

Review

Not peer-reviewed version

Explainability of Automated Fact Verification Systems: A Comprehensive Review

[Manju Vallayil](#)*, [Parma Nand](#)*, [Wei Qi Yan](#), [Héctor Allende-Cid](#)

Posted Date: 9 October 2023

doi: 10.20944/preprints202310.0366.v1

Keywords: automated fact verification; AFV; explainable artificial intelligence; XAI; explainable AFV




Preprints.org is a free multidiscipline platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This is an open access article distributed under the Creative Commons Attribution License which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Review

Explainability of Automated Fact Verification Systems: A Comprehensive Review

Manju Vallayil ^{1,*}, Parma Nand ^{1,*}, Wei Qi Yan ¹ and Héctor Allende-Cid ^{2,3}

¹ Auckland University of Technology, New Zealand; weiqi.yan@aut.ac.nz

² Pontificia Universidad Católica de Valparaíso, Chile; hector.allende@pucv.cl

³ Knowledge Discovery, Fraunhofer-Institute of Intelligent Analysis and Information Systems (IAIS), Sankt Augustin, Germany

* manju.vallayil.vijayalekshmi@aut.ac.nz (M.V.); parma.nand@aut.ac.nz (P.N.)

† Current address: Institute of Robotics and Vision (IoRV), WZ801A WZ Building, Level 8, 6-24 St. Paul St., Auckland, New Zealand.

Abstract: The rapid growth of Artificial Intelligence (AI) has led to considerable progress in Automated Fact Verification (AFV). This process involves collecting evidence for a statement, assessing its relevance, and predicting its accuracy. Recently, research has begun to explore automatic explanations as an integral part of the accuracy analysis process. However, the explainability within AFV is lagging compared to the wider field of explainable AI (XAI), which aims at making AI decisions more transparent. This study looks at the notion of explainability as a topic in the field of XAI, with a focus on how it applies to the specific task of Automatic Fact Verification. It examines the explainability of AFV, taking into account architectural, methodological, and dataset-related elements, with the aim of making AI more comprehensible and acceptable to the general society. Although there is a general consensus on the need for AI systems to be explainable, there is a dearth of systems and processes to achieve it. In this research, we investigate the concept of explainable AI in general and demonstrate its various aspects through the particular task of Automated Fact Verification. We explore the topic of faithfulness in the context of local and global explainability and how these correspond with architectural, methodological and data-based ways of achieving it. We examine these concepts for the specific case of AFV and analyze the current datasets used for AFV and how they can be adapted to further the identified aims on XAI. The paper concludes by highlighting the gaps and limitations in current data science practices and possible recommendations for modifications in architectural and data curation processes that would further the goals of XAI.

Keywords: automated fact verification; AFV; explainable artificial intelligence; XAI; explainable AFV

1. Introduction

Advances in Artificial Intelligence (AI), particularly transformer architecture [1], and the sustained success it has brought in transferring learning approaches in natural language processing, have led to advances in Automated Fact Verification (AFV). An AFV pipeline involves the subtasks for collecting evidence related to a claim, sorting most relevant evidence sentences, and predicting the veracity of the claim. Some systems such as [2] follow an additional step in the prelim to detect whether a claim is check-worthy or not before commencing on the other subtasks in the pipeline. Besides these subtasks, recent studies like [3] have started exploring how to generate automatic explanations as the reason for veracity prediction. However, not as much effort has been put into the explanation functionality of AFV compared to the strong progress made over the past few years both in fact checking technology and data sets [4]. The lack of focus on explanation is behind the growing interest in explainable AI

research [4]. Explainable AI¹ aims to provide the reasoning behind the decision (prediction) made, in contrast to the 'black box' impression² of machine learning where even the AI practitioners fail to explain the reason behind a particular decision made by an AI system they designed. Similarly, the goal of explainable AFV systems is to go beyond simple fact verification by generating interpretations that are grounded in facts and that communicate in a way that is easily understood and accepted by humans. While there is broad agreement in the research community on the importance of the explainability of AI systems [5–7], there is much less agreement on the current state of explainable AFV. The latest studies on the verification of facts [8–10] do not cohere around an aligned view on the subject. While researchers like [9] state that "Modern fact verification systems have distanced themselves from the black-box paradigm", [10] contradict this by stating, modern AFV systems estimate the truthfulness "using numerical scores which are not human-interpretable". We perceive the same impression as the latter on the basis of the literature review on the state-of-the-art AFV systems. Another of the most recent arguments supporting this view is [8]. They assert that, despite being a "nontrivial task", explainability of AFV is "mostly unexplored" and "needs to evolve" compared to the developments in explainable NLP.

This issue is further exacerbated by the fact that providing justifications for claim verdicts has always been a crucial aspect of human fact checking [4,11]. Therefore, it becomes evident that the transition from manual to automated fact checking falls short of achieving its intended human aspect to the functionality for 'Automated Fact Verification' unless there is a clear incorporation of explainability.

In this paper, we start by exploring the concepts of explainability in XAI in Section 2, followed by a specific focus on its implementation in AFV in Section 3. By defining explainability within the context of AFV and introducing the architectural, methodological, and dataset-based aspects for discussing interpretations, our aim is to support and inspire research and implementations that can initiate the process of bridging the current explainability gap in AFV. We emphasize the importance of datasets in achieving global explainability in AFV systems, suggesting that they should be a major focus of future research.

2. Explainable Artificial Intelligence

The field of interpretability in artificial intelligence is experiencing rapid growth, with numerous studies [7,12,13] exploring different facets of interpretation. These investigations are often conducted under the umbrella of Explainable AI (XAI), which encompasses various approaches and methodologies aimed at providing explanations for the decision-making process of black-box AI models, providing information on how they generate their outcomes. In this section, we review pertinent literature on XAI, with three main objectives: What is explainability? Why is it needed? and How can it be implemented?

The primary objective of XAI is to build models that humans can interpret effectively, especially in sensitive sectors like military, banking, and healthcare. These domains rely on the expertise of specialists to solve problems more efficiently, while also seeking meaningful outputs to understand and trust the solutions provided [14]. Additionally, it benefits both domain specialists and developers when appropriate outputs are available, as it encourages investigation into the system when discrepancies occur. However, many AI systems that support decision making have been developed as opaque structures that conceal their internal logic from the user, as identified by researchers [6,15]. The absence of an explanation for these black-box AI systems raises both practical and ethical concerns. Moreover, there exists an inherent tension between the performance of machine learning (such as predictive accuracy) and the explainability of the system [12]. For example, white-box models are

¹ (Also known as Interpretable AI or Explainable Machine Learning. Although used interchangeably the the subtle difference between explainability and interpretability, as is that the latter is not necessarily easily understood by those with little experience in the field unlike the former.

² https://en.wikipedia.org/wiki/Explainable_artificial_intelligence

deliberately designed to be interpretable, which makes their outputs easier to understand, but the drawback of these models is the compromise on accuracy. However, gray-box models strike a balance between interpretability and accuracy, offering a favorable trade-off [13]. However, black-box models, while more accurate, lack interpretability. Figure 1 demonstrates a comparison of these models and briefly conveys the idea of explainability in terms of AI systems. However, it is worth acknowledging that in certain scenarios, such as those involving structured data with naturally meaningful features, simpler classifiers, such as logistic regression or decision lists, may produce competitive results after appropriate preprocessing, as emphasized in the work by [16].

