of thought. *arXiv:2304.13007*, 2023.
- Gabriele Sarti, Nils Feldhus, Ludwig Sickert, and Oskar van der Wal. Inseq: An interpretability toolkit for sequence generation models. *arXiv:2302.13942*, 2023.
- Zhengxuan Wu, Atticus Geiger, Christopher Potts, and Noah D Goodman. Interpretability at

- scale: Identifying causal mechanisms in alpaca. *arXiv:2305.08809*, 2023b.
- Zongxia Li, Paiheng Xu, Fuxiao Liu, and Hyemi Song. Towards understanding in-context learning with contrastive demonstrations and saliency maps. *arXiv:2307.05052*, 2023.
- Zichen Chen, Ambuj K Singh, and Misha Sra. Lm-explainer: a knowledge-enhanced explainer for language models. *arXiv:2303.16537*, 2023.
- Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. Chatrec: Towards interactive and explainable llms-augmented recommender system. *arXiv:2303.14524*, 2023b.
- Hanlin Zhang, Ziyang Li, Jiani Huang, Mayur Naik, and Eric Xing. Improved logical reasoning of language models via differentiable symbolic programming. In *First Workshop on Pre-training: Perspectives, Pitfalls, and Paths Forward at ICML 2022*, 2022.
- Chandan Singh, Aliyah R Hsu, Richard Antonello, Shailee Jain, Alexander G Huth, Bin Yu, and Jianfeng Gao. Explaining black box text modules in natural language with language models. *arXiv:2305.09863*, 2023.
- Shiyang Li, Jianshu Chen, Yelong Shen, Zhiyu Chen, Xinlu Zhang, Zekun Li, Hong Wang, Jing Qian, Baolin Peng, Yi Mao, et al. Explanations from large language models make small reasoners better. *arXiv:2210.06726*, 2022.
- Xi Ye and Greg Durrett. The unreliability of explanations in few-shot prompting for textual reasoning. *NeurIPS*, 35:30378–30392, 2022.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. *arXiv:2305.04388*, 2023.
- Sungmin Kang, Bei Chen, Shin Yoo, and Jian-Guang Lou. Explainable automated debugging via large language model-driven scientific debugging. *arXiv:2304.02195*, 2023.
- Satyapriya Krishna, Jiaqi Ma, Dylan Slack, Asma Ghandeharioun, Sameer Singh, and Himabindu Lakkaraju. Post hoc explanations of language models can improve language models. *arXiv:2305.11426*, 2023.
- Yue Yang, Artemis Panagopoulou, Shenghao Zhou, Daniel Jin, Chris Callison-Burch, and Mark Yatskar. Language in a bottle: Language model guided concept bottlenecks for interpretable image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19187–19197, 2023.
- Nitzan Bitton-Guetta, Yonatan Bitton, Jack Hessel, Ludwig Schmidt, Yuval Elovici, Gabriel Stanovsky, and Roy Schwartz. Breaking common sense: Whoops! a vision-and-language benchmark of synthetic and compositional images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2616–2627, 2023.
- Yucheng Shi, Hehuan Ma, Wenliang Zhong, Gengchen Mai, Xiang Li, Tianming Liu, and Junzhou Huang. Chatgraph: Interpretable text classification by converting chatgpt knowledge to graphs. *arXiv:2305.03513*, 2023.
- Qika Lin, Jun Liu, Rui Mao, Fangzhi Xu, and Erik Cambria. Techs: Temporal logical graph networks for explainable extrapolation reasoning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1281–1293, 2023.
- Amirhossein Aminimehr, Pouya Khani, Amirali Molaei, Amirmohammad Kazemeini, and Erik Cambria. Tbxplain: A text-based explanation method for scene classification models with the statistical prediction correction. *arXiv preprint arXiv:2307.10003*, 2023.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- Samuel R Bowman. Eight things to know about large language models. *arXiv:2304.00612*, 2023.
- Yang Liu, Yuanshun Yao, Jean-Francois Ton, Xiaoying Zhang, Ruocheng Guo, Hao Cheng, Yegor Klochkov, Muhammad Faaiz Taufiq, and Hang Li. Trustworthy llms: a survey and guideline

- for evaluating large language models' alignment. In *Socially Responsible Language Modelling Research*, 2023.
- Q Vera Liao and Jennifer Wortman Vaughan. Ai transparency in the age of llms: A human-centered research roadmap. *arXiv:2306.01941*, 2023.
- Qianqian Xie, Weiguang Han, Yanzhao Lai, Min Peng, and Jimin Huang. The wall street neophyte: A zero-shot analysis of chatgpt over multimodal stock movement prediction challenges. *arXiv:2304.05351*, 2023.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *NeurIPS*, 35:2507–2521, 2022.
- Olga Golovneva, Moya Peng Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. Roscoe: A suite of metrics for scoring step-by-step reasoning. In *The Eleventh International Conference on Learning Representations*, 2022.
- Haiyan Zhao, Hanjie Chen, Fan Yang, Ninghao Liu, Huiqi Deng, Hengyi Cai, Shuaiqiang Wang, Dawei Yin, and Mengnan Du. Explainability for large language models: A survey. *ACM TIST*, 2023b.