



OPEN ACCESS

EDITED BY

David Tomás,
University of Alicante, Spain

REVIEWED BY

Zenun Kastrati,
Linnaeus University, Sweden
Banafsheh Rekabdar,
Portland State University, United States

*CORRESPONDENCE

Feras A. Batarseh
✉ batarseh@vt.edu

RECEIVED 18 May 2023

ACCEPTED 04 September 2023

PUBLISHED 25 September 2023

CITATION

Gurrapu S, Kulkarni A, Huang L, Lourentzou I and Batarseh FA (2023) Rationalization for explainable NLP: a survey.
Front. Artif. Intell. 6:1225093.
doi: 10.3389/frai.2023.1225093

COPYRIGHT

© 2023 Gurrapu, Kulkarni, Huang, Lourentzou and Batarseh. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Rationalization for explainable NLP: a survey

Sai Gurrapu¹, Ajay Kulkarni², Lifu Huang¹, Ismini Lourentzou¹ and Feras A. Batarseh^{2,3*}

¹Department of Computer Science, Virginia Tech, Blacksburg, VA, United States, ²Commonwealth Cyber Initiative, Virginia Tech, Arlington, VA, United States, ³Department of Biological Systems Engineering, Virginia Tech, Blacksburg, VA, United States

Recent advances in deep learning have improved the performance of many Natural Language Processing (NLP) tasks such as translation, question-answering, and text classification. However, this improvement comes at the expense of model explainability. Black-box models make it difficult to understand the internals of a system and the process it takes to arrive at an output. Numerical (LIME, Shapley) and visualization (saliency heatmap) explainability techniques are helpful; however, they are insufficient because they require specialized knowledge. These factors led rationalization to emerge as a more accessible explainable technique in NLP. Rationalization justifies a model's output by providing a natural language explanation (rationale). Recent improvements in natural language generation have made rationalization an attractive technique because it is intuitive, human-comprehensible, and accessible to non-technical users. Since rationalization is a relatively new field, it is disorganized. As the first survey, rationalization literature in NLP from 2007 to 2022 is analyzed. This survey presents available methods, explainable evaluations, code, and datasets used across various NLP tasks that use rationalization. Further, a new subfield in Explainable AI (XAI), namely, Rational AI (RAI), is introduced to advance the current state of rationalization. A discussion on observed insights, challenges, and future directions is provided to point to promising research opportunities.

KEYWORDS

rationalization, explainable NLP, rationales, abstractive rationale, extractive rationale, large language models, natural language generation, Natural Language Processing

1. Introduction

The commercialization of NLP has grown significantly in the past decade. Text has a ubiquitous nature which enables many practical NLP use cases and applications, including but not limited to text classification, fact-checking, machine translation, text2speech, and others, which significantly impact our society. Despite its diverse and practical applications, NLP faces many challenges; an important one is explainability (Madsen et al., 2021).

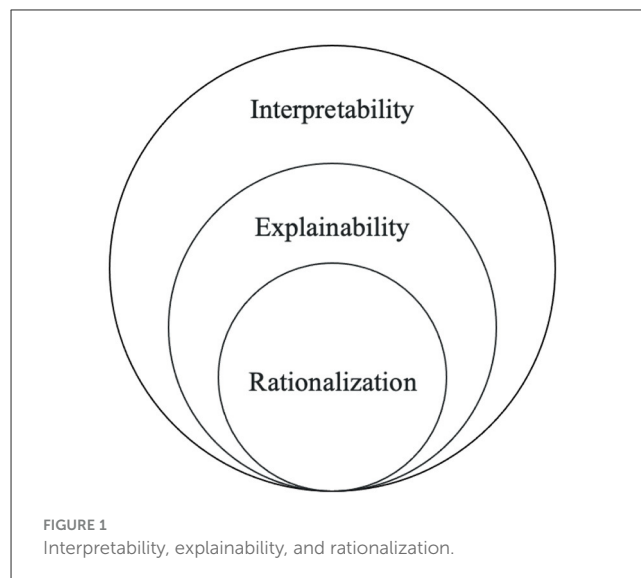
In the past, NLP systems have traditionally relied on white-box techniques. These techniques—rules, decision trees, hidden Markov models, and logistic regression—are inherently explainable (Danilevsky et al., 2020). The recent developments in deep learning have contributed to the emergence of black-box architectures that improve task performance at the expense of model explainability. Such black-box predictions make understanding how a model arrives at a decision challenging. This lack of explainability is a significant cause of concern for critical applications. For example, directly applying natural language generation methods to automatically generate radiology reports from chest X-ray images only guarantees that the produced reports will look natural rather than contain correct anatomically-aware information (Liu et al., 2019a).

Similarly, Visual Question Answering (VQA) systems are known to learn heavy language priors (Agrawal et al., 2016). Thus, a lack of transparency can affect the decision-making process and may lead to the erosion of trust between humans and Artificial Intelligence (AI) systems. This can further jeopardize users' safety, ethics, and accountability if such a system is deployed publicly (Madsen et al., 2021). Considering the utilization of NLP in healthcare, finance, and law domains, all of which can directly affect human lives, it can be dangerous to blindly follow machine predictions without fully understanding them. For instance, a physician following a medical recommendation or an operation procedure for a patient without full knowledge of the system can do more harm than good. In addition, systems employing Machine Learning (ML), such as most current NLP methods, are prone to adversarial attacks where small, carefully crafted local perturbations can maximally alter model predictions, essentially misguiding the model to predict incorrectly but with high confidence (Finlayson et al., 2019). The bar for the ethical standards and the accountability required to maintain NLP systems continue to increase as these systems become more opaque with increasingly complex networks and algorithms.

There has been significant research focus on enabling models to be more interpretable, i.e., allowing humans to understand the internals of a model (Gilpin et al., 2018). However, due to the lack of completeness, interpretability alone is not enough for humans to trust black-box models. *Completeness* is the ability to accurately describe the operations of a system that allows humans to anticipate its behavior better. Gilpin et al. (2018) argue that *explainability* improves on interpretability as a technique to describe the model's decision-making process of arriving at a prediction and the ability to be verified and audited. Therefore, models with explainability are interpretable and complete. In this survey, the focus is on explainability and mainly on the **outcome explanation problem** where Guidotti et al. (2018) describe explainability as "the perspective of an end-user whose goal is to understand how a model arrives at its result".

In NLP, there exist various explainable techniques such as LIME (Local Interpretable Model-Agnostic Explanations) (Ribeiro et al., 2016), Integrated Gradients (Sundararajan et al., 2017), and SHAP (Shapley Additive Explanations) (Lundberg and Lee, 2017a). Despite the availability of these methods, many require specialized knowledge to understand their underlying processes, which makes them indecipherable and inaccessible for the general users or audience, which we refer to as the nonexperts, hence limiting usability. In the following sections, we share how rationalization addresses these problems and helps improve explainability for nonexpert users.

The structure of this literature survey is as follows. In Section 2, we share the background and intuition for rationalization. In Section 3, we explain our paper collection methodology. We identify and describe the most commonly used rationalization techniques and point to available papers adopting them in Section 4. In Section 5, we compare and contrast abstractive and extractive rationalization techniques. We conclude with a discussion of the open challenges and promising future directions in Section 6.

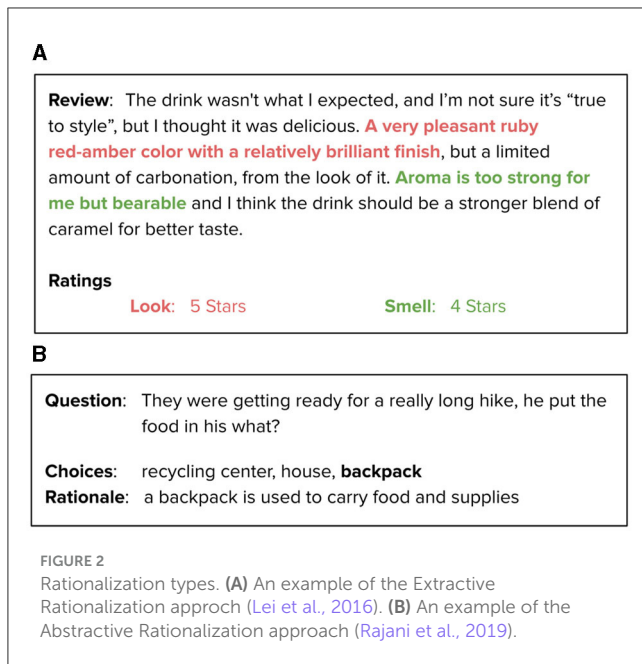


2. Background

One of the emerging explainable techniques for NLP applications is rationalization (Atanasova et al., 2020). **Rationalization** provides explanations in natural language to justify a model's prediction. These explanations are rationales, which present the input features influencing the model's prediction. The reasoning behind the prediction could be understood simply by reading the explanation/rationale, thereby revealing the model's decision-making process. Rationalization can be an attractive technique because it is human-comprehensible and allows individuals without domain knowledge to understand how a model arrived at a prediction. It essentially allows the model to "talk for themselves" (Bastings et al., 2019; Luo et al., 2021). This technique is a part of a subset of explainability because it enables models to be interpretable and complete, as shown in Figure 1. Specifically, rationalization provides a local explanation since each prediction has a unique explanation rather than one for the entire model. Local explanations can be categorized into two groups: local *post-hoc* and local self-explaining. Danilevsky et al. (2020) present local *post-hoc* methods as explaining a single prediction after the model predicts and local self-explaining methods as simultaneously explaining and predicting.

Rationalization in NLP was first introduced in 2007 (Zaidan et al., 2007). As described in Section 4.6, the objective was to use annotator rationales to improve task performance for text categorization. Interestingly, explainability was not the core objective. However, explainability is an advantage of rationalization because it makes the model inherently explainable even if used in the context of task improvement (Gurrapu et al., 2022).

Our literature review found that rationalization can be further divided into two major groups: abstractive and extractive (El-Kassas et al., 2021). In extractive rationalization, important features or sentences from the input data are extracted as rationales to support the prediction. In contrast, abstractive rationalization is a generative task in which novel sentences are



generated using new words or paraphrasing existing sentences. This is typically accomplished through the use of a language model such as BERT (Bidirectional Encoder Representations from Transformers) (Devlin et al., 2019), T5 (Raffel et al., 2020), or GPT (Generative Pre-trained Transformer) (Radford and Narasimhan, 2018). Figure 2 demonstrate the usage of the two explanation types with examples.

Recent advances in explainable NLP have led to a significant increase in rationalization research. Further, at present, the field of rationalization is disorganized. Thus, the motivations for this survey are—(a) formally define rationalization, (b) present and categorize the well-cited techniques based on NLP tasks, and (c) discuss current trends and future insights on the field. Thus, our primary goal is to provide future researchers with a comprehensive understanding of the previously scattered state of rationalization. The key contributions of this paper are as follows.

1. First literature to survey the field of rationalization in NLP.
2. Introduction of a new subfield called Rational AI (RAI) within Explainable AI (XAI).
3. A comprehensive list of details on available rationalization models, XAI evaluations, datasets, and code are provided to guide future researchers.
4. Presents NLP Assurance as an important method for developing more trustworthy and reliable NLP systems.

2.1. Related surveys

Table 1 indicates the related survey papers published in recent years. Danilevsky et al. (2020) note that previous surveys in XAI are broadly focused on AI without a specific narrow domain focus. Their work primarily focuses on surrogate-based explainability methods. NLP publications in recent years further demonstrate that this distinction is less relevant and valuable in the NLP domain

because “the same neural network can be used not only to make predictions but also to derive explanations.” Therefore, surveying the field of Explainable AI (XAI) in NLP requires NLP-specific methods that are different from the standard XAI methods that are widely known.

Thayaparan et al. (2020) survey the use of explanations specifically in Machine Reading Comprehension (MRC). The authors describe MRC papers that support explanations and provide a detailed overview of available benchmarks. Further, Madsen et al. (2021) briefly discuss rationalization and natural language explanations using a question and answer approach, CAGE (Commonsense Auto-Generated Explanations) (Rajani et al., 2019) as an example. Thus, this raises the question - *how can rationalization be generalized and applied to other tasks in the NLP domain?* However, until now, no comprehensive literature review on rationalization has been available for the prominent NLP tasks. Thus, through this survey paper, we attempt to address this need.