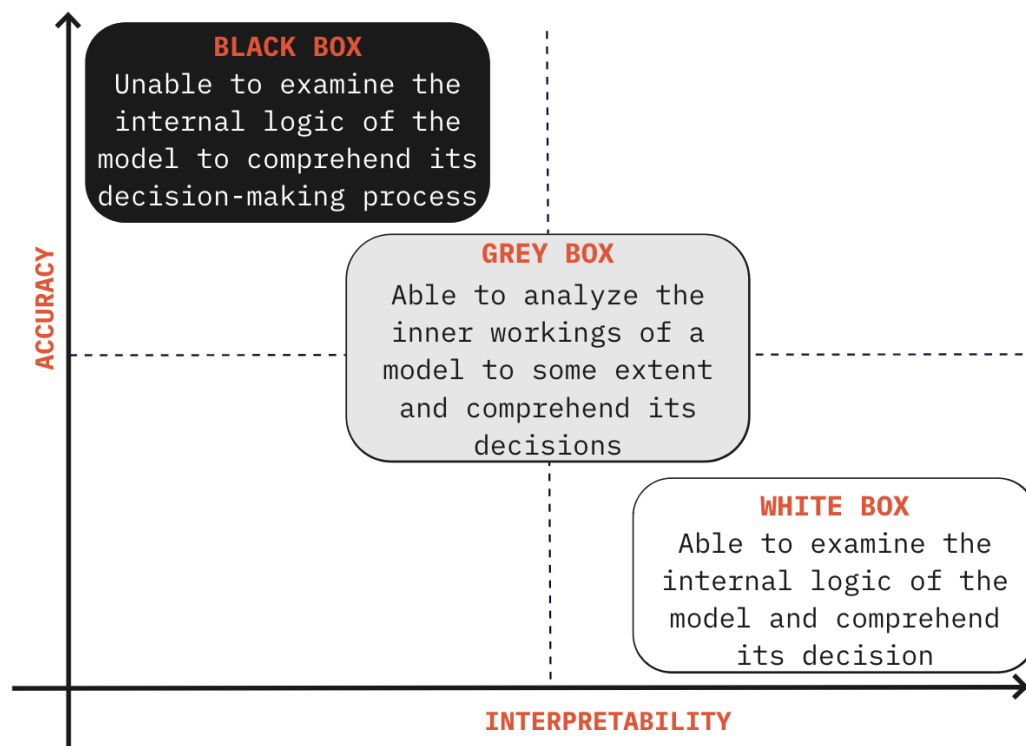


Figure 1. Comparative Analysis of Black-Box, Gray-Box, and White-Box Machine Learning Models, Highlighting an Apparent Trade-off between Accuracy (Y-Axis) and Interpretability (X-Axis).

Upon investigating the reasons for the increase in popularity of this research field, it is evident that XAI has received increasing attention from both academia and industry [4,5,7,13] attributing an inflection point in the middle of the last decade [12]. We provide a brief analysis of the factors contributing to this surge in research interest in XAI, in an attempt to resolute why Explainability is important and continues to be a pressing requirement in, AFV or AI in general. According to studies [13,17], as AI becomes more widely implemented, concerns about AI's black-box working paradigm have also become prevalent among governments and the general public. This has led to the need for regulatory authorities to act towards a push for some form of explainability. An initial step towards this AI regulation was taken by the European Parliament in 2016, when it adopted the General Data Protection Regulation (GDPR)³. With GDPR policy requiring citizens to receive explanations for algorithmic decisions, explainability has become a significant aspect of algorithm design thereafter [12,13,18]. Another authoritative supervision on XAI practices was by the R&D

³ Regulation (EU) 2016/679 of the European Parliament and of the Council, of 27 April 2016., uri=CELEX:32016R0679

Agency of the United States Department of Defense⁴, the Defense Advanced Research Projects Agency (DARPA) [19]. DARPA regulated, an XAI research program, and funded 11 research groups from the USA; they worked in the direction of a common conception; AI systems need to be more explainable to be better understood, trusted, and controlled. The main contribution of this DARPA XAI initiative is the creation of an XAI Toolkit, consolidating the diverse artifacts of the program (such as code, papers, reports, etc.) and the lessons learned from the 4-year program into a centralized publicly accessible repository⁵. This DARPA XAI program and the GDPR policy of the EU Parliament, along with the introduction of the EU AI Act (proposed European law on Artificial Intelligence)⁶ contributed majorly to the explainable AI movement we see today [13]. The drive for greater transparency and accountability in AI is not limited to the global stage; it is also reflected in national initiatives, including those spearheaded by the New Zealand government. For example, various initiatives and intergovernmental standards have been adopted by the New Zealand government, to address the transparency and accountability of AI algorithms. The G20 AI Principles, endorsed by Leaders in June 2019 and based on the OECD AI Policy Observatory, serve as a framework to promote responsible AI use. Similarly, the 'Algorithm Charter for Aotearoa New Zealand' aligns with the OECD principles and aims to improve transparency and accountability in AI algorithm usage. The country also actively contributes to global efforts through the OECD portal⁷, which showcases AI policy initiatives worldwide and emphasizes the importance of AI transparency. The AI Forum NZ has released its own set of principles, including a focus on AI transparency, while other measures such as the New Zealand Pilot Project from the World Economic Forum⁸ further support the objective of improving AI transparency and explainability.

2.1. XAI implementation: Objectives and Approaches

2.1.1. Objectives

The current body of work on XAI [6,13,14] highlights several essential conditions that must be considered when implementing the XAI model. Firstly, 'interpretability', which refers to the degree to which the model and its predictions can be understood by humans. The complexity of a predictive model (often measured in its size) is widely considered as a component for measuring interpretability. Another crucial condition is 'accuracy', which denotes the extent to which a model can accurately forecast outcomes for unseen instances. Evaluating the accuracy of a model involves using metrics such as the accuracy score, the F1-score, and other relevant evaluation measures. Finally, 'fidelity', which refers to the extent to which an interpretable model can accurately replicate the behavior of a corresponding black-box system. Fidelity measures how well the interpretable model imitates the performance of the black box. Similar to accuracy, fidelity is evaluated using metrics such as the accuracy score, F1-score, and other relevant measures, but with respect to the outcomes produced by the black box.

These three objectives serve as guidelines to develop an interpretable model; in the following sub section we review the standard approaches suggested in literature in order to achieve these objectives.

2.1.2. Approaches

Based on our review of the relevant literature, including studies by various researchers [6,13,17,20] the practical approaches to achieving the aforementioned XAI objectives are generally categorized into model explainability and post hoc explainability; the latter is further appropriated into interpretations

⁴ <https://www.darpa.mil/program/explainable-artificial-intelligence>

⁵ <https://xaitk.org/>

⁶ <https://artificialintelligenceact.eu/>

⁷ <https://oecd.ai/en/dashboards/overview>

⁸ https://www3.weforum.org/docs/WEF_Reimagining_Regulation_Age_AI_2020.pdf

at the prediction level and the data set [20]. It is important to note that the model explainability is interchangeably referred to as interpretability in certain works, such as [21].

Model-based explainability methods involve the creation of simple and transparent AI models that can be easily understood and interpreted. Such methods are particularly useful when the underlying data relationships are not highly complex, allowing simple models to effectively capture data patterns [15]. Models like decision trees and linear regression (unlike Deep Neural Networks), inherently possess model explainability [12]. When we opt for models that are inherently interpretable, they provide their own faithful explanations, faithfully representing what the model actually computes [16]. However, when dealing with data that exhibit higher degrees of complexity or non-linearity, more intricate black-box models are designed and implemented. In such cases, post hoc explainability techniques are used to extract information about the relationships learned by the model [13,15,22].

A post hoc explainability method operates on a trained and/or tested AI model, generating approximations of the model’s internal workings and decision logic [15,20]. Post hoc methods aim to reveal relationships between feature values and the model’s predictions, without requiring access to its internal mechanisms. This enables users to identify the most crucial features in a machine learning task, quantify their importance, reproduce decisions made by the black-box model, and uncover potential biases in the model or the underlying data. The prediction-level interpretation methods revolve around explaining the rationale behind the individual predictions made by the models. These methods delve into identifying the specific features and interactions that contributed to a particular prediction, and hence also named local interpretations [14] or local explanations [21] or local fidelity [23]. On the other hand, the approaches at the data set level focus on comprehending the broader associations and patterns that the model has learned, with the aim of discerning the patterns related to the predicted responses on a global scale [14,20].

2.2. XAI Taxonomy

Taxonomy in XAI provides a structured framework to categorize and organize methods, techniques, and approaches for explaining AI and machine learning models, facilitating systematic discussions of model explainability and interpretability. Figure 2 provides a hierarchical visualization of XAI by offering a structured view of the XAI approaches mentioned in Section 2.1.2.

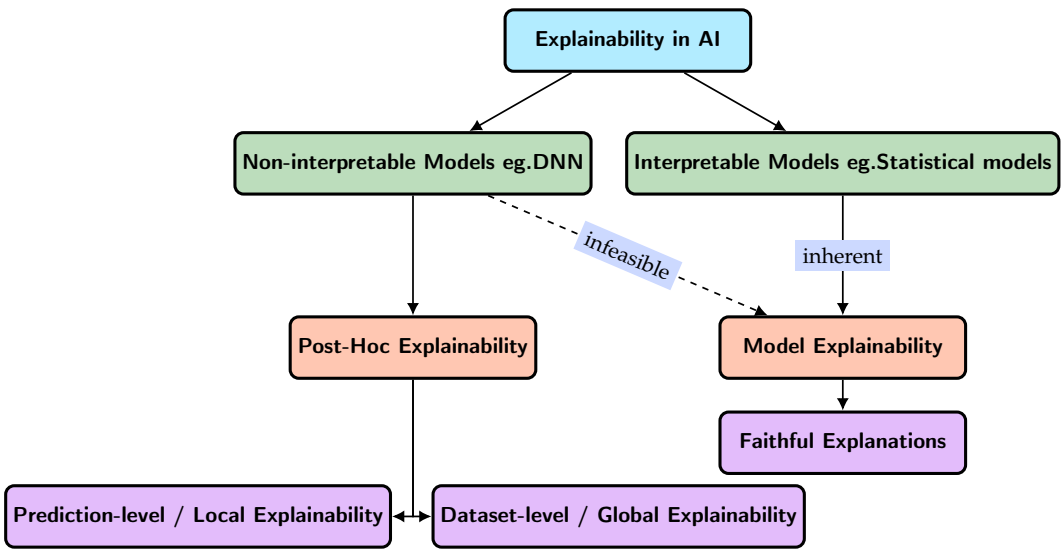


Figure 2. Hierarchical overview of approaches to achieving explainability in AI, focusing on black-box and white-box strategies. The arrows labeled ‘inherent’ and ‘infeasible’ respectively signify the natural transparency of white-box models such as Statistical models, and the difficulties in attaining model explainability in black-box models such as DNNs.

