



# Going Beyond XAI: A Systematic Survey for Explanation-Guided Learning

YUYANG GAO, Emory University, Atlanta, USA

SIYI GU, Emory University, Atlanta, USA

JUNJI JIANG, Emory University, Atlanta, USA

SUNGSOO RAY HONG, George Mason University, Fairfax, USA

DAZHOU YU, Emory University, Atlanta, USA

LIANG ZHAO, Emory University, Atlanta, USA

---

As the societal impact of Deep Neural Networks (DNNs) grows, the goals for advancing DNNs become more complex and diverse, ranging from improving a conventional model accuracy metric to infusing advanced human virtues such as fairness, accountability, transparency, and unbiasedness. Recently, techniques in Explainable Artificial Intelligence (XAI) have been attracting considerable attention and have tremendously helped Machine Learning (ML) engineers in understanding AI models. However, at the same time, we started to witness the emerging need beyond XAI among AI communities; based on the insights learned from XAI, how can we better empower ML engineers in steering their DNNs so that the model's reasonableness and performance can be improved as intended? This article provides a timely and extensive literature overview of the field Explanation-Guided Learning (EGL), a domain of techniques that steer the DNNs' reasoning process by adding regularization, supervision, or intervention on model explanations. In doing so, we first provide a formal definition of EGL and its general learning paradigm. Second, an overview of the key factors for EGL evaluation, as well as summarization and categorization of existing evaluation procedures and metrics for EGL are provided. Finally, the current and potential future application areas and directions of EGL are discussed, and an extensive experimental study is presented aiming at providing comprehensive comparative studies among existing EGL models in various popular application domains, such as Computer Vision and Natural Language Processing domains. Additional resources related to event prediction are included in the article website: <https://kugaoyang.github.io/EGL/>

CCS Concepts: • Computing methodologies → Machine learning approaches; Natural language processing; Computer vision;

Additional Key Words and Phrases: Explainable AI (XAI), faithfulness, trustworthiness, deep learning, Explanation-Guided Learning (EGL), explanation supervision, attention supervision, explanation alignment

## ACM Reference Format:

Yuyang Gao, Siyi Gu, Junji Jiang, Sungsoo Ray Hong, Dazhou Yu, and Liang Zhao. 2024. Going Beyond XAI: A Systematic Survey for Explanation-Guided Learning. *ACM Comput. Surv.* 56, 7, Article 188 (April 2024), 39 pages. <https://doi.org/10.1145/3644073>

---

This work was supported by the National Science Foundation (NSF) Grant No. 2318831, No. 1755850, No. 1841520, No. 2007716, No. 2007976, No. 1942594, No. 1907805, Cisco Faculty Research Award, Oracle for Research Grant Award, Amazon Research Award, NVIDIA GPU Grant, and Design Knowledge Company (subcontract number: 10827.002.120.04).

Authors' addresses: Y. Gao, S. Gu, J. Jiang, D. Yu, and L. Zhao (Corresponding author), Emory University, Atlanta, Georgia, USA, 30322; e-mails: [ygao13.personal@gmail.com](mailto:ygao13.personal@gmail.com), [carrie.gu@emory.edu](mailto:carrie.gu@emory.edu), [jjian50@emory.edu](mailto:jjian50@emory.edu), [dazhou.yu@emory.edu](mailto:dazhou.yu@emory.edu), [liang.zhao@emory.edu](mailto:liang.zhao@emory.edu); S. R. Hong, George Mason University, Fairfax, VA, USA; e-mail: [shong31@gmu.edu](mailto:shong31@gmu.edu).

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2024 Copyright held by the owner/author(s).

ACM 0360-0300/2024/04-ART188

<https://doi.org/10.1145/3644073>

## 1 INTRODUCTION

In recent years, techniques in **Explainable Artificial Intelligence (XAI)** have been attracting considerable attention [2, 10, 61] and have gradually become the dominating ways that connect the way **Deep Neural Networks (DNNs)** work and human reasoning [66, 91]. As DNNs cannot provide human sensible “global structure” of how the model works unlike white-box models, XAI has become an imperative tool that **Machine Learning (ML)** engineers always use to “make sense” of the way their models work [54]. In recent years, many XAI techniques have been proposed in an effort to open the “black box” of DNNs [61], such as techniques that provide saliency maps for understanding which sub-parts (i.e., features) in an instance are most responsible for the model prediction [12, 105, 106, 131, 181]. Despite the recent fast progress on XAI techniques for DNNs, the majority of the research body in XAI put focus on handling “how to generate the explanations” while showing less attention to advanced questions like “whether the explanations are reasonable/accurate,” “what if the explanations are unreasonable/inaccurate,” and, most importantly, “how to adjust the model to generate more reasonable/accurate explanations in the future.” We are starting to witness the emerging need beyond XAI; based on the insights learned from XAI, how can we better steer DNNs such that their future behavior can be improved from the insights learned from XAI techniques? We argue that understanding how to convert insights learned from XAI-driven techniques to steer DNNs would be the key to realizing the DNNs to be more powerful, fair, accountable, transparent, unbiased, and trustworthy, unraveling many real-world application areas.

In recent years, several new areas have emerged that aim at gaining a thorough grasp of the model behavior through the model explanation. Explanatory Debugging is one area of research that has gained popularity [79, 88, 158]. Interactive techniques and systems were developed to enable human users to interactively select features of interest and then investigate how the model behaves in the resulting subspaces for debugging purposes. Another interesting area of research compared the explanation provided by DNNs and the explanation provided by humans to gain a better understanding of the models’ behavior [35, 147]. Although the aforementioned studies are capable of providing more insights about whether the explanations are accurate or reasonable, they are yet to be sufficient for further handling how we can learn from those mistakes, and consequently adjust the model to get better quality explanations and enhance the model performance.

In recent years, a new research direction has emerged, focusing on leveraging XAI techniques to intervene in the behavior of machine learning models. This is achieved by introducing additional supervision signals or prior knowledge obtained from *human explanations* into the model’s reasoning process. Human explanation, in the context of machine learning, involves understanding the rationale behind a person’s decision when performing a task at which this person is competent. For example, in the diagnostic imaging domain, a radiologist or clinician can use **computed tomography (CT)** images to diagnose cancer and annotate suspicious lesion areas to support the rationale of their diagnostic result. The lesion areas provided can be considered a form of explanation for decision-making in terms of salient areas [8, 31, 60, 173]. Hence, the human insights into decision-making are often represented through *human annotations* that help explain the decision-making, which usually take a form similar to attention mechanisms, saliency maps, or other forms commonly used in model explanations. In many applications, such as the aforementioned diagnostic imaging domain, human explanation annotations are available and provide much more informative guidance than conventional prediction labels for machine learning model training. This motivates the goal of incorporating human explanation annotations into machine learning models, e.g., by aligning model explanations with human explanation annotations. This ultimately enhances machine learning models’ interpretability and generalizability, because they better grasp the correct rationale and become more robust to artifacts and noises that interfere

with model predictions. This research direction is commonly referred to as **Explanation Guided Learning (EGL)** [69, 89, 127, 141]. In addition to EGL, several related terms are frequently used within this context, such as *Explanation Supervision* [52–54], *Attention Supervision* [120, 176], *Explanation Alignment* [123, 170], and *Learning from Explanation* [24, 128, 167]. These terms collectively describe the effort to utilize insights from human explanation annotations to enhance the explainability and performance of machine learning models.

Recently, there has been a surge of research that both proposes and applies new approaches in numerous application areas, including **Computer Vision (CV)**, **Natural Language Processing (NLP)**, and **Visual Question Answering (VQA)**. Despite the fact that EGL techniques are generally still in their early stage, the majority of existing studies have produced encouraging results, showing that the main DNNs can generally benefit from the additional explanation objective in terms of both model explainability and generalizability to unseen data across various application domains. However, developing EGL frameworks can be difficult due to significant technical challenges caused by its unique characteristics, including the following:

- (1) *Gap between the pattern of model explanation and human explanation annotation*: The explanation generated by model explainers is typically continuous values, whereas human explanation annotations are typically binary. Therefore, it is difficult to align the human explanation annotation directly with the model explanation without significant efforts to fill the gap between the data domain and distributions.
- (2) *Ensuring the Alignment of Model Explanations in EGL Evaluation*: In contrast to conventional models, where task performance takes precedence, evaluating the quality of EGL outcomes necessitates sophisticated and thoughtfully designed evaluation procedures that often inherently involve subjectivity. For instance, human participants may be engaged in the evaluation process to assess the accuracy and alignment of model explanations. Moreover, beyond the realm of XAI explanations, EGL introduces the additional complexity of jointly evaluating the accuracy of predictions, explanations, and their interplay in determining model generalizability. Consequently, a significant challenge lies in the absence of systematic standardization and comprehensive summarization methods that can effectively evaluate the diverse EGL methodologies proposed to date.
- (3) *Noisiness in human explanation annotation labels*: Unlike predictive task labels, it is much more likely for human annotators to unintentionally create noisy annotation labels where either the real important features are missed or irrelevant features are mistakenly included in the explanation annotation. For instance, when annotation the image data, some important object parts or even the entire objects may be missed by the coarsely drawn boundary from human annotators. Thus, applying naive supervision directly to train the model can lead to falsely excluded non-trivial features from the input space that are important to the prediction [52].
- (4) *Difficulty in explicitly measuring the faithfulness of the explanation quality with respect to the model generalizability*: Due to the fact that EGL techniques are generally still in their infancy, most existing works still primarily focus on merely evaluating the explanation quality of the EGL model independently of the model task performance. The faithfulness of the improved explanation quality with respect to the model prediction is yet to be explored explicitly and can be a key research question to be answered for EGL techniques to further advance and enhance the model performance and generalizability.

## 1.1 Contributions

As the majority of existing EGL approaches were built for a specific application domain, cross-referencing these techniques across application domains serving different communities is

problematic and challenging. Moreover, the lack of a comprehensive review and taxonomy of existing techniques and applications in EGL creates substantial challenges for researchers working in the related field, since they lack clear information on existing bottlenecks, pitfalls, open-ended questions, and potentially fruitful future research directions.

To this end, this article provides a systematic survey of EGL models across various application domains, including CV [39, 45, 89, 103, 123, 125, 143, 157], NLP [9, 11, 17, 23, 24, 26, 28, 29, 38, 55, 57, 70, 87, 94, 95, 141, 142, 144, 167, 168, 175, 177, 180], VQA [27, 50, 113, 120, 160, 176], and more in Section 4. The goal of the survey is to help interdisciplinary researchers build a better understanding of the existing EGL techniques and develop appropriate frameworks to solve the problems in their applications domains. In addition, this survey aims at helping researchers outside the AI communities to understand the basic principles as well as identify interdisciplinary open research opportunities in the EGL domain. As far as we know, this is the first comprehensive survey on Explanation-Guided Learning. This work’s contributions are as follows:

- We summarize a general learning paradigm of EGL based on existing works in this field to provide overall guidance on identifying and designing new EGL techniques.
- We identify the key factors of comprehensively evaluating the EGL model performance and provide a summarization and categorization of the existing evaluation procedures and metrics.
- We propose a taxonomy of Explanation-Guided Learning categorized by the level of guidance and methodologies. The advantages, drawbacks, as well as relations among different subcategories of EGL techniques, are also introduced and compared.
- We introduce the broader application of EGL and detail the unique benefits and future opportunities for each application domain.
- We conduct a comprehensive experimental analysis and comparative study among existing EGL models in CV and NLP domains.
- We summarize the existing literature on EGL at the current stage and then provide a set of open problems and potential promising future research directions of EGL.

## 1.2 Relationship with Related Surveys

This section outlines previously published surveys that have some relevance to Explanation-Guided Learning. These surveys can be classified into three topics: (1) XAI technique and evaluation, (2) AI ethics, and (3) interactive machine learning, as introduced in detail below.

**Explainability Technique and Evaluation:** The related surveys of interpretability techniques provide a technical review and categorization of existing explanation techniques that can explain the machine learning model, especially for the sophisticated “black box” DNN models. Several related surveys provide an in-depth classification of machine learning interpretability methods in general [10, 61, 93, 124], while others focus on more specific fields of study. Specifically, Burkart et al. [22] review the explainability methods of supervised machine learning models. Montavon et al. [107] provide a survey that specifically focuses on the interpretability techniques designed for explaining DNNs. Zhang et al. [171] research interpretability techniques for **Convolutional Neural Networks (CNN)** and visual explanation. Tjoa et al. [151] summarized the XAI techniques that have been adopted for explaining medical data. Along the line of interpretability techniques, many recent surveys also review the methods and metrics for comprehensively validating the quality of the explanation generated by the XAI techniques [65, 104, 183].

**AI Ethics:** As the societal impact of AI grows, the goals for revising AI become more complex and diverse, ranging from improving a conventional model accuracy metric to infusing advanced human virtues such as **fairness, accountability, transparency (FaccT)**, and unbiasedness [96].

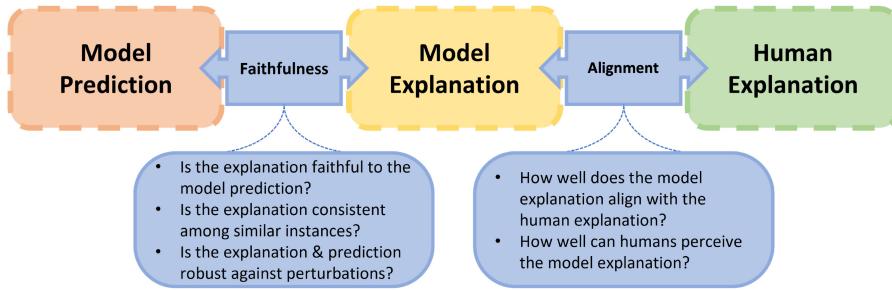


Fig. 1. The gaps in Explanation-Guided Learning performance evaluations.

Table 1. Detailed Evaluation Measures Categorized by the Gaps

Category	Evaluation Measure
Explanation Faithfulness	Perturbation based [11, 29, 38, 40, 69, 103, 140, 170] Explanation Consistency based [11, 140]
Explanation Alignment	Case study based [28, 40, 113, 175] Human annotation based <sup>1</sup> [24, 27, 38, 40, 50, 57, 63, 70, 89, 95, 103, 113, 140, 164] User study based (user-perceived understandability) [29, 70, 116, 149]

Aligning to such direction, recent surveys started to collect, synthesize, and structuralize the existing approaches meant to be designed to handle several types of bias in AI [25, 99]. The most noteworthy finding in our survey for the landmark surveys is that the approaches for detecting bias in ML are more than the ways to mitigate the bias [19, 41]. The second important finding is that even though several studies focus on showing the ways to detect bias, they also present a hint of how we can mitigate them by showing some typical bias cases [99, 110]. Last, the existing survey also provides a pressing field needs explaining why we need to improve the ways to steer models in the case of witnessing the evidence of bias [66].

**Interactive Machine Learning:** Since Fails et al. [47] proposed the idea of interactive ML, the HCI community has put a high priority on applying XAI techniques in developing interactive techniques and systems meant to help ML engineers to better understand their models' weaknesses and strengths. Landmark surveys related to human factor and interaction can be categorized into (1) the interactive design, emphasizing how to design the feedback loop between humans and ML models through system [42, 72] that are widely proposed in the human factor research communities, such as SIGCHI, CSCW, and UIST, and (2) visual analytic, focusing on how to apply visualization techniques to help ML engineers understanding complex ML model behavior [44, 166].

### 1.3 Outline of the Survey

The remaining part of the survey is organized as follows. In Section 2, we delve into the problem formulation and performance evaluations of EGL models, addressing the challenges of ensuring the alignment between model explanations and human explanation annotations (Challenge 2) and explicitly measuring the faithfulness of explanation quality with respect to model generalizability (Challenge 4), as illustrated in Figure 1 and Table 1. In Section 3, we provide a taxonomy of EGL, categorized by the level of guidance and methodologies, which directly addresses the challenges of the gap between the pattern of model explanation and human explanation annotations (Challenge 1) and the noisiness in human explanation annotation labels (Challenge 3), as illustrated in Figure 2. Furthermore, we present detailed insights into each EGL technique, including their respective advantages, drawbacks, and relationships to other techniques within the same subcategories.

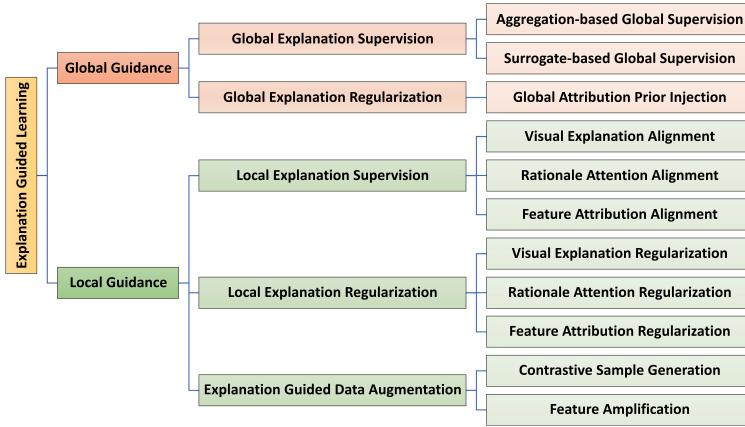


Fig. 2. Taxonomy of Explanation-Guided Learning problems and techniques.

In Sections 4 and 5, we explore the broader applications of EGL and conduct a comprehensive experimental analysis and comparative study among existing EGL models in both CV and NLP domains. Finally, in Section 6, we conclude by summarizing the current state of development in EGL techniques and propose several open problems and potential future research directions.

## 2 PROBLEM FORMULATION AND PERFORMANCE EVALUATIONS

This section begins by introducing the generic denotation and formulation of the Explanation-Guided Learning problem (Section 2.1) and then considers ways to categorize the performance evaluation measures of Explanation-Guided Learning (Section 2.2).

