

# Causal Reasoning Meets Visual Representation Learning: A Prospective Study

Yang Liu    Yu-Shen Wei    Hong Yan    Guan-Bin Li    Liang Lin

School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China

**Abstract:** Visual representation learning is ubiquitous in various real-world applications, including visual comprehension, video understanding, multi-modal analysis, human-computer interaction, and urban computing. Due to the emergence of huge amounts of multi-modal heterogeneous spatial/temporal/spatial-temporal data in the big data era, the lack of interpretability, robustness, and out-of-distribution generalization are becoming the challenges of the existing visual models. The majority of the existing methods tend to fit the original data/variable distributions and ignore the essential causal relations behind the multi-modal knowledge, which lacks unified guidance and analysis about why modern visual representation learning methods easily collapse into data bias and have limited generalization and cognitive abilities. Inspired by the strong inference ability of human-level agents, recent years have therefore witnessed great effort in developing causal reasoning paradigms to realize robust representation and model learning with good cognitive ability. In this paper, we conduct a comprehensive review of existing causal reasoning methods for visual representation learning, covering fundamental theories, models, and datasets. The limitations of current methods and datasets are also discussed. Moreover, we propose some prospective challenges, opportunities, and future research directions for benchmarking causal reasoning algorithms in visual representation learning. This paper aims to provide a comprehensive overview of this emerging field, attract attention, encourage discussions, bring to the forefront the urgency of developing novel causal reasoning methods, publicly available benchmarks, and consensus-building standards for reliable visual representation learning and related real-world applications more efficiently.

**Keywords:** Causal reasoning, visual representation learning, reliable artificial intelligence, spatial-temporal data, multi-modal analysis.

**Citation:** Y. Liu, Y. S. Wei, H. Yan, G. B. Li, L. Lin. Causal reasoning meets visual representation learning: A prospective study. *Machine Intelligence Research*, vol.19, no.6, pp.485-511, 2022. <http://doi.org/10.1007/s11633-022-1362-z>

## 1 Introduction

With the emergence of huge amounts of heterogeneous multi-modal data, including images<sup>[1-3]</sup>, videos<sup>[4-7]</sup>, texts/languages<sup>[8-10]</sup>, audios<sup>[11-14]</sup>, and multi-sensor<sup>[15-19]</sup> data, deep learning based methods have shown promising performance for various computer vision and machine learning tasks, e.g., the visual comprehension<sup>[20-23]</sup>, video understanding<sup>[24-27]</sup>, visual-linguistic analysis<sup>[28-30]</sup>, and multi-modal fusion<sup>[31-33]</sup>, etc. However, the existing methods rely heavily upon fitting the data distributions and tend to capture the spurious correlations from different modalities, and thus fail to learn the essential causal relations behind the multi-modal knowledge that have a good generalization and cognitive abilities. Inspired by the fact that most of the data in computer vision society are independent and identically distributed (i.i.d.), a substantial body of literature<sup>[34-37]</sup> adopted data augmentation, pre-training, self-supervision, and novel architectures to improve the robustness of the state-of-the-art

deep neural network architectures. However, it has been argued that such strategies only learn correlation-based patterns (statistical dependencies) from data and may not generalize well without the guarantee of the i.i.d. setting<sup>[38]</sup>.

Due to the powerful ability of to uncover the underlying structural knowledge about data generating processes that allow interventions and generalize well across different tasks and environments, causal reasoning<sup>[39-41]</sup> offers a promising alternative to correlation learning. Recently, causal reasoning has attracted increasing attention in a myriad of high-impact domains of computer vision and machine learning, such as interpretable deep learning<sup>[42-47]</sup>, causal feature selection<sup>[48-60]</sup>, visual comprehension<sup>[61-70]</sup>, visual robustness<sup>[71-78]</sup>, visual question answering<sup>[79-84]</sup>, and video understanding<sup>[85-92]</sup>. A common challenge of these causal methods is how to build a strong cognitive model that can fully discover causality and spatial-temporal relations.

In this paper, we aim to provide a comprehensive overview of causal reasoning for visual representation learning, attract attention, encourage discussions, and bring to the forefront the urgency of developing novel causality-guided visual representation learning methods.

Review  
Manuscript received May 9, 2022; accepted August 1, 2022;  
published online November 3, 2022  
Recommended by Associate Editor Friedhelm Schwenker  
© The Author(s) 2022

Although there are some surveys<sup>[40, 41, 93–95]</sup> about causal reasoning, these works are intended for general representation learning tasks such as deconfounding, out-of-distribution (OOD) generalization, and debiasing. Differently, our paper focuses on the systematic and comprehensive survey of related works, datasets, insights, future challenges and opportunities for causal reasoning, visual representation learning, and their integration. To present the review more concisely and clearly, this paper selects and cites related work by considering their sources, publication years, impact, and the cover of different aspects of the topic surveyed in this paper. The overview of the structure of this paper is shown in Fig. 1. Overall, the main contributions of this paper are given as follows.

Firstly, this paper presents the basic concepts of causality, the structural causal model (SCM), the independent causal mechanism (ICM) principle, causal inference, and causal intervention. Then, based on the analysis, this paper further gives some directions for conducting causal reasoning on visual representation learning tasks. Note that to the best of our knowledge, this paper is the first that proposes the potential research directions for causal visual representation learning.

Secondly, a prospective review is introduced to systematically and structurally review the existing works according to their efforts in the above-pointed directions for conducting causal visual representation learning more efficiently. We focus on the relation between visual representation learning and causal reasoning and provide a better understanding of why and how existing causal reasoning methods can be helpful in visual representation learning, as well as providing inspiration for future research and studies.

Thirdly, this paper explores and discusses future research areas and open problems related to using causal reasoning methods to tackle visual representation learning. This can encourage and support the broadening and deepening of research in the related fields.

The remainder of this paper is organized as follows. Section 2 provides the preliminaries, including the basic concepts of causality, the SCM, the ICM principle, caus-

al inference, and causal intervention. Section 3 discusses the ways to use causal reasoning to learn robust features, which are the key techniques for visual representation learning. Section 4 reviews some recent visual learning tasks, including visual understanding, action detection and recognition, and visual question answering, including the discussions about the existing challenges of these visual learning methods. Section 5 reviews the related causality-based visual representation learning works systematically. Section 6 provides a review of existing causal datasets for visual learning. Section 7 proposes and discusses some future research directions and finally Section 8 gives the conclusions.