On local explainability, it is worth noting that while an explanation may not achieve complete faithfulness unless it provides a full model description, it is imperative for it to at least achieve local faithfulness. This means that the explanation must accurately correspond to the model's behavior in (at least) the proximity of the instance being predicted, ensuring its meaningfulness. However, it's crucial to highlight that while global fidelity (globally faithful explanations) would encompass local fidelity, local fidelity does not imply global fidelity. Features that have global importance may not necessarily be significant in the local context, and hence the search for globally faithful, but interpretable explanations remains a challenging endeavor, especially when dealing with complex models [23].

In Section 2, we have presented an overview of explainability concepts, their significance, and associated terminology. The following section will delve into their application within the context of AFV models.

3. Explainable AFV

Despite notable progress in the development of explainable AI techniques, achieving comprehensive global explainability in AFV models remains a challenging task. However, this issue encompasses multiple aspects that pose significant obstacles to research in the field of explainable AFV. First, only a relatively small number of automated fact-checking systems include explainability components. Second, Explainable AFV systems currently do not possess the capability of global explainability. Finally, the existing datasets for AFV suffer from a lack of explanations. We address these factors through three main perspectives: architectural, methodological, and data-based, and we examine them along the objectives and approaches of explainability discussed in previous section. The emphasis is placed on the data perspective (Section 3.3) due to its crucial role in achieving global explainability within the context of AFV. The significance of this perspective is derived from the training dataset's influence on AI model behavior and the critical necessity to achieve a high level of data explainability in AFV.

3.1. Architectural Perspective

Majority of AFV systems broadly adopt a three-stage pipeline architecture similar to the Fact Extraction and VERification (FEVER) shared task [24], as identified and commented on by many researchers [24–30]. These three stages (also called sub-tasks) are, document retrieval (evidence retrieval), sentence selection (evidence selection) and Recognizing Textual Entailment or RTE (label/veracity prediction). The document retrieval component is responsible for gathering relevant documents from a knowledge base, such as Wikipedia, based on a given query. The sentence-retrieval component then selects the most pertinent evidence sentences from the retrieved documents. Lastly, the RTE component predicts the entailment relationship between the query and the retrieved evidence. While the above framework is generally followed in AFV, alternative approaches incorporate additional distinct components to identify credible claims and provide justifications for label predictions, as shown in Figure 3. The inclusion of a justification component in such alternative approaches contributes to the system's capacity for explainability within the AFV paradigm.

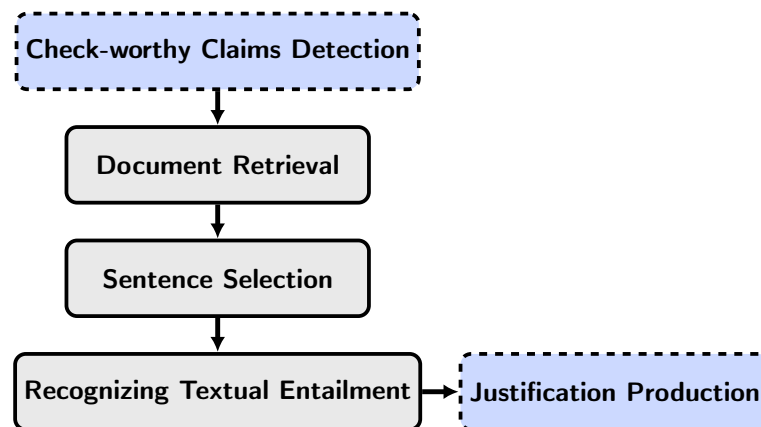


Figure 3. The figure illustrates the main stages involved in Automated Fact Verification: Document Retrieval, Sentence Selection, and Recognition of Textual Entailment. Optional components are also shown to identify verified claims and provide justifications. The inclusion of a Justification component specifically contributes to enhancing the system’s capability for explainability within the AFV framework.

The majority of AFV systems are highly dependent on deep neural networks (DNNs) for the label prediction task [4]. Furthermore, in recent years, deep learning-based approaches have demonstrated exceptional performance in detecting fake news [31]. As we mentioned in Section 2, there is, however, an inherent conflict between the performance of AI models and their ability to explain how they make decisions. However, although existing AFV systems lack inherent explainability [4], it would be foolish to overlook the potential to use these less interpretable deep models for AFV, as these models possess the ability to achieve state-of-the-art results with a remarkable level of prediction accuracy. However, this also indicates that model-based interpretation approaches may not be a suitable solution for AFV systems. The reason being that these methods require the involvement of simple and transparent AI models that can be easily understood and interpreted.

Therefore, considering the architectural characteristics of state-of-the-art AFV systems, a potential trade-off solution to achieve explainability may involve incorporating post hoc measures of explainability, either at the prediction-level or dataset-level, while still leveraging the capabilities of less interpretable deep transformer models. The subsequent subsections delve into the attempts made in the literature to incorporate post-hoc explainability in terms of methods and input within the context of AFV.

3.2. Methodological Perspective

The methodological aspect looks at the different approaches utilized in existing literature to develop explainable AFV systems.

3.2.1. Summarization Approach

In AFV, extractive and abstractive explanations serve as two types of summarization methodologies, providing a summary along with the predicted label as a form of justification or explanation. Extractive explanations involve directly extracting relevant information or components from the input data that contribute to the prediction or fact-checking outcome. These explanations typically rely on the emphasis of specific words, phrases, or evidence within the input. On the other hand, abstractive explanations involve generating novel explanations that may not be explicitly present in the input data. These explanations focus on capturing the essence or key points of the prediction or fact-checking decision by generating new text that conveys the rationale or reasoning behind the outcome. It is important to note that terminology can vary across fields. For instance, in the Explainable Natural Language Processing (Explainable NLP) literature, [32] refers to extractive explanations as ‘Highlights’ and abstractive explanations as ‘Free-text explanations’.

The approaches to explainability employed by existing explainable AFV systems are primarily extractive. For example, the work of [21] presents the first investigation on the generation of explanations automatically based on the available claim context and utilized the transformer model architecture for extraction summarization purposes. Two models are trained with the intention of addressing this issue. One model focuses on generating post hoc explanations, where the predictive and explanation models are trained independently; while the other model is trained jointly to handle both tasks simultaneously. The model that trains the explainer separately tends to slightly outperform the model trained jointly. [33] also approach the task of explanation generation as a form of summarization. However, their methodology differs from that of [21]. Specifically, the explanation models of [33] are fine-tuned for extractive and abstractive summarization, with the aim of generating novel explanations that go beyond mere extractive summaries. By training the models on a combination of extractive and abstractive summarization tasks, they enabled the models to generate more comprehensive and insightful explanations by leveraging both existing information in the input and generating new text to convey the reasoning behind the fact-checking outcomes.

A potential concern is that these models (both extractive and abstractive) may generate explanations that, while plausible in relation to the decision, do not accurately reflect the actual veracity prediction process. This issue is particularly problematic in the case of abstractive models, as they can generate misleading justifications due to the possibility of hallucinations [11].

3.2.2. Logic-based Approach

In logic-based explainability, the focus is on capturing the logical relationships and dependencies between various pieces of information involved in fact verification. This includes representing knowledge in the form of logical axioms, rules, and constraints to provide justifications for the verification results. [3] and [30] are examples of recent studies that focus on the explainability of fact verification using logic-based approaches.

[3] propose a logic-regularized reasoning framework, LOREN, for fact verification. By incorporating logical rules and constraints, LOREN ensures that the reasoning process adheres to logical principles, improving the transparency and interpretability of the fact-verification system. The experimental results demonstrate the effectiveness of LOREN in achieving an explainable fact verification. Similarly, [30] highlights the potential of natural logic theorem proving, as a promising approach for explainable fact verification systems. The system named ProoFVer applies logical inference rules to derive conclusions based on given premises, providing transparent explanations for the verification process. The experimental evaluation shows the efficacy of ProoFVer in accurately verifying factual claims while also offering interpretable justifications through the logical reasoning steps.