### 2.1 Problem Formulation

Consider a differentiable model  $f$  parameterized by  $\theta$  that learns to fit inputs  $X \in \mathbb{R}^{N \times D}$  and the corresponding one-hot class labels  $Y \in \mathbb{R}^{N \times K}$ , where  $N$  denotes the total number of data samples,  $D$  denotes the input dimension, and  $K$  denotes the number of classes. An explainer  $g$  is considered to extract the explanation  $M$  from the model  $f$  given its parameter  $\theta$  and a set of data points  $\langle X, Y \rangle$ . Generally speaking, the model explanation  $M$  represents the marginal contribution of each input feature to the model's decision after all possible combinations have been considered. Notice that in this article we use the terms *rationale*, *attention*, and *saliency maps* interchangeably as the specific form of  $M$  that is frequently used by the corresponding application domains. Depending on the way the explanation is calculated,  $M$  can be generally represented by either local explanation  $M^{(L)}$ , where  $M_i^{(L)}$  is the local explanation of model  $f$  with respect to sample  $\langle X_i, Y_i \rangle$ , or a single global explanation  $M^{(G)}$  of the model  $f$ .

**The EGL paradigm.** The general goal for Explanation-Guided Learning is to boost both the task performance as well as the interpretability of the backbone model by jointly optimizing model prediction as well as the explanation. Based on the earlier exploration of explanation supervision frameworks design [52, 53, 103, 126], we introduce the key objective function of Explanation-Guided Learning as follows:

$$\min \underbrace{\mathcal{L}_{\text{Pred}}(f(X), Y)}_{\text{task supervision}} + \underbrace{\alpha \mathcal{L}_{\text{Exp}}(g(f, \langle X, Y \rangle), \hat{M})}_{\text{explanation supervision}} + \underbrace{\beta \Omega(g(f, \langle X, Y \rangle))}_{\text{explanation regularization}} , \quad (1)$$

where  $\hat{M}$  explicitly incorporates the “right” explanation, which can be typically realized by human annotation masks [38, 54]. The function  $g$  is typically realized by existing differentiable XAI explainers. Common explainers that are widely used for this purpose include GradCAM [131] and attention mechanisms [13, 152]. These explainers offer established methods for interpreting the model’s decision-making process.

As shown in Equation (1), the key objective function of Explanation-Guided Learning mainly consists of three terms, namely (1) task supervision term for the typical prediction loss (such as the cross-entropy loss), (2) explanation supervision term for supervising the model explanation with some explicit knowledge of what the “right” explanation should be, and (3) explanation regularization term for enforcing some general properties about the “right” explanation (such as maintaining the sparsity nature of the explanation). Notice that all three terms above can be defined and implemented differently depending on each particular Explanation-Guided Learning method.

## 2.2 Performance Evaluations

Unlike the evaluation of conventional machine learning models, which primarily centers on assessing the model’s performance, and the evaluation of traditional explainable AI models, which concentrates on the quality of generated model explanations, EGL introduces a holistic assessment that simultaneously considers model prediction performance, the quality of model explanation, and their interplay. As depicted in Figure 1, we discern and categorize two pivotal dimensions of evaluation that are fundamental for gauging the effectiveness of EGL models: *faithfulness* and *alignment* of the model explanation.

*Explanation Faithfulness*, in the context of EGL, revolves around the concept of whether the model explanation remains true to the underlying model’s reasoning. It encompasses questions such as “Is the explanation faithful to the model’s true reasoning process?”, “Is the explanation consistent across similar instances?”, and “Is the explanation and prediction robust against perturbation?” For example, in a medical diagnosis EGL model, faithfulness entails assessing whether the model’s explanation aligns with the medical knowledge and diagnostic reasoning used by human experts. A faithful explanation should consistently reflect similar diagnostic results for patients with similar symptoms, and it should remain robust when presented with minor variations in patient data.

*Explanation Alignment*, however, focuses on the accuracy and perceptibility of the model explanation from a human perspective. It delves into questions like “How well does the model explanation align with human explanation?” and “How well can humans perceive the model explanation?” Alignment would encompass assessing whether the explanations generated by the model accurately align with the nuances and interpretations of language used by humans, thereby ensuring the *correctness* of the AI explanation by aligning it with human explanation annotations. It also gauges how easily humans can comprehend and trust the explanations provided by the model.

These two dimensions, faithfulness and alignment, are crucial for comprehensively evaluating EGL models, because they ensure that the model not only provides faithful, accurate explanations but also that these explanations are presented in a manner that is understandable and trustable to humans. This broader perspective on evaluation acknowledges the unique challenges and complexities of EGL, where the alignment and faithfulness of model explanations are integral to achieving the model’s ultimate goal: improving transparency and decision-making in complex, real-world applications.

Here we summarize the existing evaluation metrics into the two categories in Table 1 and introduce each type of metric in great detail in the following two subsections.

**2.2.1 Metrics on Evaluating Explanation Faithfulness.** Here we introduce the metrics for explanation faithfulness (model prediction vs. model explanation) evaluation, which aims at evaluating how the model-generated explanation influences the corresponding model's prediction.

**Perturbation-based evaluations:** To evaluate the faithfulness of the model explanation, the study of how different types of perturbations on the input space influence the model prediction has become a very common and well-received approach in the literature [11, 29, 38, 40, 69, 103, 140, 170]. Existing measures can be mainly categorized into three groups, depending on the type of perturbation, as follows:

- *Occlusion-based perturbation:* These metrics basically study how much influence on the model's prediction if the important feature or rationale identified by the model explanation are occluded or masked from the original sample [29, 38, 69, 103, 170]. One commonly used occlusion-based metric is *comprehensiveness* [38], where the difference of the predicted probability from the model  $f(\cdot)$  for the same class  $Y_i$  is compared between the original input  $X_i$  and  $X_i \setminus g(f, \langle X_i, Y_i \rangle)$ , where the operation “\” represents the exclusion of the supporting rationales  $g(f, \langle X_i, Y_i \rangle)$  from input  $X_i$ . Mathematically, Comprehensiveness can be defined as follows:

$$\text{Comprehensiveness} = f(X_i)_{Y_i} - f(X_i \setminus g(f, \langle X_i, Y_i \rangle))_{Y_i}. \quad (2)$$

Besides the comprehensiveness score, many other intuitive methods are also used to evaluate the quality of the explanation. Inspired by previous works [108, 132], a common intuitive strategy to measure the faithfulness of the explanation used by existing works [29, 103, 170] is to track the degradation of model performance by removing importance features (often in decreasing order) from the input.

- *Insertion-based perturbation:* These metrics study how well the prediction aligns between the original sample and an artificially generated sample where only the important feature/rationales are included [11, 38, 103, 140]. One popular metric is *Sufficiency* [38], which captures the degree to which the snippets within the extracted rationales  $g(f, \langle X_i, Y_i \rangle)$  are adequate for a model to make a prediction. Concretely, it can be defined as follows:

$$\text{Sufficiency} = f(X_i)_{Y_i} - f(g(f, \langle X_i, Y_i \rangle))_{Y_i}. \quad (3)$$

Similarly, many other intuitive methods are also used following the insertion idea. A common strategy used by existing works [11, 103] is to track the increase in model performance by gradually inserting the important features (often in decreasing order) from the input.

- *Adversarial perturbation:* These metrics in general check whether the model explanation is still faithful to the model prediction under adversarial attacks [40, 170]. For instance, Reference [170] leveraged the sanity check method originally proposed in Reference [3] to check if attribution maps look different when the deep network being explained is extremely perturbed or under adversarial attacks. The intuition behind this measure is that a faithful attribution method should yield different explanations for the randomized model.

**Consistency-based evaluations:** Besides the perturbation-based metrics that only focus on evaluating each instance locally at a time, existing works also propose consistency-based evaluation, where more global evaluation metrics have been proposed to validate how well the explanation aligns across similar instances [11, 140]. More specifically, Reference [11] proposed a metric called *Data Consistency* that measures how similar the explanations for similar instances are. Although the specific equation of the measurement in the article is specifically designed for NLP and generative explanation, the basic idea can be generally expressed as follows:

$$\text{Data Consistency} = |g(f, \langle X_i, Y_i \rangle) - g(f, \langle X_i \setminus M, Y_i \rangle)|, \quad (4)$$

where  $M$  is a random mask that masks out  $K$  input features from  $X_i$ ,  $K$  can be treated as a hyper-parameter dependent on the dataset characteristics. The underlying assumption here is that model explanations between highly similar samples should exhibit proximity, and, therefore, higher Data Consistency values indicate superior performance in preserving this similarity.

It is worth noting that the authors of this metric emphasized the use of the L1 loss function instead of L2 [11]. This choice serves as a strategy to mitigate the penalization in the rare situation where the random mask  $M$  unintentionally masks out important features, potentially leading to significantly different outlier predictions. Furthermore, in practical scenarios where ground-truth explanation annotations are available, this information can be employed to guide the selection of  $M$ .

Similarly to the above idea, another work employed **Intersection over Union (IoU)** score to measure explanation stability across similar instances [140]. Specifically, they proposed to find similar instances by searching for the nearest neighbors of  $X_i$  in the dataset based on both the semantic similarity (the cosine of their BERT representations) and the lexical similarity (the ratio of overlapping  $n$ -grams).

**2.2.2 Metrics on Evaluating Explanation Alignment.** Here we introduce the metrics for explanation alignment evaluation, which aims at evaluating how well the model-generated explanation aligns with the human explanation annotation or how well can humans perceive the model-generated explanation.

**Case study:** Case study has been widely used as a conventional method for qualitatively evaluating the explanation generated by the model [28, 40, 113, 175], where a set of instances and their corresponding model explanations are selected and investigated qualitatively. Although making qualitative assessments and detailed analyses of just a few samples can be easily achieved, it is in general less scientifically rigorous, and the claims or conclusions are pruned to be biased due to the author's subjectivity.

**User study (user-perceived understandability):** User study, specifically user-perceived understandability, has been commonly used as a qualitative evaluation method to assess how humans can understand the explanation generated by the model [26, 70, 116, 149]. The user-perceived understandability methods are typically achieved by developing a user interface to show the model explanations to human subjects, and collecting the rating of how likely the important features identified by the model explanation can lead to the correct prediction of the underlying ground-truth label.

**Human annotation-based evaluation:** Explanation alignment is a unique yet commonly used quantitative metric in Explanation-Guided Learning which measures how the human-annotated ground-truth explanation is aligned with the model generated explanation [24, 27, 38, 40, 50, 57, 63, 70, 89, 95, 103, 113, 140, 164]. The distance is commonly measured by the IoU score [38, 95, 140], precision, recall, and F-1 scores [57, 140].

**2.2.3 Other General Metrics.** Besides measuring the faithfulness and alignment of model explanation, most of the papers also included the conventional model task performance metrics to verify if the Explanation-Guided Learning actually helped the generalizability of the backbone DNN models. Like most papers working on classification tasks, the common metrics used to evaluate model performance are accuracy, **Area Under the ROC Curve (AUC)** score, and F1 score.

### 3 EXPLANATION-GUIDED LEARNING TECHNIQUES

This section focuses on the taxonomy and representative techniques utilized for each category and subcategory. According to the level at which the model explanation is obtained and supervised, the technique types for EGL can be divided into global guidance and local guidance, as shown in

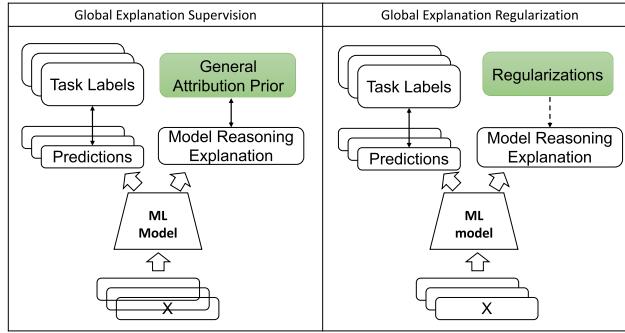


Fig. 3. Illustration of Global Guidance techniques. Specifically, global explanation supervision techniques (left) aim at providing supervision in terms of model attribution, while global explanation regularization techniques (right) aim at confining the model reasoning process with prior knowledge.

Figure 2. Specifically, global guidance focuses on the model’s global explanation and refines the model’s overall decision-making process, while local guidance guides the model with each sample-specific explanation. The aforementioned techniques are then further categorized in terms of the way explanation guidance is injected during the course of model training.

### 3.1 Global Guidance

Global explanation guidance focuses on injecting prior knowledge or adding supervision signals to improve the model’s global explanation that explains the decision-making process of the model in general. Based on the way explanation guidance is injected, global explanation guidance methods can be categorized into two types: (1) Global Explanation Supervision, where the ground-truth explanation labels are provided as an additional supervision signal to train the featurewise explanation of the model, and (2) Global Explanation Regularization, in which some regularization terms that represent some general prior knowledge about the model explanation are added to regularize the featurewise explanation of the model, as illustrated in Figure 3.

**3.1.1 Global Explanation Supervision.** The techniques proposed in global explanation supervision [45, 94, 157] aim at providing a single featurewise explanation of the model globally. Compared with instance-level local explanation supervision where the explanation ground truth is provided for each instance [52, 54, 103], global explanation aims to provide a more effective global guide to the model’s behavior as a whole. Depending on the strategies to compute the global explanation of the model, current literature can be mainly categorized in two directions: (1) aggregation based [45, 94] and (2) surrogate based [34, 118, 153].

**Aggregation-based Global Supervision:** This type of method typically achieves explanation supervision by first estimating the global feature attribution via aggregating local feature attribution of each sample and aligning it with a single ground-truth feature attribution vector  $\hat{m}$  as the additional supervision signal to train the model jointly with the conventional task loss. The common techniques used to calculate each sample’s feature attribution are integrated gradient [94] and the expected gradient proposed by Erion et al. [45]. Specifically, the objective function for aggregation-based global supervision can be summarized as follows:

$$\min \mathcal{L}_{\text{Pred}}(f(X), Y) + \alpha \mathcal{L}_{\text{Exp}} \left( \frac{1}{N} \sum_{i=1}^N g(f, \langle X_i, Y_i \rangle), \hat{M} \right). \quad (5)$$

This type of method has been utilized and shown promising results in many application domains, such as image classification tasks [45] and text classification [94]. For image data use cases [45], the global explanation  $\hat{M}$  represents the global pixel-level saliency map calculated using expected gradients. As demonstrated in Figure 1 in Reference [45], the authors illustrate the significant improvements in global explanation quality achieved by combining higher-level attribution priors with expected gradients attributions. This enhancement is consistent across various data domains, including image, gene expression, and healthcare datasets. In the context of text data applications [94], the global explanation  $\hat{M}$  denotes the global attribution values assigned to each potential token in the text dataset, which are calculated using Integrated Gradients attributions. The authors applied this approach to an NLP task aimed at identifying toxic conversations. In Table 8 in Reference [94], the authors showcase the successful highlighting of toxic words with higher attribution values using their proposed method, even across different data ratios, as compared to the baseline.

The advantage of the aggregation-based global supervision methods is that they can easily adapt existing techniques developed for local explanation with little to no extra effort. Besides, the acquisition of only one single feature attribution vector as a classwise explanation signal is much more affordable, as compared with instancewise supervision methods that require much more labor from human annotators. However, the drawbacks of this type of technique also come from the aggregation of local explanation, as the aggregated explanation is sensitive to the samples used to calculate, and thus could bring the sample bias into the global explanation of the model estimated.

**Surrogate-based Global Supervision:** This branch of work achieves explanation supervision by first estimating the global explanation of the target model via a surrogate model where the model-level explanation is easy to obtain, and then human knowledge can be leveraged to guide the global explanation and consequently supervise the model behavior. In this branch, the rule-based explanation is commonly used as it can be easily edited by practitioners [34, 118, 153].

Rule-based explanation supervision can be achieved from many different angles. For instance, Vojíř et al. [153] proposed the editable rule-based models that enable the users to edit rules and replace the underlying machine learning model, and Popordanoska et al. [118] proposed the Explanatory Guided Learning framework that creates simple rules capturing the prediction of the target model and allows the user to correct instances that are incorrect and the model is retrained. The rule-based explanation is also used as a mechanism for feedback that supports user adjustments without retraining the model [34]. Cor nec et al. [33] developed the AI Model Explorer and Editor tool, which provides visualization of model decision boundaries using interpretable surrogates and allows for the real-time modification of the decision boundaries. More recently, Lee et al. [86] proposed SELOR, a framework for upgrading a deep model with a Self-Explainable version with LLogic rule Reasoning capability, inspired by neuro-symbolic reasoning [36] that integrates deep learning with logic rule reasoning to inherit advantages from both. SELOR provides high human precision by explaining logic rules while also maintaining high prediction performance and does not require predefined rule sets and can be learned in a differentiable way.

**3.1.2 Global Explanation Regularization.** Global explanation regularization is the method where some regularization terms that incorporate general prior knowledge about the global explanation are applied to the model. A good example of a preferred property of the model explanation is the sparseness, as it can provide a better understanding of the model behavior by humans, and in the meantime, serve as a regularizer of the explanation space to enhance model generalizability [51, 115, 133]. Concretely, the objective function for global explanation regularization can be summarized as

$$\min \mathcal{L}_{\text{Pred}}(f(X), Y) + \beta \Omega(M^{(G)}), \quad (6)$$

where the function  $\Omega(\cdot)$  represents the specific regularization function for regulating the model's global explanation and  $M^{(G)}$  represents the model's global explanation vector calculated based on intrinsic parameters of the model  $f$ .