2.2. Definitions

To provide clarity and distinguish terms that are typically used interchangeably in published literature, a list of definitions is provided in Table 2. These terms are used throughout the paper.

3. Methodology

The following are the inclusion-exclusion criteria for our publications collection methodology. The first known use of rationalization in NLP was in the year 2007. Our survey focuses on the domain of NLP from 2007 to early 2022. We have included peer-reviewed publications within this range that include a significant rationalization component as a method to provide explainability. We defined *significance* as rationalization being the main component of their research methodology and approach. We have eliminated a number of publications that are either not entirely in the NLP domain or do not contain a significant rationalization component.

For identifying and selecting articles, the range of keywords and topics was limited to the following in the NLP domain: rationalization, explanation, justification, and explainable NLP. Thus, this survey includes reviews of the articles from journals, books, industry research, dissertations, and conference proceedings from commonplace AI/NLP venues such as ACL, EMNLP, NAACL, AAAI, NeurIPS, and others. Finally, these articles are categorized by important NLP tasks, as shown in Table 3. In recent years, there has been an increase in focus on explainability in NLP after a rise in deep learning techniques (Danilevsky et al., 2020). Due to this, a majority of the papers collected were from recent years (2016 and onwards), as illustrated in Figure 3.

The availability of relevant articles was limited. After following the above approach, 33 articles were downselected to be the primary focus of this paper's discussion. Instead of providing a broad yet surface-level understanding of the work, we focus on demonstrating in-depth the most important approaches and

TABLE 1 Three related survey papers on explainability and interpretability in NLP.

No.	Survey title	Year	Papers reviewed
1	A Survey on the state of Explainable AI for Natural Language Processing (Danilevsky et al., 2020)	2020	50
2	A survey on explainability in machine reading comprehension (Thayaparan et al., 2020)	2020	69
3	Post-hoc interpretability for neural NLP: a survey (Madsen et al., 2021)	2021	27

TABLE 2 List of common terms that are used interchangeably in published literature.

Term	Definition
Black-box Model	A “machine-learning obscure model, whose [architecture] internals are either unknown to the observer, or they are known but uninterpretable by humans” (Guidotti et al., 2018).
Interpretability	Interpretability “aims at developing tools to understand and investigate the behavior of an AI system” (Thayaparan et al., 2020). Doshi-Velez and Kim (2017) augments this definition by adding that interpretability tools allow us to “explain or to present in understandable terms to a human” what the AI system is performing.
Explainability	There is no consensus on the nature of explanations since they are entirely task-dependent, and AI embraces a wide variety of tasks (Miller, 2019a). We treat explainability as a specialization of interpretability where the aim of explainability is to design inherently interpretable models capable of performing transparent inference through the generation of an explanation for the final prediction (Thayaparan et al., 2020). In this survey, we focus on black-box explainability, specifically as the outcome explanation problem.
Rationalization	The term rationalization is interchangeable with explanation or justification. Rationalization has rarely been formally defined in the context of NLP. Therefore, we propose the following definition: Rationalization justifies a model’s output by providing a natural language explanation. This is accomplished by either extracting text fragments from the input (extractive rationalization) or by generating a novel explanation (abstractive rationalization).
NLP Assurance	A process that is applied at all stages of the NLP development lifecycle to ensure that all outcomes are valid, verified, trustworthy, and explainable to a non-expert, ethical in the context of its deployment, unbiased in its learning, and fair to its users. This definition is adopted from Batarseh et al. (2021) and modified to fit the scope of the NLP domain.

progress made in each NLP task. Overall, we selected six articles in multiple NLP domains, five on Machine Reading Comprehension and Sentiment Analysis, four on Text Classification, Fact-Checking and Commonsense Reasoning, and three on Natural Languages Inference, and two articles on Neural Machine Translation (NMT).

4. Rationalization techniques

In this section, we discuss relevant papers and their rationalization techniques categorized by the NLP tasks listed in Table 3. Tables with important information on the papers for each subsection are presented at the beginning.

TABLE 3 Details on seven NLP tasks and their definitions which are surveyed in this paper.

NLP task	Definition
Machine Reading Comprehension (MRC)	Enabling a model to answer questions regarding a given context (Baradaran et al., 2022).
Commonsense Reasoning	Going beyond pattern recognition to make inferences using world knowledge (Apperly, 2011; Sap et al., 2020).
Natural Language Inference	Determining if a hypothesis entails or contradicts a premise (MacCartney, 2009).
Fact Checking	Classifying if a claim is either true or false based on evidence (Vargo et al., 2018; Gurrapu et al., 2022).
Sentiment Analysis	Quantifying whether the textual data has a positive, negative, or neutral emotion (Zhang et al., 2018).
Text Classification	Categorizing textual data by automatically assigning labels (Minaee et al., 2021).
Neural Machine Translation	Translating languages using deep neural networks (Sutskever et al., 2014).

4.1. Machine reading comprehension

MRC enables a model to answer questions regarding a given context (Baradaran et al., 2022). For this reason, it also frequently referred to as Question Answering (QA) Systems. As shown in Table 4, for MRC applications, we found five recent articles from which three articles provide novel datasets (Ling et al., 2017; Mihaylov et al., 2018; Xie et al., 2020) and the remaining articles Sharp et al. (2017) and Lakhota et al. (2021) each propose a new MRC framework.

The first article, published in 2018, presented a new question-answering dataset based on the open book exam environment for elementary-level science—OpenBookQA (Mihaylov et al., 2018). This dataset consists of two components—(i) Questions (Q): a set of 5,958 multiple choice questions and (ii) Facts (F): a set of 1,326 diverse facts about elementary level science. This dataset was further tested for evaluating the performance of existing QA systems and then compared with the human performance. The results indicated that human performance was close to 92%, but many existing QA systems showed poor performance close to the random guessing baseline of 25%. Additionally, the authors found that simple neural networks achieved an accuracy of about 50%, but it is still not close to the human performance, about 92%. Recently an extension of the WorldTree project (Jansen, 2018), i.e., WorldTree V2 (Xie et al., 2020), is presented. The main goal of this project is to generate a science domain explanation with a supporting semi-structured knowledge base. The WorldTree project is a part of

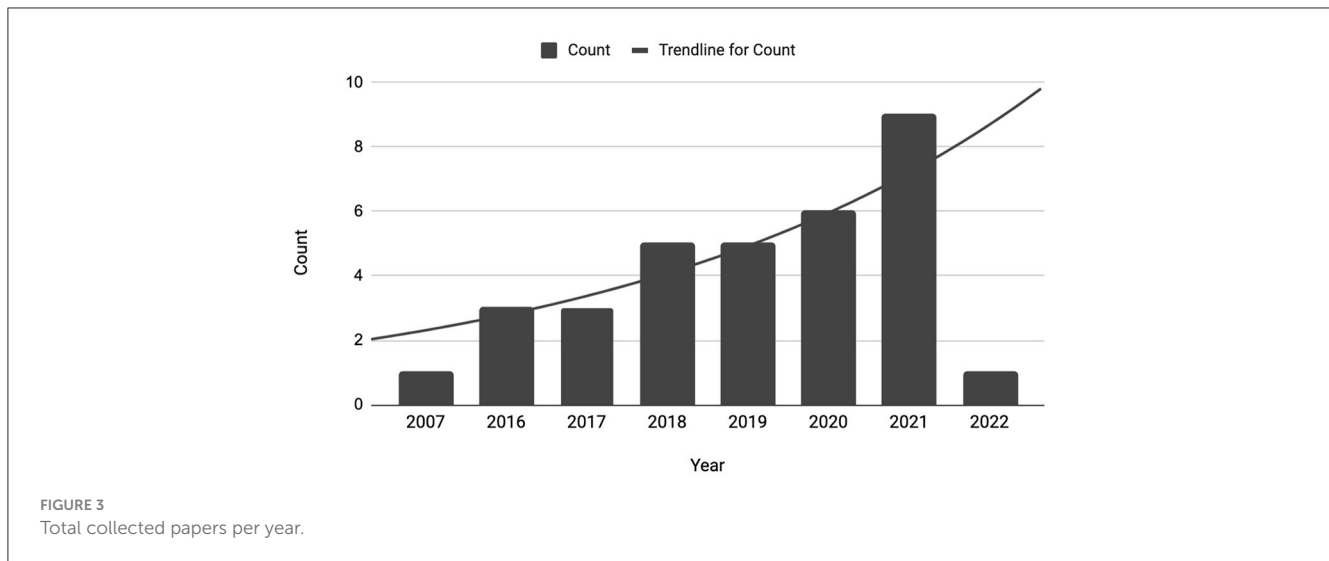


TABLE 4 Selected Machine Reading Comprehension papers.

References	Name	Year	Explanation	Models	XAI Metric	Dataset	Code
Sharp et al. (2017)	-	2017	Extractive	TF-IDF, FFNN	-	AI2 science, Aristo Mini	-
Ling et al. (2017)	-	2017	Extractive	LSTM, Seq2Seq	-	AQuA	✓
Mihaylov et al. (2018)	OpenBookQA	2018	Abstractive	BiLSTM Max-out	-	OpenBookQA,	✓
Xie et al. (2020)	WorldTree V2	2018	Abstractive	TF-IDF, BERT	-	WorldTree V2	✓
Lakhotia et al. (2021)	FiD-Ex	2021	Extractive	T5, BERT-to-BERT	-	Natural Questions	-

explainable question answering tasks that provide answers to natural language questions and their human-readable explanations for why the answers are correct. Xie et al. (2020) notes that most multi-hop inference models could not demonstrate combining more than two or three facts to perform inference. However, here the authors merge, on average, six facts from a semi-structured knowledge base of 9216 facts. Thus, this resulted in the WorldTree V2 corpus for standardized science questions. This corpus consists of 5100 detailed explanations to support training and instrumenting multi-hop inference question answering systems.

Lakhotia et al. (2021) demonstrates a new MRC framework called FiD-Ex (Extractive Fusion-in-Decoder). It has been noted that seq2seq (Sequence to Sequence) models work well at generating explanations and predictions together. However, these models require a large-labeled dataset for training and bring a host of challenges such as fabricating explanations for incorrect predictions, the difficulty of adapting to long input documents, etc. Thus, to tackle these challenges, the FiD-Ex framework includes sentence markers to encourage extractive explanations and intermediate fine-tuning for improving few-shot performance on open-domain QA datasets. This new framework is tested on ERASER (Evaluating Rationales And Simple English Reasoning) datasets and their benchmarks for evaluations (DeYoung et al., 2020). This experiment concludes that FiD-Ex significantly improves upon prior work on the explanation metrics and task accuracy on supervised and few-shot settings.

Sharp et al. (2017) proposes a new neural network architecture that re-ranks answer justifications as an intermediate step in answer selection. This new approach alternates between a max pooling layer and a shallow neural network (with ten nodes, gloriot uniform initializations, tanh activation, and L2 regularization of 0.1) for providing a justification. This approach contains three components: 1) retrieval component, which retrieves a pool of candidates' answer justification, 2) extractor, which extracts the features and 3) scores, which perform the scoring of the answer candidate based on the pool of justifications. The authors used 8th-grade science questions provided by Allen Institute for Artificial Intelligence (AI2) for evaluations. The training set includes 2,500 questions with four options, and the test set consists of 800 publicly released questions. Further, a pool of candidate justifications corpora containing 700k sentences from StudyStack and 25k sentences from Quizlet is used. The top 50 sentences were retrieved as a set of candidate justification. For model tuning, the authors used five-fold cross-validation, and during testing, the model architecture and hyperparameters were frozen. The authors compared results using two baselines: IR baseline and IR++. They concluded that this new approach showed better accuracy and justification quality while maintaining near state-of-the-art performance for the answer selection task.

Ling et al. (2017) presented a dataset and an approach that provides answer rationales, sequences of natural language, and human-readable mathematical expressions for solving algebraic word problems. The authors proposed a sequence-to-sequence model which generates a sequence of instructions and provides

the rationales after selecting the answer. For this purpose, a two-layer LSTM (Long Short-Term Memory) with a hidden size of 200 and word embedding with a size of 200 is utilized. Further, the authors also built a dataset containing 100,000 problems in which each question is decomposed into four parts - two inputs and two outputs. This new dataset is used for generating rationales for math problems and for understanding the quality of rationales as well as the ability to obtain a correct answer. Further, the authors used an attention-based sequence to sequence model as a baseline and compared results based on average sentence level perplexity and BLEU-4 (Bilingual Evaluation Understudy). The authors noted that this new approach could outperform the existing neural models in the ability to solve problems and the fluency of the generated rationales.