It is important to acknowledge certain limitations and drawbacks associated with this logic-based approach. First, the complexity and computational cost of logic-based reasoning can limit its scalability and practical applicability in real-world fact verification scenarios. Furthermore, while logic provides a structured and interpretable framework for reasoning, it may not capture all the nuances and complexities of natural language and real-world information. That means, the effectiveness of these approaches heavily relies on the adequacy and comprehensiveness of the predefined logical rules, which may not cover all possible scenarios and domains. Lastly, interpretability of the generated explanations may still be challenging for nonexpert users. They may involve complex logical steps that require expertise to fully understand and interpret.

3.2.3. Attention-based Approach

Different from the summarization and the logic-based techniques, explainable AFV systems such as [22,34] use visualizations to illustrate important features or evidence utilized by AFV models for predictions. This provides users with a means to understand the relationships that influence the decision-making process. For example, the AFV model proposed by [34] introduces an attention

mechanism that directs the focus towards the salient words in the article in relation to the claim. This enables the generation of the most significant words in the article as evidence (words with more weights are highlighted with darker shades in the verdict) and [34] claim that this strategy enhances the transparency and interpretability of the model. The explanation module of the fact checking framework by [22] also utilizes the attention mechanism to generate explanations for the model's predictions, highlighting the important features and evidence used for classification.

However, [11] illustrated several critical concerns associated with the reliability of attention as an explanatory method, citing pertinent studies [35–37] to reinforce the argument. The authors point out that the removal of tokens assigned high attention scores does not invariably affect the model's predictions, illustrating that some tokens, despite their high scores, may not be pivotal. On the contrary, certain tokens with lower scores have been found to be crucial for accurate model predictions. These observations collectively indicate a possible 'fidelity' issue in the explanations yielded by attention mechanisms questioning the reliability and interpretability of attention mechanisms in models. Furthermore, [11] argue that the complexity of these attention-based explanations can pose substantial challenges for people lacking an in-depth understanding of the model architecture, compromising readability and overall comprehension. This scrutiny of the limitations inherent to attention-based explainability methods highlights the pressing need to reevaluate their applicability and reliability within the realm of AFV.

3.2.4. Counterfactual Approach

Counterfactual explanations, also known as inverse classification, describe minimal changes to input variables that would lead to an opposite prediction, offering the potential for recourse in decision-making processes [16]. These explanations allow users to understand what modifications are needed to reverse a prediction made by a model. In the context of AFV, counterfactual explanations have been explored. The study by [38], for example, explicitly focuses on the interpretability aspect of counterfactual explanations, in order to help users understand why a specific piece of news was identified as fake. The comprehensive method introduced in that work involves question-answering and entailment reasoning to generate counterfactual explanations, which could enhance users' understanding of model predictions in AFV. In a recent study [39] exploring debiasing for fact verification, the researchers propose a method called CLEVER that operates from a counterfactual perspective to mitigate biases in predicting the veracity. CLEVER stands out by training separate models for claim-evidence fusion and claim-only prediction, allowing the unbiased aspects of predictions to be highlighted. This method could be explored further in the context of explainability in AFV, as it allows users to discern the factors that lead to specific predictions, even if the main emphasis of the cited work was on bias mitigation.

Nevertheless, counterfactual explanations in AFV, while providing valuable insights into why a model makes specific predictions, may also confront challenges in their practical application. One notable limitation lies in the potential complexity and difficulty of interpreting minimal changes in input variables, especially in cases involving complex facts and evidence. This complexity could pose challenges to users in grasping the precise factors that influence the predictions of the models, which is a key aspect in achieving a broader interpretability in AI systems, as discussed in Section 2.

Table 1 categorizes the existing approaches to develop explainable AFV systems into four methodological aspects discussed: Summarization, logic-based, attention-based, and counterfactual. Each category is illustrated with examples of studies that employ these methods, highlighting their unique contributions as well as potential drawbacks. The table serves as a comprehensive overview, aiding in understanding the various techniques used to enhance the explainability and interpretability in state-of-the-art AFV systems.

Additionally, it is important to note the inherent complexity in typical DNN-based AFV systems. When considered alongside the objectives of XAI outlined in Section 2.1.1, which emphasize that the interpretability of a predictive model is often assessed through its complexity (commonly measured by

its size), this factor adds another layer of complexity to the already challenging task of achieving model explainability in state-of-the-art AFVs. However, the situation is equally challenging when it comes to post hoc explainability, especially in terms of achieving global explainability. None of the explainable AFV systems we discussed provides global explainability; they mainly focus on prediction-level or local explainability by explaining the model’s decision-making process for specific instances or cases. On the other hand, global interpretability at the data set level aims to uncover more general relationships learned by the model and provides a greater understanding of how the model learns and generalizes across different examples [20]. The following section explores the extent of dataset-level explainability in AFV, leading to an examination of its current state.

Table 1. Comparative Analysis of Diverse Methodologies Employed in Enhancing Explainability in Automated Fact Verification Systems.

Methodological Aspect	Examples	Drawbacks
Summarization Approach	[21]Utilizes the transformer model for extractive summarization. Two models trained separately and jointly. [33]Fine-tuned for both extractive and abstractive summarization.	May generate misleading or inaccurate explanations. Particularly problematic for abstractive models.
Logic-based Approach	[3]LOREN framework uses logical rules for transparency. [30]ProofVer uses natural logic theorem proving.	Complexity and computational cost limit scalability. May not capture all nuances of natural language.
Attention-based Approach	[34] Uses attention mechanism to focus on salient words. [22]Utilizes attention for feature and evidence highlighting.	Relies on human experts for visualizations, diverging from the principles of XAI.
Counterfactual Approach	[38]Focuses on interpretability by generating counterfactual explanations in AFV through question-answering and entailment reasoning. [39]Proposes the CLEVER method, which operates from a counterfactual perspective to mitigate biases in veracity prediction within AFV.	May face complexity in interpreting minimal input changes, particularly in intricate factual claims and evidence scenarios, potentially hindering broader interpretability.

3.3. Data Perspective

The potential of data explainability lies in its ability to provide deep insights that enhance the explainability of AI systems (which rely heavily on data for knowledge acquisition) [6,13]. Data explainability methods encompass a collection of techniques aimed at better comprehending the data sets used in the training and design of AI models [13]. The importance of a training data set in shaping the behavior of AI models highlights the need to achieve a high level of data explainability. Therefore, it is crucial to note that constructing a high-performing and explainable model requires a high-quality training dataset. In AFV, the nature of this dataset, also known as the source of evidence, has evolved over time. Initially, the evidence was primarily based on claims, where information directly related to the claim was used for verification. Subsequently, knowledge-base-based approaches were introduced, utilizing structured knowledge sources to support the verification process. Further advances led to the adoption of text-based evidence, where relevant textual sources were used for verification. In recent developments, there has been a shift towards dynamically retrieved sentences, where the system dynamically retrieves and selects sentences that are most relevant to the claim for verification purposes. We will explore these changes through the lens of explainability.

Systems such as [40] that process the claim itself, using no other source of information as evidence, can be termed as ‘knowledge-free’ or ‘retrieval-free’ systems. In these systems, the linguistic characteristics of the claim are considered the deciding factor. For example, claims that contain

a misleading phrase are labeled 'Mostly False'. [41] also employ a similar approach, focusing on linguistic patterns, but incorporate a hybrid methodology by including claim-related metadata with the input text to the deep learning model. These additional data include information such as the claim reporter's profile and the media source where the claim is published. These knowledge-free systems face limitations in their performance, as they depend only on the information inherent in the claim and do not consider the current state of affairs [42]. The absence of contextual understanding and the inability to incorporate external information make dataset-level explainability infeasible in these systems.

In knowledge-base-based fact-verification systems [43–45], a claim is verified against the RDF triples present in a knowledge graph. The veracity of the claim is calculated by assessing the error between the claim and the triples based on different approaches such as rule-based, subgraph-based, or an embedding-based one. The drawback of such systems is the likelihood of a claim being verified as false, based on the assumption that the supporting facts of a true claim are already present in the graph, which is not always feasible. This limited scalability and the inability to capture nuanced information hinder the achievement of explainability in these type of fact verification models.