## 2 Preliminaries

### 2.1 Causal learning and reasoning

As the sentence “correlation is not causation” says, two variables are correlated does not mean that one of them causes the other. Actually, statistical learning models the correlations of data. By observing a sufficient amount of i.i.d. data, the statistical learning method can perform considerably well under i.i.d. settings. However, when facing problems that do not satisfy i.i.d. assumptions, the performance of these methods often seems poor (e.g., image recognition models tend to predict “bird” when seeing “sky” in the image, since bird and sky usually appear simultaneously in the dataset). Causal learning<sup>[39]</sup> is different from statistical learning, which aims to discover causal relationships beyond statistical relations. Learning causality requires machine learning methods not only to predict the outcome of i.i.d. experiments but also to reason from a causal perspective. Causal reasoning can be divided into three levels. The first level is association. The statistical machine learning methods mentioned above belong to this level. A typical question of association is “How would the weather change when the sky is turning grey”, which asks about the association between “weather” and “the appearance of the sky”. The second

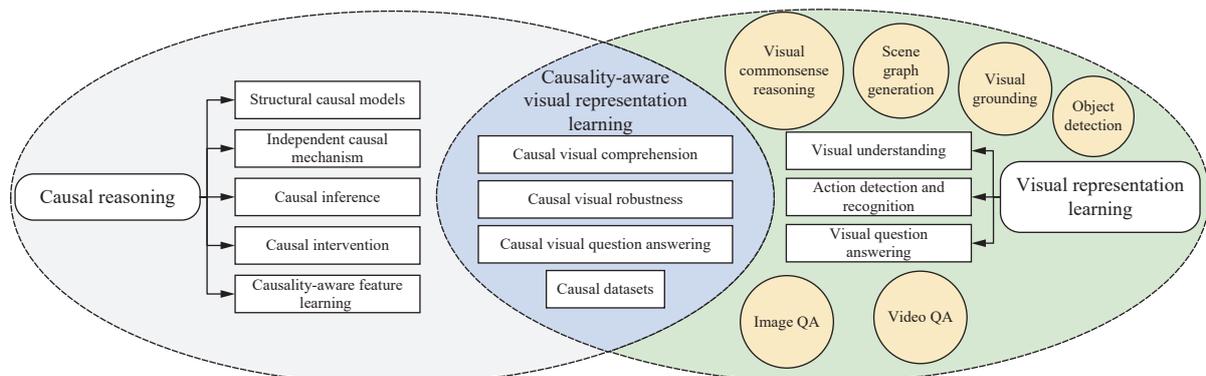


Fig. 1 Overview of the structure of this paper, including the discussion of related methods, datasets, challenges, and the relations among causal reasoning, visual representation learning, and their integration

level is intervention. An intervention-based question asks about the effect of the intervention (e.g., “Would I become stronger if I go to the gym every day?”). Intervention-based questions require us to answer the outcome when taking specific treatment, which can not be answered by only learning data associations (e.g., If we only learn the associations, then if we observe that a man who goes to the gym every day may not be stronger than a professional athlete, we may conclude that going to the gym not always makes you stronger). The third level is counterfactual. A typical form of a counterfactual question is “What if I had ...”, which focuses on the outcome when the condition is not realized. Counterfactual inference aims to compare different outcomes under the same condition, but the antecedent of the counterfactual question is not real.

### 2.2 Structural causal models

The SCM considers the formulation of a causality style. Assume that we have a set of variables  $X_1, X_2, \dots, X_n$ , each variable is a vertex of a causal graph (i.e., a directed acyclic graph (DAG) describes causal relations of variables). Then, those variables could be written as the outcome of a function:

$$X_i = f_i(\mathbf{PA}_i, U_i) \approx P(X_i|\mathbf{PA}_i), \quad i = 1, \dots, n \quad (1)$$

where  $\mathbf{PA}_i$  indicates the parents of  $X_i$  in the causal graph, and  $U_i$  refers to unmeasured factors such as noise. The deterministic function gives a mathematical form of the effect of direct causes of  $X_i$  on the variable  $X_i$ . Using the graphical causal model and SCM language, we can express joint distributions as follows:

$$P(X_1, X_2, \dots, X_n) = \prod_{i=1}^n P(X_i|\mathbf{PA}_i). \quad (2)$$

Equation (2) is called a product decomposition of the joint distributions. After the decomposition and graphical modeling, the causal relations and effects of a dataset can be represented as the causal graph and the joint distribution.

### 2.3 Independent causal mechanism

The independent causal mechanism principle<sup>[40]</sup> can be expressed as follows:

**ICM principle.** The causal generative process of a system’s variables is composed of autonomous modules that do not inform or influence each other. In the probabilistic case, this means that the conditional distribution of each variable given its causes (i.e., its mechanism) does not inform or influence the other conditional distributions.

The ICM principle describes the independence of caus-

al mechanisms. If we conceive that the real world is composed of modules in variable styles, then the modules could represent the physically independent mechanisms of the world. When applying the ICM principle to the disentangled factorization (2), it can be written as<sup>[40]</sup>:

1) Changing (or performing an intervention upon) one mechanism  $P(X_i|\mathbf{PA}_i)$  does not change any of the other mechanisms  $P(X_j|\mathbf{PA}_j)$  ( $i \neq j$ ).

2) Knowing some other mechanisms  $P(X_i|\mathbf{PA}_i)$  does not give us information about a mechanism  $P(X_j|\mathbf{PA}_j)$  ( $i \neq j$ ).

The ICM principle guarantees that our intervention on one mechanism does not affect others, which further reveals the possibility of transferring knowledge across domains that have the same modules.

### 2.4 Causal inference

The purpose of causal inference is to estimate the outcome shift (or effect) of different treatments. Let symbol  $A$  denote a treatment that refers to an action that applies to a unit. For example, if we have a medicine  $A$ , let  $A = 1$  denote applying medicine  $A$  and  $A = 0$  denote not applying medicine  $A$ , then  $A = 1$  is a treatment, and the recovery of the patient is the outcome of the treatment  $A = 1$ . Under this condition, the aim of causal inference is to uncover the effect of applying treatment  $A$ . A counterfactual outcome is the potential outcome of an action that has not been taken. For example, if we take treatment  $A = 1$ , then the outcome of  $A = 0$  is counterfactual. Then the average treatment effect (ATE) of treatment  $A = 1$  could be written as

$$ATE = E[Y(A = 1) - Y(A = 0)] \quad (3)$$

where  $Y(A = a)$  denotes the potential outcome of treatment  $A = a$ . If we have taken treatment  $A = 1$ , then  $Y(A = 0)$  is the counterfactual outcome.

The goal of causal inference is to estimate the treatment effects given the observational data, which is usually incomplete in real-world scenarios due to the cost and moral problems. From a counterfactual perspective, we cannot always obtain a no-treatment outcome if we apply the treatment. Thus, we need to adopt causal inference to analyze the effect of a certain treatment.

### 2.5 Causal intervention

Causal intervention for machine learning aims to capture the causal effects of interventions (i.e., variables), and take advantage of causal relations in datasets to improve model performance and generalization ability. The basic idea of causal intervention is to use an adjustment strategy that modifies the graphical model and manipulates conditional probabilities to discover causal relationships among variables. In this section, we review two ad-

justment strategies: back-door adjustment and front-door adjustment.