4.2. Commonsense reasoning

Commonsense knowledge helps humans navigate everyday situations. Similarly, commonsense reasoning in NLP is the ability for a model to go beyond pattern recognition and use world knowledge to make inferences (Apperly, 2011; Sap et al., 2020). As shown in Table 5 on commonsense reasoning, we found four articles, and all of them provide unique solutions that contribute to the development of commonsense reasoning frameworks.

Sap et al. (2020) demonstrates a solution for commonsense reasoning using LSTM encoder and decoder. The main goal was to convert the actions of an autonomous agent into natural language using neural machine translation. For this purpose, the authors built a corpus of thoughts of people as they complete tasks in the Frogger game which are then stored as states and actions. In the next step, LSTM encoder and decoder are used to translate actions as well as states into natural language. Lastly, the authors used the BLEU score to calculate sentence similarity and assessed the accuracy for selecting the best rationale. The authors also conducted a survey to evaluate the rationales based on human satisfaction. The Frogger experiment is concluded with Encoder-Decoder framework outperforming the baselines and demonstrates that the use of game theory approaches for generating rationales is a promising technique (Chang et al., 2019, 2020; Yu et al., 2019; Li et al., 2022).

Further, it is noted that deep learning model performance is poor when used in tasks that require commonsense reasoning due to limitations with available datasets. To tackle this problem, Rajani et al. (2019) developed the Commonsense Auto-Generated Explanations (CAGE) framework for generating explanations for Commonsense Question Answering (CQA). The authors also created a new dataset—Common Sense Explanations (CoS-E)—by collecting human explanations for commonsense reasoning and highlighting annotations. From this paper, the authors concluded that CAGE could be effectively used with pre-trained language models to increase commonsense reasoning performance.

Recently, Majumder et al. (2021) and Tang et al. (2021) presented novel solutions for commonsense reasoning. Majumder et al. (2021) focused on the Natural Language Explanations (NLEs), which are more detailed than Extractive rationales but fall short in terms of commonsense knowledge. In this solution, the

authors proposed a self-rationalizing framework RExC (Rationales, Explanations, and Commonsense). RExC first extracts rationales that act as features for the prediction then expands the extractive rationales using commonsense resources. In the last step, RExC selects the best suitable commonsense knowledge for generating NLEs and a final prediction. The authors tested RExC on five tasks - three natural language understanding tasks and two vision language understanding tasks. Overall, the results indicated improvement in the quality of extractive rationales and NLEs that bridges the gap between task performance and explainability. On the other hand, Tang et al. (2021) focused on Visual Commonsense Reasoning (VCR). They focused on a problem when a question with a corresponding input image is given to the system, and it attempts to predict an answer with a rationale statement as the justification. To explore this, author presented a multi-model approach by combining Computer Vision (CV) and NLP. Their approach leverages BERT and ResNet50 (Residual neural network) as the feature representation layer and BiLSTM (Bidirectional LSTM) and Attention for the multimodal feature fusion layer. These layers are then concatenated into an LSTM network for the encoder layer before passing into the classifier for the prediction layer. This was tested on the benchmark VCR dataset and it indicated significant improvements over existing methods and it also provided a more interpretable intuition into visual commonsense reasoning.

4.3. Natural Language Inference

Natural Language Inference (NLI) task helps with identifying a natural language hypothesis from a natural language premise (MacCartney, 2009). For this application, as shown in Table 6, we found three articles. The first article presents a new dataset—e-SNL (explanation-augmented Stanford Natural Language Inference) (Camburu et al., 2018)—and the other two articles discuss approaches that can improve NLI.

Camburu et al. (2018) extended the Stanford NLI (SNLI) (Bowman et al., 2015a) dataset by providing human-annotated explanations for the entailment relations. This new dataset—e-SNLI—is used in a series of classification experiments involving LSTM-based networks for understanding its usefulness for providing human-interpretable full-sentence explanations. The authors also evaluated these explanations as an additional training signal for improving sentence representation and transfer capabilities of out-of-domain NLI datasets. Thus, from these experiments, the authors conclude that e-SNLI can be used for various goals mentioned above and also be utilized for improving models as well as asserting their trust.

Another issue with NLI is the faithfulness of the generated explanations, tackled by Kumar and Talukdar (2020) and Wiegrefe et al. (2021). Kumar and Talukdar (2020) mentioned that existing methods do not provide a solution for understanding correlations of the explanations with the model's decision-making and this can affect the faithfulness of the generated explanations. Considering this problem, the authors proposed and presented a new framework - NILE (Natural language Inference over Label-specific Explanations). The NILE framework can generate natural language explanations for each possible decision and process these

TABLE 5 Selected commonsense reasoning papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Ehsan et al. (2018)	-	2018	Extractive	LSTM, Seq2Seq	-	-	-
Rajani et al. (2019)	CAGE	2019	Abstractive	GPT, BERT	-	CoS-E, CommonsenseQA	✓
Majumder et al. (2021)	RExC	2021	Extractive	Transformer	-	ComVE, e-SNLI, COSe, e-SNLI-VE, VCR	-
Tang et al. (2021)	DMVCR	2021	Extractive	LSTM, BERT	-	VCR	✓

TABLE 6 Selected Natural Language Inference papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Camburu et al. (2018)	e-SNLI	2018	Abstractive	BiLSTM, Seq2Seq	-	e-SNLI	✓
Kumar and Talukdar (2020)	NILE	2020	Abstractive	GPT-2, RoBERTa	-	e-SNLI	✓
Wiegrefe et al. (2021)	-	2020	Abstractive	T5	-	CoS-E, SNLI	✓

explanations to produce a final decision for the classification problems. To test this approach, the authors used two datasets - SNLI and e-SNLI - and compared NILE with baseline and other existing approaches based on explanation accuracy, in-domain evolution sets (SNLI), and on out-of-domain examples (train on SNLI and test on MNLI) ([Williams et al., 2018](#)). Based on the first 100 SNLI test samples, the results indicated that NILE variants are comparable with the ETPA (Explain Then Predict Attention) baseline, and NILE explanations generalize significantly better on out-of-domain examples. For out-of-domain examples (MNLI), results showed that the percentage of correct explanations in the subset of correct label predictions was significantly better for all the NILE variants. Thus, the authors concluded that NILE is an effective approach for accurately providing both labels and explanations. Further, [Kumar and Talukdar \(2020\)](#) also focused on the need for faithfulness for denoting the model's decision-making process by investigating abstractive rationales. The author proposed two measurements - robustness equivalence and feature importance agreement - to investigate the association of the labels and predicted rationales, which are required for a faithful explanation. This investigation was performed on CommonsenseQA ([Talmor et al., 2019](#)) and SNLI dataset using T5-based models ([Narang et al., 2020](#)). The results indicated that state-of-the-art T5-based join models demonstrate desirable properties and potential for producing faithful abstractive rationales.

4.4. Fact-checking

Fact-checking has become a popular application of NLP in recent years given its impact on assisting with misinformation and a majority of the work has been with claim verification ([Vargo et al., 2018](#); [Gurrapu et al., 2022](#)). Based on a paper published in 2016, there are 113 active fact-checking groups and 90 of which were established after 2010 ([Graves and Cherubini, 2016](#)). This indicates the growth of the fact-checking application. Considering the scope of this literature review, as shown in [Table 7](#), we found four articles on fact-checking. Two of the studies in this section present novel datasets, and the remaining two provide new techniques to improve fact-checking.

In 2017, a large dataset for the fact-checking community called LIAR ([Wang, 2017](#)) was introduced, including POLITIFACT data. Most works on this data were focused on using the claim and its speaker-related metadata to classify whether a verdict is true or false. The evidence—an integral part of any fact-checking process—was not part of the LIAR and was overlooked. Thus, in [Alhindi et al. \(2018\)](#) extended the LIAR dataset to LIAR-plus by including the evidence/justification. The authors treated the justification as a rationale for supporting and explaining the verdict. Further, they used Feature-based Machine Learning models (Logistic Regression and Support Vector Machine) and deep learning models (Bi-Directional Long Short-term Memory (BiLSTM) and Parallel-BiLSTM) for binary classification tasks to test the data. The results demonstrated a significant performance improvement in using the justification in conjunction with the claims and metadata. Further, [Hanselowski et al. \(2019\)](#) introduced a new corpus for training machine learning models for automated fact-checking. This new corpus is based on different sources (blogs, social media, news, etc.) and includes two granularity levels—the sources of the evidence and the stance of the evidence toward the claim—for claim identification. Authors then used this corpus to perform stance detection, evidence extraction and claim validation experiments. In these experiments, a combination of LSTMs, baseline NN, pre-trained models have been used, and their results are compared based on precision, recall, and F1 macro. The results indicated that fact-checking using heterogeneous data is challenging to classify claims correctly. Further, the author claims that the fact-checking problem defined by this new corpus is more difficult compared to other datasets and needs more elaborate approaches to achieve higher performance.

It has been noted that the fact-checking systems need appropriate explainability for the verdicts they predict. The justifications that are human-written can help to support and provide context for the verdicts, but they are tedious, unscalable, and expensive to produce ([Atanasova et al., 2020](#)). Considering this issue, [Atanasova et al. \(2020\)](#) proposed that the creation of the justifications needs to be automated to utilize them in a large-scale fact-checking system. The authors presented a novel method that automatically generates the justification from the claim's context and jointly models with veracity prediction. Further,

TABLE 7 Selected fact-checking papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Alhindi et al. (2018)	LIAR-PLUS	2018	Extractive	SVM, BiLSTM	-	LIAR-PLUS	✓
Hanselowski et al. (2019)	-	2019	Extractive	BERT	-	FEVER	✓
Atanasova et al. (2020)	-	2020	Extractive	DistilBERT	-	LIAR-PLUS	-
Rana et al. (2022)	RERRFACT	2022	Extractive	RoBERTa, BioBERT	-	SCIFACT	-

this new method is then tested on the LIAR dataset (Wang, 2017) for generating veracity explanations. The results indicated that this new method could combine predictions with veracity explanations, and manual evaluations reflected the improvement in the coverage and quality of the explanations. Another important domain in which fact-checking is useful is Science. Researching and providing substantial evidence to support or refute a scientific claim is not a straightforward task. It has been seen that scientific claim verification requires in-depth domain expertise along with tedious manual labor from experts to evaluate the credibility of a scientific claim. Considering this problem, Rana et al. (2022) proposed a new framework called RERRFACT (Reduced Evidence Retrieval Stage Representation) for classifying scientific claims by retrieving relevant abstracts and training a rationale-selection model. RERRFACT includes a two-step stance prediction that differentiates non-relevant rationales then identifies a claim's supporting and refuting rationales. This framework was tested on the SCI-FACT dataset (Wadden et al., 2020) and performed competitively against other language model benchmarks on the dataset leaderboard.

4.5. Sentiment analysis

Sentiment Analysis is a subset of the text classification field (Minaee et al., 2021). It focuses specifically on the "computational study of people's opinions, sentiments, emotions, appraisals, and attitudes toward entities such as products, services, organizations, individuals, issues, events, topics and their attributes" (Zhang et al., 2018). The use of rationales to support sentiment analysis models in NLP is widely used compared to other NLP tasks. For this task, as shown in Table 8, we identified five papers in this field.

In 2016, Lei et al. (2016) pioneered rationalization in sentiment analysis by proposing a problem: "prediction without justification has limited applicability". To make NLP outcomes more transparent, the authors propose an approach to extract input text which serves as justifications or rationales for a prediction. These are fragments from the input text which themselves are sufficient to make the same prediction. Their implementation approach includes a generator and an encoder architecture. The generator determines which can be potential candidates for a rationale from the input text. Those candidates are fed into the encoder to determine the prediction and the rationales are not provided during training. They employ an RCNN (Region-based Convolutional Neural Network) and an LSTM architecture and when compared with each other the RCNN performed better. The experiment was conducted on the BeerAdvocate dataset. The paper's approach outperforms attention-based baseline models.