Unlike the latter two approaches; in the evidence retrieval approach, supporting pieces of evidence for the claim verdict have to be fetched from a relevant source using an information retrieval method. While the benefits of such systems outweigh the limitations of static approaches mentioned earlier, there are certain significant constraints that can also affect the explainability of these models. While the quality of the source (biased or unreliable), availability of the source (geographical or language restrictions), and resources for the retrieval process (time-consuming, and expensive human and computational resources) can have a significant impact on the evidence retrieval and limit the scope of evidence; a deep understanding of claim context is critical to avoid misinterpreted and incomplete evidence which lead to erroneous verdicts. Nevertheless, these limitations suggest that the evidence retrieval approach might not be entirely consistent with key XAI principles such as 'Accuracy' and 'Fidelity'. This, in turn, casts doubt on the effectiveness of any post hoc explainability measures attempted within this data aspect.

An alternative approach is using text from verified sources of information as evidence; Encyclopedia articles, journals, Wikipedia, and fact-checked databases are some examples. Since Wikipedia is an open source web-based encyclopedia and contains articles on a wide range of topics, it is consistently considered an important source of information for many applications, including economic development [46], education [47], data mining [48], and AFV. For example, the FEVER task [24], an application in AFV, relies on the retrieval of evidence from Wikipedia pages. In the FEVER dataset, each SUPPORTED/REFUTED claim is annotated with evidence from Wikipedia. This evidence could be a single sentence, multiple sentences, or a composition of evidence from multiple sentences, sourced from the same page or multiple pages of Wikipedia. This approach aligns well with the XAI principle of 'Interpretability', as Wikipedia is a widely accessible and easily understandable source of information. However, it is crucial to note that Wikipedia also comes with limitations that could impact the 'Accuracy' and 'Fidelity' principles of XAI, which can potentially impact the interpretability of models relying on Wikipedia as a primary data source. Firstly, like any other source, Wikipedia pages can contain biased and inaccurate content, and these can remain undetected for a longer period (same with outdated information); this compromises the 'Accuracy' of any AFV model trained on these data. Secondly, despite covering a wide range of topics, Wikipedia suffers deficiencies in comprehensiveness⁹, limiting a model's ability to understand contextual information fully, thereby affecting 'Interpretability'. Lastly, models trained predominantly on Wikipedia's textual content can develop biases and limitations inherent to the nature and scope of Wikipedia's content, impacting

⁹ https://en.wikipedia.org/wiki/Reliability_of_Wikipedia#Coverage

both ‘Fidelity’ and ‘Interpretability’ when applied to diverse real-world scenarios and varied types of unstructured data.

Given these considerations and their misalignment with the XAI objectives of ‘Interpretability,’ ‘Accuracy,’ and ‘Fidelity,’ it becomes evident that relying solely on Wikipedia as a training dataset may not be the most effective pathway toward explainable AFV.

Alternatively, Wikipedia can be used as an elementary corpus to train the AI model to achieve a general understanding of various knowledge domains for AFV, and this background or prior knowledge can then be harnessed further with additional domain data to gain a deeper context (which helps the model to attain information on global relationship and thus increase explainability). Being the largest Wikipedia-based benchmark dataset for fact verification [26,49], the FEVER dataset can unarguably be considered as this elementary corpus for AFV tasks, and Transformers and Transfer Learning is the most pragmatic technology choice for AFV according to state-of-the-art systems [29,30,50].

The quality of the data set we use or create for the application is a major factor in determining the explainability of a transformer-based AFV model and its ability to comprehend the underlying context. For example; [51] developed the SCIFACT data set in order to expand the ideas of FEVER to COVID-19 applications. SCIFACT comprises 1.4K expert-written scientific claims along with 5K+ abstracts (from different scientific articles) that either support or refute each claim and is annotated with rationales, which consists of a minimal collection of sentences from the abstract that imply the claim. The study demonstrated the obvious advantages of using such a domain-specific dataset (can also be called subdomain here since scientific claim verification is a subtask of claim verification) as opposed to just using a Wikipedia-based evidence dataset. [51] argues that the inclusion of rationales in the training data set “facilitates the development of interpretable models” that not only label predictions but also identify the specific sentences necessary to support their decisions. However, the limited scale of the dataset, consisting of only 1.4K claims, necessitates caution in interpreting assessments of system performance and underscores the need for more expansive datasets to propel advancements in explainable fact-checking research.

Building on this perspective of improving the quality and diversity of the data set, [52] critically evaluated the FEVER corpus, emphasizing its reliance on synthetic claims from Wikipedia and advocating for a corpus that incorporates natural claims from a variety of web sources. In response to this identified need, they introduced a new, mixed-domain corpus, which includes domains like blogs, news, and social media, the mediums often responsible for the spread of unreliable information. This corpus, which encompasses 6,422 validated claims and over 14,000 documents annotated with evidence, addresses the prevalent limitations in existing corpora, including restricted sizes, lack of detailed annotations, and domain confinement. However, through meticulous error analysis, [52] discovered inherent challenges and biases in claim classification, attributed to the heterogeneous nature of the data and the incorporation of Fine-Grained Evidence (FGE) from unreliable sources. We infer that these findings illustrate substantial barriers to realizing the fundamental goals of XAI, particularly accuracy and fidelity. Moreover, [52]’s focus on diligently modeling meta-information related to evidence and claims could be understood as their implicit recognition of the crucial role of explainability in the realm of automated fact-checking. By suggesting the integration of diverse forms of contextual information and reliability assessments of sources, they highlight the necessity of developing models that are not only more accurate but also capable of providing reasoned and understandable decisions, a pivotal step towards fostering explainability in automated fact-checking systems.

Table 2 offers a comprehensive categorization of the datasets used in Fact Verification systems, highlighting a variety of dataset types, each highlighting distinctive attributes and challenges. The datasets are categorized meticulously based on their inherent nature and source, such as ‘Knowledge-free Systems’, ‘Knowledge-Base-Based’, ‘Wikipedia-Based’, ‘Domain(Single)-Specific-Corpus’, and ‘Mixed-domain-Corpus (non-Wikipedia-based)’. Each type is represented with illustrative studies and remarks to provide insight into the inherent limitations or challenges in relation to enhancing explainability in AFV systems. The categorization is enriched with

subclassifications under ‘Knowledge Type’, ‘Text Type’, and ‘Domain Type’. ‘Knowledge-free systems’ are denoted with dashes (-) under ‘Text Type’ and ‘Domain Type’, indicating the inherent absence of these attributes. This underscores the retrieval-free nature of such systems, which predominantly rely on the intrinsic linguistic features of the claims, thus lacking contextual understanding aKnowledge-Base-Based explainability infeasible. The ‘Knowledge-Base-Based’ type can be either single-dosubcategoriesdomain, represented by checkmarks in both sub-categories under ‘Domain Type’. This illustrates the versatility of knowledge-based systems in utilizing structured information from a specialized domain or amalgamating insights from multiple domains. The ability to cater to varied domains accentuates the expansive applicability of such systems, though it also brings forth challenges related to scalability and capturing nuanced information. ‘Wikipedia-Based’ datasets, inherently multi-domain, are highlighted separately to focus on the specific challenges of using Wikipedia as the main information source, such as dealing with potential biases and inaccuracies. The ‘Domain(Single)-Specific-Corpus’ is distinguished as it focuses on a specialized or singular domain, providing depth and specificity. While this focus allows for a detailed exploration of a particular domain, it also poses limitations due to the restricted scope and potential biases inherent to the selected domain, thereby affecting the overall evaluation and applicability of the system. Additionally, the ‘Mixed-domain Corpus’ type emphasizes the inclusion of diverse domains, especially those not solely reliant on Wikipedia, addressing the challenges arising from data heterogeneity and reliability.

Table 2. Comparative Analysis of Dataset Types and Their Impact on Explainability in AFV Systems.

Fact_Verification Dataset Type	Example_Studies	Knowledge_Type		Text_Type		Domain_Type		Remarks
		Knowledge-Free	External-knowledge	Structured-Data	Free-Text	Single-Domain	Multi-Domain	
Knowledge-free Systems	[40], [41]	✓	×	—	—	—	—	Lack of contextual understanding, dataset-level explainability infeasible
Knowledge-Base-Based	[43], [44]	×	×	✓	×	✓	✓	Limited scalability, inability to capture nuanced information
Wikipedia-Based	[53]	×	✓	×	✓	×	✓	Biased and inaccurate content, limited comprehensiveness
Domain(Single)-Specific-Corpus	[51]	×	✓	×	✓	✓	×	Limited size, potential for biased evaluation
Mixed-domain-Corpus (non-Wikipedia-based)	[52]	×	✓	×	✓	×	✓	Challenges in classification due to heterogeneous data (impact accuracy); evidence from unreliable sources (impact fidelity)

The categorization in Table 2, coupled with associated remarks, is intended to act as a resource, providing information on the various challenges and possibilities to improve explainability within AFV systems. This categorization can guide researchers and practitioners in making informed decisions regarding dataset selection and utilization, providing a clearer understanding of the implications and limitations of different dataset types in the context of Automated Fact Verification.