### 2.5.1 Back-door adjustment

Assume that we want to gauge the causal effect between  $X$  and  $Y$  by Bayes' rules; we can have

$$P(Y|X) = \sum_z P(Y|X, z)P(z|X). \tag{4}$$

This conditional distribution could not represent the true causal effect of  $X$  on  $Y$ , due to the existence of back-door path  $X \leftarrow Z \rightarrow Y$ . Variable  $Z$  here is a confounder that not only affects pre-intervention  $X$  but also the outcome  $Y$ , which would make the conditional distribution a collective effect of  $X$  and  $Z$ , and thus leads to spurious correlation. To eliminate the spurious correlation introduced by the back-door path, the back-door adjustment uses do-operator to calculate the intervened probability  $P(Y|do(X))$  instead of the conditional probability  $P(Y|X)$ :

$$P(Y|do(X)) = \sum_z P(Y|X, z)P(z). \tag{5}$$

Compared with  $P(Y|X)$ ,  $P(Y|do(X))$  is replaced with the conditional distribution  $P(z|X)$  with the marginal distribution  $P(z)$ . Fig. 2 is a graphical view of the do-operator. The edge from  $Z$  to  $X$  is deleted in the intervened causal graph to block the back-door path  $X \leftarrow Z \rightarrow Y$ ; thus,  $X$  and  $Z$  become independent after the intervention. After the back-door adjustment, the intervened distribution  $P(Y|do(X))$  can remove the spurious correlation between  $X$  and  $Z$  and calculate the true causal effect of variable  $X$ . Backdoor adjustment measures the causal effect of a variable by finding and blocking back-door paths points to it.

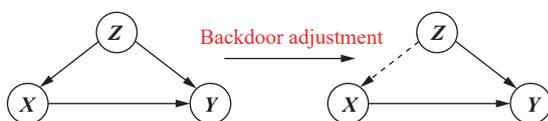


Fig. 2 An example of back-door adjustment, the back-door path from  $X$  to  $Y$  is blocked by cutting off the edge from  $Z$  to  $X$ .

### 2.5.2 Front-door adjustment

The back-door criterion may not be satisfied in some causal graphical patterns (e.g., no back-door paths exist in causal graphs, or variables that block the back-door paths are unobserved). In such a case, the front-door adjustment pattern can be applied to estimate causal effects. As Fig. 3 shows, assuming that the variable  $Z$  is an unobserved variable, the back-door adjustment becomes invalid because the marginal distribution  $P(z)$  is not observed. However, if we have an observed mediator variable  $W$  on the front-door path  $X \rightarrow W \rightarrow Y$ , then we can identify the effect of  $X$  on  $W$  directly since the back-door path from  $X$  to  $W$  is blocked by the collider at  $Y$ :

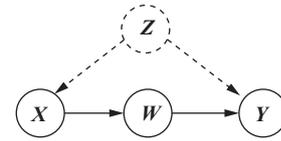


Fig. 3 Back-door criterion is not satisfied since  $Z$  is an unobserved variable.

$$P(W|do(X)) = P(W|X). \tag{6}$$

Note that there is a back-door path from  $W$  to  $Y$ :  $W \leftarrow X \leftarrow Z \rightarrow Y$ , which can be blocked by applying back-door adjustment on  $X$ :

$$P(Y|do(W)) = \sum_x P(Y|W, x)P(x). \tag{7}$$

And the total effect of  $X$  on  $Y$  could be written by summing on  $W$ :

$$P(Y|do(X)) = \sum_w P(Y|do(W = w))P(W = w|do(X)). \tag{8}$$

Then, the front-door adjustment formulation is obtained by applying (6)–(8):

$$P(Y|do(X)) = \sum_w \sum_x P(Y|W = w, x)P(W = w|x)P(x). \tag{9}$$

The front-door adjustment identifies the effect of  $X$  on  $Y$  by applying the do-operator twice, one at the mediator variable  $W$  and the other at variable  $X$  that blocks the back-door path. In this way, the unobserved variable  $Z$  can be bypassed in intervention.

### 2.5.3 Back-door or front-door?

The back-door adjustment requires us to determine what the confounder is in advance. Thus, the back-door adjustment is effective when the confounder is observable. However, in visual domains, data biases are complex, and it is hard to know and disentangle different types of confounders. Especially for some challenging tasks like the visual-linguistic question reasoning where the confounders in visual and linguistic modalities are not always observable. Therefore, the front-door causal intervention gives a feasible way to calculate  $P(Y|do(X))$  when we cannot explicitly represent the confounder.

## 3 Causality-aware feature learning

Traditional feature learning methods usually learn the spurious correlation introduced by confounders. This will reduce the robustness of models and make models hard to generalize across domains. Causal reasoning, a learning paradigm that reveals the real causality from the outcome, overcomes the essential defect of correlation learning and learns robust, reusable, and reliable features. In

this chapter, we review the recent representative causal reasoning methods for general feature learning, which mainly consist of three main paradigms: 1) structural causal model (SCM) embedded, 2) applying causal intervention/counterfactual, and 3) Markov boundary (MB) based feature selection.

For embedding the SCM, Mitrovic et al.<sup>[96]</sup> proposed representation learning via invariant causal mechanisms (RELIC) to address self-supervised learning problems and achieved competitive performance in terms of robustness and out-of-distribution generalization on ImageNet. Shen et al.<sup>[97]</sup> proposed a disentangled generative causal representation (DEAR) learning method for causal controllable generation on both synthesized and real datasets.

To apply causal intervention or counterfactual inference for feature learning, Huang et al.<sup>[63]</sup> proposed a causal intervention-based deconfounded visual grounding method to eliminate the confounding bias. Zhang et al.<sup>[65]</sup> present a causal inference based weakly-supervised semantic segmentation framework. Tang et al.<sup>[67]</sup> present a causal inference framework that disentangles the paradoxical effects of the momentum to remove the confounder of long-tail classification. Chen et al.<sup>[84]</sup> proposed a counterfactual critic MultiAgent training (CMAT) approach to learn the visual context properly.