They also demonstrate their approach on a Q&A retrieval task indicating that leveraging rationales for sentiment analysis tasks is very beneficial.

Similarly, Du et al. (2019) claim that explainability alone is not sufficient for a DNN (Deep Neural Network) to be viewed as credible unless the explanations align with established domain knowledge. In essence, only the correct evidences are to be used by the networks to justify predictions. In this paper, the authors define credible DNNs as models that provide explanations consistent with established knowledge. Their strategy is to use domain knowledge to improve DNNs credibility. The authors explore a specific type of domain knowledge called a rationale which are the salient features of the data. They propose an approach called CREX (Credible Explanation), which regularizes DNNs to use the appropriate evidence when making a decision for improved credibility and generalization capability. During training, instances are coupled with expert rationales and the DNN model is required to generate local explanations that conform to the rationales. They demonstrate it on three types of DNNs (CNN, LSTM, and self-attention model) and various datasets for testing. Results show that the CREX approach allows DNNs to look at the correct evidences rather than the specific bias in training dataset. Interestingly, they point that incorporating human knowledge does not always improve neural network performance unless the knowledge is very high quality.

Many papers published in the rationalization field indicate that a machine learning system learning with human provided explanations or "rationales" can improve its predictive accuracy (Zaidan et al., 2007). Strout et al. (2019) claim that this work hasn't been connected to the XAI field where machines attempt to explain their reasoning to humans. The authors attempt to show in their paper that rationales can improve machine explanations as evaluated by human judges. Although automated evaluation works, Strout et al. (2019) believe that since the explanations are for users, therefore humans should directly evaluate them. The experiment is done by using the movie reviews dataset and by having a supervised and an unsupervised CNN model for a text classification task. They use attention mechanism and treat the rationales as supervision in one of the CNN models. Results indicate that a supervised model trained on human rationales outperforms the unsupervised on predictions. The unsupervised is the model where the rationales/explanations are learned without any human annotations.

The selective rationalization mechanism is commonly used in complex neural networks which consist of two components—rationale generator and a predictor. This approach has a problem of model interlocking which arises when the predictor overfits to the features selected by the generator. To tackle this problem this paper proposes a new framework A2R which introduces

TABLE 8 Selected sentiment analysis papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Lei et al. (2016)	-	2016	Extractive	LSTM, RCNN	-	BeerAdvocate, AskUbuntu	✓
Du et al. (2019)	CREX	2019	Extractive	CNN, LSTM	-	BeerAdvocate, MovieReview	-
Strout et al. (2019)	-	2019	Extractive	RA-CNN, AT-CNN	-	MovieReview	-
Yu et al. (2021)	A2R	2021	Extractive	BiGRU	-	BeerAdvocate, MovieReview	✓
Antognini and Faltings (2021)	ConRAT	2021	Extractive	CNN, BiGRU	-	AmazonReviews, BeerAdvocate	-

a third component for soft attention into the architecture (Yu et al., 2021). The authors have used BeerAdvocate and MovieReview for understanding the effectiveness of the framework. The authors compared results from A2R with the original rationalization technique RNP (Rationalizing Neural Predictions) along with 3PLAYER, HARD-KUMA and BERT-RNP. For implementation authors have used bidirectional Gated Recurrent Units (GRU) in the generators and the predictors. Furthermore, they performed two synthetic experiments using BeerAdvocate dataset by deliberately inducing interlocking dynamics and then they performed experiments in real-world setting with BeerAdvocate and MovieReview. From the results they made two conclusions—(1) A2R showed consistent performance compared to other baselines on both the experiments, (2) A2R helps to promote trust and interpretable AI. In the future, the authors would like to improve A2R framework for generating casually corrected rationales to overcome the lack of inherent interpretability in the rationalization models.

Existing methods in rationalization compute an overall selection of input features without any specificity and this does not provide a complete explanation to support a prediction. Antognini and Faltings (2021) introduce ConRAT (Concept-based RATionalizer), a self-interpretable model which is inspired by human decision-making where key concepts are focused using the attention mechanism. The authors use the BeerReviews dataset to not only predict the review sentiment but also predict the rationales for key concepts in the review such as Mouthfeel, Aroma, Appearance, Taste, and Overall. ConRAT is divided into three submodels, a Concept Generator which finds the concepts in the review, a Concept Selector that determines the presence or absence of a concept, and a Predictor for final review predictions. ConRAT outperforms state-of-the-art methods while using only the overall sentiment label. However, Antognini et al. (2021) have further demonstrated that attention mechanism usage can contribute to a tradeoff between noisy rationales and a decrease in prediction accuracy.

4.6. Text classification

Text classification, also commonly known as text categorization, is the process of assigning labels or tags to textual data such as sentences, queries, paragraphs, and documents (Minaee et al., 2021). Classifying text and extracting insights can lead to a richer understanding of the data but due to their unstructured nature, it is challenging and tedious. NLP techniques

in text classification enable automatic annotation and labeling of data to make it easier to obtain those deeper insights of the data. For this task, as shown in Table 9, we identified four papers in this field.

Traditionally, rationales provide well-defined kinds of data to nudge the model on why a prediction is the way it is given the data. Moreover, they require little additional effort for annotators and yield a better predictive model. When classifying documents, it is beneficial to obtain sentence-level supervision in addition to document-level supervision when training new classifications systems (Zhang et al., 2016). Previous work relied on linear models such as SVMs (Support Vector Machines), therefore, Zhang et al. (2016) propose a novel CNN model for text classification that exploit associated rationales of documents. Their work claims to be the “first to incorporate rationales into neural models for text classification”. The authors propose a sentence-level CNN to estimate the probability that a sentence in a given document can be a rationale. They demonstrate that their technique outperforms baselines and CNN variants on five classification datasets. Their experimentation task uses Movie Reviews and the Risk of Bias (RoB) datasets. On the movie review dataset, their technique performs with a 90.43% accuracy with the RA-CNN (Recurrent Attention Convolutional Neural Network) model and similar strong results are also indicated on the RoB datasets.

It seems intuitive that more data or information can lead to better decision-making by the neural networks. Zaidan et al. (2007) propose a new framework to improve performance for supervised machine learning by using richer “kinds” of data. Their approach is called the “annotator rationales” technique and it is to leverage a training dataset with annotated rationales. The rationales highlight the evidence supporting the prediction. Zaidan et al. (2007) test their approach on text categorization tasks, specifically, sentiment classification of movie reviews and they claim that these rationales enable the machine to learn why the prediction is the way it is. Rationales help the model learn the signal from the noise. ML algorithms face the “credit-assignment problem” which means that many features in the data (X) could have affected the predicted result (Y). Rationales provide a “shortcut” to simplifying this problem since they provide hints on which features of X were important. Zaidan et al. (2007) used a discriminative SVM for experimentation and the results indicate that this technique significantly improves results for the sentiment classification and they hypothesize that leveraging rationales might be more useful than providing more training examples.

Recently, rationales have been a popular method in NLP to provide interpretability in the form of extracted subsets of texts. It is common to have spurious patterns and co-varying aspects in the

TABLE 9 Selected text classification papers.

References	Name	Year	Explanation	Models	XAI Metric	Dataset	Code
Zaidan et al. (2007)	-	2007	Extractive	SVM	-	MovieReview	-
Zhang et al. (2016)	-	2016	Extractive	SVM, RA-CNN	-	Risk of Bias	-
Liu et al. (2019b)	GEF	2019	Extractive	CNN, LSTM	-	PCMag Reviews, Skytrax User Reviews	✓
Plyler et al. (2022)	CDA	2021	Extractive	RL, RNN	-	TripAdvisor Reviews, RateBeer	✓

dataset due to which rationale selectors do not capture the desired relationship between input text and target labels. Considering this problem this paper proposes CDA (Counterfactual Data Augmentation) framework to aid rational models trained with Maximum Mutual Information (MMI) criteria (Plyler et al., 2022). CDA consists of transforms—for rational and classifications—because of their effectiveness over RNNs in NLP. The authors used [TripAdvisor.com](#) and RateBeer datasets for testing CDA with three baselines - MMI, FDA (Factual Data Augmentation), and ANT (simple substitution using antonyms). The results of the rational models were compared using precision and the accuracy of the classifier is reported based on the development set. From the results, authors concluded that the models trained using the CDA framework learn higher quality rationales and it doesn't need human intervention. In the future, the authors would like to explore more on counterfactual predictors and on CDA framework that could connect with other rationalization strategies. Similarly, Liu et al. (2019b) proposed a novel Generative Explanation Framework (GEF) for classification problems that can generate fine-grained explanations. The motivation behind this explanation framework is to provide human-readable explanations without ignoring fine-grained information such as textual explanations for the label. For understanding the accuracy of explanations, the authors conducted experiments on two datasets—PCMag and Skytrax User Reviews—which were processed by the Stanford Tokenizer. Further, the authors used Encoder-Predictor architecture in which they used Conditional Variational Autoencoder (CVAR) as a base model for text explanations and Long Short-Term Memory (LSTM) for numerical explanations. The experimental results indicated that after combining base models with GEF the performance of the base model was enhanced along with improving the quality of explanations. Further, the authors also used human evaluation for evaluating the explainability of the generated text explanations. The authors noted that for 57.62% of the tested items GEF provided better or equal explanations compared with the basic model.

4.7. Neural machine translation

With the advent of deep learning, Neural Machine Translation (NMT) became the successor to traditional translation methods such as Rule-based or Phrase-Based Statistical Machine Translation (PBSMT) (Yang et al., 2020). NMT models leverage Deep Neural Networks architecture to train the model end-to-end to improve translation quality and only require a fraction of the storage memory needed by PBSMT models (Sutskever et al., 2014). The

use of explanations to support NMT model's prediction is relatively new, however, there has been some pioneering work to provide more explainability. For this task, as shown in Table 10, we identified two relevant papers in this area.

Quality Estimation (QE) models perform well at analyzing the overall quality of translated sentences. However, determining translation errors is still a difficult task such as identifying which words are incorrect due to the limited amounts available training data. The authors explore the idea that since QE models depend on translation errors to predict the quality, using explanations or rationales extracted from these models can be used to better detect translation errors (Fomicheva et al., 2021). They propose a novel semi-supervised technique for word-level QE and demonstrate the QE task as a new benchmark for evaluating feature attribution (the interpretability of model explanations to humans). Instead of natural language explanations, their technique employs various feature attribution methods such as LIME, Integrated Gradients, Information Bottleneck, causal, and Attention. It was shown that explanations are useful and help improve model performance and provide better explainability.

Deep learning models are black-boxes because they involve a large number of parameters and complex architectures which makes them uninterpretable. Considering this problem and to bring interpretability in deep learning models (Alvarez-Melis and Jaakkola, 2017) propose a model-agnostic method for providing explanations. The explanations provided by this method consist of sets of inputs and output tokens that are causally related in the black-box model. Further, these causal relations are inferred by performing perturbations on the inputs from the black-box models, generating a graph of tokens, and then solving a partitioning problem to select the most relevant components. To test the method authors used a symmetric encoder-decoders consisting of recurrent neural networks with an intermediate variational layer. This method was tested for three applications - simple mappings, machine translation, and a dialogue system. For simple mapping, the authors used the CMU (Carnegie Mellon University) Dictionary of word pronunciations and evaluated inferred dependencies by randomly selecting 100 key-value pairs. For Machine Translation the authors used three black-boxes - Azure's Machine Translation system, Neural MY model, and human - for translating English to German. Finally, for the dialogue system, the authors used OpenSubtitle. From the results, the authors concluded that this model-agnostic method can produce reasonable, coherent, and often insightful expatiations. Additionally in future work, the authors noted that for Machine Translation and dialogue system applications potential improvements are needed for questioning seemingly correct predictions and explaining those that are not.

TABLE 10 Selected neural machine translation papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Alvarez-Melis and Jaakkola (2017)	SOCRAT	2017	Extractive	RNN	Attention Score	WMT14	✓
Fomicheva et al. (2021)	-	2021	Extractive	Transformer	LIME, Integrated Gradients	MLQE-PE	-

TABLE 11 Selected multiple domain papers.