We acknowledge the extensive investigations conducted by [32] in Explainable NLP and by [4] in Explainable AFV, which provide meticulous lists and insightful analyses of prevalent datasets in their respective fields. It is crucial to clarify that our endeavor in this section (Section 3.3) does not aim to perform an exhaustive review of datasets, a task diligently undertaken by the aforementioned studies. Instead, our work is uniquely positioned to illuminate the distinctive attributes and inherent diversity within various dataset types in AFV. We hope that our attempt to examine the impact of different data types on explainability serves as a thoughtful addition to ongoing discussions and reflections on the subject, offering a new perspective on the multifaceted interactions between data diversity and explainability in AFV.

4. Discussion

While fact-checking datasets commonly support the standard 3-stage pipeline of fact verification, there is currently a lack of datasets that specifically facilitate explanation learning aligned with government and intergovernmental standards on XAI. This is of paramount importance towards Explainable AFV in as much as we expect an AI system to produce explanations, it should also have the ability/opportunity to consume explanations. To achieve this, it is necessary to train the model network using an explanation-learn-friendly (ELF) dataset. However, prominent large-scale datasets like FEVER [53] and MultiFC [54] lack this aspect of the fact verification task. Furthermore, currently there are no alternative resources available to address this limitation, as commented on by [55]. To create an ELF dataset, it is essential to analyze the data set practices of fact verification systems with a focus on explainability. This paper has undertaken this crucial initial step and found that the absence of an explanation-based fact verification corpus presents a significant obstacle to advancing research in the field of explainable fact-checking.

In addition to the lack of suitable ELF datasets in AFV, another significant challenge to the growth of the explainable AFV field is the ambiguity and discrepancies surrounding the concepts of local and global explainability. Global interpretability refers to the ability to comprehend the overall logic and reasoning of a model, including all possible outcomes. It involves understanding the complete decision-making process and the underlying principles of the model. On the other hand, local interpretability refers to the ability to understand the specific reasons or factors that contribute to a particular decision or prediction. It focuses on the interpretability of individual predictions or decisions rather than the entire model. These terms are not consistently understood and implemented in different research communities, leading to confusion and causing impeding progress in the field. While XAI researchers [6,14,15,20] generally adhere to a consistent understanding of local and global explainability, Explainable AFV researchers have different interpretations and perspectives, contributing to the ambiguity surrounding explainable AFV. For example, [33] focuses on local coherence and global coherence, evaluating sentence cohesion and the appropriateness of explanations in relation to the claim and associated label, both at the prediction level. On the other hand, [21] discuss the explainability as providing local explanations for individual data points, without specifically addressing local or global aspects. As a result, the definitions of local coherence, global coherence, and explainability in AFV studies predominantly refer to prediction-level explainability, leaving the concept of global explainability in AFV insufficiently defined.

The lack of recognition of the importance of global explainability is evident in the implementations as well. Existing systems primarily focus on local explainability, which hampers adequate understanding of the system's decision-making process at the data set level. In an extensive survey conducted by [4] on explainable AFV systems, it was found that all the examined systems focused primarily on providing explanations for individual predictions rather than offering explanations about the underlying fact-checking model itself. This indicates a prevalent trend in the field of explainable AFV, where the emphasis is on local explainability. However, this local explainability is not sufficient for AFV systems because it only provides insights into individual predictions without offering a holistic view of the system's overall behavior. Furthermore, global explainability is crucial for AFV systems, as it provides a comprehensive understanding of how the system arrives at its predictions and decisions. This global approach also allows AFV systems to align with advances in XAI research and comply with the XAI principles, enabling transparency and accountability.

In addition to the ambiguity surrounding local and global concepts, the field of explainable AFV is further complicated due to variations in how explainability concepts are categorized, suggesting a lack of consensus on taxonomy. For example, while [22] distinguishes between intrinsic explainability and post hoc explainability, other researchers in explainable AFV, such as [21], propose a categorization that broadly divides XAI into interpretability and explainability. [22] describe intrinsic explainability as the process of creating self-explanatory models that inherently incorporate explainability. This suggests that their definition of 'intrinsic explainability' closely aligns with the general notion of

`interpretability' related to model-level reasoning, as discussed in Section 2.2. However, the choice of the term `intrinsic' by [22] adds a distinct nuance to this categorization. On the other hand, their view on post hoc explainability is in line with standard XAI. In contrast, while [21] aligns `interpretability' with the mainstream XAI taxonomy, they adopt a narrower view for `explainability', reserving it for local explanations of individual instances, which is a subset of post-hoc explainability. This deviates from the wider view where `explainability' usually refers to model-level justifications.

These disparities in taxonomy demonstrate that the ambiguity extends beyond the local and global dimensions, contributing to the overall ambiguity within the field of explainable AFV. This disagreement and discrepancy among the relatively few existing explainable AFV systems pose significant challenges for the growth and advancement of research in this field and highlights the need for a more standardized approach to explainability in AFV systems.

4.1. Limitations

We concentrated our efforts on exploring the explainability of DNN based AFV models, and thus neglected other explainability approaches such as rule discovery [56,57]. This research gap provides an opportunity for future studies to investigate the model explainability of DNNs, particularly transformer models, in the context of AFV.

Similarly, to limit the scope of this paper, we did not address the absence of a clear and established link between the various interpretation methods proposed in the literature and the evaluation criteria for measuring explainability; the lack of clarity regarding how to measure explainability is a significant challenge in this field of the research. This aspect warrants further investigation in future research to enhance the assessment of explainable AFV systems.

4.2. Future Research Directions

In addition to the future plans outlined in the limitations of this study, we propose the following directions for exploration and research.

- *Direction 1: Exploring a Balanced Approach to Explainability in AFV:* Researchers should explore the development of techniques and methodologies aimed at achieving a balanced approach to explainability, integrating both global and local perspectives in AFV systems. This involves understanding the broader relationships and patterns that underlie AFV model decisions across diverse factual claims and evidence (global explainability), while also addressing the specific concerns related to individual instances (local explainability). For instance, a nuanced exploration of gray-explainability could involve refining gray-box models to optimize the trade-off between interpretability and accuracy, ensuring that the explanations provided are as meaningful and understandable as possible without incurring a substantial loss in predictive accuracy. Although examples such as dispute resolution and individual patient treatment decisions illustrate the broader applicability and importance of this approach beyond the realm of fact verification, they underscore the universal need for tailored and comprehensible explanations in individual cases. In fact verification systems, a balanced approach is particularly crucial for gaining both a localized understanding of individual claims and a broader insight that can inform strategies in handling diverse types of information and evidence. By investigating methods that provide insights into AFV model behavior and reasoning patterns on both the macro and micro levels, researchers can work towards achieving a holistic understanding of explainability in AFV systems. Following the principles of XAI, one potential starting point, as suggested by the principles of XAI, such as [23], could be to explain multiple representative individual predictions (local) as a means to gain insights towards a more comprehensive understanding (insight). This nuanced exploration, which aligns with the overarching goal of achieving explainability in AFV systems, ensures that the insights gained are as widely applicable as they are individually relevant, potentially leading to more informed and equitable decision-making processes across different domains.

- *Direction 2: Comprehensive Investigation and Comparative Analysis of AFV Datasets:* Future research endeavors could benefit from undertaking a meticulous and comprehensive review of the applicable data sets for AFV, informed by the insights provided in our tables 1 and 2. Table 1 outlines a comparative analysis of various methodologies used to improve explainability in AFV systems, while Table 2 delves into the distinctive attributes and inherent diversity within various types of data set in AFV. A focused study in this direction could reveal deeper insights into the suitability and compatibility of various datasets with different AFV models and explainability techniques, providing a more nuanced understanding of the interactions between dataset characteristics and explainability. Such an investigation would not only enrich the understanding of the influence of diverse datasets on the explainability of AFV models but also reveal untapped potential in utilizing underexplored types of dataset to enhance model transparency and interpretability. By synergizing the diverse techniques for explainability and the variety of dataset types highlighted in our tables, this research direction has substantial potential to lessen the gap in the field of explainable AFV.
- *Direction 3: Development of an Explainability-Focused, Explanation Learning-Friendly (ELF) Dataset:* As a logical progression from Direction 2, researchers should prioritize developing an ELF dataset to address the lack of explanations in existing AFV datasets, enabling more nuanced studies in explainability in AFV. This customized data set would serve as a benchmark to assess the effectiveness of various AFV models in generating meaningful explanations, thereby fostering advancements in creating explainable AFV systems. Such a focused endeavor would be pivotal in bridging existing gaps and furthering research in explainable AFV, allowing for an exploration of the interplay between dataset attributes, model structures, and explainability methodologies.