Causal feature selection aims to find a subset of features from a large number of predictive features to reduce computational cost and build predictive models for variables of interest. Recent causality-based feature selection methods use Bayesian network (BN) and Markov boundary (MB) to identify potential causal features. BN is used as a DAG representing the causal relations between variables, and MB implies the local causal relationships between the class variable and the features in its MB. Since the BN of variables may be very large and hard to compute, current causal-based methods focus on identifying the MB as a variable or a subset of the MB. For example, Wu et al.<sup>[59]</sup> introduced the PCMasking concept to explain a type of incorrect CI tests in MB discovery and proposed a cross-check and complement MB discovery (CCMB) algorithm to solve the incorrect test problem. Yu et al.<sup>[60]</sup> presented theoretical analyses on the conditions for MB discovery in multiple interventional datasets and designed an algorithm for learning MBs from multiple interventional datasets. Yu et al.<sup>[58]</sup> formulated the causal feature selection problem with multiple datasets as a search problem and gave the upper and lower bounds of the invariant set, then proposed a multi-source feature selection algorithm. Yang et al.<sup>[55]</sup> proposed the concept of N-structures and then designed an MB discovering subroutine to integrate MB learning with N-structures to discover MB while distinguishing direct causes from direct effects. Yu et al.<sup>[53]</sup> proposed a multi-label feature selection algorithm, multi-label feature selection to causal structure learning (M2LC), which learns the causal mechanism behind the data and is able to select causally informative features and visualize common

features. Guo et al.<sup>[51]</sup> proposed an error-aware Markov blanket learning algorithm to solve the conditional independence test error in causal feature selection. Ling et al.<sup>[57]</sup> proposed an efficient local causal structure learning algorithm, local causal structure learning by feature selection (LCS-FS), which speeds up parent and children discovery by employing feature selection without searching for conditioning sets. Yu et al.<sup>[50]</sup> proposed a multiple imputation MB framework MimMB for causal feature selection with missing data. MimMB integrates data imputation with MB learning in a unified framework to enable the two key components to engage with each other.

Finding causal features improves the explanatory capability and robustness of models. Causal feature selection methods can provide a more convincing explanation for prediction than correlation-based methods. As the ICM principle implies, the underlying mechanism of the class variable can be learned from causal relations and thus can be transferred across different settings or environments. Although the existing causal feature learning methods achieve promising performance, most of them focus on general feature learning without considering a more specific problem, visual representation learning.

## 4 Visual representation learning: State-of-the-art

Visual representation learning has made great progress in recent years, which can utilize spatial or/and temporal information to complete specific tasks, including visual understanding (object detection, scene graph generation, visual grounding, visual commonsense reasoning), action detection and recognition, and visual question answering, etc. In this section, we introduce these representative visual learning tasks and discuss the existing challenges and necessity of applying causal reasoning to visual representation learning.

Object detection aims to determine where objects are located in a given image (object localization) and to which category each object belongs to (object classification) and label them with rectangular bounding boxes (BBs) to show their confidence in existence. In image object detection, deep learning frameworks for object detection are divided into two types. The first type is to follow the traditional object detection process, generating region proposals firstly and then classifying each proposal into a different object class. The other type is to treat object detection as a regression or classification problem and adopt a unified framework to directly obtain the final predictions (category and location). Region proposal-based methods mainly include R-CNN<sup>[98]</sup>, spatial pyramid pooling (SPP-Net)<sup>[99]</sup>, Fast R-CNN<sup>[100]</sup>, Faster R-CNN<sup>[101]</sup>, feature pyramid network (FPN)<sup>[102]</sup>, region-based fully convolutional Network (R-FCN)<sup>[103]</sup>, and Mask R-CNN<sup>[104]</sup>, some of which are interrelated (e.g., SPP-net modifies R-CNN with an SPP layer). Based on regression/classification, the methods mainly include

MultiBox<sup>[105]</sup>, AttentionNet<sup>[106]</sup>, G-CNN<sup>[107]</sup>, YOLO<sup>[108]</sup>, single shot MultiBox detector (SSD)<sup>[109]</sup>, YOLOv2<sup>[110]</sup>, deeply supervised object detector (DSOD)<sup>[111]</sup> and deconvolution single shot detector (DSSD)<sup>[112]</sup>. The correlations between these two pipelines are connected by anchors introduced in Faster R-CNN. In video saliency object detection, extending state-of-the-art saliency detectors from images to videos is challenging. Li et al.<sup>[113]</sup> presented a flow-guided recurrent neural encoder (FGRNE), which works by enhancing the temporal coherence of the per-frame feature by exploiting both motion information in terms of optical flow and sequential feature evolution encoding in terms of long short-term memory (LSTM) networks. Li et al.<sup>[114]</sup> developed a multi-task motion-guided video salient object detection network, which learns to accomplish two sub-tasks using two sub-networks, one sub-network is for salient object detection in still images and the other one is for motion saliency detection in optical flow images. Yan et al.<sup>[115]</sup> presented an effective video saliency detector that consists of a spatial refinement network and a spatiotemporal module. By utilizing the generated pseudo-labels together with a part of manual annotations, the detector can learn spatial and temporal cues for both contrast inference and coherence enhancement. For video salient object detection, how to effectively take object motion into consideration and obtain robust spatial-temporal information is crucial in video salient object detection. However, some non-object, occlusion, motion blur, and lens movement make the model hard to concentrate on the true interesting object area.

Scene graph generation (SGG) aims to describe object instances and relations between objects in a scene. With its powerful representation ability, SGG can encode images<sup>[116, 117]</sup> and videos<sup>[118, 119]</sup> as its abstract semantic elements without any restrictions on the attributes, types, and relations between objects. Therefore, the task of SGG is to build a graph structure that associates its nodes and edges well with objects in the scene and their relations, where the key challenge task is to detect/recognize relations between objects. Currently, SGG can be divided into two classes: 1) with facts alone and 2) introducing prior information. Besides, these SGG methods pay more attention to the methods with facts alone, including CRF-based (conditional random field) SGG<sup>[117, 120]</sup>, VTransE-based (visual translation embedding) SGG<sup>[121, 122]</sup>, RNN/LSTM-based SGG<sup>[123, 124]</sup>, Faster RCNN-based SGG<sup>[125, 126]</sup>, graph neural network (GNN)<sup>[127, 128]</sup>, etc. Furthermore, SGG adds different types of prior information, such as language priors<sup>[129]</sup>, knowledge priors<sup>[130, 131]</sup>, visual contextual information<sup>[132]</sup>, visual cue<sup>[133]</sup>, etc. Fig. 4 shows the related work on SGG, and it can be clearly seen that most of the methods use the GNN model or introduce relevant prior information when conducting SGG. Existing SGG methods are still far from building a practical knowledge base. There exists a serious conditional distribution bias of the relation-

ship in SGG methods. For example, knowing that the subject and object are person and head it is easy to guess that the relationship is that a person has a head.