A commonly used prior knowledge to define  $\Omega(\cdot)$  is to ensure the sparseness of the explanation, where a regularization term is proposed to penalize small magnitude weights of  $f$  that connect to the input features [51, 115, 133]. The existing studies suggest that this can result in a feature selection effect and greatly enhance the model's computational efficiency as well as generalizability. In addition, Burkart et al. [21] proposed a batchwise regularization technique to enhance the interpretability of DNN models by means of a global surrogate rule list with a novel regularization approach that yields a differentiable penalty term. In Wu et al. [161], the authors proposed the regional tree regularization that encourages a DNN model to be well approximated by several separate decision trees specific to predefined regions of the input space, yielding simpler explanations without compromising model accuracy.

### 3.2 Local Guidance

Local explanation guidance focuses on applying supervision signals or regularization terms to the model explanation of each local sample to guide the model learning. As shown in Figure 4, compared with the global explanation guidance, local guidance is more commonly used and explored in the current research communities thanks to the development of local explanation techniques, such as GradCAM [131] and attention mechanism [13, 152]. Based on the way explanation guidance is injected, local explanation guidance techniques can be categorized into three types: (1) Local Explanation Supervision, where the ground-truth explanation labels for each individual sample are provided as additional supervision signals to train the corresponding model explanation, (2) Local Explanation Regularization, in which some regularization terms that represent some general prior knowledge about the local model explanation are added to regularize all the local explanation of the model, and (3) Explanation Guided Data Augmentation, where the local model explanations are used to construct additional data samples for model training.

**3.2.1 Local Explanation Supervision.** Just as we supervise the model prediction via ground-truth labels, local explanation supervision methods add additional supervision signals to align the model explanation with ground-truth explanation labels (e.g., human annotation masks) during model training. The explanation loss and the conventional prediction loss are typically jointly optimized during model training. The general assumption behind this approach is that the model can benefit from the explanation labels by learning to focus on the right features and consequently lead to better generalizability to unseen instances. Depending on the data representation and application domains, we further narrow down the techniques into three subcategories: (1) visual explanation alignment, (2) rationale attention alignment, and (3) feature attribution alignment.

**Visual Explanation Alignment:** The visual explanation of image data is typically represented by a heat map overlaid on top of the original image, and the ground-truth explanation labels  $\hat{M}$  are typically obtained by human annotation in the form of bounding boxes or fine-grained contours.

The first framework that can be applied to visual explanation alignment was proposed by Ross et al. [126], where the authors defined a very generic Explanation-Guided Learning loss called “**Right for the Right Reasons**” loss (RRR) as follows:

$$\min \sum_{i=1}^N -Y_i \log(f(X_i)) + \alpha \sum_{n=1}^N (\hat{M}_i \frac{\partial}{\partial X_i} \log(f(X_i)))^2 + \beta \|\theta\|_2^2, \quad (7)$$

where  $\hat{M}_i$  denotes the ground-truth explanation mask of a sample  $i$ ; the task supervision loss is implemented as the conventional cross-entropy loss, and the explanation supervision loss is

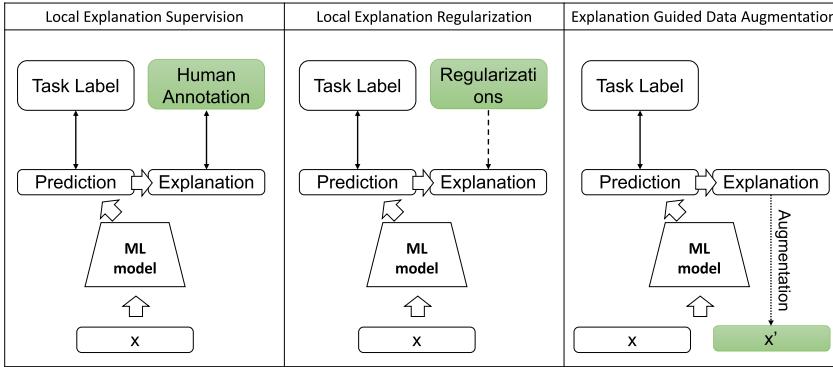


Fig. 4. Illustration of Local Guidance techniques. (Left) Local Explanation Supervision, where the ground-truth explanation labels for each individual sample are provided as additional supervision signals to train the corresponding model explanation; (middle) Local Explanation Regularization, in which some regularization terms that represent some general prior knowledge about the local model explanation are added to regularize all the local explanation of the model; and (right) Explanation Guided Data Augmentation, where the local model explanations are used to construct additional data samples for model training.

designed to enforce the alignment of the ground-truth explanation mask  $\hat{M}$  and the gradient maps via inner product operations.

Later, the RRR loss was further extended by Schramowski et al. [129] and Dhama et al. [39] regarding the definition of the explanation losses. Specifically, instead of regularizing the gradients with respect to input  $X$ , Schramowski et al. [129] proposed to regularize the gradients of the final convolutional layer of the model that corresponds to GradCAM explanation and add a rescaling weight  $c_k$  to each class  $k$  for handling the unbalanced dataset issue. In Dhama et al. [39], the explanation loss is broken down into two terms to characterize the sensitivity of the gradient maps differently based on the relationship between each pixel of input and the ground-truth mask as follows:

$$\mathcal{L}_{\text{Exp}} = \alpha_1 \sum_{j \in \hat{M}_i} \frac{\partial \mathcal{L}_{\text{Pred}}(f(X_i), Y_i)}{\partial X_{i,j}} + \alpha_2 \sum_{j \in [d] \setminus \hat{M}_i} \frac{\partial \mathcal{L}_{\text{Pred}}(f(X_i), Y_i)}{\partial X_{i,j}}, \quad (8)$$

where  $[d] \setminus \hat{M}_i$  represent the complement subset of the explanation  $\hat{M}_i$  of the whole feature set.

Many more models that are designed for visual explanation alignment have been proposed [52, 54, 103, 143, 165]. Specifically, Stammer et al. [143] proposed a visual explanation alignment model based on symbolic (concept) alignment, where the symbolic (concept) explanation is modeled by a set transformer module. Mitsuhasha et al. [103] proposed a visual explanation alignment objective specifically designed for the **Attention Branch Network (ABN)** [49], where the attention branch outputs are used as the model explanation. The limitation of this work is that it can only work under ABN architecture. Ying et al. [165] proposed the Visual Feature Importance Supervision framework that optimizes four key model objectives: (1) accurate predictions given limited but sufficient information (Sufficiency), (2) max-entropy predictions given no important information (Uncertainty), (3) invariance of predictions to changes in unimportant features (Invariance), and (4) alignment between model explanations and human explanation annotations (Plausibility) to improve model accuracy as well as performance. Nguyen et al. [109] proposed two novel architectures of self-interpretable image classifiers that first explain, and then predict by harnessing the visual correspondences between a query image and exemplars and demonstrated the improvement on out-of-distribution dataset scenarios. Gao et al. [54] proposed a more generic visual

explanation alignment framework called GRADIA based on the GradCAM explanation. In addition, they proposed the Reasonability Matrix that can better determine what samples need to be adjusted to improve the model performance and explanation quality. More recently, Gao et al. [52] proposed a robust visual explanation alignment framework that can better handle the nosiness of human annotation on image data. Specifically, the explanation loss is defined as follows:

$$\mathcal{L}_{\text{Exp}} = \min_{\theta, \lambda, \phi} \sum_i^N \max\{0, \|\tilde{M}_i - \hat{M}_i\| - \lambda\} + (g(f, \langle X_i, Y_i \rangle) - h_\phi(\hat{M}_i))^2, \quad (9)$$

where  $\phi$  is the parameter set of the imputation function  $h_\phi(\cdot)$ . The imputation function can be realized by applying multiple layers of convolution operations with learnable kernels over the raw annotation labels;  $M_i$  is a binary projection of the explanation  $g(f, \langle X_i, Y_i \rangle)$  by a threshold  $\lambda$ , as

$$\tilde{M}_i = \begin{cases} 1 & g(f, \langle X_i, Y_i \rangle) \geq \lambda \\ -1 & g(f, \langle X_i, Y_i \rangle) < \lambda \end{cases}. \quad (10)$$

Besides the application to general images, visual explanation alignment techniques have also been applied to medical image domains [136, 137, 184]. Please refer to the Applications Section 4 for more details.

**Rationale Attention Alignment:** The explanation of natural language data is typically represented by the rationales (e.g., word tokens) that highlight the most significant part of the data for making specific task predictions. The ground-truth explanation labels are typically obtained by human annotation in the form of rationales or natural language format (sentences). In this domain, the datasets collected by the ERASER benchmark [38] are commonly used as the datasets come with ground-truth rationales obtained from human annotators.

Many existing works have proposed to supervise the rationale attention of the model to improve the model performance and quality of attention [24, 28, 57, 74, 142, 175, 180]. The explanation loss is commonly realized by conventional losses, such as cross-entropy loss [28, 74], Mean Squared Error [141, 142], and KL-divergence loss [180]. In addition, Atanasova et al. [11] proposed several novel ways to enforce the alignment, such as Data Consistency, Confidence Indication, and Faithfulness.

In addition to leveraging the model attention value itself, many existing works have also proposed to directly generate rationales [17, 70, 87, 141, 167, 180] or natural language [23, 95] as the “explanation” of the model to be aligned with ground-truth labels via additional decoders, such as **Conditional Random Field (CRF)** [81, 167], **Gated Recurrent Unit (GRU)** [30, 180], and Transformer-based models [37, 95, 141, 152].

**Feature Attribution Alignment:** Besides the specific domain of applications, local explanation supervision can be generally applied to any dataset where the input feature importance can be computed. Such feature importance is typically referred to as “Feature Attribution” and can be also treated as the model explanation to be aligned with human explanation annotation ground truth. For instance, in Balayan et al. [15] the feature attribution is computed by a specific designed semantic layer (an intermediate output that is the importance of each feature) and is aligned with human-labeled feature masks by Cross-Entropy loss; in Singh et al. [138], feature attribution is calculated by Contextual Decomposition and is aligned with the ground-truth human labels by  $\ell_1$  distance [123].

Overall, the idea of local explanation supervision has been explored extensively in many application domains in recent few years, primarily due to the fact that (1) it is straightforward for human annotators to provide an instancewise explanation with necessary domain knowledge and (2) the development and popularity of local explanation techniques, such as GradCAM [131] and attention mechanism [13, 152] to explain complex DNNs in high-dimensional problem space (such as image and text data). So far the results seem to be promising, as most existing works suggest

that applying the local explanation supervision during training can greatly enhance both the quality of the explanation as well as the performance of the backbone DNNs model. However, as also pointed out by several existing works, the scalability remains the biggest challenge for this kind of approach, as the additional instancewise human explanation annotations may not be easily accessible and require non-trial effort from human annotators [52, 53]. Designing effective semi-supervised or weakly-supervised explanation supervision frameworks, or even adapting the idea from active learning, can be promising future directions to further overcome this limitation. Nevertheless, existing works also demonstrated the effectiveness of local explanation supervision under very limited training sample sizes [52, 54], which could suggest the potential benefit of applying current techniques to the domains where data samples are limited and hard to acquire, yet both model performance and the explainability are on-demand, such as in medical domains.

**3.2.2 Local Explanation Regularization.** Local explanation regularization methods add additional regularization terms to regularize each local explanation to ensure the generated model explanations follow some general properties (such as smoothness, stability, and sparsity) or follow the knowledge from other existing well-trained models. The additional explanation of regularization loss is typically jointly optimized with the prediction loss during model training. Depending on the type of regularization terms, we break down the existing techniques into two subcategories: (1) Property-based Regularization and (2) Explanation Distillation Regularization.

**Property-based Regularization:** The additional regularization terms are injected into the model explanation to enforce some general properties (such as smoothness, stability, and sparsity). Specifically, in Lei et al. [87] the authors proposed the *continuity* and *sparsity* regularization terms. In Erion et al. [45] the authors proposed the *smoothness* Regularization (i.e., Laplace 0-mean prior) on the model Explanation computed by expected gradient. In Halliwell et al. [62] the authors proposed the Prediction-guided sparsity regularization (in Equations (5) and (6)) to penalize the model to have small values in saliency maps (computed by GradCAM and guided BP) if the prediction is incorrect. Alvarez et al. [5] proposed a gradient regularization approach for enforcing explanation *robustness/stability*. Plumb et al. [117] apply the *fidelity* and *stability* regularization on the explanation. Specifically, the explainer  $g(\cdot)$  is realized by Local Interpretable Model-Agnostic Explanations [122] with a linear function  $l(\cdot)$ , and the authors applied two regularization terms on the model explanation, (1) neighborhood-fidelity and (2) stability based on the neighborhood of input, as follows:

$$\Omega = \underbrace{\mathbb{E}_{X' \sim \mathcal{N}_{X_i}} [(l(X') - f(X'))^2]}_{\text{neighborhood-fidelity}} + \underbrace{\mathbb{E}_{X' \sim \mathcal{N}_{X_i}} [\|g(f, \langle X_i, Y_i \rangle) - g(f, \langle X', Y_i \rangle)\|_2^2]}_{\text{stability}}, \quad (11)$$

where  $\mathcal{N}_{X_i}$  is a neighborhood of sample  $X_i$  in the space of probability distributions over the whole input data distribution  $X$ , and  $X'$  is sampled from the neighborhood  $\mathcal{N}_{X_i}$ . Intuitively, the fidelity regularization enhanced the explanation to accurately convey which patterns the model used to make this prediction, while the stability regularization will lead to more stable explanations, which will improve the model's trustworthiness [5, 6].

**Explanation Distillation Regularization:** Besides enforcing predefined properties of the explanation, this line of work tries to distill explanation knowledge from other well-trained models to guide the explanation of the target model. In Zeng et al. [170], the authors proposed to align the explanation of a target model with another pre-trained adversarially counterpart model generated explanation using  $\ell_2$  distance loss. Singh et al. [139] proposed to align the target model's **Class Activation Maps (CAM)** [181] explanation with a pre-trained model's explanation by minimizing the overlap between each classes explanation. More specifically, the explanation loss consists of two terms, (1) regularization loss, which measures the distance between the target models' and

the corresponding pre-trained model's explanation of class  $i$ , and (2) overlapping loss, which calculates the similarity between the target model's explanation of different classes. As a result, the model can be trained under the constraints whereby (1) the target model explanation should be as close as possible to a pre-trained model and (2) the model explanation of different classes should be as different as possible, leading to a better explanation quality and higher accuracy. More recently, Fernandes et al. [48] proposed the **Scaffold-Maximizing Training (SMaT)** framework for directly optimizing explanations of the model's predictions to improve the training of a student simulating the said model. The authors found that, across tasks and domains, explanations learned with SMaT both lead to students that simulate the original model more accurately and are more aligned with how people explain similar decisions.

While using the idea from model distillation to extract the knowledge to guide the model explanation is an interesting direction, the potential positive effect is largely dependent on the choice and quality of the pre-trained model and is prone to negative transfer, such as contextual bias in pre-trained model explanation, that can hurt the target model performance. Thus additional validation and guidelines are on demand for real-world problems.

**3.2.3 Explanation Guided Data Augmentation.** Explanation Guided Data Augmentation is an emerging subdomain in the data augmentation domain, where the ground-truth explanation (i.e., rationale) of the prediction task is taken into account when building up additional augmented samples for model training. The general formulation for generating explanation-guided data augmentation samples can be summarized as follows:

$$X'_i = \text{aug}(X_i, g(f, \langle X_i, Y_i \rangle)), \quad (12)$$

where  $\text{aug}(\cdot)$  denotes the specific augmentation function based on the original input sample  $X_i$  and the model's explanation for the given input-output pair  $\langle X_i, Y_i \rangle$ .

The underlying assumption is that training the model with the augmented samples  $X'$  will encourage the model to better learn to pay attention to the right rationales for the prediction tasks and thus naturally enhance both the explainability as well as the generalizability of the model. Based on the way explanation is used for the data augmentation, existing techniques can be categorized into two directions: (1) Rationale Inclusion/Amplification and (2) Rationale Exclusion/Masking.

**Rationale Inclusion/Amplification:** This line of works typically emphasizes the right rationales and de-emphasize other irrelevant features. The inclusion/amplification-based augmentation function can be generally defined as follows:

$$\text{aug}_{in}(X_i, g(f, \langle X_i, Y_i \rangle)) = X_i \times (\gamma + \lambda g(f, \langle X_i, Y_i \rangle))), \quad (13)$$

where  $\gamma$  is used to set a default offset value to preserve all the feature values regardless of the importance;  $\lambda$  is the scale factor that controls the degree of amplification of the important features.

Specifically, Sharma et al. [135] proposed to amplify the feature values of the right rationales relatively higher by a certain degree. In their experiments,  $\gamma$  is set to 0.01 and  $\lambda$  is set to 1 to emphasize the rationale features in the augmented samples. The results demonstrated the general effectiveness of the proposed method on several conventional ML models, such as Naive Bayes, logistic regression, and SVM. In Saha et al. [127], only the important part of the image for network prediction is selected using saliency-based explanations and stored in the episodic memory with the corner coordinate for continual learning. Ismail et al. [69] proposed to minimize the KL divergence between  $f(X)$  and  $f(X')$ , where  $X'$  is augmented by masking the features with low gradient values. These types of methods can be seen as a special case of Equation (13), where  $\gamma$  is set to 0 and  $\lambda$  is set to 1 to only include the rationale features in the augmented samples.

Besides the simple augmentation of the feature values, several other works have also proposed some novel and unique ways to augment data to best leverage the extra information from the

model explanation. In Pillai et al. [116], the saliency map explanation of the original sample and other samples as a composed image is aligned with the model explanation of the original input sample. In Teso et al. [149], the irrelevant features are perturbed, while the rationales and task labels are preserved as new samples to guide the model attending the ground-truth rationales. In Schneider et al. [128], the model explanation (generated by GradCAM) is treated as additional input for model prediction and requires a specific change of the model architecture. In Gu et al. [60], a Saliency-guided Data Augmentation (ESSA) framework is proposed for conducting explanation supervision and adversarial-trained image data augmentation. It operates through a synergized iterative loop that handles the translation from annotation to sophisticated images and the generation of synthetic image-annotation pairs, employing an alternating training strategy. In summary, the general idea stays the same, which is to build additional samples and inform the model to better learn which features are the right rationales to make the right prediction of the downstream tasks.