5. Conclusion

This study serves as a basis for future research in addressing the global challenge of explainability in AFV. By adopting a discrete, yet interconnected approach in discussing explainable AI and the explainability of AFV, this paper also provides valuable information for researchers interested in exploring alternative strategies to solve the 'black box' problem in AFV. Moreover, as emphasized by [16], addressing the challenge of interpretability requires domain-specific definitions, which means that future research should closely examine the distinct technical hurdles associated with key domains within the field. To the best of our knowledge, there have been no other studies that have investigated the explainability of AFV models according to the principles, objectives, and approaches defined in the field of XAI.

In alignment with insights from various research studies, including the work of [14], our study advocates the notion that current AFV systems, which primarily emphasize local explainability, need to expand their horizons to address the broader objectives of XAI. The core objective of attaining global explainability in AFV is to discover overarching relationships and patterns that AFV models incorporate and apply across diverse factual claims and evidence, which extend beyond local interpretations. To initiate this transition, our study suggests that, the development of a specialized training dataset tailored for this purpose can serve as a pivotal step. In addition to these points, our study provides an analysis of data practices in current AFV, offering valuable insights into the existing gaps and limitations in achieving explainability. Furthermore, our study offers a brief overview of the ambiguity and discrepancies present among AFV systems regarding the concepts and perspectives of explainability.

We are of the opinion that following the research directions we have identified would be a promising beginning for further progress in the area of AFV explainability. However, it is essential to recognize that there is still a long way to go to accomplish a genuine transition from manual to automated fact verification, as was pointed out in the problem statement in the Introduction.

Author Contributions: Conceptualization, M.V. and P.N.; methodology, M.V. and P.N.; writing—original draft preparation, M.V.; writing—review and editing, P.N., W.Q.Y. and H.A-C.; visualization, M.V.; supervision, P.N. and W.Q.Y. All authors have read and agreed to the published version of the manuscript.

Institutional Review Board Statement: Not applicable

Informed Consent Statement: Not applicable

Data Availability Statement: Not applicable

Acknowledgments: We would like to express our profound gratitude to Dr. Amrith Krishna, for his invaluable insights and meticulous feedback on this work. Dr. Amrith Krishna's extensive expertise and groundbreaking work in Explainable Fact Verification have significantly enhanced the depth and rigor of our research. We are deeply appreciative of the time and effort he devoted to reviewing our work and providing critical insights that have been instrumental in refining this paper. However, any errors or omissions in this work are solely our responsibility.

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

Abbreviations

The following abbreviations are used in this manuscript:

AI	Artificial Intelligence
XAI	Explainable Artificial Intelligence
AFV	Automated Fact Verification
ELF	Explanation Learn-Friendly
DNN	Deep Neural Network
RTE	Recognizing Textual Entailment

References

1. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is all you need. *Advances in Neural Information Processing Systems*. Neural information processing systems foundation, 2017, Vol. 2017-Decem, pp. 5999–6009. doi:10.48550/arxiv.1706.03762.
2. Hassan, N.; Zhang, G.; Arslan, F.; Caraballo, J.; Jimenez, D.; Gawsane, S.; Hasan, S.; Joseph, M.; Kulkarni, A.; Nayak, A.K.; Sable, V.; Li, C.; Tremayne, M. ClaimBuster: The First-Ever End-to-End Fact-Checking System. *Proc. VLDB Endow.* **2017**, *10*, 1945–1948. doi:10.14778/3137765.3137815.
3. Chen, J.; Bao, Q.; Sun, C.; Zhang, X.; Chen, J.; Zhou, H.; Xiao, Y.; Li, L. Loren: Logic-regularized reasoning for interpretable fact verification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 2022, Vol. 36, pp. 10482–10491.
4. Kotonya, N.; Toni, F. Explainable Automated Fact-Checking: A Survey. *COLING 2020 - 28th International Conference on Computational Linguistics, Proceedings of the Conference*. Online, 2020, pp. 5430–5443, [2011.03870]. doi:10.18653/v1/2020.coling-main.474.
5. Došilović, F.K.; Brčić, M.; Hlupić, N. Explainable artificial intelligence: A survey. *41st International convention on information and communication technology, electronics and microelectronics (MIPRO)*. IEEE, 2018, pp. 0210–0215.
6. Guidotti, R.; Monreale, A.; Ruggieri, S.; Turini, F.; Giannotti, F.; Pedreschi, D. A survey of methods for explaining black box models. *ACM Computing Surveys (CSUR)* **2018**, *51*, 1–42.
7. Kim, T.W. Explainable artificial intelligence (XAI), the goodness criteria and the grasp-ability test. *ArXiv* **2018**, abs/1810.09598.
8. Das, A.; Liu, H.; Kovatchev, V.; Lease, M. The state of human-centered NLP technology for fact-checking. *Information Processing & Management* **2023**, *60*, 103219.
9. Olivares, D.G.; Quijano, L.; Liberatore, F. Enhancing Information Retrieval in Fact Extraction and Verification. *Proceedings of the Sixth Fact Extraction and VERification Workshop (FEVER)*, 2023, pp. 38–48.
10. Rani, A.; Tonmoy, S.M.T.I.; Dalal, D.; Gautam, S.; Chakraborty, M.; Chadha, A.; Sheth, A.; Das, A. FACTIFY-5WQA: 5W Aspect-based Fact Verification through Question Answering, 2023, [arXiv:cs.CL/2305.04329].
11. Guo, Z.; Schlichtkrull, M.; Vlachos, A. A Survey on Automated Fact-Checking. *Transactions of the Association for Computational Linguistics* **2022**, *10*, 178–206, [2108.11896]. doi:10.1162/tacl_a_00454.

12. Gunning, D.; Vorm, E.; Wang, J.Y.; Turek, M. DARPA's explainable AI (XAI) program: A retrospective. *Applied AI Letters* **2021**, *2*, e61. doi:10.1002/AIL2.61.
13. Ali, S.; Abuhmed, T.; El-Sappagh, S.; Muhammad, K.; Alonso-Moral, J.M.; Confalonieri, R.; Guidotti, R.; Del Ser, J.; Díaz-Rodríguez, N.; Herrera, F. Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. *Information Fusion* **2023**, p. 101805.
14. Doshi-Velez, F.; Kim, B. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv: Machine Learning* **2017**.
15. Moradi, M.; Samwald, M. Post-hoc explanation of black-box classifiers using confident itemsets. *Expert Systems with Applications* **2021**, *165*, 113941.
16. Rudin, C. Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead. *Nature Machine Intelligence* **2019**, *1*, 206–215.
17. Mueller, S.T.; Hoffman, R.R.; Clancey, W.; Emrey, A.; Klein, G. Explanation in Human-AI Systems: A Literature Meta-Review Synopsis of Key Ideas and Publications and Bibliography for Explainable AI. Technical report, DARPA XAI Program, 40 S. Alcaniz St., Pensacola, FL 32502, 2019.
18. Goodman, B.; Flaxman, S. European Union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine* **2017**, *38*, 50–57.
19. Gunning, D. Broad agency announcement explainable artificial intelligence (XAI). Technical report, Technical report, Defense Advanced Research Projects Agency Information Innovation Office 675 North Randolph Street Arlington, VA 22203-2114, 2016.
20. Murdoch, W.J.; Singh, C.; Kumbier, K.; Abbasi-Asl, R.; Yu, B. Definitions, methods, and applications in interpretable machine learning. *Proceedings of the National Academy of Sciences* **2019**, *116*, 22071–22080. doi:10.1073/pnas.1900654116.
21. Atanasova, P.; Simonsen, J.G.; Lioma, C.; Augenstein, I. Generating Fact Checking Explanations. Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics; Association for Computational Linguistics: Online, 2020; pp. 7352–7364. doi:10.18653/v1/2020.acl-main.656.
22. Shu, K.; Cui, L.; Wang, S.; Lee, D.; Liu, H. defend: Explainable fake news detection. Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, 2019, pp. 395–405.
23. Ribeiro, M.T.; Singh, S.; Guestrin, C. “Why Should I Trust You?” Explaining the Predictions of Any Classifier. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2016, pp. 1135–1144.
24. Thorne, J.; Vlachos, A.; Cocarascu, O.; Christodoulopoulos, C.; Mittal, A. The Fact Extraction and VERification (FEVER) Shared Task. Proceedings of the First Workshop on Fact Extraction and VERification (FEVER); Association for Computational Linguistics: Brussels, Belgium, 2018; pp. 1–9. doi:10.18653/v1/W18-5501.
25. Soleimani, A.; Monz, C.; Worring, M. BERT for evidence retrieval and claim verification. Advances in Information Retrieval. ECIR 2020. Lecture Notes in Computer Science, 2019, Vol. 12036 LNCS, pp. 359–366, [1910.02655]. doi:{10.1007/978-3-030-45442-5_45}.
26. Zhong, W.; Xu, J.; Tang, D.; Xu, Z.; Duan, N.; Zhou, M.; Wang, J.; Yin, J. Reasoning over semantic-level graph for fact checking. Proceedings of the Annual Meeting of the Association for Computational Linguistics. Association for Computational Linguistics (ACL), 2020, pp. 6170–6180. doi:10.18653/v1/2020.acl-main.549.
27. Jiang, K.; Pradeep, R.; Lin, J. Exploring Listwise Evidence Reasoning with T5 for Fact Verification. ACL-IJCNLP 2021 - 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, Proceedings of the Conference, 2021, Vol. 2, pp. 402–410. doi:10.18653/v1/2021.acl-short.51.
28. Chen, J.; Zhang, R.; Guo, J.; Fan, Y.; Cheng, X. GERE: Generative Evidence Retrieval for Fact Verification. SIGIR 2022 - Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. Association for Computing Machinery, Inc, 2022, pp. 2184–2189, [2204.05511]. doi:10.1145/3477495.3531827.
29. DeHaven, M.; Scott, S. BEVERS: A General, Simple, and Performant Framework for Automatic Fact Verification. Proceedings of the Sixth Fact Extraction and VERification Workshop (FEVER); Association for Computational Linguistics: Dubrovnik, Croatia, 2023; pp. 58–65.
30. Krishna, A.; Riedel, S.; Vlachos, A. ProofVer: Natural Logic Theorem Proving for Fact Verification. *Transactions of the Association for Computational Linguistics* **2022**, *10*, 1013–1030, [2108.11357].