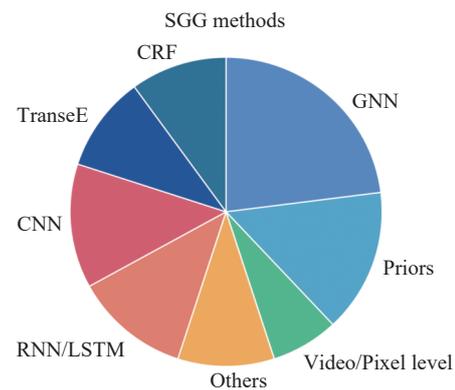


Fig. 4 Classification and statistics of SGG methods

Visual grounding usually involves two modalities, visual and linguistic data. This task aims to locate the target object in the image according to the corresponding object description (title or description) and the given image. When locating the target object, it is necessary to understand the input description information, and integrate the information of the visual modality for localization prediction. Currently, visual grounding methods can be classified into three types: fully supervised<sup>[134–141]</sup>, weakly supervised<sup>[142]</sup>, and unsupervised<sup>[143]</sup>. First, the fully supervised methods contain box annotations with object-phrase information. This method can be further divided into two-stage methods<sup>[135, 136, 138, 141]</sup> and one-stage method<sup>[139]</sup>. The two-stage approach is to extract candidate proposals and their features in the first stage through a region proposal network (RPN)<sup>[100]</sup> or traditional algorithms (Edgebox<sup>[144]</sup>, Selective Search<sup>[145]</sup>). Second, weakly supervised methods<sup>[146, 147]</sup> only have images and corresponding sentences, and no box annotations for object-phrases in the sentences. Due to the lack of mapping between phrases and boxes, weak supervision will additionally design many loss functions, such as designing reconstruction loss, introducing external knowledge, and designing loss functions based on image-caption matching. Third, there is no image-sentence information in the unsupervised method. Wang and Specia<sup>[148]</sup> used off-the-shelf approaches to detect objects, scenes and colors in images and explore different approaches to measuring semantic similarity between the categories of detected visual elements and words in phrases. To locate the object instance described by a natural language referring expression in an image, some referring expression comprehension methods are proposed. Yang et al.<sup>[149]</sup> proposed a dynamic graph attention network to perform multi-step reasoning by modeling the relationships among the objects in the image and the linguistic structure of the expression. Yang et al.<sup>[134]</sup> proposed a cross-modal relationship extractor (CMRE) to adaptively highlight objects

and relationships with a cross-modal attention mechanism, and represented the extracted information as a language-guided visual relation graph. Furthermore, Yang et al.<sup>[22]</sup> proposed a cross-modal relationship extractor to adaptively highlight objects and relationships (spatial and semantic relations) related to the given expression with a cross-modal attention mechanism, and represent the extracted information as a language-guided visual relation graph. Yang et al.<sup>[150]</sup> proposed a scene graph-guided modular network (SGMN), which performed reasoning over a semantic graph and a scene graph with neural modules under the guidance of the linguistic structure of the expression. However, due to the existence of linguistic and visual biases, most visual grounding models are heavily dependent on specific datasets, without good transfer ability and generalization performance.

Due to the success of BERT-related models in the field of NLP, researchers have begun to focus on more a challenging multi-modal reasoning task, visual commonsense reasoning (VCR). The VCR task needs to combine image information with the understanding of questions, and obtain the correct answer as well as the reasoning process based on the commonsense. Given an image, the image contains a series of bounding boxes with labels. In general, VCR can be divided into two sub-tasks:  $Q \rightarrow A$  task is choosing an answer based on the question; and  $QA \rightarrow R$  task is reasoning based on the question and the answer, explaining why the answer was chosen. Due to the challenging nature of VCR, there are actually relatively few existing studies. Some of them resort to designing specific model architectures<sup>[151–155]</sup>. Recognition to cognition networks (R2C)<sup>[151]</sup> implemented this task with a three-step approach, associating text with objects involved, linking answers with corresponding questions and objects, and finally, reasoning about shared representations. Inspired by brain neuron connectivity, CCN<sup>[152]</sup> dynamically modeled the visual neuron connectivity, which is contextualized by the queries and responses. Heterogeneous graph learning (HGL)<sup>[153]</sup> leveraged visual answering and dual question answering heterogeneous graphs to seamlessly connect vision and language. Zhang et al.<sup>[155]</sup> proposed a multi-level counterfactual contrastive learning network for VCR by jointly modeling the hierarchical visual contents and the inter-modality relationships between the visual and linguistic domains. Recently, BERT-based pre-training methods have been extensively explored in vision and language domains. In general, most of them adopt a pre-training-then-transfer scheme and achieve significant performance improvements on the VCR<sup>[156–158]</sup> benchmarks. These models are usually pre-trained on large-scale multi-modal datasets (e.g., concept captioning<sup>[159]</sup>) and then fine-tuned on VCR. At present, the promising performance of VCR is generally attributed to the pre-trained big model and the prior external knowledge. Compared with simple vision-linguistic domain tasks, the introduction of external knowledge brings

new challenges: 1) How to retrieve limited supporting knowledge from external knowledge bases that contain massive data. 2) How to effectively integrate external knowledge with visual and linguistic features. 3) The reasoning process gives interpretability needs supporting facts, which depends heavily on the language structure design.

The task of action detection and recognition includes two aspects, one is to identify all action instances in the video, and the other one is to localize actions spatially and temporally. Nowadays, spatial-temporal action detection or recognition models can be divided into two categories, the first one<sup>[6, 7, 160–167]</sup> is to model spatial-temporal relationships based on convolutional neural networks (CNNs), and the other one<sup>[168–171]</sup> is based on video transformer structures. Besides, the skeleton-based models<sup>[172–175]</sup> have recently attracted great attention. Sun et al.<sup>[161]</sup> proposed an actor-centric relational network (ACRN), which used two-stream to extract the central character feature and global background information from the input clip, and then performed feature fusion for action classification. Feichtenhofer et al.<sup>[165]</sup> proposed a two-stream model named SlowFast networks that contains a Slow pathway and a Fast pathway. Bertasius et al.<sup>[169]</sup> simply extended the ViT<sup>[176]</sup> design to video by proposing several scalable schemes for space-time self-attention. Arnab et al.<sup>[177]</sup> proposed pure-transformer architectures for video classification, including several variants of the model by factorizing the spatial and temporal dimensions of the input video. Although great progress has been made in spatial-temporal action detection and recognition based on the CNNs or transformer models, there exist some critical problems in terms of the robustness and the transferability of the models. The existing action detection and recognition models rely heavily on scenes and objects. When a model is well-trained in one dataset, it is hard to be generalized to another dataset with different scenes. Additionally, the methods are easily focused on some static appearance or background information rather than the true motion area due to the essential correlation learning in most of the models. This may be harmful to the reliability of the model, as well as the robustness of the learned spatial-temporal representations. Causal reasoning has the powerful ability to uncover the underlying structural knowledge about human actions that build a strong cognitive model that can fully discover causality and spatial-temporal relations.