**Rationale Exclusion/Masking:** Conversely to inclusion/amplification, this line of works typically teaches the model not to attend irrelevant rationales by excluding/masking out the right rationales, as summarized by the following equation:

$$\text{aug}_{ex}(X_i, g(f, \langle X_i, Y_i \rangle)) = X_i \times (\gamma - g(f, \langle X_i, Y_i \rangle))), \quad (14)$$

where  $\gamma$  is typically set to be the maximum possible value of the importance, e.g., 1, to exclude the value of the important features from  $X$  and thus serve as a masking function for the data augmentation.

Specifically, Zaidan et al. [168] propose to construct some additional samples by masking out those important features of some existing samples to simulate the loss of confidence (uncertainty should raise) in predicting the right answer. A similar idea can be also found in Li et al. [89], where the proposed self-guidance is basically using the model explanation as a mask to augment the original image and thus construct an unsupervised loss based on the augmented image.

Overall, the unique advantages of explanation-guided data augmentation techniques can be summarized as follows: (1) It takes the model behavior (i.e., rationale for the prediction) into consideration; (2) it can be model agnostic with respect to the specific explainability techniques used for calculating the model explanation; and (3) it can be used in combination with other conventional data augmentation techniques and in parallel with other EGL techniques for model training. However, the effectiveness of the existing works is mainly supported by intuitions and empirical observations. Thus further development of quantitative evaluation metrics as well as theoretical analysis and justification of the techniques can be essential to further advance this field of research.

## 4 APPLICATIONS

### 4.1 Computer Vision

Applying EGL to solve image classification problems has become a hot and attractive research area in recent years [52, 54, 103], largely thanks to the popularity and advancement of visual explanation techniques [131, 171]. Depending on the nature of the image source, existing works can be further categorized into (1) general image prediction and (2) medical image analysis.

**4.1.1 General Image Prediction.** The application of EGL on general images typically involves image classification tasks on natural image data such as ImageNet [78], Caltech-UCSD Birds [154], Microsoft COCO [92], and Places365 [182], and some synthetic image data such as ToyColor [126], MNIST [83], and many MNIST variants, including Fashion-MNIST [163], Decoy-MNIST [126], and Color-MNIST [90]. The typical EGL technique used in this application domain is local explanation supervision and regularization, where the sample level visual explanation of the model is jointly

optimized together with the conventional prediction loss [39, 45, 52, 54, 80, 89, 103, 123, 125, 143, 146, 157]. When applying the explanation supervision techniques, the ground-truth explanation labels are typically collected from human annotators, and the additional attention loss is typically realized by a distance loss between the ground-truth and the model visual explanation at the sample level. For model explanation assessment, case studies are most commonly used for qualitative analysis [52, 54, 103], while IoU score is for quantitative evaluation [52, 54, 89].

**4.1.2 Medical Image Analysis.** Besides generic image applications, EGL has also been widely studied in the medical domain, thanks to the availability of domain-expert annotation on many medical image datasets [31, 82, 156]. In general, we observed a variety of datasets studied by existing works, including but not limited to the ISIC Skin Cancer dataset [31], Iris-Cancer dataset, scaphoid fracture detection dataset [82], Fundus image dataset [119], and the pneumonia detection X-ray dataset [156] for disease identification task [184]. Similarly to most EGL frameworks on generic image data, an additional explanation loss is added to the model objective and is typically realized by a distance loss between the ground-truth annotation collected from domain experts and the model visual explanation [179]. However, compared with generic image data, several unique challenges have been identified by existing works when applying EGL to medical images, such as (1) difficulty in assessing the quality of the model explanation, and (2) the scalability of the sample size of the annotation labels of the datasets.

## 4.2 Natural Language Processing

Interest has recently grown in applying EGL to designing NLP systems. Based on how the explanation is acquired, we have two categories of the application: (1) using the attention mechanism as the explanation and (2) using a generative model to generate the explanation.

**4.2.1 Attention Mechanism as the Explanation.** NLP systems generally use variants of attention mechanisms to get explanations. To evaluate the explanation, the ground-truth explanation labels are typically collected from human annotators (Stacey et al. [142] use TextRank to get ground-truth labels), and the evaluation metric can be the F1 score and IoU score based on token or snippet level. In addition to the agreement with human rationales, a faithful explanation is related to the downstream task performance, so rationale-level supervision is widely applied [28, 55, 142, 144, 175, 180]. Comprehensiveness and sufficiency are two main metrics regarding the influence of the explanation on the downstream task, Faithfulness, Data Consistency, and Confidence Indication are other diagnostic properties [11]. Attention mechanisms can learn to assign soft weights to token representations so that one can extract highly weighted tokens as rationales [38]. While this is intuitive for most of the NLP systems, the weights can be useless to give a faithful explanation because of the complex interaction of tokens. Another strand of works [24, 57] hard-select tokens or snippets from the input and only uses the selected part for the downstream task to get untangled explanation. This strand can be further divided into pipeline approaches and reinforcement learning approaches according to how the models are trained. Aligning the explanation with human annotators is not necessarily the optimal objective for improving model accuracy, the various loss strategies are proposed [24]. By, for example, masking out important explanation features of existing samples [168], one can augment the data. Liu et al. [94] propose global supervision by adding feature attribution prior to the total loss.

**4.2.2 Generative Model to Generate the Explanation.** In addition to giving explanations directly by the attention mechanism of NLP systems, a lot of works apply additional generative models to generate natural language explanations [17, 70, 87]. Although the rationales acquired from attention mechanisms provide concise and quick explanations, they may not have the means to provide

important details of the reasoning of a model. By using an additional natural language decoder, one can generate a comprehensive description of the decision-making process behind a prediction, some examples of the generative module include a CRF [167], a natural language decoder [23, 95], a GRU following an MLP [177], BiLSTM, and Transformer [141]. The commonly used datasets and corresponding tasks are ComVE [155] for commonsense validation, e-SNLI [23] for natural language inference, COSe [121] for commonsense question answering, e-SNLI-VE [75] for visual entailment, and VCR [169] for visual commonsense reasoning. The ground truth is usually also natural language explanations provided by humans. To evaluate the quality of natural language explanations, one can either use automatic metrics like METEOR [16], BERTScore [172], and BLEURT [130] or use human evaluation with metrics like e-ViL score [75], confidence, and readability. In terms of the faithfulness of the natural language explanations, Wiegreffe et al. [159] provides two necessary conditions: feature importance agreement and robustness equivalence.

### 4.3 Visual Question Answering

Attention and reasoning are two intertwined mechanisms underlying VQA tasks. Thanks to the widely used attention mechanism in VQA, applying EGL to help improve both the interpretability and performance of VQA tasks have become an important research area in recent years. The typical EGL technique used in VQA tasks is local explanation supervision and regularization, where the sample-level visual explanation of the model is jointly optimized together with the conventional prediction loss. When applying the explanation supervision techniques, the ground-truth explanation labels can be collected from human annotators [27, 50, 160] or generated by another model [120, 176]. There have been a lot of VQA datasets with annotations; some are annotated with human-generated questions and answers like MovieQA [148] and the VQA v1.0 dataset in Reference [7], while others are developed with synthetic scenes and rule-based templates like GQA [68], Clevr [73], and VCR [169]. VQA-2.0 [58] includes complementary images that lead to different answers, reducing language bias by forcing the model to use visual information. The AiR-D [27] is the first dataset of eye-tracking data collected from humans performing the VQA tasks. The VQA-HAT dataset [35] is a visual explanation dataset that collects human attention maps by giving human experts blurred images and asking them to determine where to deblur to answer a given visual question. VQA-CP [4] contains QA pairs whose distribution is significantly different between the training and test set. VQA-X [112] offers human textual explanations that can be used to determine important objects and then are grounded to important regions in the image as the explanation. The additional attention loss can be attention accuracy [27], false sensitivity rate [160], rank correlation loss [120, 176], and IoU loss [50].

### 4.4 Healthcare

EGL techniques have also been well explored in general healthcare applications, such as on gene interaction graph [59], Adult Changes in Thought [102], Mount Sinai Brain Bank, Religious Orders Study/Memory and Aging Project [1], and healthcare mortality prediction [101]. Specifically, Erion et al. [46] studied the tissue-specific gene interaction graph for the tissue most closely related to acute myeloid leukemia (AML, a type of blood cancer) in the HumanBase database [59] on how penalizing differences between the attributions of neighbors in an arbitrary graph connecting the features can be used to incorporate prior biological knowledge about the relationships between genes, yield more biologically plausible explanations of drug response predictions, and improve test error. They tested the model performance on a healthcare mortality prediction dataset [101], where the model inputs are 35 features representing patients' demographic information and medical data. Erion et al. [45] then further extended their previous study and proposed to add a graph attribution prior regularization on explanation to a two-layer neural network. In addition,

Weinberger et al. [157] extracted prior information from multiple gene expression datasets of the Accelerating Medicines Partnership Alzheimer’s Disease Project portal, incorporated meta-features in a gene–gene interaction graph and proposed a deep attribution prior framework to Alzheimer’s disease biomarker prediction. More recently, Zhao et al. [179] proposed an EGL framework for pulmonary nodule detection, utilizing public thoracic CT image datasets. This framework integrates the robust explanation supervision technique to ensure the performance of nodule classification and morphology. The method aims to reduce the workload of radiologists, enabling them to focus on the diagnosis and prognosis of potential cancerous pulmonary nodules at the early stage to improve outcomes for lung cancer patients. Following a similar direction, Zhang et al. [173] proposed a Multi-annotated Explanation-Guided Learning framework for explanation supervision, incorporating comprehensive and high-quality generated annotations by learning the characteristics of each annotator. Their experimental results demonstrate that the proposed method can significantly outperform all other methods.

#### 4.5 Chemistry

EGL has also started to see emerging applications in the chemistry domain, especially for molecular puppetry prediction tasks [97, 98, 145]. For instance, one recent work proposed an EGL framework for **Graph Neural Networks (GNNs)** by supervising their node- and edge-level explanation to align with domain expert annotation labels [53]. In this work, the authors studied three binary classification molecular datasets, namely (1) the **Blood-brain barrier penetration (BBBP)** dataset comes from a recent study [97] on the modeling and prediction of barrier permeability, (2) the BACE dataset provides quantitative and qualitative (binary label) binding results for a set of inhibitors of human b-secretase 1 [145], and the “**Toxicology in the 21st Century**” (TOX21) initiative created a public database measuring the toxicity of compounds [162]. The general goal for each dataset is identifying functional groups on organic molecules for biological molecular properties. Each dataset contains binary classifications of small organic molecules as determined by the experiment [67]. The experimental results suggest that the proposed GNES framework can effectively improve the reasonability of the explanation while still keeping or even improving the backbone GNNs model performance.

#### 4.6 Crime

EGL has also been studied in the application of risk and crime-related applications, where it is important to check if the model is leveraging reasonable features when predicting crime incidences or assessing the future risk of crime suspects. For instance, several works have studied the Propublica’s COMPAS Recidivism Risk Score datasets,<sup>2</sup> which contains data for predicting recidivism (i.e., whether a person commits a crime/a violent crime within 2 years) from many attributes [5, 123]. COMPAS dataset is designed for checking whether there exist biases in the mode explanation, such as the model’s treatment of the person’s race attribute when making the prediction. Specifically, Rieger et al. [123] proposed **contextual decomposition explanation penalization (CDEP)**, a method that enables practitioners to leverage explanations to improve the performance of a deep learning model. In particular, CDEP enables inserting domain knowledge into a model to ignore spurious correlations, and correct errors, and demonstrates the ability to increase performance on real datasets; Alvarez et al. [5] proposed an EGL framework by explicitly enforcing three basic desiderata for interpretability—explicitness, faithfulness, and stability—during training to enhance the robustness and interpretability of model explanations. Besides, Balayan et al. [15] studied a private online retailer fraud detection dataset with the proposed JOEL framework, a neural

<sup>2</sup>Available online at: [github.com/propublica/compas-analysis/](https://github.com/propublica/compas-analysis/)

network-based framework to jointly learn a decision-making task and associated explanations that convey domain knowledge. Specifically, JOEL is tailored to human-in-the-loop domain experts that lack deep technical ML knowledge, providing high-level insights about the model’s predictions that very much resemble the experts’ own reasoning. Moreover, they collect the domain feedback from a pool of certified experts and use it to ameliorate the model (human teaching), hence promoting seamless and better-suited explanations.

#### 4.7 Potential Future Domains of Applications

Despite the recent attention and major advance of EGL in the aforementioned popular application domains, there are still a number of open problems and potentially fruitful directions for future research and application of EGL, as follows:

**4.7.1 FaccT.** FaccT are becoming as important as—or depending on application areas—more important than model accuracy as an evaluation metric. Since it is nearly not feasible to prepare an impeccable dataset that can equally represent every possible feature related to a model’s task, blindly pursuing a model’s accuracy cannot exclude the chance of causing “catastrophic consequences” in critical circumstances [66]. One of EGL’s crucial application areas is to realize the balance between the model accuracy and FaccT by allowing human users to elicit their perspectives on steering the model. In shaping the balance, one crucial research direction is to understand how to maximize the case where reasonable human reasoning can also cause accurate prediction. There are several arguments discussing when human reasoning can cause a beneficial or detrimental effect on model prediction. While the debate is ongoing, we are gradually seeing more evidence where human involvement can result in a positive effect [32, 56]. For example, Shao et al. find that humans “arguing against” unreasonable explanation can benefit the model [134]. At the end of the day, from the perspective of model accuracy and FaccT, a railroad should not be the reason for predicting a train [85], a snowboard cannot be a male class [64], and a shopping cart should not only belong to a woman class [178].

**4.7.2 Adversarial Learning.** Adversarial perturbations can significantly drop the model’s accuracy. In a dramatic situation, it can reach nearly to 0%. Current ML models are vulnerable to adversarial attacks. Since the majority of adversarial attack shift model’s attention, applying EGL in detecting unusual shifts could be one of the solutions for developing a more robust ML model against adversarial attacks. However, in pursuing such a direction, the change of the attention map after the attack can be subtle from human eyes [20]. To apply EGL in the area of adversarial learning, we see devising better solutions in the following areas to be crucial. First, providing additional signals other than model attention can help human users effortlessly detect the attacked cases. Second, devising an advanced EGL mechanism that can (1) guide the users to generate effective input (2) and applying such input to improve the model’s robustness would be essential. Following this line of thought, very recently, Jeong et al. [71] proposed Generative Noise Injector for Model Explanations, a novel defense framework that perturbs model explanations to minimize the risk of model inversion attacks while preserving the interpretabilities of the generated explanations. Thus, we believe future studies on model explanation defense and attack can be one of the key research sub-areas of EGL domain.

**4.7.3 Continual and Active Learning.** EGL’s core principle is motivating ML engineers’ iterative training, such as continual learning [43, 127] and active learning [26, 74]; helping them to figure out the vulnerability through explanation and fixing the issue by providing a human-level guideline. In supporting such an iterative training, we believe one of the promising areas is “data iteration,” a design that can help ML engineers to fortify the dataset by adding more examples based on

Table 2. A List of Publicly Available Datasets for EGL with Human Explanation Annotation Labels

Dataset	Type	Link	Annotation Type
Gender Classification	Vision	<a href="https://github.com/YuyangGao/RES">https://github.com/YuyangGao/RES</a>	Pixel level
Scene Recognition	Vision	<a href="https://github.com/YuyangGao/RES">https://github.com/YuyangGao/RES</a>	Pixel level
Face Glasses Recognition	Vision	<a href="https://github.com/carriegu0818/EGL_benchmark">https://github.com/carriegu0818/EGL_benchmark</a>	Pixel level
Prohibited Item Detection	Vision	<a href="https://github.com/carriegu0818/EGL_benchmark">https://github.com/carriegu0818/EGL_benchmark</a>	Pixel level
ACT-X	Vision	<a href="https://github.com/Seth-Park/MultimodalExplanations">https://github.com/Seth-Park/MultimodalExplanations</a>	Pixel level and Textual
Caltech-UCSD Birds	Vision	<a href="https://authors.library.caltech.edu/27452/">https://authors.library.caltech.edu/27452/</a>	Pixel level(bounding box)
The PASCAL VOC Challenge 2007	Vision	<a href="http://host.robots.ox.ac.uk/pascal/VOC/voc2007/">http://host.robots.ox.ac.uk/pascal/VOC/voc2007/</a>	Pixel level
The PASCAL VOC Challenge 2012	Vision	<a href="http://host.robots.ox.ac.uk/pascal/VOC/voc2012/">http://host.robots.ox.ac.uk/pascal/VOC/voc2012/</a>	Pixel level
ISIC2018 Challenge	Vision	<a href="https://challenge.isic-archive.com/landing/2018/">https://challenge.isic-archive.com/landing/2018/</a>	Pixel level(bounding box)
Pneumonia Detection	Vision	<a href="https://www.kaggle.com/c/rsna-pneumonia-detection-challenge">https://www.kaggle.com/c/rsna-pneumonia-detection-challenge</a>	Pixel level(bounding box)
Movie Review	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Span-level rationale
MultiRC	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Single sentence-level rationale
FEVER	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Sentence-level rationale
BoolQ	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Token-level rationale
Evidence inference	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Sentence-level rationale
e-SNLI	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Token-level rationale
Commonsense Explanations (CoS-E)	NLP	<a href="https://github.com/jayded/eraserbenchmark">https://github.com/jayded/eraserbenchmark</a>	Sentence-level rationale
VQA-HAT	VQA	<a href="https://computing.ece.vt.edu/~abhsikdd/vqa-hat/">https://computing.ece.vt.edu/~abhsikdd/vqa-hat/</a>	Pixel level
GQA	VQA	<a href="https://cs.stanford.edu/people/dorarad/gqa/about.html">https://cs.stanford.edu/people/dorarad/gqa/about.html</a>	Pixel level
VQA-X	VQA	<a href="https://github.com/Seth-Park/MultimodalExplanations">https://github.com/Seth-Park/MultimodalExplanations</a>	Pixel level and Textual
VQS	VQA	<a href="https://github.com/Cold-Winter/vqs">https://github.com/Cold-Winter/vqs</a>	Pixel level
BBBP	Graph	<a href="https://github.com/YuyangGao/GNES">https://github.com/YuyangGao/GNES</a>	Node and edge level
BACE	Graph	<a href="https://github.com/YuyangGao/GNES">https://github.com/YuyangGao/GNES</a>	Node and edge level
TOX21	Graph	<a href="https://github.com/YuyangGao/GNES">https://github.com/YuyangGao/GNES</a>	Node and edge level

detected vulnerabilities through explanation. In such a direction, we believe understanding the pros and cons of retraining and continual learning can be crucial. For example, there can be a case where newly found data points can be stacked up on an existing dataset and be used in retraining [14]. Another case can be to iteratively update the last model through some of the existing techniques in continual learning [111]. In general, in the world of EGL, understanding when to apply retraining or continual learning and what are the pros and cons of each training strategy are not well understood. Understanding which strategy can yield what strengths and weaknesses in the scenario of data iteration would be one of the core future applications of EGL.