References	Name	Year	Explanation	Models	XAI metric	Dataset	Code
Ribeiro et al. (2016)	LIME	2016	Extractive	SP-LIME, Parzen	-	Product Reviews	✓
Lakhotia et al. (2021)	ERASER	2020	Extractive	LSTM, BERT-to-BERT	Sufficiency	Movie Reviews, e-SNLI	✓
Das and Chernova (2020)	RGA	2020	Extractive	Utility Function	-	Stockfish	-
Sharma et al. (2020)	EPITOME	2020	Extractive	RoBERTa, Attention	-	Reddit, TalkLife	✓
Chan et al. (2021)	UNIREX	2021	Extractive	BigBird-Base	Comprehensiveness	SST, Movie Reviews, CoS-E	-
Zhang et al. (2021)	ExPred	2021	Extractive	MLP, GRU, BERT	-	FEVER, Movie Reviews	✓

4.8. Multiple domains

To demonstrate the effectiveness and generalizability of rationalization, many papers have attempted to demonstrate the use of rationales in multiple NLP tasks ([DeYoung et al., 2020](#); [Sharma et al., 2020](#)) or in conjunction with other disciplines such as [Sharma et al. \(2020\)](#). In this section, as shown in [Table 11](#), we present six papers with work in more than one NLP task or if the work was in another discipline but leveraged rationalization.

Currently in NLP many state-of-the-art tasks use deep neural networks and [DeYoung et al. \(2020\)](#) claim they are opaque in terms of their interpretability, or the way they make predictions. Lots of work has been conducted in this area, however, there is no standardization. The work has been with different datasets, NLP techniques and tasks which all have different aims and success metrics and this creates a challenge in this field of research in terms of tracking progress. To mitigate this, the authors propose a new benchmark called Evaluating Rationales And Simple English Reasoning (ERASER). There are multiple datasets (seven total) for various NLP tasks included in the benchmark. Datasets include human annotations of rationales which are the supporting evidence for a task's prediction. This is an extractive rationalization technique. For example, ERASER includes the Movie Reviews dataset for sentiment classification and each review has a rationale or an annotated sentence that supports the prediction for that review. In addition, metrics are also provided as a baseline benchmark to evaluate the extract rationales quality. The authors believe that this benchmark will facilitate in creating better interpretable NLP architectures.

It is important to understand the reasons behind the predictions for assessing trust which is important for making decision or deploying a new model. Considering this problem, [Ribeiro et al. \(2016\)](#) have proposed a novel model-agnostic approach LIME for providing explanations from any classifier about a local prediction and SP-LIME for providing global view of the model. For understanding the effectiveness of these methods simulated user experiment and evaluations with human subjects are performed. In the simulated user experiment two sentiment analysis datasets (books and DVDs) were used and

different classifiers such as - decision trees, logistic regression (L2 regularization), nearest neighbors, and SVM with RBF (Radial Basis Function) kernel were also used. To explain individual predictions authors compared LIME with parzen, greedy and random procedures. Further, evaluation with human subjects is performed using Amazon Mechanical Turk to estimate the real-world performance by creating a new religion dataset. Based on the results the authors concluded that LIME explanations are faithful to the models and helps in assessing trust in individual predictions. Further, SP-selected LIME explanations are good indicators of generalization which is validated via evaluations with human subjects. In future authors would like to explore how to perform a pick step for images and would like to explore different applications in speech, video, and medical domains. Additionally, they also would like to explore theoretical properties and computational optimizations for providing accurate, real-time explanations which are needed in any human-in-the-loop systems.

Rational extractions should be faithful, plausible, data-efficient, and fast with maintaining good performance but existing rational extractors are ignoring one or more of these aspects. Considering this challenge ([Chan et al., 2021](#)) propose UNIREX (Unified Learning Framework for Rationale Extraction) an end-to-end rational extractor that accounts for all these mentioned aspects. For understating UNIREX's effectiveness three text classification detests—SST (Stanford Sentiment Treebank), CoS-E (Commonsense Explanations), Movies—were used and compared with several baselines such as Vanilla, SGT (Sequence Graph Transform), SGT+P, FRESH etc. The authors used BigBird-base model in all their experiments which was pre-trained from the Hugging Face Transformers library. For training the authors used a learning rate of $2e-5$ and batch size of 32 with maximum 10 epochs. For evaluating faithfulness, the authors used comprehensiveness (Comp) and sufficiency (Suff) for $K = [1, 5, 10, 20, 50]$ which combined into a single metric Comp-Suff-Difference (CSD). For plausibility authors measured similarity to gold rationales using AUPRC (Area Under the Precision-Recall Curve), AP (Average Precision) and TF1 (Token F1). For data efficiency measurement how, performance varies with the percentage of train instances with gold rationale supervision is used. For speed authors compared

methods via runtime complexity analysis with respect to the number of LM (Linear Model) forward/backward passes. Finally for performance standard dataset-specific metrics accuracy (for SST and CoS-E) and macro-averaged F1 (for Movies) are used. The results indicated that UNIREX allows effective trade-off between performance, faithfulness, and plausibility to identify better rationale extractions. Further, authors also mention UNIREX trained rationale extractors can generalize to unseen datasets and tasks.

With many ML systems demonstrating performance beyond that of human across many applications, the field of XAI is advancing techniques to improve transparency and interpretability. [Das and Chernova \(2020\)](#) explores XAI in the context of a question previously unexplored in ML and XAI communities: “Given a computational system whose performance exceeds that of its human user, can explainable AI capabilities be leveraged to improve the performance of the human?”. The authors investigate this question through the game of Chess where computational game engines performance surpass the average player. They present an automated technique for generating rationales for utility-based computational methods called the Rationale-Generating Algorithm. They evaluate this with a user study against two baselines and their findings show that the machine generated rationales can lead to significant improvement in human task performance. They demonstrate that rationales can not only be used to explain the system’s actions but also instruct the user to improve their performance.

[Sharma et al. \(2020\)](#) explores an application of rationalization in the mental health support field and understanding empathy. Empathy is important for mental health support and with the rise of text-based internet platforms, it becomes crucial to understanding empathy in only communication. The paper presents a computational approach to understanding empathy by developing a corpus of 10,000 pairs of posts and responses with supporting rationales as evidence. They use a multi-task RoBERTa-based bi-encoder model to identify empathy in conversations and extract rationales for predictions. Their results demonstrate that their approach can effectively identify empathic conversations.

To improve interpretability for NLP tasks, recent rationalization techniques include Explain-then-Predict models. In this technique, an extractive explanation from the input text is generated and then a prediction is generated. However, these models do not use the rationales appropriately and consider the task input as simply a signal to learn and extract rationales. [Zhang et al. \(2021\)](#) propose a novel technique to prevent this problem with their approach called ExPred where they leverage multi-task learning on the explanation phase and embed a prediction network on the extracted explanations to improve task performance. They experiment with three datasets [Movie Reviews, FEVER (Fact Extraction and VERification), MultiRC] and conclude that their model significantly outperforms existing methods.

5. Extractive and abstractive methods

This section compares extractive and abstractive rationalization techniques. It can be observed from [Figure 4](#) that there is more interest and focus on extractive rationalization techniques compared to abstractive. There are multiple reasons for this, and

the progress in the Automatic Text Summarization (ATS) domain can help explain.

5.1. Extractive

In most extractive rationalization approaches, generating a rationale is similar to text summarization. These rationales contain the salient features of the input text, which users need to understand as the most influenced features of the model’s prediction.

Next, two steps are implemented while performing the task—(i) irrelevant information is ignored, and (ii) most crucial information is selected based on a scoring algorithm. This approach is a common foundation of summarization techniques. In extractive summarization, meaningful sentences are extracted to form a summary of the original text while still retaining the overall subject matter ([El-Kassas et al., 2021](#)). The critical difference with rationalization is that it is able to justify a neural network’s prediction with evidence. In a way, extractive rationalization uses extractive summarization’s fundamentals and takes it further. It frames the task as can we rationalize the output prediction where rationalize means to understand the prediction process and reason with supporting evidence. This introduces an interdependent relationship between the rationale and the prediction. This process is close to how humans rationalize with a sequence of reasons to justify a decision. This can be implemented in the NLP process to make models more explainable.

As interest in ATS systems grew in the past few decades, researchers have mainly focused on extractive summarization due to its simplicity, and reliability ([El-Kassas et al., 2021](#)). The abstractive summarization needed reliable natural language generation; thus, it was in its infancy from the 2000s to the early 2010s. Therefore, an increasing body of knowledge on extractive techniques is available, which researchers interested in rationalization could leverage and build on. This intuition behind extractive summarization paves the way for extractive rationalization. The stark difference between extractive and abstractive in [Figure 3](#) is expected and reasonable, and the fields of summarization and rationalization follow similar paths. However, summarization approaches should purely be used for inspiration - following the identical methods for rationalization would be insufficient, and it does not provide reliable model explainability. [Chan et al. \(2021\)](#) notes that for appropriate explainability, the desiderata for the rationale is that—(i) it must reflect the model’s reasoning process (faithfulness), (ii) be convincing to the nonexpert (plausibility), and (iii) the rationale extraction should not hurt task performance. Thus, there is more work than simply extracting sentences as rationales. Moreover, extractive rationalization is insufficient because extracted sentences themselves are insufficient to provide full explainability. Humans do not fully understand without context and a coherent and logical explanation.

5.2. Abstractive

The extensive research in extractive summarization reached its maturity, has peaked in terms of performance, and now the progress is stagnated ([Gupta and Gupta, 2019](#)). Recent advances

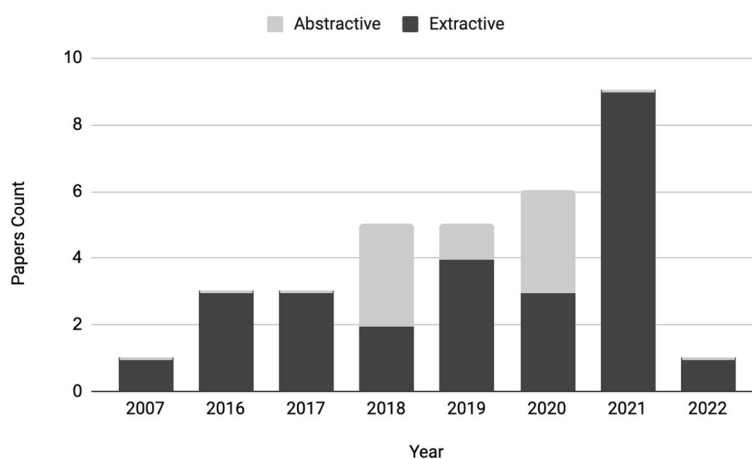


FIGURE 4
Total paper count for extractive and abstractive methods.

in deep learning and the advent of the Transformer architecture in 2017 have led to more reliable and influential language models (Vaswani et al., 2017; Devlin et al., 2019). In 2019, Liu and Lapata (2019) demonstrated a BERT-based abstractive summarization model that outperforms most non-Transformer-based models. Their model achieved state-of-the-art (SOTA) in automatic and human-based evaluations for summarization. Abstractive techniques allowed novel words and phrases to be generated instead of extracting spans from the input. Due to these advances, the research focuses gradually shifted from extractive to abstractive summarization. It is expected that rationalization will follow a similar trend.

Abstractive rationalization is still relatively new, with limited research available. However, there have been promising and pioneering approaches such as Rajani et al. (2019) and Kumar and Talukdar (2020). Almost every paper discussed with an abstractive rationalization technique in Section 4 leveraged some implementation of the Transformer architecture, such as BERT, GPT-2 (Radford et al., 2019), and T5, amongst others. BERT was the most frequently used language model. When BERT was released in 2018, it achieved SOTA results on many NLP tasks and surpassed human performance on tasks such as question answering and commonsense reasoning (Devlin et al., 2019). It made a giant leap in terms of performance compared to other language models of its time. This led to wide adoption and variations of BERT for the tasks where the Transformer-based model was required. Recently introduced models such as BART (Lewis et al., 2020), GPT-3 (Brown et al., 2020), and T5 demonstrate promising results and surpass BERT in some tasks. This is due to language models growing exponentially, and they continue to improve and perform incredibly well at natural language generation (Sanh et al., 2019). For example, in some cases, text produced by GPT-3 is almost on par if not better, than human-written text. This enables more opportunities for research in abstractive rationalization, which is needed. By leveraging SOTA language models, explanations can become more comprehensive and convincing when illustrating a model's decision-making process. As mentioned in Section 2, it is

almost as if the models are “talking for themselves”. We believe that significant progress can be made in rationalization by focusing more on improving abstractive techniques.