Visual question answering (VQA) is a vision-language task that has received much attention recently. The objective of VQA is: Given the image/video and a related question, a machine needs to reason over visual elements and general knowledge to infer the correct answer. The attention mechanism is widely used in VQA models, which aim to focus on the critical part of the image and question, and find cross-modality correlations. The UpDn<sup>[178]</sup> framework is a typical conventional VQA meth-

od based on attention, which uses a top-down attention LSTM<sup>[179]</sup> for the fusion of visual and linguistic features. Besides using LSTM, the transformer<sup>[180]</sup> can also be adapted to the VQA task, thanks to its powerful scaled dot-product attention block. Visual-language pre-training (VLP) models based on BERT<sup>[181]</sup> show remarkable performance in the VQA task. ViLBERT<sup>[156]</sup> is a BERT-based visual and language pre-training framework, which uses a self-attention transformer block<sup>[180]</sup> to model in-modality relation and develop a co-attention transformer block to compute cross-modality attention score, and it finally achieves a state-of-the-art on four visual-language tasks including VQA at that time. Compared with the image QA<sup>[29, 30, 178, 182, 183]</sup>, the video question answering (VideoQA) task is much more challenging due to the existence of extra temporal information. To accomplish the VideoQA problem, the model needs to capture spatial, temporal, visual, and linguistic relations to reason about the answer. To explore relational reasoning in VideoQA, Xu et al.<sup>[184]</sup> proposed an attention mechanism to exploit the appearance and motion knowledge with the question as guidance. Later on, some hierarchical attention and co-attention-based methods are proposed to learn appearance-motion and question-related multi-modal interactions. Le et al.<sup>[185]</sup> proposed a hierarchical conditional relation network (HCRN) to construct sophisticated structures for representation and reasoning over videos. Jiang and Han<sup>[186]</sup> introduced a heterogeneous graph alignment (HGA) network. Huang et al.<sup>[187]</sup> proposed a location-aware graph convolutional network to reason over detected objects. Lei et al.<sup>[188]</sup> employed sparse sampling to build a transformer-based model named CLIPBERT and achieve end-to-end video-and-language understanding. Liu et al.<sup>[189]</sup> proposed a hierarchical visual-semantic relational reasoning (HAIR) framework to perform hierarchical relational reasoning. Although hierarchical attention mechanisms successfully improve the visual-language task performance, these models remain with a strong reliance on modality bias<sup>[28, 190]</sup> and tend to capture the spurious linguistic or visual correlations within the images/videos, and thus fail to learn the multi-modal knowledge with good generalization ability and interpretability.

## 5 Causality-aware visual representation learning

According to the above-discussed visual representation learning methods, the current machine learning, especially representation learning, faces several challenges: 1) lack of interpretability, 2) poor generalization ability, and 3) over-reliance on correlations of data distribution. Causal reasoning offers a promising alternative to address these challenges. The discovery of causality helps to uncover the causal mechanism behind the data, allowing the machine to understand better why and to make decisions through intervention or counterfactual reasoning.

Since Section 3 has reviewed the recent causal reasoning methods for general feature learning, it provides a good theoretical basis for further research on causal reasoning with specific visual representation learning tasks. In this section, we summarize some recent approaches for causal visual representation learning, as shown in Table 1. The causal visual representation learning is an emerging research topic and has appeared since the 2020s. The related tasks can be roughly categorized into several main aspects: 1) causal visual understanding, 2) causal visual robustness, and 3) causal visual question answering. In this section, we discuss these three representative causal visual representation learning tasks.

### 5.1 Causal visual understanding

Visual understanding contains several tasks, such as object detection, scene graph generation, visual grounding, visual commonsense reasoning, etc. However, some challenges exist in these tasks: 1) For image/video object saliency detection, some non-object, occlusion, motion blur, and lens movement make the model hard to concentrate on the true interesting object area. To this end, causal reasoning can make the model focus on the essential interesting object area by learning robust and reliable visual representations. 2) For the SGG problem that contains superficial bias and insufficient generalization ability, causal reasoning can be introduced to mitigate these problems well. For example, an item such as towel is used to bathe in the bathroom, but is used to wash the face in the office. Introducing causal reasoning into SGG can generalize the functionality of an item to different scenarios. 3) Due to the existence of linguistic and visual biases, most visual grounding models are heavily dependent on specific datasets without good transfer ability and generalization performance. This problem can be mitigated by causal reasoning methods, which learn robust and transferable features to mitigate the visual and linguistic biases. 4) For visual commonsense reasoning, linguistic biases may directly affect the reasoning performance. Generally, the superficial correlations captured by the existing VCR models can be mitigated by introducing causality that integrates external knowledge and visual and linguistic features into a robust and discriminative representation space. Non-causal visual understanding methods are easily affected by confounders in visual content. Illumination, position, backgrounds, co-occurrence of objects, and other visual factors are confounders that are inevitable in common settings. With traditional correlation learning, spurious correlations introduced by the confounders degrade the robustness of representation learning. For example, since the co-occurrence of “bird” and “sky” are high, the model would learn a strong correlation between them. Thus, when seeing a picture of a floating balloon that also contains “sky”, it would also make a confident prediction that it is a picture of a bird.

Table 1 Recent causal visual representation learning methods, tasks, basic models, and causal reasoning types

Author & Year	Tasks	Basic models	Types of causal reasoning
Agarwal et al. <sup>[191]</sup> 2020	Visual question answering	GAN <sup>[192]</sup>	Counterfactual sample synthesising
Chen et al. <sup>[84]</sup> 2020	Visual question answering	CCS <sup>[84]</sup>	Counterfactual sample synthesising
Zhang et al. <sup>[65]</sup> 2020	Weakly-supervised semantic segmentation	Pseudo-mask generation	Backdoor adjustment
Tang et al. <sup>[66]</sup> 2020	Scene graph generation	Unbiased training	Causal inference
Tang et al. <sup>[67]</sup> 2020	Long-tailed classification	De-confounded training	Causal inference
Wang et al. <sup>[68]</sup> 2020	IC & VQA & VCR	R-CNN	Causal intervention
Yue et al. <sup>[74]</sup> 2020	Few-shot learning	–	Backdoor adjustment
Hu et al. <sup>[72]</sup> 2021	Class incremental learning	–	Causal inference
Yue et al. <sup>[73]</sup> 2021	Unsupervised domain adaptation	GAN & VAE	Causal inference
Mao et al. <sup>[77]</sup> 2021	Domain adaptation	GAN	Causal inference
Tang et al. <sup>[71]</sup> 2021	Adversarial defense	Instrumental variable intervention	Causal intervention
Wang et al. <sup>[61]</sup> 2021	Visual recognition	Adversarial training	Causal inference
Yue et al. <sup>[62]</sup> 2021	Zero-shot & Open-set visual recognition	–	Counterfactual inference
Huang et al. <sup>[63]</sup> 2022	Visual grounding	CNN, BERT <sup>[181]</sup> , Attention <sup>[180]</sup>	Causal inference
Zhang et al. <sup>[193]</sup> 2020	Visual question answering	BERT <sup>[181]</sup>	Backdoor adjustment
Yang et al. <sup>[82]</sup> 2021	Visual question answering	Attention <sup>[180]</sup>	Front-door adjustment
Niu et al. <sup>[81]</sup> 2021	Visual question answering	CF-VQA <sup>[81]</sup>	Counterfactual inference
Li et al. <sup>[194]</sup> 2022	Video question answering	IGV <sup>[194]</sup>	Invariant grounding
Liu et al. <sup>[195]</sup> 2022	Visual recognition	CNN/Transformer, CCD <sup>[195]</sup>	Causal context debiasing
Liu et al. <sup>[196]</sup> 2022	Motion forecasting	Encoder-decoder <sup>[196]</sup>	Causal invariant learning
Lv et al. <sup>[197]</sup> 2022	Domain generalization	CIRL <sup>[197]</sup>	Causal intervention
Lin et al. <sup>[198]</sup> 2022	Video anomaly detection	UVAD <sup>[198]</sup>	Causal intervention, counterfactual inference
Lin et al. <sup>[199]</sup> 2022	Salient object detection	USOD <sup>[199]</sup>	Causal intervention
Liu et al. <sup>[200]</sup> 2022	Event-level video question answering	Transformer, CMCIR <sup>[200]</sup>	Causal intervention