**4.7.4 Contrastive Learning.** Contrastive learning is a powerful self-supervised learning strategy that encourages augmentations of the same input to have more similar representations compared to augmentations of different inputs. In the field of EGL, we have started to see several works that apply the contrastive objective to the model explanation between similar/dissimilar samples to build up the explanation objective [40, 114, 139, 168]. The most significant advantage of leveraging the contrastive learning paradigm for explanation guidance is that no ground-truth explanation annotation labels are required for model training. However, designing an appropriate contrastive framework for EGL can be more challenging due to the lack of a standard form of model explanation under different application domains. Besides, how to define and formulate the positive and negative explanation samples to contrast with the anchor sample’s explanation can be challenging without knowing the ground-truth labels. Thus, we believe the further development of the contrastive EGL framework can be one of the core future directions in EGL, and it can lead to a significant leap in the application of EGL to the domains where ground-truth explanation labels are generally difficult to obtain in large scale.

## 5 EXPERIMENTS

This section aims at providing an extensive and comprehensive experimental study among existing EGL models in various popular application domains. Specifically, the comparative studies of four datasets from the CV domain [174], namely (1) Gender Classification, (2) Scene Recognition, (3) Face Glasses Recognition, and (4) Prohibited Item Detection, and three datasets from the NLP, namely (1) Movie Review, (2) MultiRC, and (3) FEVER, are provided. The details about each dataset are included in Table 2, where a full list of publicly available datasets for EGL is provided.

## 5.1 Visual Explanation Guided Learning

5.1.1 *Gender Classification* [52]. The gender classification task is derived from the Microsoft COCO dataset<sup>3</sup> [92]. Images containing the keywords “men” or “women” in their captions were initially selected. Subsequently, the dataset underwent a rigorous filtering process to ensure its suitability for a single person gender classification task. Images were excluded if they exhibited one or more of the following characteristics: (1) captions containing references to both genders, (2) images featuring multiple individuals, or (3) images where humans were not clearly recognizable [54]. To maintain data integrity and prepare it for human explanation annotation, a subset of the images underwent manual annotation by human annotators. The resulting dataset comprises a total of 1,736 images with human explanation annotations, demonstrating an even distribution of female and male subjects. To simulate a scenario where human explanation annotation labels are limited, only a random sample of 100 images was used for the training set, while the validation and test sets each consists of 700 images.

5.1.2 *Scene Recognition Dataset* [52]. The scene recognition dataset is originally derived from the Places365 dataset<sup>4</sup> [182] and manually annotated by Gao et al. [52]. The task for this dataset is a binary classification of scene recognition: nature vs. urban. Specifically, the categories used to sample the data are as follows:

- *Nature*: mountain, pond, waterfall, field wild, forest broadleaf, rainforest
- *Urban*: house, bridge, campus, tower, street, driveway

The dataset consists of a total of 2,086 images with human explanation annotation labels. Similarly, we split the data randomly with a sample size of 100/700/700 for training, validation, and testing.

5.1.3 *Face Glasses Recognition*. We construct the glasses recognition dataset from the CelebAMask-HQ dataset<sup>5</sup> [84] by categorizing face images with and without glasses. In CelebAMask-HQ, masks were manually annotated with 19 classes including all facial components and accessories. The rationale of the task is that we are able to obtain factual annotation labels by the segmentation of eyes and glasses directly. While the original dataset is highly imbalanced in the ratio between faces with and without glasses, we randomly select an equal number of images in both classes, with a total of 100/393/392 images for training/validation/testing, respectively.

5.1.4 *Prohibited Item Detection*. The task is constructed from the Sixray dataset<sup>6</sup> [100] by splitting images based on the presence of prohibited items. Sixray is highly imbalanced with 1,059,231 X-ray images, including six classes of 8,929 prohibited items. Merging the six prohibited classes, the task of the new dataset is a binary prohibited item detection. Bounding boxes of prohibited items are included in all images. Due to data imbalance, the dataset is further filtered into 100/5296/5298 images for training, validation, and testing, respectively.

5.1.5 *Evaluation Metrics*. We evaluate the model in terms of prediction performance as well as in terms of explanation performance. For prediction performance, we use AUC and accuracy as evaluation metrics. To evaluate explanation faithfulness, we employ the Matrix for comprehensiveness and sufficiency by ERASER [38]. For explanation alignment assessment, we compare the saliency map generated by GradCAM with ground-truth annotation masks. Specifically, we use the IoU score [18], the bitwise intersection and union operations between the ground-truth explanation and the binarized model explanation. We further evaluate explanation performance with

<sup>3</sup>Available at: <https://cocodataset.org/>

<sup>4</sup>Available at: <http://places2.csail.mit.edu/index.html>

<sup>5</sup>Available at: [http://mmlab.ie.cuhk.edu.hk/projects/CelebA/CelebAMask\\_HQ.html](http://mmlab.ie.cuhk.edu.hk/projects/CelebA/CelebAMask_HQ.html)

<sup>6</sup>Available online at: <https://github.com/MeioJane/SIXray>

Explanatory F1, precision, and recall by bitwise comparison between ground-truth explanation and model explanation.

**5.1.6 Comparison Methods.** We compare the performance of several models as listed below:

- *Baselines*: Baseline 1 and 2 are pre-trained ResNet50 and VGG16 model that trains only on prediction loss without explanation loss.
- *Supervised EGL models*:
  - *GRADIA* [54]: A framework that trains the DNN model with both the prediction loss as well as a conventional L1 loss that directly minimizes the distance between the continuous model explanation and the binary positive explanation labels.
  - *RES* [52]: A framework that trains the DNN with both factual and counterfactual annotations with two imputation functions:  $g(\cdot)$  as a fixed value Gaussian convolution filter and learnable imputation function  $g_\phi(\cdot)$  via multiple layers of learnable kernels.
  - *CDEP* [123]: A framework that incorporates Contextual Decomposition to penalize spurious correlations and therefore correct errors.
  - *RRR* [129]: A framework that was initially introduced in Reference [126] and altered in Reference [129], which aims to regularize the model to be right for the right reasons.
- *Unsupervised EGL models*:
  - *SGT* [69]: A framework that introduces saliency-guided training for neural networks to reduce noisy gradients in predictions.
  - *SENN* [5]: A framework that applied two regularization terms on the model explanation: (1) neighborhood-fidelity and (2) stability based on the neighborhood of input.

**5.1.7 Implementation Details.** All models are trained for 50 epochs with the same train/val/test split as mentioned above. We use the ADAM optimizer with a learning rate of 0.0001 [77]. The architecture of each model is listed in Table 3. To better compare the performance on explainability, the model explanations are generated by GradCAM [131]. When calculating the explanation evaluation metrics, the explanation maps were further binarized by a fixed threshold of 0.5. We use a batch size of 32 for training and 100 for testing. For GRADIA and RES, we set the slack variable  $\alpha$  to 0.1 and 0.01, respectively, and the regularization factor to 0. For RRR, we set the regularization parameter to 1. For CDEP, we set the regularizer rate to 0, 0.1, and 10. For SENN, we set the robust regularization, sparsity regularization, and concept regularization hyperparameters to 0.0001, 0.00002, and 1, respectively. For SGT, we set features dropped to 0.1 and 0.3.

**5.1.8 Quantitative Analysis.** Model prediction performance and explanation quality in the domain of computer vision are presented in Table 3. We evaluate six models through gender classification, scene recognition, glasses identification, and prohibited item discovery. We evaluate two baseline model performances: ResNet50 and VGG16, as they are employed by the selected paper. Overall, supervised models demonstrate better prediction power and higher explanation quality than unsupervised models. ResNet50 seems to perform slightly better in prediction, and significantly better in explanation quality compared with VGG16.

For gender classification, GRADIA generally has the best prediction performance and presents the highest explanation quality, with the best scores in all metrics besides comprehensiveness and Exp recall, and minimal differences of 0.9% and 2.1% in comprehensiveness and explanatory recall. For models with ResNet50 as the backbone architecture, GRADIA and RES significantly outperform SGT, since SGT is unsupervised. SGT presents a lower accuracy, comprehensiveness, IoU, and explanatory F1 than the baseline model, which implies that the unsupervised model is not improving model performance and explanation quality. Yet SGT achieves the lowest sufficiency, which measures how well the prediction aligns between the original input and an explanation-generated

Table 3. The Classification Performance and Explanation Evaluation on Visual Explanation Guided Learning

Model	Architecture	Prediction		Exp Faithfulness		Exp Alignment			
		Acc. $\uparrow$	AUC $\uparrow$	Comp. $\uparrow$	Suff. $\downarrow$	IoU $\uparrow$	F1 $\uparrow$	Precision $\uparrow$	Recall $\uparrow$
Gender Classification									
Baseline 1	ResNet50	0.680	0.664	0.048	0.105	0.147	0.470	0.554	0.525
Baseline 2	VGG16	0.637	0.671	0.048	0.051	0.058	0.155	0.340	0.132
GRADIA [54]	ResNet50	<b>0.695</b>	<b>0.764</b>	0.107	0.090	<b>0.243</b>	<b>0.625</b>	<b>0.787</b>	0.608
RES [52]	ResNet50	0.690	0.744	<b>0.108</b>	0.097	0.240	0.614	0.742	<b>0.621</b>
RRR [129]	VGG16	0.624	0.628	0.027	0.015	0.091	0.270	0.424	0.257
CDEP [123]	VGG16	0.628	0.635	0.014	0.010	0.100	0.254	0.419	0.231
SENN [5]	SENN(CNN)	0.589	0.627	0.005	0.027	0.062	0.186	0.300	0.186
SGT [69]	ResNet50	0.645	0.687	0.012	<b>0.002</b>	0.044	0.352	0.502	0.257
Scene Recognition									
Baseline 1	ResNet50	0.947	0.965	0.068	0.255	0.397	0.702	0.906	0.628
Baseline 2	VGG16	0.953	<b>0.988</b>	0.134	0.117	0.191	0.324	0.890	0.226
GRADIA [54]	ResNet50	0.952	0.987	<b>0.255</b>	0.073	0.378	0.606	0.912	0.501
RES [52]	ResNet50	<b>0.956</b>	<b>0.988</b>	0.189	<b>0.002</b>	<b>0.435</b>	<b>0.722</b>	<b>0.909</b>	<b>0.647</b>
RRR [129]	VGG16	0.953	0.987	0.014	0.019	0.224	0.364	0.925	0.250
CDEP [69]	VGG16	0.934	0.952	0.026	0.039	0.127	0.232	0.807	0.153
SENN [123]	SENN(CNN)	0.733	0.798	0.022	0.042	0.082	0.183	0.721	0.108
SGT [5]	ResNet50	0.937	0.985	0.164	0.039	0.056	0.301	0.796	0.213
Face Glasses Recognition									
Baseline 1	ResNet50	0.991	<b>0.999</b>	0.302	0.163	0.134	0.971	0.998	0.954
Baseline 2	VGG16	<b>0.996</b>	0.864	0.183	0.039	0.299	0.873	0.987	0.804
GRADIA [54]	ResNet50	0.990	<b>0.999</b>	0.368	0.262	<b>0.375</b>	<b>0.949</b>	0.993	<b>0.917</b>
RES [52]	ResNet50	0.991	<b>0.999</b>	<b>0.396</b>	0.128	0.302	0.932	<b>0.997</b>	0.887
RRR [129]	VGG16	0.994	<b>0.999</b>	0.384	0.160	0.332	0.909	0.992	0.864
CDEP [123]	VGG16	<b>0.996</b>	<b>0.999</b>	0.004	<b>0.003</b>	0.042	0.203	0.540	0.248
SENN [5]	SENN(CNN)	0.797	0.873	0.002	0.023	0.045	0.202	0.639	0.134
SGT [69]	ResNet50	<b>0.996</b>	<b>0.999</b>	0.083	0.106	0.292	0.671	0.974	0.556
Prohibited Item Detection									
Baseline 1	ResNet50	0.961	0.992	0.161	0.053	0.195	0.823	0.870	0.784
Baseline 2	VGG16	0.917	0.988	-0.026	<b>-0.010</b>	0.147	0.330	0.788	0.241
GRADIA [54]	ResNet50	<b>0.974</b>	<b>0.997</b>	<b>0.176</b>	0.272	0.213	0.703	0.928	0.610
RES [52]	ResNet50	0.962	0.995	0.152	0.295	<b>0.235</b>	<b>0.837</b>	<b>0.964</b>	<b>0.776</b>
RRR [129]	VGG16	0.950	0.995	0.055	0.128	0.155	0.343	0.448	0.320
CDEP [123]	VGG16	0.959	0.992	0.038	0.021	0.061	0.403	0.615	0.298
SENN [5]	SENN(CNN)	0.754	0.839	-0.026	0.095	0.042	0.152	0.584	0.098
SGT [69]	ResNet50	0.962	0.993	0.027	0.066	0.062	0.552	0.648	0.481

The results are obtained from 3 individual runs and the best results of each metric are highlighted in bold.

input. Models with VGG16 as the backbone report lower sufficiency than those with ResNet50, while models with ResNet50 generally hold higher accuracy, comprehensiveness, IoU, and explanatory F1. SENN, which develops its own architecture with a set of conceptizer, parametrizer, and aggregator, underperforms in all metrics, since the backbone is a simple CNN model.

For scene recognition, the baseline VGG16 achieves better accuracy, comprehensiveness, and sufficiency, compared with the baseline ResNet50. VGG16 has a significantly low explanatory recall and IoU, which results in 53.8% worse performance in explanatory F1. For the selected models, RES yields the best performance on all metrics, slightly improving prediction performance and

boosting explanation quality significantly, with 1.0%, 2.4%, 177.9%, -99.2%, 9.6%, 2.8% changes in accuracy, AUC, comprehensiveness, sufficiency, IoU, and explanatory F1, respectively. SENN consistently underperforms in all metrics. Among models with VGG16, RRR is able to maintain a similar accuracy as the baseline while improving explanation alignment (IoU and Exp F1) by 17.3% and 12.3%, while CDEP results in worse explanation alignment. RRR and CDEP both obtain a lower score in comprehensiveness and sufficiency, which suggests that even stripping off the model-generated explanations, the models are able to generate similar predictions. The stripped model-generated explanations are useful in terms of prediction, but not all useful information is covered by the saliency map. Moreover, baseline ResNet50 under-performs in terms of explanation faithfulness but outperforms in terms of explanation alignment. This implies that while ResNet50 is able to generate explanations that exhibit the pattern of human explanation annotation, the generated explanations are not useful for the model in terms of prediction.

In the task of the glasses identification, all models achieve high performance besides SENN, which consistently results in worse performance in all metrics. In terms of explanation comprehensiveness, and faithfulness, GRADIA, RES, RRR, and CDEP show (21.9%, 60.7%), (31.1%, -21.5%), (109.8%, 310.3%), and (-97.8%, -92.3%) changes with respect to the baseline. RES holds the highest comprehensiveness score and CDEP holds the lowest sufficiency score. This implies that GRADIA, RES, and RRR are successful at extracting all useful attention for prediction, while GRADIA and RRR sacrifice the prediction power if solely using generated attention as the input. Among supervised models, GRADIA has the highest IoU as well as the highest percentage increase, which implies that the model successfully learns the pattern of human explanation annotation and is able to produce saliency maps that are most aligned with human explanation annotation. However, since GRADIA also has the highest sufficiency score, it further implies that learning from the human explanation annotation may not be sufficient for the model to make correct predictions. Meanwhile, SGT shows high accuracy and IoU even as an unsupervised model. Yet it suffers from low Explanatory recall, low comprehensiveness, and high sufficiency score, indicating that the model-generated attention is not successful in terms of prediction.