6. Discussions

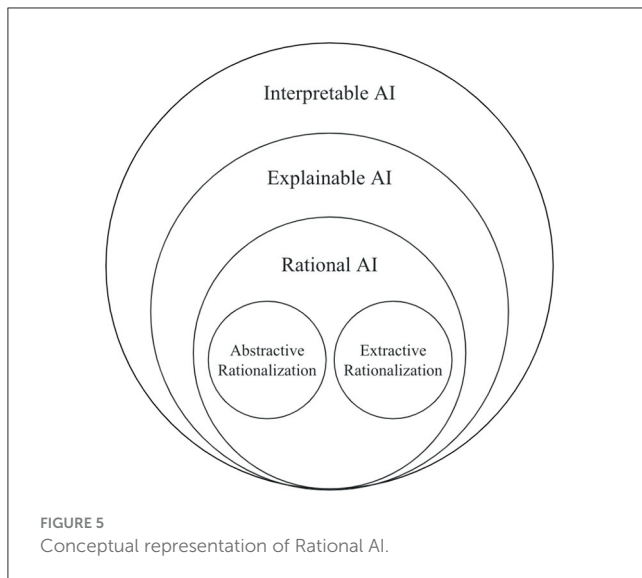
In this section, we discuss insights from the literature reviewed, challenges, and potential future directions to propel progress on rationalization. Most importantly, we introduce a new XAI subfield called Rational AI.

6.1. Introducing Rational AI

In Section 1, we have seen the need for explainability and the available methods in NLP. The numerical methods, such as SHAP values (Lundberg and Lee, 2017b) or Attention scores, visualization methods, such as LIME (Ribeiro et al., 2016), and saliency heatmaps, all require specialized domain knowledge to understand. At the same time, with increasing interactions with NLP-based systems, the nonexpert also deserves to know and understand how these black-box systems work because it has some degree of influence on their lives. This is formally called the *right to an explanation*, a right to receive an explanation for an algorithm's output (Miller, 2019b). A classic example is a bank system with an NLP model that automatically denies a loan application. In this situation, providing the loan applicant with SHAP values or saliency heatmaps to justify the bank's algorithms is not very meaningful. Thus, explainability methods are truly explainable and helpful if the nonexpert can understand them (Mittelstadt et al., 2019). We introduce Rational AI (RAI) as a potential solution.

6.1.1. Rational AI

Rationalization techniques come the closest to this goal because they are built on natural language explanations (NLEs). NLEs



are intuitive and human comprehensible because they are simply descriptive text. The textual information can be easily understood and translated into other languages if needed. Across all of the NLP tasks discussed in Section 4, we have seen the benefits of NLEs and the accessibility it provides to the nonexpert. We believe there is a critical need to focus on explainability techniques with NLEs. Considering these factors, we propose a new subfield in Explainable AI called Rational AI as shown in Figure 5. We define Rational AI as follows.

Rational AI: A field of methods that enable a black-box system to rationalize and produce a natural language explanation (rationale) to justify its output.

Rationality is the process of applying RAI to make models more explainable through an NLE. This is similar to the relationship between explainability and XAI. Further, rationality should not be confused or used interchangeably with the general AI term of a rational agent Russell and Norvig (2002). These are distinct topics with similar names. In this survey, RAI and rationality are purely considered in the context of model explainability. We also have not seen any usage or previous definitions of RAI within this context.

We compare rationality to the other fields shown in Figure 5. Models with interpretability are interpretable, while those with explainability are interpretable and complete, as described in Section 1. Models with rationality are interpretable and complete and can rationalize their behavior through an NLE.

The explainability methods described earlier in this subsection explain, but they do not justify in a way that is accessible and comprehensible to the nonexpert. In recent years, language models have become powerful and incredibly good at language generation tasks, but we have yet to see their full potential. As they continue to grow exponentially, we predict this is the beginning of explainability techniques using NLEs. The intuition behind RAI is that rationalization is one such technique, and many are yet to be developed. This calls for a more organized field to improve research focus and the need for RAI to exist.

6.1.2. Generalizing RAI

Although RAI arises from the need for better explainability for NLP tasks, it is potentially applicable in general AI and other fields in AI. Other fields, such as Computer Vision, Speech, and Robotics, could leverage rationalization methods to improve their model explainability. For example, rationalization in Computer Vision can help explain through an NLE which visual features contributed the most to an image classifier prediction in place of complex explainable techniques (Sundararajan et al., 2017; Tjoa and Guan, 2021). Many promising opportunities exist for researchers to apply rationalization in other disciplines.

6.2. Challenges

We have seen that rationalization is a relatively new technique, and with it, various challenges exist. In this subsection, we share challenges and potential solutions to improve the current state.

6.2.1. Statistical evaluations

No standard statistical evaluations exist currently for rationalization. There is a wide variety of metrics that are in use, such as Mean Squared Error (Lei et al., 2016), Accuracy (Zaidan et al., 2007; Du et al., 2019; Rajani et al., 2019), F1 Score (Alhindi et al., 2018; Rana et al., 2022), ANOVA (Analysis of variance) (Das and Chernova, 2020), and Precision (Plyler et al., 2022). We have observed that the most preferred statistical metric is accuracy. It is reasonable for evaluation metrics to be task-dependent and focused on the prediction. However, those alone are insufficient because the accuracy of the NLE also needs to be considered. For example, if the task prediction had high accuracy, but the NLE was unclear and incomprehensible, then it is not helpful. Metrics such as the BLEU (BiLingual Evaluation Understudy) score by Papineni et al. (2002) and the ROUGE-N (Recall-Oriented Understudy for Gisting Evaluation) score by Lin (2004) exist for evaluating open-ended machine-generated texts. However, we have seen limited use in the literature review, such as Camburu et al. (2018). The scores work by comparing the generated text with a set of ground-truth reference texts, and often these are human-written references. These scores are helpful, especially for abstractive rationalization, where explanations can be open-ended; however, they come with limitations since the evaluation is effectively token-level matching. Since an NLE is the typical outcome of systems with rationalization, adopting a standard evaluation metric can help improve research progress. Consistent evaluations also make it easier to compare different experiments and approaches.

6.2.2. Data

The availability and the need for more diversity of appropriate datasets is also a problem hindering progress.

Availability: Data collection is an expensive and time-consuming task. It is possible to repurpose existing datasets, but modifying them requires manual human labor. Thus, researchers often build their datasets for a specific task they are working on. Camburu et al. (2018) developed the e-SNLI dataset by modifying the SNLI dataset from Bowman et al. (2015b). Camburu et al.

TABLE 12 Human-Centered Evaluations (HCE) are performed in 15 out of 33 papers surveyed in this review.

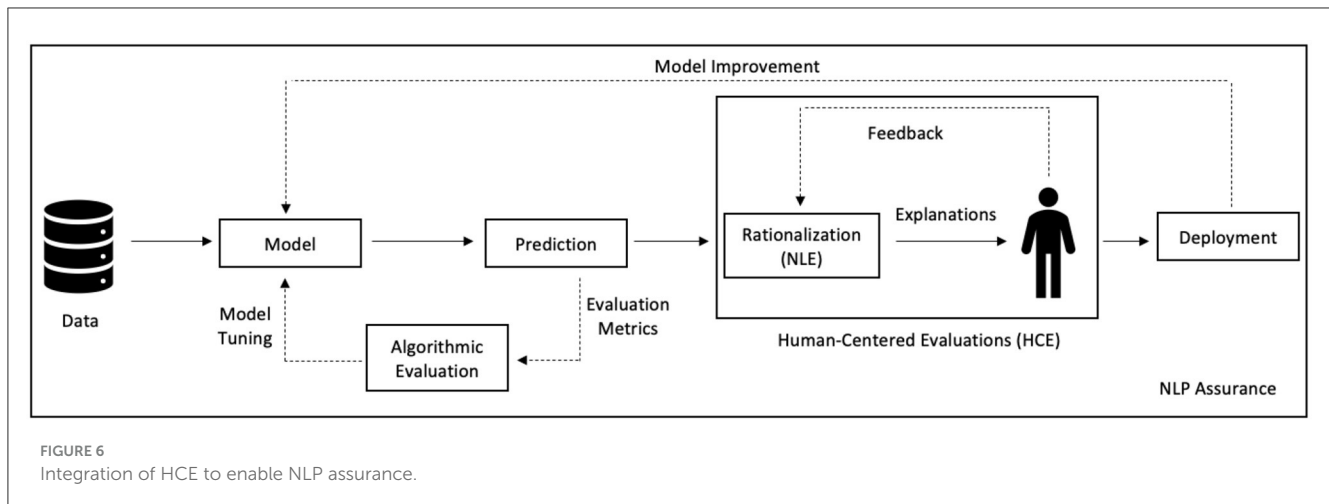
Year	References	Method name	Venue	NLP domain	HCE performed?
2007	Zaidan et al. (2007)	-	NAACL	Text Classification	No
2016	Ribeiro et al. (2016)	LIME	NAACL	Multiple Domain	Yes
2016	Lei et al. (2016)	-	EMNLP	Sentiment Analysis	No
2016	Zhang et al. (2016)	-	EMNLP	Text Classification	No
2017	Sharp et al. (2017)	-	CoNLL	Machine Reading Comprehension	No
2017	Ling et al. (2017)	-	ACL	Machine Reading Comprehension	No
2017	Alvarez-Melis and Jaakkola (2017)	SOCRAT	EMNLP	Neural Machine Translation	Yes
2018	Ehsan et al. (2018)	-	AIES	Commonsense Reasoning	Yes
2018	Alhindi et al. (2018)	LIAR-PLUS	FEVER	Fact-Checking	No
2018	Mihaylov et al. (2018)	OpenBookQA	EMNLP	Machine Reading Comprehension	Yes
2018	Xie et al. (2020)	WorldTree V2	LREC	Machine Reading Comprehension	Yes
2018	Camburu et al. (2018)	e-SNLI	NeurIPS	Natural Language Inference	No
2019	Rajani et al. (2019)	CAGE	ACL	Commonsense Reasoning	Yes
2019	Hanselowski et al. (2019)	-	CoNLL	Fact-Checking	No
2019	Du et al. (2019)	CREX	ICDM	Sentiment Analysis	No
2019	Strout et al. (2019)	-	BlackboxNLP	Sentiment Analysis	Yes
2019	Liu et al. (2019b)	GEF	ACL	Text Classification	Yes
2020	Atanasova et al. (2020)	-	ACL	Fact-Checking	Yes
2020	Lakhotia et al. (2021)	ERASER	ACL	Multiple Domain	Yes
2020	Das and Chernova (2020)	RGA	IUI	Multiple Domain	Yes
2020	Sharma et al. (2020)	EPITOME	EMNLP	Multiple Domain	No
2020	Kumar and Talukdar (2020)	NILE	ACL	Natural Language Inference	Yes
2020	Wiegrefe et al. (2021)	-	EMNLP	Natural Language Inference	No
2021	Majumder et al. (2021)	RExC	ICLR	Commonsense Reasoning	Yes
2021	Tang et al. (2021)	DMVCR	DaWaK	Commonsense Reasoning	No
2021	Tang et al. (2021)	FiD-Ex	EMNLP	Machine Reading Comprehension	No
2021	Chan et al. (2021)	UNIREX	ICML	Multiple Domain	No
2021	Zhang et al. (2021)	EXPred	WSDM	Multiple Domain	No
2021	Fomicheva et al. (2021)	-	ACL	Neural Machine Translation	No
2021	Yu et al. (2021)	A2R	NeurIPS	Sentiment Analysis	Yes
2021	Antognini and Faltings (2021)	ConRAT	ACL	Sentiment Analysis	Yes
2021	Plyler et al. (2022)	CDA	NeurIPS	Text Classification	No
2022	Atanasova et al. (2020)	RERRFACT	AAAI	Fact-Checking	No

(2018) achieved promising results on their task, demonstrating how their dataset can enable a wide range of new research directions by altering and repurposing existing datasets.