Causal reasoning provides a good solution to address the above problem. By replacing the conditional distribution with the intervened distribution, the spurious correlation can be eliminated, and the machine can learn the real causality. Applying intervention in the training procedure is a widely used implementation of the causal intervention. In the visual recognition task, Wang et al.<sup>[61]</sup> combined adversarial training with causal intervention, modeled the different causal effects of mediators and confounders, and designed an adversarial training pipeline to improve the effect of mediators while suppressing the effect of confounders. Yue et al.<sup>[62]</sup> applied counterfactual inference to zero-shot and open-set visual recognition by proposing a generative causal model to generate counterfactual samples. A confounding effect also exists in the visual grounding task, Huang et al.<sup>[63]</sup> proposed a deconfounded visual grounding framework by conducting interventions on linguistic features. For the weakly-supervised semantic segmentation task, Zhang et al.<sup>[65]</sup> used the structural causal model to formulate the causalities

between components, then constructed a confounder set and removed confounders by back-door adjustment. The prior bias is also a non-trivial problem in the scene graph generation (SGG) task. To reduce the negative impact of training bias in scene graph generation, Tang et al.<sup>[66]</sup> built a causal graph and extracted counterfactual causality from the trained graph to infer the causal effects of training bias and then remove the negative bias. Besides, causal reasoning can replace the traditional re-weighting and re-sampling methods in resolving long-tailed distribution problems. Tang et al.<sup>[67]</sup> analyzed that the momentum in stochastic gradient descent (SGD) introduces the unbalanced sample distribution and then proposed to use counterfactual inference in the test stage to detect and remove the causal effect of the momentum item. Wang et al.<sup>[68]</sup> proposed an unsupervised commonsense learning framework to learn intervened visual features by back-door adjustment, which can be used in the downstream task as image captioning, visual question answering, and visual commonsense reasoning. Liu et al.<sup>[195]</sup> es-

tablished a SCM to uncover the causal relevance among contextual priori, object feature, contextual bias, and final prediction in multi-target visual tasks. Liu et al.<sup>[196]</sup> introduced a causal formalism of motion forecasting, which casts the problem as a dynamic process with three groups of latent variables, namely invariant variables, style confounders, and spurious features. Lin et al.<sup>[198]</sup> proposed a causal graph to analyze the confounding effect of the pseudo label generation process for unsupervised video anomaly detection. Lin et al.<sup>[199]</sup> proposed a causal-based debiasing framework to disentangle the unsupervised salient object detection from the impact of contrast distribution bias and spatial distribution bias.

## 5.2 Causal visual robustness

The ubiquitous spurious correlation learned by deep learning models reduces the model robustness, which is a potential vulnerability of the conventional deep learning paradigm. In this perspective, the causal learning paradigm can be introduced to avoid the presence of confounding effects and make the model more robust<sup>[201]</sup>.

Confounders are widespread in visual robustness problems, including few-shot learning, class-incremental learning, domain adaptation, generative model, etc. Yue et al.<sup>[74]</sup> uncovered that pre-trained knowledge is a confounder in fewshot learning and developed a few-shot learning paradigm by introducing back-door adjustment to control the pre-trained knowledge. The confounding effect can be leveraged by attackers, Tang et al.<sup>[71]</sup> proposed an instrumental variable<sup>[202]</sup> estimation-based causal regularization method for adversarial defense. Hu et al.<sup>[72]</sup> explained the catastrophic forgetting effect in class-incremental learning in terms of causality: the causal effect of old data is zero, and then proposed distilling the causal effect of old data by controlling the collider effect of the causal graph. As ICM inferred, causal mechanisms could be invariant across domains; hence, learning invariant causal knowledge is likely to be superior in robustness. To learn cross-domain knowledge, Yue et al.<sup>[73]</sup> disentangled semantic attributes in images into causal factors and used CycleGAN<sup>[203]</sup> to generate counterfactual samples in the counterpart domain, then exploited the counterfactual sample and a latent variable encoded by variational auto encoder (VAE)<sup>[204]</sup> as proxy variables of an unobserved attribute for intervention. Apart from generating counterfactual samples, intervention can also be implemented by the generative method. Mao et al.<sup>[77]</sup> argued that conventional randomized control trials and intervention approaches could hardly be used in naturally collected images, then introduced a framework performing interventions on realistic images by steering generative models to generate intervened distribution. Lv et al.<sup>[197]</sup> introduced a causality inspired representation learning (CIRL) algorithm that enforced the representations to satisfy three properties, and then used them to simulate causal factors.

## 5.3 Causal visual question answering

For visual question answering, the real causality behind the visual-linguistic modalities and the interaction between the appearance-motion and language knowledge are neglected in most of the existing methods. In recent works, the purpose of introducing causality into visual question-answering tasks is to reduce language bias in VQA tasks. Strong correlations between the question and the answer will make VQA models rely on spurious correlations without concerning visual knowledge. For example, since the answer to the question “What is the color of the apple?” is “red” in most cases, the VQA model will easily learn the correlation between the word “apple” and the word “red”. Thus, when given an image of a green apple, the model still predicts the answer “red” with strong confidence. Although simply balancing the dataset<sup>[28, 205]</sup> can partly mitigate the linguistic bias, the spurious correlation still exists in the model. From this perspective, the causality-based solution is better than simply balancing the data, since the causal reasoning cuts off the superficial correlations and makes the VQA models focus on the real causality.