In the Prohibited Item Detection task, models generally exhibit a pattern of maintaining high accuracy, better explanation alignment, and comprehensiveness but worse sufficiency. ResNet50 baseline achieves higher prediction and explanation robustness compared with VGG16. While most models' accuracy ranges from 0.917 to 0.974, SENN yields an accuracy of 0.754, which implies that a CNN model is insufficient for the dataset. All models with VGG16 as the backbone performs well in terms of low sufficiency but poorly in comprehensiveness and explanatory recall. GRADIA yields the highest accuracy (0.974), AUC (0.997), and comprehensiveness score (0.176) but poor sufficiency score of 0.272 when the baseline has a sufficiency of 0.053. RES shows the worst good performance on sufficiency but second to best score on comprehensiveness among EGL models. The baseline model with VGG16 achieves the best sufficiency score of -0.010. This is because the model sufficiency score is highly influenced by the size of generated annotation map, and the VGG16 baseline model sacrifices the sparseness of explanation to achieve a higher sufficiency and consequently leads to the worst comprehensiveness score. To validate the above statement, we further compute the proportion of attention map generated respectively to the entire image of the VGG16 baseline, RES, and GRADIA models. We find that the average explanation map sizes of the VGG16 baseline model are on average 84% and 33% greater than those of RES and GRADIA, respectively. This additional observation provides additional support for our assumption that the baseline model tends to generate larger explanation maps, leading to a much higher sufficiency score but a much worse comprehensiveness score. In terms of explanation alignment, RES and RRR are able to improve IoU from 0.195 to 0.101 and from 0.147 to 0.155, the best among each architecture. RES improves explanatory F1 from 0.827 to 0.837 and CDEP improves explanatory F1

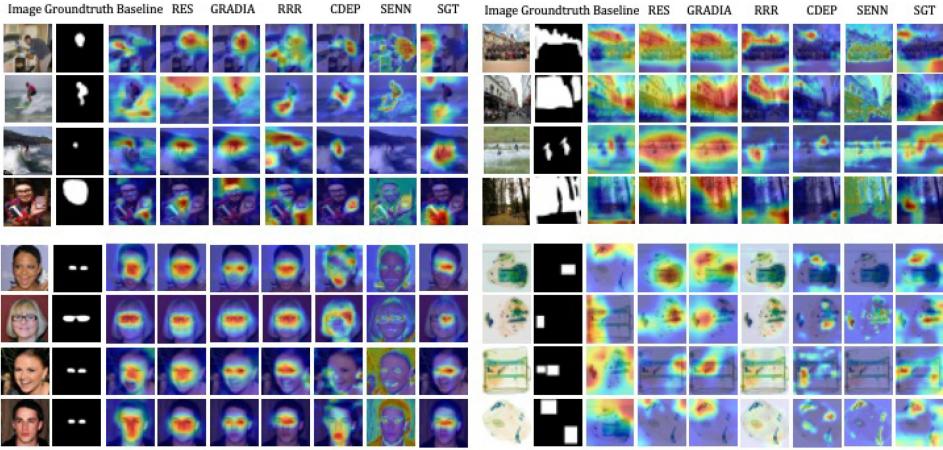


Fig. 5. Selected explanation visualization results on four vision datasets: gender classification, scene recognition, face glasses detection, and prohibited item identification. The model-generated explanations are highlighted.

from 0.330 to 0.403. Overall, GRADIA achieves the best prediction accuracy and faithfulness while RES achieves the best explanatory alignment among all models.

**5.1.9 Qualitative Case Study.** Figure 5 displays the visualization results of the four vision tasks: gender classification, scene recognition, face glasses detection, and prohibited item identification. Overall, RES and GRADIA have more overlap with the ground truth. CDEP is more fine-grained. SENN generates a saliency map that highlights all areas instead of focusing on a particular place. In the gender classification task, RES and GRADIA perform well in identifying the human body, even with distractions. RRR, CDEP, and SGT sometimes highlight areas that are disruptive, such as phones and kitchen stoves. In terms of saliency map size, CDEP generates a saliency map that focuses only on a small area, whereas SENN generates a thin layer of attention all over the image, which explains their low performance in IoU and explanatory F1. In the scene recognition task, RES, GRADIA, and SENN have the most overlapping areas with the ground-truth label. RRR and CDEP sometimes are biased. For example, in the last image, RRR considers the floor key elements in scene recognition, whereas the ground-truth label is the trees. SGT mostly attends to areas other than the ground truth, which explains its low IoU compared with other models. For face glasses detection, all models generate saliency maps focused on the face area. While the baseline model focus on the entire face, RES, GRADIA, RRR, and SGT generate maps that are more specific to the eye areas. CDEP focuses more on the lower half of the face and SENN attends to all face parts, such as the eyes, mouth, and so on. While all models are highly accurate at identifying the presence of prohibited items, model-generated maps do not align with ground-truth labels. Most models are able to focus on some small objects, not necessarily the prohibited items. SGT focuses on the white area outside of the baggage, which accounts for its low accuracy and low explanation performance.

## 5.2 Rationale Attention Guided Learning

To evaluate the performance of EGL models on NLP tasks, three datasets with explanation rationales are selected for the experimental study [38], and the details about each dataset are presented as follows.

5.2.1 *Movie Review* [52]. The movie review dataset includes binary sentiment labels as well as rationale annotations at the span level. The task is to classify movies with positive sentiments from those with negative sentiments. We randomly split the data with a sample size of 1600/150/200 for training, validation, and testing. The explanation label is the sentiment  $\in \{\text{positive}, \text{negative}\}$ .

5.2.2 *MultiRC* [76]. MultiRC is a reading comprehension dataset originally composed of a series of rationale/question/answer triplets. This is also a binary classification task where the prediction label is True or False. The dataset is divided into 24029/3214/4848 for training, validation, and testing. The ground truth for explanation indicates if the answer is correct.

5.2.3 *FEVER* [150]. FEVER is a fact verification dataset, where each claim can be classified into supported, refuted, or not enough information. DeYoung et al. [38] further took a subset of the dataset and included only support and refuted claims. The dataset is further separated into 97957/6122/6111 images for training, validation, and testing respectively. For explanation rationales, the model has to predict the veracity of a claim  $\in \{\text{support}, \text{refuse}\}$ .

5.2.4 *Evaluation Metrics*. We evaluate the model in two categories: (1) prediction performance and (2) explanation performance. For prediction performance, accuracy and AUC are computed to evaluate the predictive power of the model. For explanation evaluation, we incorporated six matrices to fully examine the explanation robustness. Matrix for comprehensiveness and sufficiency are derived from ERASER [38]. In addition, we measure the token level IoU [18] between ground-truth rationale and predicted rationale through IoU, explanatory F1, precision, and recall.

5.2.5 *Comparison Methods*. We compare the performance of several models as follows:

- *Baseline*: Baselines 1, 2, and 3 are pre-trained models that train only using the prediction loss without explanation loss. The pre-trained architecture for baselines 1, 2, and 3 are BERT+MLP, BERT+LSTM, and BERT+BERT, respectively.
- *ERASER* [38]: A pipeline model that first trains the encoder to extract rationales and then trains the decoder to perform prediction using only rationales.
- *Glockner et al.* [57]: A differentiable training framework that aims to output faithful rationales on a sentence level.
- *Carton et al.* [24]: A model that applies sentence-level rationale supervision, non-occluding “importance embeddings” on selective rationales with high sufficiency-accuracy.
- *Expred* [177]: A novel explanation generation framework work using multi-task learning that is task aware and can exploit rationales data for effective explanations.
- *FRESH* [70]: A model that aims to produce faithful rationales for neural text classification by defining independent snippet extraction and prediction modules.

5.2.6 *Implementation Details*. The data preprocessing follows the setting of ERASER [38]. We train all the models equally for 20 epochs, and Adam is used for optimization with a learning rate of  $2e-5$ . To evaluate the explanation performance, the threshold for the calculated rationales is set to be 0.5. We follow the hyperparameter settings reported in the papers of the above methods.

5.2.7 *Quantitative Analysis*. Table 4 presents the model prediction performance and explanation quality of the Movie Review, MultiRC, and FEVER datasets. The best results for each dataset are highlighted with bold. In general, when comparing with the baseline, all models achieve better accuracy and explanation alignment. The sufficiency score also decreases compared with the baseline model, which implies that the model-generated rationale is representative of the entire document.

Table 4. The Classification Performance and Explanation Evaluation on NLP Tasks

Model	Architecture	Prediction		Exp. Faithfulness		Exp. alignment		
		Acc. $\uparrow$	AUC $\uparrow$	Comp. $\uparrow$	Suff. $\downarrow$	IoU $\uparrow$	F1 $\uparrow$	Precision $\uparrow$
Movie Review								
Baseline 1	BERT+MLP	0.516	0.478	0.086	0.145	0.242	0.365	0.441
Baseline 2	BERT+LSTM	0.622	0.591	0.027	0.126	0.043	0.112	0.462
Baseline 3	BERT+BERT	0.756	0.703	0.112	0.113	0.085	0.188	0.411
ERASER [38]	BERT+LSTM	0.826	0.805	0.128	0.093	0.598	0.749	0.734
Glockner et al. [57]	BERT+MLP	0.564	0.511	0.114	0.103	0.541	0.702	0.693
Carton et al. [24]	BERT+BERT	<b>0.834</b>	<b>0.812</b>	0.138	0.084	0.585	0.738	0.726
Expred [177]	BERT+GRU+MLP	0.794	0.779	0.094	<b>0.076</b>	<b>0.639</b>	<b>0.779</b>	<b>0.781</b>
FRESH [70]	BERT+LSTM	0.678	0.653	<b>0.144</b>	0.093	0.569	0.726	0.745
MultiRC								
Baseline 1	BERT+MLP	0.564	0.511	0.012	0.188	0.235	0.459	0.534
Baseline 2	BERT+LSTM	0.593	0.573	0.081	0.205	0.106	0.280	0.471
Baseline 3	BERT+BERT	0.627	0.580	0.054	0.154	0.076	0.234	0.485
ERASER [38]	BERT+LSTM	0.639	0.615	0.039	0.132	0.448	0.618	0.615
Glockner et al. [57]	BERT+MLP	0.587	0.547	0.065	0.136	0.409	0.580	0.576
Carton et al. [24]	BERT+BERT	<b>0.647</b>	0.613	0.074	0.076	<b>0.473</b>	<b>0.642</b>	<b>0.633</b>
Expred [177]	BERT+GRU+MLP	0.638	<b>0.622</b>	0.032	<b>0.061</b>	0.447	0.618	0.602
FRESH [70]	BERT+LSTM	0.607	0.586	<b>0.096</b>	0.113	0.437	0.608	0.613
Fever								
Baseline 1	BERT+MLP	0.822	0.803	0.075	0.126	0.103	0.319	0.513
Baseline 2	BERT+LSTM	0.851	0.822	0.022	0.099	0.157	0.391	0.454
Baseline 3	BERT+BERT	0.872	0.856	0.017	0.117	0.036	0.145	0.612
ERASER [38]	BERT+LSTM	0.874	0.867	0.036	0.053	0.679	0.808	0.805
Glockner et al. [57]	BERT+MLP	0.835	0.813	<b>0.122</b>	0.066	0.672	0.803	<b>0.833</b>
Carton et al. [24]	BERT+BERT	0.893	0.876	0.084	0.048	<b>0.707</b>	<b>0.828</b>	0.831
Expred [177]	BERT+GRU+MLP	<b>0.903</b>	<b>0.889</b>	0.043	<b>0.027</b>	0.696	0.820	0.817
FRESH [70]	BERT+LSTM	0.862	0.832	0.106	0.053	0.627	0.771	0.732

The best results of each metric are highlighted in bold.

For the Movie Review dataset, Carton et al. [24] yields the highest classification accuracy and Expred [177] generates the explanations with the highest quality. Compared with the baseline architecture BERT+LSTM, ERASER [38] improve the model accuracy and AUC by 32.8% and 36.2% and boost the explanation quality by 374.1%, -26.2%, 1290.7%, and 568.8% in terms of comprehensiveness, sufficiency, IoU, and explanatory F1 scores, respectively, while FRESH [70] improves model accuracy, AUC, Sufficiency, and Exp F1 by 9.0%, 10.5%, 433.3%, -26.2%, 1223.3%, and 548.2% respectively. ERASER [38] has better performance in terms of both model performance as well as explanation quality. Carton et al. [24], which employs the BERT+BERT architecture, increases accuracy and AUC by 10.3% and 15.5%, and ERASER [38] achieves the second-best result with an architecture of BERT+LSTM. Expred [177] obtain the highest explanation alignment and lowest sufficiency with a model architecture of BERT+GRU+MLP, and FRESH [70](BERT+LSTM) holds the highest comprehensiveness score among the selected models.

For the MultiRC dataset, Carton et al. [24] achieves the highest classification accuracy as well as the highest explanation faithfulness and alignment. It improves the model accuracy and AUC by 3.2% and 15.5%, lowers sufficiency score by 50.6%, and boosts IoU and explanation F1 by 522.4% and 174.4%, respectively, compared with the baseline. For all models with the architecture BERT+LSTM, while they consistently obtain better results than baseline except for comprehensiveness ERASER [38], outperforms FRESH [70] by 5.3%, 4.9%, 2.5%, and 1.6% in terms of model accuracy, AUC, IoU, and explanatory F1. FRESH [70] is more accurate when assessing explanation faithfulness through the sufficiency and comprehensiveness score, with a 14.4% decrease and 146.2% boost compared with ERASER [38].

The performance varies for the FEVER dataset, as FRESH [70] achieves the highest accuracy, AUC, and sufficiency scores, and Expred [177] yields the highest comprehensiveness, IoU, Explanatory F1, and Explanatory Recall. All the models perform generally well in the fact verification

Ground Truth	Baseline	ERASER	Glockner et al.	Carton et al.	Expred	FRESH
... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .	... the canadians can make a good movie . the world is coming to an end . we do n't know why or how , but apparently there is no way to stop it . the world has had this information for months , as most of the rioting and other assorted chaos has passed and governments have shut down operations ... these people 's lives however all intersect during their final six hours . writer - director - star don mckellar has crafted a highly unique and emotional film .
Prediction: Positive	Prediction: Negative	Prediction: Positive	Prediction: Positive	Prediction: Positive	Prediction: Positive	Prediction: Negative

Fig. 6. Selected explanation visualization results on FEVER dataset. The model explanations are highlighted.

Ground Truth	Baseline	ERASER	Glockner et al.	Carton et al.	Expred	FRESH
... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died	... fossils in younger rocks look like animals and plants that are living today . fossils in older rocks are less like living organisms . fossils can tell us about where the organism lived . was it land or marine ? fossils can even tell us if the water was shallow or deep . fossils can even provide clues to ancient climates . what can we tell about former living organisms from fossils ?    how they died
Prediction: False	Prediction: True	Prediction: False	Prediction: True	Prediction: False	Prediction: False	Prediction: True

Fig. 7. Selected explanation visualization on Movie Review dataset. The model explanations are highlighted.

task in terms of accuracy, with a range of 0.835 to 0.903. In terms of explanatory faithfulness, FRESH [70] performs worse than the baseline in comprehensiveness but reduces sufficiency by 46.5%. Expred [177] obtains the greatest boost in comprehensive and sufficiency, with a change of 394.1% and -59.0% respectively. The baseline models generally show poor performance in explanation faithfulness and alignment, which are improved significantly across all five models. Expred [177] is able to improve IoU by 1863.9% and Explanatory F1 by 471.0%.

**5.2.8 Qualitative Case Study.** Figures 6, 7, and 8 provide examples of visualization results on the FEVER, Movie Review, and MultiRC datasets. The model-generated explanations are highlighted. In general, the baseline model highlights areas that are scattered all around the corpus, whereas trained models generate explanation rationales that are more aggregated. In Figure 6, ERASER [38] and Glockner et al. [57] are highly aligned with ground truth, aligned with their high performance in IoU. While Expred [177] obtains the highest accuracy and comprehensiveness, its generated-explanation does not align with the ground-truth annotations, which implies that the ground-truth labels may not be sufficient for the model to learn the prediction. FRESH [70] generates explanations that are poorly aligned with the ground truth and outputs a wrong prediction label.

In Figure 7, while Carton et al. [24] aligns well with the ground truth, it focuses on a higher percentage of tokens, which explains why it slightly underperforms in explanatory precision and comprehensiveness but outperforms in sufficiency. There exhibits a compromise between high

Ground Truth	Baseline	ERASER	Glockner et al.	Carton et al.	ExpreD	FRESH
Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.	Mount Hood , called Wy'east by the Multnomah tribe , is a potentially active stratovolcano in the Cascade Volcanic Arc of northern Oregon . It was formed by a subduction zone on the Pacific coast and rests in the Pacific Northwest region of the United States . It is located about 50 mi east-southeast of Portland , on the border between Clackamas and Hood River counties ... Mount Hood is in the Pyrenees.
Prediction: REFUTES	Prediction: SUPPORTS	Prediction: REFUTES	Prediction: REFUTES	Prediction: REFUTES	Prediction: REFUTES	Prediction: SUPPORTS

Fig. 8. Selected explanation visualization results on MultiRC dataset. The model explanations are highlighted.

accuracy and high explanation quality, as Carton et al. [24] achieves the highest accuracy but lowest comprehensiveness among the selected models. This example shows how the amount of attention may manipulate the result of explanation faithfulness. If a high amount of tokens are considered important, then sufficiency will be close to 0 and comprehensiveness will be relatively high. Therefore, it is necessary to consider both explanation faithfulness and alignment when analyzing the explanation quality. This example also reveals the importance of a case study to visualize the quantitative results and understand how attention performs in terms of alignment and faithfulness.