Diversity: Without enough datasets, new research in rationalization will be limited. Researchers will be constrained to the existing datasets to make new progress. This trend is evident in the literature reviewed in MRC and Sentiment Analysis compared to NMT. In MRC, the datasets are very diverse. In sentiment analysis, most papers rely on either the BeerAdvocate (McAuley et al., 2012) or MovieReviews (Maas et al., 2011) datasets

to perform their experiments. In both domains, we discovered five publications each. For a domain such as NMT, progress seems limited, and we found only two publications. The lack of appropriate rationalization datasets for NMT tasks is a possible reason for this.

As we observed in our literature review, there is a direct relationship between dataset availability and the progress made. More work in creating new datasets for rationalization can help improve diversity and the progress of certain domains lagging behind, such as NMT. New datasets across all domains, in general,



will increase the interest and work in rationalization because researchers will have more flexibility in designing new techniques and experimenting with a wide variety of data. [Stamper et al. \(2010\)](#) has organized the largest repository of learning science datasets called DataShop, and it led to improvements in research progress. Similarly, an organized central repository for rationalization supporting datasets can be beneficial. Without a centralized model evaluation and development system, reproducibility and accessibility will remain low.

6.3. Human-centered evaluations and assurance

NLP has direct applications in many disciplines. For example, MRC and commonsense reasoning are helpful in the education discipline. Our literature review indicates using Q&A tools and commonsense injection to generate explanations for educational needs ([Mihaylov et al., 2018](#); [Li et al., 2019](#)). Further, NLP has also been used to enhance human task performance, as we saw in [Das and Chernova \(2020\)](#), and to provide support for mental health ([Sharma et al., 2020](#)). Additionally, fact-checking is another application, and it is crucial in social media, fake news detection, and law ([Alhindi et al., 2018](#)). It has become common to interact with these systems, and they may have a significant influence on all aspects of our society. Due to this, the European Union recently passed a regulation that requires algorithms to provide explanations that can significantly affect users based on their user-level predictions ([Doshi-Velez and Kim, 2017](#)).

6.3.1. Human-centered evaluations (HCE)

The explanations provided by the NLP systems must provide enough information to the user to help them understand its decision-making process ([Putnam and Conati, 2019](#)). Considering these aspects, the human-machine partnership is essential for evaluating and generating accurate explanations. This calls for better methods to evaluate the explanations generated. The field of HCE addresses this problem, and [Sperrle et al. \(2021\)](#) defines

it as a “field of research that considers humans and machines as equally important actors in the design, training, and evaluation of co-adaptive machine learning scenarios.”

In this literature survey, we found 15 out of 33 papers in which HCE is performed, and a summary is provided in [Table 12](#). [Sperrle et al. \(2021\)](#) shares the increasing trend of HCE since 2017 compared to the previous years. While conducting this literature survey, this trend was not observed in the rationalization domain. Overall, we found that HCE is incorporated in most of the papers on Machine Reading Comprehension (2 out of 5), Commonsense Reasoning (3 out of 4), Fact-Checking (1 out of 4), Natural Language Inference (2 out of 5), Neural Machine Translation (1 out of 2), Sentiment Analysis (3 out of 5) and Multiple Domain (3 out of 6). From our observations, researchers give more attention to performance while evaluating AI algorithms and ignore human factors such as usability, user intentions, and user experience. Thus, along with the accuracy of AI algorithms, it is also essential to focus on the interpretability and reliability of the explanations generated by AI algorithms. The articles in which HCE is used are primarily performed via crowdsourcing using Amazon Mechanical Turk, and the focus is on user-based evaluations or annotations. This pattern necessitates conducting expert evaluations to understand users’ needs better because it can help improve trust in AI algorithms.

HCE is a subset of the Human-Computer Interaction (HCI) field, which is integrated with the AI paradigm after the algorithmic performance evaluations as shown in [Figure 6](#). This integration can be regarded as human-centered AI, and [Riedl \(2019\)](#) claims this as an AI/ML perspective that intelligent systems are part of a more extensive system that also includes human stakeholders. The literature ([Hwang and Salvendy, 2010](#)) on usability, and user experience testing demonstrated three widely used methods to perform HCE - Think Aloud (TA), Heuristic Evaluation (HE), and Cognitive Walkthrough (CW). The TA method is a standard method and can be more effective considering the evaluations of explanations in the NLP domain. In the TA method, evaluators are asked to “think aloud” about their experience while an experimenter observes them and listens to their thoughts ([Fox, 2015](#)). This way, an HCE method can be used in the final step to understand usability, user intentions, and user experience. This may lead to a better understanding of the interpretability and reliability of the

explanations generated by rationalization. Therefore, in addition to statistical evaluation techniques, we strongly encourage researchers to integrate HCE as part of their evaluations.

6.3.2. Assurance

It is critical to perform rigorous testing and validation of NLP systems at all stages before their deployment. For example, it should be ensured that the data is unbiased, models are interpretable, and the process of arriving at the outcome is explainable to a nonexpert. In the last step of this process, it would be beneficial to use RAI techniques. Integrating rationalization with human-centered evaluations and elements of NLP Assurance can invoke human-AI trust and safety with the systems - with the recent rise of chatbots such as ChatGPT, the need for more rigorous validation is more important than any other time. This process may also transform black-box systems into white-box systems and make NLP models more comprehensible and accessible for nonexpert users.

Author contributions

SG developed the research work, managed the writing, reviewed, and designed the paper. AK provided reviews of multiple works in the field of NLP and helped with writing and editing the paper. LH and IL reviewed and edited the paper. FB led the overall design of the paper and edited and evaluated the manuscript's scientific merit and overall quality. All authors contributed to the article and approved the submitted version.

References

- Agrawal, A., Batra, D., and Parikh, D. (2016). "Analyzing the behavior of visual question answering models," in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, eds. J. Su, X. Carreras, and K. Duh. Austin, Texas: The Association for Computational Linguistics, 1955–1960.
- Alhindi, T., Petridis, S., and Muresan, S. (2018). "Where is your evidence: Improving fact-checking by justification modeling," in *Proceedings of the First Workshop on Fact Extraction and VERification (FEVER)*. Brussels, Belgium: Association for Computational Linguistics, 85–90.
- Alvarez-Melis, D., and Jaakkola, T. (2017). "causal framework for explaining the predictions of black-box sequence-to-sequence models," in *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Copenhagen, Denmark: Association for Computational Linguistics, 412–421.
- Antognini, D., and Faltings, B. (2021). "Rationalization through concepts," in *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. Copenhagen, Denmark: Association for Computational Linguistics, 761–775.
- Antognini, D., Musat, C., and Faltings, B. (2021). "Multi-dimensional explanation of target variables from documents." *Proc. AAAI Conf. Artif. Intell.* 35, 12507–12515. doi: 10.1609/aaai.v35i14.17483
- Apperly, I. (2011). *Mindreaders the Cognitive Basis of Theory of Mind*. London: Psychology Press.
- Atanasova, P., Simonsen, J. G., Lioma, C., and Augenstein, I. (2020). "Generating fact checking explanations," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Copenhagen, Denmark: Association for Computational Linguistics, 7352–7364.
- Baradaran, R., Ghiasi, R., and Amirkhani, H. (2022). A survey on machine reading comprehension systems. *Nat. Lang. Eng.* 2022, 1–50. doi: 10.1017/S1351324921000395
- Bastings, J., Aziz, W., and Titov, I. (2019). "Interpretable neural predictions with differentiable binary variables," in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 2963–2977.
- Batarseh, F. A., Freeman, L., and Huang, C.-H. (2021). A survey on artificial intelligence assurance. *J. Big Data* 8, 7. doi: 10.1186/s40537-021-00445-7
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015a). "A large annotated corpus for learning natural language inference," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, 632–642.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. (2015b). A large annotated corpus for learning natural language inference. *arXiv*. doi: 10.18653/v1/D15-1075
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., et al. (2020). "Language models are few-shot learners," in *Advances in Neural Information Processing Systems, volume 33*, eds. H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin. New York: Curran Associates, Inc, 1877–1901.
- Camburu, O.-M., Rocktäschel, T., Lukasiewicz, T., and Blunsom, P. (2018). "e-snl: Natural language inference with natural language explanations," in *Advances in Neural Information Processing Systems*, eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett. New York: Curran Associates, Inc, 9539–9549.
- Chan, A., Sanjabi, M., Mathias, L., Tan, L., Nie, S., Peng, X., et al. (2021). Unirex: a unified learning framework for language model rationale extraction. *CoRR*. doi: 10.18653/v1/2022.bigsience-1.5
- Chang, S., Zhang, Y., Yu, M., and Jaakkola, T. (2020). "Invariant rationalization," in *Proceedings of the 37th International Conference on Machine Learning*. New York City: PMLR, 1448–1458.
- Chang, S., Zhang, Y., Yu, M., and Jaakkola, T. S. (2019). "A game theoretic approach to class-wise selective rationalization," in *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019*, eds. H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alch Buc, E. A. Fox, and R. Garnett. Vancouver, BC, Canada: NeurIPS, 10055–10065.
- Danilevsky, M., Qian, K., Aharonov, R., Katsis, Y., Kawas, B., and Sen, P. (2020). "A survey of the state of explainable AI for natural language processing," in *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational*

Funding

This work was funded by the Commonwealth Cyber Initiative (CCI).