31. Huang, Y.; Gao, M.; Wang, J.; Shu, K. DAFD: Domain Adaptation Framework for Fake News Detection. *Neural Information Processing*; Mantoro, T.; Lee, M.; Ayu, M.A.; Wong, K.W.; Hidayanto, A.N., Eds.; Springer International Publishing: Cham, 2021; pp. 305–316.
32. Wiegrefe, S.; Marasovic, A. Teach Me to Explain: A Review of Datasets for Explainable Natural Language Processing. *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021.
33. Kotonya, N.; Toni, F. Explainable automated fact-checking for public health claims. *EMNLP 2020 - 2020 Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 2020*, pp. 7740–7754, [2010.09926]. doi:10.18653/v1/2020.emnlp-main.623.
34. Popat, K.; Mukherjee, S.; Yates, A.; Weikum, G. Declare: Debunking fake news and false claims using evidence-aware deep learning. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018, 2018*, pp. 22–32, [1809.06416]. doi:10.18653/v1/d18-1003.
35. Jain, S.; Wallace, B.C. Attention is not Explanation. *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*; Association for Computational Linguistics, , 2019; pp. 3543–3556.
36. Serrano, S.; Smith, N.A. Is attention interpretable? *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*; Association for Computational Linguistics, , 2019; pp. 2931–2951.
37. Pruthi, D.; Gupta, M.; Dhingra, B.; Neubig, G.; Lipton, Z.C. Learning to deceive with attention-based explanations. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*; Association for Computational Linguistics, , 2020; pp. 4782–4793.
38. Dai, S.C.; Hsu, Y.L.; Xiong, A.; Ku, L.W. Ask to Know More: Generating Counterfactual Explanations for Fake Claims. *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2022*, pp. 2800–2810.
39. Xu, W.; Liu, Q.; Wu, S.; Wang, L. Counterfactual Debiasing for Fact Verification. *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2023*, pp. 6777–6789.
40. Rashkin, H.; Choi, E.; Jang, J.Y.; Volkova, S.; Choi, Y. Truth of varying shades: Analyzing language in fake news and political fact-checking. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, 2017*, pp. 2931–2937. doi:10.18653/v1/d17-1317.
41. Wang, W.Y. “Liar, Liar Pants on Fire”: A New Benchmark Dataset for Fake News Detection. *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*; Association for Computational Linguistics: Vancouver, Canada, 2017; pp. 422–426. doi:10.18653/v1/P17-2067.
42. Thorne, J.; Vlachos, A. Automated Fact Checking: Task Formulations, Methods and Future Directions. *Proceedings of the 27th International Conference on Computational Linguistics*; Association for Computational Linguistics: Santa Fe, New Mexico, USA, 2018; pp. 3346–3359.
43. Shi, B.; Weninger, T. Discriminative predicate path mining for fact checking in knowledge graphs. *Knowledge-Based Systems* **2016**, *104*, 123–133, [1510.05911]. doi:10.1016/j.knosys.2016.04.015.
44. Gardner, M.; Mitchell, T. Efficient and expressive knowledge base completion using subgraph feature extraction. *Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing, 2015*, pp. 1488–1498. doi:10.18653/v1/d15-1173.
45. Bordes, A.; Usunier, N.; Garcia-Durán, A.; Weston, J.; Yakhnenko, O. Translating embeddings for modeling multi-relational data. *Advances in Neural Information Processing Systems*, 2013.
46. Sheehan, E.; Meng, C.; Tan, M.; Uzgent, B.; Jean, N.; Burke, M.; Lobell, D.; Ermon, S. Predicting economic development using geolocated wikipedia articles. *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining, 2019*, pp. 2698–2706.
47. Brailas, A.; Koskinas, K.; Dafermos, M.; Alexias, G. Wikipedia in Education: Acculturation and learning in virtual communities. *Learning, Culture and Social Interaction* **2015**, *7*, 59–70. doi:https://doi.org/10.1016/j.lcsi.2015.07.002.
48. Schwenk, H.; Chaudhary, V.; Sun, S.; Gong, H.; Guzmán, F. WikiMatrix: Mining 135M Parallel Sentences in 1620 Language Pairs from Wikipedia. *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*; Association for Computational Linguistics: Online, 2021; pp. 1351–1361. doi:10.18653/v1/2021.eacl-main.115.

49. Shorten, C.; Khoshgoftaar, T.M.; Furht, B. Deep Learning applications for COVID-19. *Journal of Big Data* **2021**, *8*, 1–54.
50. Stambach, D. Evidence Selection as a Token-Level Prediction Task. FEVER 2021 - Fact Extraction and VERification, Proceedings of the 4th Workshop. Association for Computational Linguistics (ACL), 2021, pp. 14–20. doi:10.18653/v1/2021.fever-1.2.
51. Wadden, D.; Lin, S.; Lo, K.; Wang, L.L.; van Zuylen, M.; Cohan, A.; Hajishirzi, H. Fact or Fiction: Verifying Scientific Claims. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP); Association for Computational Linguistics: Online, 2020; pp. 7534–7550. doi:10.18653/v1/2020.emnlp-main.609.
52. Hanselowski, A.; Stab, C.; Schulz, C.; Li, Z.; Gurevych, I. A richly annotated corpus for different tasks in automated fact-checking. Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL); Association for Computational Linguistics: Hong Kong, China, 2019; pp. 493–503. doi:10.18653/v1/K19-1046.
53. Thorne, J.; Vlachos, A.; Christodoulopoulos, C.; Mittal, A. FEVER: a Large-scale Dataset for Fact Extraction and VERification. Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies; Association for Computational Linguistics: New Orleans, Louisiana, 2018; pp. 809–819. doi:10.18653/v1/N18-1074.
54. Augenstein, I.; Lioma, C.; Wang, D.; Chaves Lima, L.; Hansen, C.; Hansen, C.; Simonsen, J.G. MultiFC: A Real-World Multi-Domain Dataset for Evidence-Based Fact Checking of Claims. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP); Association for Computational Linguistics: Hong Kong, China, 2019; pp. 4685–4697. doi:10.18653/v1/D19-1475.
55. Stambach, D.; Ash, E. e-FEVER: Explanations and Summaries for Automated Fact Checking. *BRISK Binary Robust Invariant Scalable Keypoints* **2020**, pp. 12–19. doi:10.3929/ETHZ-B-000453826.
56. Gad-Elrab, M.H.; Stepanova, D.; Urbani, J.; Weikum, G. Exfakt: A framework for explaining facts over knowledge graphs and text. Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, 2019, pp. 87–95.
57. Ahmadi, N.; Lee, J.; Papotti, P.; Saeed, M. Explainable Fact Checking with Probabilistic Answer Set Programming. Conference for Truth and Trust Online, 2019. doi:10.36370/tto.2019.15.

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