## 6 CONCLUSION

This survey has presented a comprehensive survey of existing methodologies developed in the field of EGL, a group of techniques that applies XAI-driven insights to steer the DNNs' behavior in realizing iterative model revision. It provides an extensive overview of the EGL challenges, techniques, applications, evaluation procedures, as well as extensive experimental comparison among existing techniques under popular application areas. It summarizes the findings of the research presented in more than 150 publications on EGL, the majority of which were released since 2018. Concretely, in this survey, the formal definition of EGL and its general learning paradigm is first given, along with an overview of the key factors for EGL evaluation, as well as summarization and categorization of existing evaluation procedures and metrics for EGL are provided. Based upon the numerous historical and state-of-the-art works discussed in this survey, the article concludes by discussing the current and potential future application areas of EGL and provides an extensive experimental study that aims at providing the first comprehensive comparative study among existing EGL models in various popular application domains, such as the CV and NLP domains.

## REFERENCES

- [1] David A. Bennett, Julie A. Schneider, Zoe Arvanitakis, and Robert S. Wilson. 2012. Overview and findings from the religious orders study. *Curr. Alzheimer Res.* 9, 6 (2012), 628–645.
- [2] Amina Adadi and Mohammed Berrada. 2018. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* 6 (2018), 52138–52160.
- [3] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. In *Advances in Neural Information Processing Systems*, Vol. 31, 9525–9536.

- [4] Aishwarya Agrawal, Dhruv Batra, Devi Parikh, and Aniruddha Kembhavi. 2018. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4971–4980.
- [5] David Alvarez Melis and Tommi Jaakkola. 2018. Towards robust interpretability with self-explaining neural networks. In *Advances in Neural Information Processing Systems*, Vol 31, 7786–7795.
- [6] David Alvarez-Melis and Tommi S. Jaakkola. 2018. On the robustness of interpretability methods. arXiv:1806.08049. Retrieved from <https://arxiv.org/abs/1806.08049>
- [7] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision*. 2425–2433.
- [8] Samuel G. Armato III, Geoffrey McLennan, Luc Bidaut, Michael F. McNitt-Gray, Charles R. Meyer, Anthony P. Reeves, Binsheng Zhao, Denise R. Aberle, Claudia I. Henschke, Eric A. Hoffman, et al. 2011. The lung image database consortium (LIDC) and image database resource initiative (IDRI): A completed reference database of lung nodules on CT scans. *Med. Phys.* 38, 2 (2011), 915–931.
- [9] Ines Arous, Ljiljana Dolamic, Jie Yang, Akansha Bhardwaj, Giuseppe Cuccu, and Philippe Cudré-Mauroux. 2021. Marta: Leveraging human rationales for explainable text classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 5868–5876.
- [10] Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, et al. 2020. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* 58 (2020), 82–115.
- [11] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2022. Diagnostics-guided explanation generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 10445–10453.
- [12] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS One* 10, 7 (2015), e0130140.
- [13] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR'15)*.
- [14] Guangji Bai, Chen Ling, Yuyang Gao, and Liang Zhao. 2023. Saliency-augmented memory completion for continual learning. In *Proceedings of the SIAM International Conference on Data Mining (SDM'23)*. SIAM, 244–252.
- [15] Vladimir Balayan, Pedro Saleiro, Catarina Belém, Ludwig Krippahl, and Pedro Bizarro. 2020. Teaching the machine to explain itself using domain knowledge. arXiv:2012.01932. Retrieved from <https://arxiv.org/abs/2012.01932>
- [16] Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*. 65–72.
- [17] Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. 2018. Deriving machine attention from human rationales. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'18)*. ACL, 1903–1913.
- [18] David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'17)*. 6541–6549.
- [19] Rachel K. E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, et al. 2019. AI Fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. *IBM J. Res. Dev.* 63, 4/5 (2019), 4–1.
- [20] Akhilan Boopathy, Sijia Liu, Gaoyuan Zhang, Cynthia Liu, Pin-Yu Chen, Shiyu Chang, and Luca Daniel. 2020. Proper network interpretability helps adversarial robustness in classification. In *Proceedings of the International Conference on Machine Learning (ICML'20)*. PMLR, 1014–1023.
- [21] Nadia Burkart, Philipp M. Faller, Elisabeth Peinsipp, and Marco F. Huber. 2020. Batch-wise regularization of deep neural networks for interpretability. In *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration (MF'20)*. IEEE, 216–222.
- [22] Nadia Burkart and Marco F. Huber. 2021. A survey on the explainability of supervised machine learning. *J. Artif. Intell. Res.* 70 (2021), 245–317.
- [23] Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. In *Advances in Neural Information Processing Systems*, Vol. 31, 9560–9572.
- [24] Samuel Carton, Surya Kanoria, and Chenhao Tan. 2022. What to learn, and how: Toward effective learning from rationales. In *Findings of the Association for Computational Linguistics (ACL'22)*. Association for Computational Linguistics, 1075–1088.
- [25] Simon Caton and Christian Haas. 2020. Fairness in machine learning: A survey. arXiv:2010.04053. Retrieved from <https://arxiv.org/abs/2010.04053>

- [26] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. 2019. A game theoretic approach to class-wise selective rationalization. In *Advances in Neural Information Processing Systems*, Vol. 32, 10055–10065.
- [27] Shi Chen, Ming Jiang, Jinhui Yang, and Qi Zhao. 2020. Air: Attention with reasoning capability. In *European Conference on Computer Vision*. Springer, 91–107.
- [28] Seungtaek Choi, Haeju Park, Jinyoung Yeo, and Seung-won Hwang. 2020. Less is more: Attention supervision with counterfactuals for text classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'20)*. 6695–6704.
- [29] George Chrysostomou and Nikolaos Aletras. 2021. Enjoy the salience: Towards better transformer-based faithful explanations with word salience. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 8189–8200.
- [30] Junyoung Chung, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. In *Neural Information Processing Systems Workshop on Deep Learning*.
- [31] Noel C. F. Codella, David Gutman, M. Emre Celebi, Brian Helba, et al. 2018. Skin lesion analysis toward melanoma detection: A challenge at the 2017 international symposium on biomedical imaging (isbi), hosted by the international skin imaging collaboration (isic). In *Proceedings of the IEEE International Symposium on Biomedical Imaging (ISBI'18)*. IEEE, 168–172.
- [32] Dennis Collaris and Jarke J. van Wijk. 2020. ExplainExplore: Visual exploration of machine learning explanations. In *Proceedings of the IEEE Pacific Visualization Symposium (PacificVis'20)*. IEEE, 26–35.
- [33] Owen Cor nec, Rahul Nair, Elizabeth Daly, Dennis Wei, and Oznur Alkan. 2021. AIMEE: Interactive model maintenance with rule-based surrogates. In *Proceedings of the Annual Conference on Neural Information Processing Systems (Vol. 176)*. PMLR, 288–291.
- [34] Elizabeth M. Daly, Massimiliano Mattetti, Oznur Alkan, and Rahul Nair. 2021. User driven model adjustment via boolean rule explanations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35, 5896–5904.
- [35] Abhishek Das, Harsh Agrawal, Larry Zitnick, et al. 2017. Human attention in visual question answering: Do humans and deep networks look at the same regions? *Comput. Vis. Image Understand.* 163 (2017), 90–100.
- [36] Luc De Raedt, Sebastian Dumancic, Robin Manhaeve, and Giuseppe Marra. 2020. From statistical relational to neuro-symbolic artificial intelligence. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'20)*. 4943–4950.
- [37] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'19)*. ACL, 4171–4186.
- [38] Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. ERASER: A benchmark to evaluate rationalized NLP models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. ACL, 4443–4458.
- [39] K. C. Dharma and Chicheng Zhang. 2021. Improving the trustworthiness of image classification models by utilizing bounding-box annotations. *CoRR* abs/2108.10131.
- [40] Mengnan Du, Ninghao Liu, Fan Yang, and Xia Hu. 2021. Learning credible DNNs via incorporating prior knowledge and model local explanation. *Knowl. Inf. Syst.* 63, 2 (2021), 305–332.
- [41] Mengnan Du, Fan Yang, Na Zou, and Xia Hu. 2020. Fairness in deep learning: A computational perspective. *IEEE Intell. Syst.* 36, 4 (2020), 25–34.
- [42] John J. Dudley and Per Ola Kristensson. 2018. A review of user interface design for interactive machine learning. *ACM Trans. Interact. Intell. Syst.* 8, 2 (2018), 1–37.
- [43] Sayna Ebrahimi, Suzanne Petryk, Akash Gokul, William Gan, et al. 2021. Remembering for the right reasons: Explanations reduce catastrophic forgetting. *Appl. AI Lett.* 2, 4 (2021), e44.
- [44] Alex Endert, William Ribarsky, Cagatay Turkay, B. L. William Wong, et al. 2017. The state of the art in integrating machine learning into visual analytics. In *Computer Graphics Forum*, Vol. 36. Wiley Online Library, 458–486.
- [45] Gabriel Erion, Joseph D. Janizek, Pascal Sturmfels, et al. 2021. Improving performance of deep learning models with axiomatic attribution priors and expected gradients. *Nat. Mach. Intell.* 3, 7 (2021), 620–631.
- [46] Gabriel G. Erion, Joseph D. Janizek, Pascal Sturmfels, Scott M. Lundberg, and Su-In Lee. 2019. Learning explainable models using attribution priors. *arXiv:1906.10670*. Retrieved from <http://arxiv.org/abs/1906.10670>.
- [47] Jerry Alan Fails and Dan R. Olsen Jr. 2003. Interactive machine learning. In *Proceedings of the 8th International Conference on Intelligent User Interfaces*. 39–45.
- [48] Patrick Fernandes, Marcos Treviso, Danish Pruthi, André F. T. Martins, and Graham Neubig. 2022. Learning to scaffold: Optimizing model explanations for teaching. In *Advances in Neural Information Processing Systems*.
- [49] Hiroshi Fukui, Tsubasa Hirakawa, Takayoshi Yamashita, and Hironobu Fujiyoshi. 2019. Attention branch network: Learning of attention mechanism for visual explanation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'19)*. 10705–10714.

[50] Chuang Gan, Yandong Li, Haoxiang Li, Chen Sun, and Boqing Gong. 2017. Vqs: Linking segmentations to questions and answers for supervised attention in vqa and question-focused semantic segmentation. In *Proceedings of the International Conference on Computer Vision (ICCV'17)*. 1811–1820.

[51] Yuyang Gao, Giorgio A. Ascoli, and Liang Zhao. 2021. Schematic memory persistence and transience for efficient and robust continual learning. *Neural Netw.* 144 (2021), 49–60.

[52] Yuyang Gao, Tong Sun, Guangji Bai, Siyi Gu, Sungsoo Ray Hong, and Liang Zhao. 2022. RES: A robust framework for guiding visual explanation. In *Proceedings of the ACM Special Interest Group on Knowledge Discovery in Data (SIGKDD'22)*. ACM, 432–442.

[53] Yuyang Gao, Tong Sun, Rishab Bhatt, Dazhou Yu, Sungsoo Hong, and Liang Zhao. 2021. GNES: Learning to explain graph neural networks. In *Proceedings of the IEEE International Conference on Data Mining (ICDM'21)*. IEEE, 131–140.

[54] Yuyang Gao, Tong Steven Sun, Liang Zhao, and Sungsoo Ray Hong. 2022. Aligning eyes between humans and deep neural network through interactive attention alignment. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2 (2022), 1–28.

[55] Reza Ghaeini, Xiaoli Fern, Hamed Shahbazi, and Prasad Tadepalli. 2019. Saliency learning: Teaching the model where to pay attention. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'19)*. ACL, 4016–4025.

[56] Yolanda Gil, James Honaker, Shikhar Gupta, Yibo Ma, Vito D’Orazio, Daniel Garjio, Shruti Gadewar, Qifan Yang, and Neda Jahanshad. 2019. Towards human-guided machine learning. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI'19)*. 614–624.

[57] Max Glockner, Ivan Habernal, and Iryna Gurevych. 2020. Why do you think that? Exploring faithful sentence-level rationales without supervision. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'20)*. Association for Computational Linguistics, 1080–1095.

[58] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'17)*. 6904–6913.

[59] Casey S. Greene, Arjun Krishnan, Aaron K. Wong, Emanuela Ricciotti, et al. 2015. Understanding multicellular function and disease with human tissue-specific networks. *Nat. Genet.* 47, 6 (2015), 569–576.

[60] Siyi Gu, Yifei Zhang, Yuyang Gao, Xiaofeng Yang, and Liang Zhao. 2023. Essa: Explanation iterative supervision via saliency-guided data augmentation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 567–576.

[61] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2018. A survey of methods for explaining black box models. *ACM Comput. Surv.* 51, 5 (2018), 1–42.

[62] Nicholas Halliwell and Freddy Lecue. 2020. Trustworthy convolutional neural networks: A gradient penalized-based approach. arXiv:2009.14260. Retrieved from <https://arxiv.org/abs/2009.14260>

[63] Xiaochuang Han and Yulia Tsvetkov. 2021. Influence tuning: Demoting spurious correlations via instance attribution and instance-driven updates. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'21)*. Association for Computational Linguistics, 4398–4409.

[64] Lisa Anne Hendricks, Kaylee Burns, Kate Saenko, Trevor Darrell, and Anna Rohrbach. 2018. Women also snowboard: Overcoming bias in captioning models. In *Proceedings of the European Conference on Computer Vision (ECCV'18)*. 771–787.

[65] Robert R. Hoffman, Shane T. Mueller, Gary Klein, and Jordan Litman. 2018. Metrics for explainable AI: Challenges and prospects. arXiv:1812.04608. Retrieved from <https://arxiv.org/abs/1812.04608>

[66] Sungsoo Ray Hong, Jessica Hullman, and Enrico Bertini. 2020. Human factors in model interpretability: Industry practices, challenges, and needs. *Proc. ACM Hum.-Comput. Interact.* 4 (2020), 1–26.

[67] Eugen Hruska, Liang Zhao, and Fang Liu. 2022. Ground truth explanation dataset for chemical property prediction on molecular graphs. ChemRxiv. 2022. DOI: [10.26434/chemrxiv-2022-96slq](https://doi.org/10.26434/chemrxiv-2022-96slq)

[68] Drew A. Hudson and Christopher D. Manning. 2019. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'19)*. 6700–6709.

[69] Aya Abdelsalam Ismail, Hector Corrada Bravo, and Soheil Feizi. 2021. Improving deep learning interpretability by saliency guided training. *Adv. Neural Inf. Process. Syst.* 34 (2021), 26726–26739.

[70] Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C. Wallace. 2020. Learning to faithfully rationalize by construction. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'20)*. ACL, 4459–4473.

[71] Hoyong Jeong, Suyoung Lee, Sung Ju Hwang, and Sooel Son. 2022. Learning to generate inversion-resistant model explanations. In *Advances in Neural Information Processing Systems*.

[72] Liu Jiang, Shixia Liu, and Changjian Chen. 2019. Recent research advances on interactive machine learning. *J. Vis.* 22, 2 (2019), 401–417.

[73] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, et al. 2017. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'17)*. 2901–2910.

[74] Teja Kanchinadam, Keith Westpfahl, Qian You, and Glenn Fung. 2020. Rationale-based human-in-the-loop via supervised attention. In *Proceedings of the Workshop on Data Science with Human in the Loop at the Conference on Knowledge Discovery and Data Mining (DaSH@ KDD'20)*.

[75] Maxime Kayser, Oana-Maria Camburu, Leonard Salewski, Cornelius Emde, et al. 2021. e-vil: A dataset and benchmark for natural language explanations in vision-language tasks. In *Proceedings of the International Conference on Computer Vision (ICCV'21)*. 1244–1254.

[76] Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. 2018. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'18)*. 252–262.

[77] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference on Learning Representations*.

[78] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.

[79] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI'15)*. 126–137.

[80] Nahyun Kwon, Tong Sun, Yuyang Gao, Liang Zhao, Xu Wang, Jeeeon Kim, and Ray Hong. 2024. 3DPFIX: Improving remote novices' 3D printing troubleshooting experience through human-AI collaboration design. *Proc. ACM Hum.-Comput. Interact.* (2024).

[81] John D. Lafferty, Andrew McCallum, and Fernando C. N. Pereira. 2001. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the International Conference on Machine Learning (ICML'01)*, Carla E. Brodley and Andrea Pohoreckyj Danyluk (Eds.). Morgan Kaufmann, 282–289.

[82] David W. G. Langerhuizen et al. 2020. Is deep learning on par with human observers for detection of radiographically visible and occult fractures of the scaphoid? *Clin. Orthopaed. Relat. Res.* 478, 11 (2020), 2653.

[83] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.

[84] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. 2020. MaskGAN: Towards diverse and interactive facial image manipulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR'20)*. 5548–5557.

[85] Seungho Lee, Minhyun Lee, Jongwuk Lee, and Hyunjung Shim. 2021. Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'21)*. 5495–5505.

[86] Seungeon Lee, Xiting Wang, Sungwon Han, Xiaoyuan Yi, Xing Xie, and Meeyoung Cha. 2022. Self-explaining deep models with logic rule reasoning. In *Advances in Neural Information Processing Systems*.

[87] Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'16)*. ACL, 107–117.

[88] Piyawat Lertvittayakumjorn and Francesca Toni. 2021. Explanation-based human debugging of nlp models: A survey. *Trans. Assoc. Comput. Ling.* 9 (2021), 1508–1528.

[89] Kunpeng Li, Ziyan Wu, Kuan-Chuan Peng, Jan Ernst, and Yun Fu. 2018. Tell me where to look: Guided attention inference network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9215–9223.

[90] Yi Li and Nuno Vasconcelos. 2019. Repair: Removing representation bias by dataset resampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 9572–9581.

[91] Zewen Li, Fan Liu, Wenjie Yang, Shouheng Peng, and Jun Zhou. 2021. A survey of convolutional neural networks: Analysis, applications, and prospects. *IEEE Trans. Neural Netw. Learn. Syst.* 33, 12 (2021), 6999–7019.

[92] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, et al. 2014. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*. Springer, 740–755.

[93] Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. 2020. Explainable ai: A review of machine learning interpretability methods. *Entropy* 23, 1 (2020), 18.