Acknowledgments

We would like to thank the members of the A3 Lab (<https://ai.bse.vt.edu/>) at Virginia Tech for their inputs to the paper, discussions, and support.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Linguistics and the 10th International Joint Conference on Natural Language Processing. Suzhou, China: Association for Computational Linguistics, 447–459.
- Das, D., and Chernova, S. (2020). “Leveraging rationales to improve human task performance,” in *Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI '20*. New York, NY, USA: Association for Computing Machinery, 510–518.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Minneapolis, Minnesota: Association for Computational Linguistics, 4171–4186.
- DeYoung, J., Jain, S., Rajani, N. F., Lehman, E., Xiong, C., Socher, R., et al. (2020). “ERASER: A benchmark to evaluate rationalized NLP models,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Minneapolis, Minnesota: Association for Computational Linguistics, 4443–4458.
- Doshi-Velez, F., and Kim, B. (2017). Towards a rigorous science of interpretable machine learning. *arXiv*.
- Du, M., Liu, N., Yang, F., and Hu, X. (2019). “Learning credible deep neural networks with rationale regularization,” in *2019 IEEE International Conference on Data Mining (ICDM)*. Los Alamitos, CA, USA: IEEE Computer Society, 150–159.
- Ehsan, U., Harrison, B., Chan, L., and Riedl, M. O. (2018). “Rationalization: A neural machine translation approach to generating natural language explanations,” in *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society, AIES '18*. New York, NY, USA: Association for Computing Machinery, 81–87.
- El-Kassab, W. S., Salama, C. R., Rafea, A. A., and Mohamed, H. K. (2021). Automatic text summarization: A comprehensive survey. *Expert Syst. Appl.* 165:113679. doi: 10.1016/j.eswa.2020.113679
- Finlayson, S. G., Bowers, J. D., Ito, J., Zittrain, J. L., Beam, A. L., and Kohane, I. S. (2019). Adversarial attacks on medical machine learning. *Science* 363, 1287–1289. doi: 10.1126/science.aaw4399
- Fomicheva Specia, Aletas (2021). *Translation Error Detection As Rationale Extraction*. Dublin, Ireland: Association for Computational Linguistics.
- Fox, J. E. (2015). “The science of usability testing,” in *Proceedings of the 2015 Federal Committee on Statistical Methodology (FCSM) Research Conference*. Washington, DC, USA: Federal Committee on Statistical Methodology, 1–3.
- Gilpin, L. H., Bau, D., Yuan, B. Z., Bajwa, A., Specter, M., and Kagal, L. (2018). “Explaining explanations: an overview of interpretability of machine learning,” in *2018 IEEE 5th International Conference on Data Science and Advanced Analytics (DSAA)*, 80–89.
- Graves, L., and Cherubini, F. (2016). *The Rise of Fact-Checking Sites in Europe*. Oxford: Oxford University: Reuters Institute.
- Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., and Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Comput. Surv.* 51, 5. doi: 10.1145/3236009
- Gupta, S., and Gupta, S. K. (2019). Abstractive summarization: an overview of the state of the art. *Expert Syst. Appl.* 121, 49–65. doi: 10.1016/j.eswa.2018.12.011
- Gurrapu, S., Huang, L., and Batarseh, F. (2022). “Exclaim: Explainable neural claim verification using rationalization,” in *2022 IEEE 29th Annual Software Technology Conference (Maryland)*.
- Hanselowski, A., Stab, C., Schulz, C., Li, Z., and Gurevych, I. (2019). “A richly annotated corpus for different tasks in automated fact-checking,” in *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*. Hong Kong, China: Association for Computational Linguistics, 493–503.
- Hwang, W., and Salvendy, G. (2010). Number of people required for usability evaluation: the 10±2 rule. *Commun. ACM* 53, 130–133. doi: 10.1145/1735223.1735255
- Jansen, P. (2018). “Multi-hop inference for sentence-level TextGraphs: How challenging is meaningfully combining information for science question answering?,” in *Proceedings of the Twelfth Workshop on Graph-Based Methods for Natural Language Processing (TextGraphs-12)*. New Orleans, Louisiana, USA: Association for Computational Linguistics, 12–17.
- Kumar, S., and Talukdar, P. (2020). “NILE: Natural language inference with faithful natural language explanations,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 8730–8742.
- Lakhotia, K., Paranjape, B., Ghoshal, A., Yih, S., Mehdad, Y., and Iyer, S. (2021). “FiD-ex: Improving sequence-to-sequence models for extractive rationale generation,” in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Punta Cana, Dominican Republic: Association for Computational Linguistics, 3712–3727.
- Lei, T., Barzilay, R., and Jaakkola, T. (2016). “Rationalizing neural predictions,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, 107–117.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., et al. (2020). “BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 7871–7880.
- Li, S., Chen, J., and Yu, D. (2019). Teaching pretrained models with commonsense reasoning: A preliminary kb-based approach. *arXiv*.
- Li, S., Majumder, B. P., and McAuley, J. (2022). “Self-supervised bot play for transcript-free conversational recommendation with rationales,” in *Proceedings of the 16th ACM Conference on Recommender Systems, RecSys '22*. New York, NY, USA: Association for Computing Machinery, 327–337.
- Lin, C.-Y. (2004). “Rouge: A package for automatic evaluation of summaries,” in *Text Summarization Branches Out*, 74–81.
- Ling, W., Yogatama, D., Dyer, C., and Blunsom, P. (2017). “Program induction by rationale generation: Learning to solve and explain algebraic word problems,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Vancouver, Canada: Association for Computational Linguistics, 158–167.
- Liu, G., Hsu, T.-M. H., McDermott, M., Boag, W., Weng, W.-H., Szolovits, P., et al. (2019a). “Clinically accurate chest x-ray report generation,” in *Proceedings of the 4th Machine Learning for Healthcare Conference*. 249–269.
- Liu, H., Yin, Q., and Wang, W. Y. (2019b). “Towards explainable NLP: A generative explanation framework for text classification,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 5570–5581.
- Liu, Y., and Lapata, M. (2019). “Text summarization with pretrained encoders,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, 3730–3740.
- Lundberg, S. M., and Lee, S.-I. (2017a). A unified approach to interpreting model predictions. *Adv. Neural Inf. Process. Syst.* 30.
- Lundberg, S. M., and Lee, S.-I. (2017b). “A unified approach to interpreting model predictions,” in *Advances in Neural Information Processing Systems*, eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. Long Beach, California: NeurIPS 2017 Proceedings.
- Luo, S., Ivison, H., Han, C., and Poon, J. (2021). *Local Interpretations for Explainable Natural Language Processing: A Survey*.
- Maas, A. L., Daly, R. E., Pham, P. T., Huang, D., Ng, A. Y., and Potts, C. (2011). “Learning word vectors for sentiment analysis,” in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*. Portland, Oregon, USA: Association for Computational Linguistics, 142–150.
- MacCartney, B. (2009). *Natural Logic and Natural Language*. Stanford, CA: Stanford University.
- Madsen, A., Reddy, S., and Chandar, S. (2021). *Post-Hoc Interpretability for Neural nlp: A Survey*.
- Majumder, B., Camburu, O.-M., Lukaszewicz, T., and McAuley, J. (2021). *Rationale-Inspired Natural Language Explanations with commonsense*.
- McAuley, J., Leskovec, J., and Jurafsky, D. (2012). “Learning attitudes and attributes from multi-aspect reviews,” in *2012 IEEE 12th International Conference on Data Mining*. IEEE, 1020–1025.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. (2018). “Can a suit of armor conduct electricity? a new dataset for open book question answering,” in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Brussels, Belgium: Association for Computational Linguistics, 2381–2391.
- Miller, T. (2019a). Explanation in artificial intelligence: Insights from the social sciences. *Artif. Intell.* 267:1–38.
- Miller, T. (2019b). Explanation in artificial intelligence: insights from the social sciences. *Artif. Intell.* 267, 1–38. doi: 10.1016/j.artint.2018.07.007
- Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., and Gao, J. (2021). Deep learning-based text classification: a comprehensive review. *ACM Comput. Surv.* 54, 3. doi: 10.1145/3439726
- Mittelstadt, B., Russell, C., and Wachter, S. (2019). “Explaining explanations in AI,” in *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 279–288.
- Narang, S., Raffel, C., Lee, K., Roberts, A., Fiedel, N., and Malkan, K. (2020). Wt5?! training text-to-text models to explain their predictions. *ArXiv*.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). “Bleu: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 311–318.
- Plyler, M., Green, M., and Chi, M. (2022). *Making a (counterfactual) Difference One Rationale at a Time*.

- Putnam, V., and Conati, C. (2019). Exploring the need for explainable artificial intelligence (xai) in intelligent tutoring systems (its). *IUI Workshops*. 19, 1–7.
- Radford, A., and Narasimhan, K. (2018). *Improving Language Understanding by Generative Pre-Training*. San Francisco: OpenAI.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. (2019). *Language Models are Unsupervised Multitask Learners*.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., et al. (2020). Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* 21, 1–67.
- Rajani, N. F., McCann, B., Xiong, C., and Socher, R. (2019). “Explain yourself! leveraging language models for commonsense reasoning,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Florence, Italy: Association for Computational Linguistics, 4932–4942.
- Rana, A., Khanna, D., Singh, M., Ghosal, T., Singh, H., and Rana, P. (2022). *Rerrfact: Reduced Evidence Retrieval Representations for Scientific Claim Verification*.
- Ribeiro, M. T., Singh, S., and Guestrin, C. (2016). “Why should i trust you?: Explaining the predictions of any classifier,” in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. New York, NY, USA: Association for Computing Machinery, 1135–1144.
- Riedl, M. O. (2019). Human-centered artificial intelligence and machine learning. *Human Behav. Emerg. Technol.* 1, 33–36. doi: 10.1002/hbe2.117
- Russell, S., and Norvig, P. (2002). *Artificial Intelligence: A Modern Approach*.
- Sanh, V., Debut, L., Chaumond, J., and Wolf, T. (2019). Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*.
- Sap, M., Shwartz, V., Bosselut, A., Choi, Y., and Roth, D. (2020). “Commonsense reasoning for natural language processing,” in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Tutorial Abstracts*. Association for Computational Linguistics, 27–33.
- Sharma, A., Miner, A., Atkins, D., and Althoff, T. (2020). “A computational approach to understanding empathy expressed in text-based mental health support,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 5263–5276.
- Sharp, R., Surdeanu, M., Jansen, P., Valenzuela-Escárcega, M. A., Clark, P., and Hammond, M. (2017). “Tell me why: Using question answering as distant supervision for answer justification,” in *Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017)*. Vancouver, Canada: Association for Computational Linguistics, 69–79.
- Sperrle, F., El-Assady, M., Guo, G., Borgo, R., Chau, D. H., Endert, A., et al. (2021). *A Survey of Human-Centered Evaluations in Human-Centered Machine Learning*. Toms River, NJ: Computer Graphics Forum. doi: 10.1111/cgf.14329
- Stamper, J., Koedinger, K., d Baker, R. S., Skogsholm, A., Leber, B., Rankin, J., et al. (2010). “Pslc datashop: a data analysis service for the learning science community,” in *International Conference on Intelligent Tutoring Systems*. Cham: Springer, 455–455.
- Strout, J., Zhang, Y., and Mooney, R. J. (2019). *Do Human Rationales Improve Machine Explanations?*, 56–62.
- Sundararajan, M., Taly, A., and Yan, Q. (2017). “Axiomatic attribution for deep networks,” in *International Conference on Machine Learning*. New York City: PMLR, 3319–3328.
- Sutskever, I., Vinyals, O., and Le, Q. V. (2014). “Sequence to sequence learning with neural networks,” in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*. Cambridge, MA, USA: MIT Press, 3104–3112.
- Talmor, A., Herzig, J., Lourie, N., and Berant, J. (2019). “CommonsenseQA: A question answering challenge targeting commonsense knowledge,” in *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)*. Minneapolis, Minnesota: Association for Computational Linguistics, 4149–4158.
- Tang, X., Huang, X., Zhang, W., Child, T. B., Hu, Q., Liu, Z., et al. (2021). Cognitive visual commonsense reasoning using dynamic working memory. *CoRR*. doi: 10.1007/978-3-030-86534-4_7
- Thayaparan, M., Valentino, M., and Freitas, A. (2020). *A Survey on Explainability in Machine Reading Comprehension*.
- Tjoa, E., and Guan, C. (2021). A survey on explainable artificial intelligence (xai): toward medical xai. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 4793–4813. doi: 10.1109/TNNLS.2020.3027314
- Vargo, C., Guo, L., and Amazeen, M. (2018). The agenda-setting power of fake news: A big data analysis of the online media landscape from 2014 to 2016. *Sage J.* 20, 5. doi: 10.1177/1461444817712086
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., et al. (2017). “Attention is all you need,” in *Advances in Neural Information Processing Systems*, eds. I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett. New York: Curran Associates, Inc.
- Wadden, D., Lin, S., Lo, K., Wang, L. L., van Zuylen, M., Cohan, A., et al. (2020). “Fact or fiction: verifying scientific claims,” in *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Wang, W. Y. (2017). “Liar, liar pants on fire: a new benchmark dataset for fake news detection,” in *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Vancouver, Canada: Association for Computational Linguistics, 422–426.
- Wiegrefe, S., Marasović, A., and Smith, N. A. (2021). “Measuring association between labels and free-text rationales,” in *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. Punta Cana, Dominican Republic: Association for Computational Linguistics, 10266–10284.
- Williams, A., Nangia, N., and Bowman, S. (2018). “A broad-coverage challenge corpus for sentence understanding through inference,” in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, 1112–1122.
- Xie, Z., Thiem, S., Martin, J., Wainwright, E., Marmorstein, S., and Jansen, P. (2020). “Worldtree v2: a corpus of science-domain structured explanations and inference patterns supporting multi-hop inference,” in *Proceedings of the 12th Language Resources and Evaluation Conference*, 5456–5473.
- Yang, S., Wang, Y., and Chu, X. (2020). A survey of deep learning techniques for neural machine translation. *CoRR*.
- Yu, M., Chang, S., Zhang, Y., and Jaakkola, T. (2019). “Rethinking cooperative rationalization: Introspective extraction and complement control,” in *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. Hong Kong, China: Association for Computational Linguistics, 4094–4103.
- Yu, M., Zhang, Y., Chang, S., and Jaakkola, T. (2021). *Understanding Interlocking Dynamics of Cooperative Rationalization*.
- Zaidan, O. F., Eisner, J., and Piatko, C. D. (2007). *Using “Annotator Rationales” to Improve Machine Learning for Text Categorization*, 260–267.
- Zhang, L., Wang, S., and Liu, B. (2018). Deep learning for sentiment analysis: a survey. *Wiley Interdisc. Rev.* 8, 1253. doi: 10.1002/widm.1253
- Zhang, Y., Marshall, I., and Wallace, B. C. (2016). “Rationale-augmented convolutional neural networks for text classification,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas: Association for Computational Linguistics, 795–804.
- Zhang, Z., Rudra, K., and Anand, A. (2021). *Explain and Predict, and Then Predict Again*. New York, NY, USA: Association for Computing Machinery, 418–426.