[94] Frederick Liu and Besim Avci. 2019. Incorporating priors with feature attribution on text classification. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. ACL, 6274–6283.

[95] Bodhisattwa Prasad Majumder, Oana-Maria Camburu, Thomas Lukasiewicz, and Julian McAuley. 2021. Rationale-inspired natural language explanations with commonsense. arXiv:2106.13876. Retrieved from <https://arxiv.org/abs/2106.13876>

[96] Yaoli Mao, Dakuo Wang, Michael Muller, et al. 2019. How data scientists work together with domain experts in scientific collaborations: To find the right answer or to ask the right question? *Proc. ACM Hum.-Comput. Interact.* 3, GROUP (2019), 1–23.

[97] Ines Filipa Martins, Ana L Teixeira, Luis Pinheiro, and Andre O. Falcao. 2012. A bayesian approach to in silico blood-brain barrier penetration modeling. *J. Chem. Inf. Model.* 52, 6 (2012), 1686–1697.

[98] Andreas Mayr, Günter Klambauer, Thomas Unterthiner, and Sepp Hochreiter. 2016. DeepTox: Toxicity prediction using deep learning. *Front. Environ. Sci.* 3 (2016), 80.

[99] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. 2021. A survey on bias and fairness in machine learning. *ACM Comput. Surv. (CSUR)* 54, 6 (2021), 1–35.

[100] Caijing Miao, Lingxi Xie, Fang Wan, Chi Su, Hongye Liu, Jianbin Jiao, and Qixiang Ye. 2019. SIXray: A large-scale security inspection x-ray benchmark for prohibited item discovery in overlapping images. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR’19)*. 2119–2128.

[101] Henry W. Miller. 1973. *Plan and Operation of the Health and Nutrition Examination Survey, United States, 1971–1973*. DHEW Publication no. (PHS), Dept. of Health, Education, and Welfare.

[102] Jeremy A. Miller, Angela Guillozet-Bongaarts, Laura E. Gibbons, Nadia Postupna, et al. 2017. Neuropathological and transcriptomic characteristics of the aged brain. *eLife* 6 (2017), e31126.

[103] Masahiro Mitsuhashi, Hiroshi Fukui, Yusuke Sakashita, Takanori Ogata, et al. 2019. Embedding human knowledge into deep neural network via attention map. arXiv:1905.03540. Retrieved from <https://arxiv.org/abs/1905.03540>

[104] Sina Mohseni, Niloofer Zarei, and Eric D. Ragan. 2021. A multidisciplinary survey and framework for design and evaluation of explainable AI systems. *ACM Trans. Interact. Intell. Syst.* 11, 3-4 (2021), 1–45.

[105] Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, et al. 2019. Layer-wise relevance propagation: An overview. *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, 193–209.

[106] Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. 2017. Explaining nonlinear classification decisions with deep taylor decomposition. *Pattern Recogn.* 65 (2017), 211–222.

[107] Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. 2018. Methods for interpreting and understanding deep neural networks. *Digit. Sign. Process.* 73 (2018), 1–15.

[108] Dong Nguyen. 2018. Comparing automatic and human evaluation of local explanations for text classification. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL’18)*. 1069–1078.

[109] Giang Nguyen, Mohammad Reza Taesiri, and Anh Nguyen. 2022. Visual correspondence-based explanations improve AI robustness and human-AI team accuracy. In *Advances in Neural Information Processing Systems*.

[110] Eirini Ntoutsi, Pavlos Fafalios, Ujwal Gadiraju, et al. 2020. Bias in data-driven artificial intelligence systems—an introductory survey. *Data Min. Knowl. Discov.* 10, 3 (2020), e1356.

[111] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. 2019. Continual lifelong learning with neural networks: A review. *Neural Netw.* 113 (2019), 54–71.

[112] Dong Huk Park, Lisa Anne Hendricks, Zeynep Akata, Anna Rohrbach, et al. 2018. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR’18)*. 8779–8788.

[113] Badri Patro, Vinay Namboodiri, et al. 2020. Explanation vs attention: A two-player game to obtain attention for vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 11848–11855.

[114] Tejaswini Pedapati, Avinash Balakrishnan, et al. 2020. Learning global transparent models consistent with local contrastive explanations. In *Advances in Neural Information Processing Systems*, Vol. 33, 3592–3602.

[115] Yan Peng, Zheng Xuefeng, Zhu Jianyong, and Xiao Yumhong. 2009. Lazy learner text categorization algorithm based on embedded feature selection. *J. Syst. Eng. Electr.* 20, 3 (2009), 651–659.

[116] Vipin Pillai and Hamed Pirsiavash. 2021. Explainable models with consistent interpretations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 35. 2431–2439.

[117] Gregory Plumb, Maruan Al-Shedivat, Ángel Alexander Cabrera, et al. 2020. Regularizing black-box models for improved interpretability. In *Advances in Neural Information Processing Systems*, Vol. 33, 10526–10536.

[118] Teodora Popordanoska, Mohit Kumar, and Stefano Teso. 2020. Machine guides, human supervises: Interactive learning with global explanations. arXiv:2009.09723.

[119] Prasanna Porwal, Samiksha Pachade, Ravi Kamble, Manesh Kokare, Girish Deshmukh, Vivek Sahasrabuddhe, and Fabrice Meriaudeau. 2018. Indian Diabetic Retinopathy Image Dataset (IDRiD). Retrieved from <https://doi.org/10.21227/H25W98>

[120] Tingting Qiao, Jianfeng Dong, and Duanqing Xu. 2018. Exploring human-like attention supervision in visual question answering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32. 7300–7307.

[121] Nazneen Fatema Rajani, Bryan McCann, Caiming Xiong, and Richard Socher. 2019. Explain Yourself! Leveraging language models for commonsense reasoning. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics. ACL*, 4932–4942.

[122] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In *Proceedings of the ACM Special Interest Group on Knowledge Discovery in Data (SIGKDD'16)*. 1135–1144.

[123] Laura Rieger, Chandan Singh, William Murdoch, and Bin Yu. 2020. Interpretations are useful: Penalizing explanations to align neural networks with prior knowledge. In *International Conference on Machine Learning*. PMLR, 8116–8126.

[124] Ribana Roscher, Bastian Bohn, Marco F. Duarte, and Jochen Garcke. 2020. Explainable machine learning for scientific insights and discoveries. *IEEE Access* 8 (2020), 42200–42216.

[125] Andrew Ross and Finale Doshi-Velez. 2018. Improving the adversarial robustness and interpretability of deep neural networks by regularizing their input gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'18)*, Vol. 32. 1660–1669.

[126] Andrew Slavin Ross, Michael C. Hughes, and Finale Doshi-Velez. 2017. Right for the right reasons: Training differentiable models by constraining their explanations. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI'17)*. 2662–2670.

[127] Gobinda Saha and Kaushik Roy. 2021. Saliency guided experience packing for replay in continual learning. arXiv:2109.04954. Retrieved from <https://arxiv.org/abs/2109.04954>

[128] Johannes Schneider and Michalis Vlachos. 2020. Reflective-net: Learning from explanations. arXiv:2011.13986. Retrieved from <https://arxiv.org/abs/2011.13986>

[129] Patrick Schramowski, Wolfgang Stammer, Stefano Teso, et al. 2020. Making deep neural networks right for the right scientific reasons by interacting with their explanations. *Nat. Mach. Intell.* 2, 8 (2020), 476–486.

[130] Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. BLEURT: Learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*. ACL, 7881–7892.

[131] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the International Conference on Computer Vision (ICCV'17)*. 618–626.

[132] Sofia Serrano and Noah A. Smith. 2019. Is attention interpretable? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. ACL, 2931–2951.

[133] Rudy Setiono and Huan Liu. 1997. Neural-network feature selector. *IEEE Trans. Neural Netw.* 8, 3 (1997), 654–662.

[134] Xiaoting Shao, Tjitzie Rienstra, Matthias Thimm, and Kristian Kersting. 2020. Towards understanding and arguing with classifiers: Recent progress. *Datenb.-Spektr.* 20, 2 (2020), 171–180.

[135] Manali Sharma, Di Zhuang, and Mustafa Bilgic. 2015. Active learning with rationales for text classification. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'15)*. 441–451.

[136] Haifeng Shen, Kewen Liao, Zhibin Liao, et al. 2021. Human-AI interactive and continuous sensemaking: A case study of image classification using scribble attention maps. In *Extended Abstracts of CHI*. 1–8.

[137] Becks Simpson, Francis Dutil, Yoshua Bengio, and Joseph Paul Cohen. 2019. Gradmask: Reduce overfitting by regularizing saliency. arXiv:1904.07478. Retrieved from <https://arxiv.org/abs/1904.07478>

[138] Chandan Singh, Wooseok Ha, and Bin Yu. 2022. Interpreting and improving deep-learning models with reality checks. In *International Workshop on Extending Explainable AI Beyond Deep Models and Classifiers*. Springer, 229–254.

[139] Krishna Kumar Singh, Dhruv Mahajan, Kristen Grauman, et al. 2020. Don't judge an object by its context: Learning to overcome contextual bias. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'20)*. 11070–11078.

[140] Xuelin Situ, Ingrid Zukerman, Cecile Paris, Sameen Maruf, and Gholamreza Haffari. 2021. Learning to explain: Generating stable explanations fast. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'21)*. 5340–5355.

[141] Ekta Sood, Simon Tannert, Philipp Müller, and Andreas Bulling. 2020. Improving natural language processing tasks with human gaze-guided neural attention. In *Advances in Neural Information Processing Systems 33* (2020), 6327–6341.

[142] Joe Stacey, Yonatan Belinkov, and Marek Rei. 2022. Supervising model attention with human explanations for robust natural language inference. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI'22)*, Vol. 36. 11349–11357.

[143] Wolfgang Stammer, Patrick Schramowski, and Kristian Kersting. 2021. Right for the right concept: Revising neuro-symbolic concepts by interacting with their explanations. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'21)*. 3619–3629.

[144] Julia Strout, Ye Zhang, and Raymond J. Mooney. 2019. Do human rationales improve machine explanations? In *ACL Workshop BlackboxNLP*.

[145] Govindan Subramanian, Bharath Ramsundar, et al. 2016. Computational modeling of  $\beta$ -secretase 1 (BACE-1) inhibitors using ligand based approaches. *J. Chem. Inf. Model.* 56, 10 (2016), 1936–1949.

[146] Tong Steven Sun, Yuyang Gao, Shubham Khaladkar, Sijia Liu, Liang Zhao, Young-Ho Kim, and Sungsoo Ray Hong. 2023. Designing a direct feedback loop between humans and convolutional neural networks through local explanations. *Proc. ACM Hum.-Comput. Interact.* 7, CSCW2 (2023), 1–32.

[147] Chenhao Tan. 2021. On the diversity and limits of human explanations. arXiv:2106.11988. Retrieved from <https://arxiv.org/abs/2106.11988>

[148] Makarand Tapaswi, Yukun Zhu, Rainer Stiefelhagen, Antonio Torralba, Raquel Urtasun, and Sanja Fidler. 2016. Movieqa: Understanding stories in movies through question-answering. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'16)*. 4631–4640.

[149] Stefano Teso and Kristian Kersting. 2019. Explanatory interactive machine learning. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*. 239–245.

[150] James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. FEVER: A large-scale dataset for fact extraction and VERification. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'18)*. ACL, 809–819.

[151] Erico Tjoa and Cuntai Guan. 2020. A survey on explainable artificial intelligence (xai): Toward medical xai. *IEEE Trans. Neural Netw. Learn. Syst.* 32, 11 (2020), 4793–4813.

[152] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, Vol. 30. 5998–6008.

[153] Stanislav Vojíř and Tomáš Kliegr. 2020. Editable machine learning models? A rule-based framework for user studies of explainability. *Adv. Data Anal. Class.* 14, 4 (2020), 785–799.

[154] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. 2011. *The Caltech-UCSD Birds-200-2011 Dataset*. California Institute of Technology.

[155] Cunxiang Wang, Shuailong Liang, Yue Zhang, Xiaonan Li, and Tian Gao. 2019. Does it make sense? And why? A pilot study for sense making and explanation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL'19)*. ACL, 4020–4026.

[156] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, et al. 2017. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR'17)*. 2097–2106.

[157] Ethan Weinberger, Joseph Janizek, and Su-In Lee. 2020. Learning deep attribution priors based on prior knowledge. *Adv. Neural Inf. Process. Syst.* 33 (2020), 14034–14045.

[158] James Wexler, Mahima Pushkarna, Tolga Bolukbasi, Martin Wattenberg, Fernanda Viégas, and Jimbo Wilson. 2019. The what-if tool: Interactive probing of machine learning models. *IEEE Trans. Vis. Comput. Graph.* 26, 1 (2019), 56–65.

[159] Sarah Wiegreffe, Ana Marasović, and Noah A. Smith. 2021. Measuring association between labels and free-text rationales. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP'21)*. ACL, 10266–10284.

[160] Jialin Wu and Raymond Mooney. 2019. Self-critical reasoning for robust visual question answering. *Adv. Neural Inf. Process. Syst.* 32 (2019), 8601–8611.

[161] Mike Wu, Sonali Parbhoo, Michael Hughes, Ryan Kindle, Leo Celi, Maurizio Zazzi, Volker Roth, and Finale Doshi-Velez. 2020. Regional tree regularization for interpretability in deep neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 6413–6421.

[162] Zhenqin Wu, Bharath Ramsundar, Evan N. Feinberg, Joseph Gomes, Caleb Geniesse, Aneesh S. Pappu, et al. 2018. MoleculeNet: A benchmark for molecular machine learning. *Chem. Sci.* 9, 2 (2018), 513–530.

[163] Han Xiao, Kashif Rasul, and Roland Vollgraf. 2017. Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms. arXiv:1708.07747. Retrieved from <https://arxiv.org/abs/1708.07747>

[164] Huihan Yao, Ying Chen, Qinyuan Ye, Xisen Jin, and Xiang Ren. 2021. Refining language models with compositional explanations. *Adv. Neural Inf. Process. Syst.* 34 (2021), 8954–8967.

[165] Zhuofan Ying, Peter Hase, and Mohit Bansal. 2022. VisFIS: Visual feature importance supervision with right-for-the-right-reason objectives. In *Advances in Neural Information Processing Systems*.

[166] Jun Yuan, Changjian Chen, Weikai Yang, Mengchen Liu, Jiazhi Xia, and Shixia Liu. 2021. A survey of visual analytics techniques for machine learning. *Comput. Vis. Media* 7, 1 (2021), 3–36.

[167] Omar Zaidan and Jason Eisner. 2008. Modeling annotators: A generative approach to learning from annotator rationales. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 31–40.

[168] Omar Zaidan, Jason Eisner, and Christine Piatko. 2007. Using “annotator rationales” to improve machine learning for text categorization. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL'07)*. 260–267.

[169] Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6720–6731.

[170] Guohang Zeng, Yousef Kowsar, Sarah Erfani, and James Bailey. 2021. Generating deep networks explanations with robust attribution alignment. In *Proceedings of the Asian Conference on Machine Learning*. PMLR, 753–768.

[171] Quan-shi Zhang and Song-Chun Zhu. 2018. Visual interpretability for deep learning: A survey. *Front. Inf. Technol. Electr. Eng.* 19, 1 (2018), 27–39.

[172] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating text generation with BERT. In *Proceedings of the 8th International Conference on Learning Representations*.

[173] Yifei Zhang, Siyi Gu, Yuyang Gao, Bo Pan, Xiaofeng Yang, and Liang Zhao. 2023. MAGI: Multi-annotated explanation-guided learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1977–1987.

[174] Yifei Zhang, Siyi Gu, James Song, Bo Pan, and Liang Zhao. 2023. XAI Benchmark for visual explanation. arXiv:2310.08537. Retrieved from <https://arxiv.org/abs/2310.08537>

[175] Ye Zhang, Iain Marshall, and Byron C. Wallace. 2016. Rationale-augmented convolutional neural networks for text classification. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP’16)*, Vol. 2016. NIH Public Access, 795.

[176] Yundong Zhang, Juan Carlos Niebles, and Alvaro Soto. 2019. Interpretable visual question answering by visual grounding from attention supervision mining. In *Proceedings of the IEEE CVF Winter Conference on Applications of Computer Vision (WACV’19)*. IEEE, 349–357.

[177] Zijian Zhang, Koustav Rudra, and Avishek Anand. 2021. Explain and predict, and then predict again. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 418–426.

[178] Jieyu Zhao, Tianlu Wang, Mark Yatskar, Vicente Ordonez, and Kai-Wei Chang. 2017. Men also like shopping: Reducing gender bias amplification using corpus-level constraints. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP’17)*. ACL.

[179] Qilong Zhao, Chih-Wei Chang, Xiaofeng Yang, and Liang Zhao. 2024. Robust explanation supervision for false positive reduction in pulmonary nodule detection. *Med. Phys.* (2024).

[180] Ruiqi Zhong, Steven Shao, and Kathleen McKeown. 2019. Fine-grained sentiment analysis with faithful attention. arXiv:1908.06870. Retrieved from <https://arxiv.org/abs/1908.06870>

[181] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR’16)*. 2921–2929.

[182] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. 2017. Places: A 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 6 (2017), 1452–1464.

[183] Jianlong Zhou, Amir H. Gandomi, Fang Chen, and Andreas Holzinger. 2021. Evaluating the quality of machine learning explanations: A survey on methods and metrics. *Electronics* 10, 5 (2021), 593.

[184] Jiaxin Zhuang, Jiabin Cai, Ruixuan Wang, et al. 2019. Care: Class attention to regions of lesion for classification on imbalanced data. In *International Conference on Medical Imaging with Deep Learning*. PMLR, 588–597.

Received 31 December 2022; revised 31 October 2023; accepted 21 December 2023