Constructing a confounder set has been commonly used in causal intervention practice. VC-RCNN<sup>[68]</sup> constructed an object level visual confounder set for performing back-door adjustment in a visual task. Following VC-RCNN, DeVLBert<sup>[193]</sup> treated nouns in linguistic modality as confounders and constructed language confounder sets using their average Bert representation vectors. Besides, DeVLBert incorporated the intervention into Bert’s<sup>[181]</sup> pre-training process and combined mask modeling objective with causal intervention. As another implementation of the intervention, Yang et al.<sup>[82]</sup> designed the In-Sample attention and Cross-Sample attention module to conduct front-door adjustment, where the In-Sample attention module approximates probability  $P(W = w|x)$ , and the Cross-Sample attention module approximates probability  $P(x)$ . Using these attention modules, a cross modality causal attention network was proposed for the VQA task by combining causal attention with the previous LXMERT<sup>[206]</sup> framework. Counterfactual-based solutions are also worth noting. Agarwal et al.<sup>[191]</sup> proposed a counterfactual sample synthesizing method based on generative adversarial network (GAN)<sup>[192]</sup>. Overcoming the complexity of the GAN based synthesizing method, Chen et al.<sup>[84]</sup> tried to replace critical objects and critical words with mask tokens and reassigned an answer to the synthesis of counterfactual QA pairs. Apart from sample synthesizing methods, Niu et al.<sup>[81]</sup> developed a counterfactual VQA framework that reduces multi-modality bias by using the total indirect effect (TIE)<sup>[39]</sup> for final inference. By blocking the direct effect of one modality, the TIE measures the total causal effect of the question and visual information, thus reducing language bias in the VQA. Li et al.<sup>[194]</sup> proposed an invariant grounding for VideoQA

(IGV) to force the VideoQA models to shield the answering process from the negative influence of spurious correlations, which significantly improves the reasoning ability. Liu et al.<sup>[200]</sup> proposed a causality-aware event-level visual question answering framework named cross-modal causal relational reasoning (CMCIR) to discover true causal structures via causal intervention on the integration of visual and linguistic modalities.

### 5.4 Comparisons and discussions

To summarize the development line and the current state of causal visual representation learning, we show the development situation of causal visual representation learning in Fig. 5, including the past, current, and future directions. Although the above-mentioned causal visual representation learning methods successfully apply causal reasoning methods to uncover causal mechanisms and achieve promising results, causal reasoning for visual representation learning is still in its infancy stage with many challenges. Firstly, the existing causal visual representation tasks are limited to several computer vision tasks without being applied to more diverse and challenging tasks such as video understanding, human-computer interaction, urban computing, etc. It should be noticed that recent large pre-trained vision-language models like contrastive language-image pre-training (CLIP)<sup>[207]</sup> have shown great potential in learning representations that are transferable across a wide range of downstream tasks. Different from traditional representation learning, which is based mostly on discretized labels, popular prompt learning adopts vision-language pre-training and aligns images/videos/texts in a common feature space, which allows zero-shot transfer to any downstream task via prompting. Therefore, how to apply causality-aware knowledge to prompt learning may be a potential direction. Secondly, causal reasoning has been burgeoned for many visual learning tasks. So far, the existing evaluation datasets are still as traditional datasets for correlation learn-

ing without proper large-scale benchmarking datasets and pipelines specified for causal reasoning. Thirdly, most of the existing methods focus on causality discovery on either visual or linguistic modality without considering both of them. Therefore, a more in-depth analysis of the relations between causal reasoning and visual representation learning is required.

## 6 Related causal datasets

Correlation-based models may perform well in existing datasets, not because these models have a strong reasoning capability, but because these datasets cannot fully support the evaluation of the models' reasoning capability. Spurious correlations in these datasets can be exploited by the model to cheat, which means that the model just concentrates on superficial correlation learning, not real causal reasoning, only approximating the distribution of the dataset. For example, in the VQA v1.0<sup>[182]</sup> dataset for the VQA task, the model simply answers "yes" when seeing the question "Do you see a ...", which will achieve nearly 90% accuracy. Due to this shortcoming in current datasets, researchers need to build benchmarks that can evaluate the true causal reasoning capability of models. In this section, we take image question answering benchmarks and video question answering benchmarks as examples to analyze the current research situation of related causal reasoning datasets and give some future directions.

### 6.1 Image question answering

Image question answering benchmarks evaluate the models' capability to answering natural language questions based on a corresponding image. Recent image question answering benchmarks try to collect or generate balanced QA pairs to make the dataset distribution more balanced in question distribution. VQA v2.0<sup>[28]</sup> collects complementary QA pairs by replacing the image and the answer in QA pairs. VQA-CP<sup>[190]</sup> resplited the VQA v1

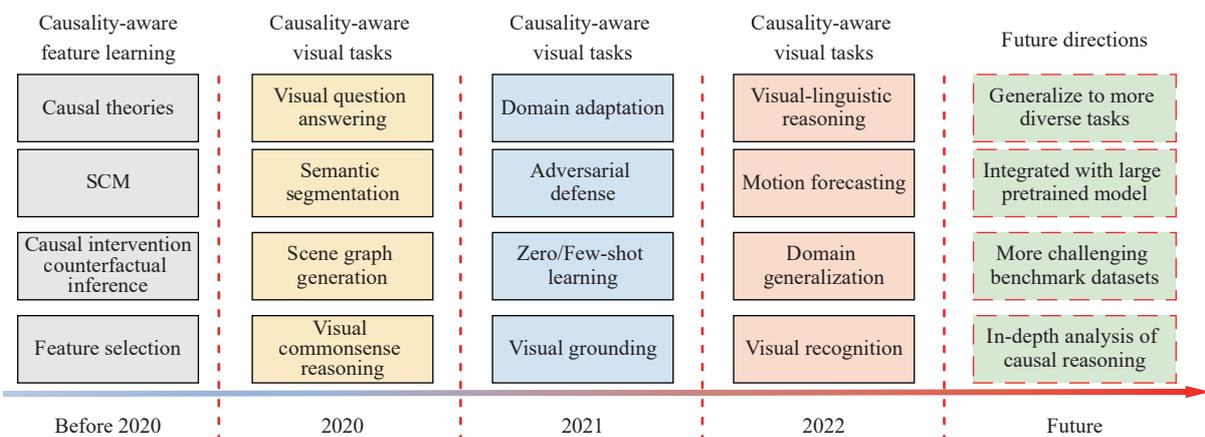


Fig. 5 Development timeline of causal visual representation learning including the past, current, and future directions

dataset and VQA v2 dataset to construct two new datasets VQA-CP v1 and VQA-CP v2. As Fig. 6 shows, Agarwal et al.<sup>[191]</sup> constructed IV-VQA and CV-VQA datasets using semantic editing to generate images and reexamine the image by a human. Li et al.<sup>[208]</sup> proposed a human-machine adversarial to collect robust QA pairs. Fig. 7 illustrates the adversarial data collection procedure. In Table 2, we summarize these datasets in terms of image source, split numbers, collected or not, and rebalanced or not.

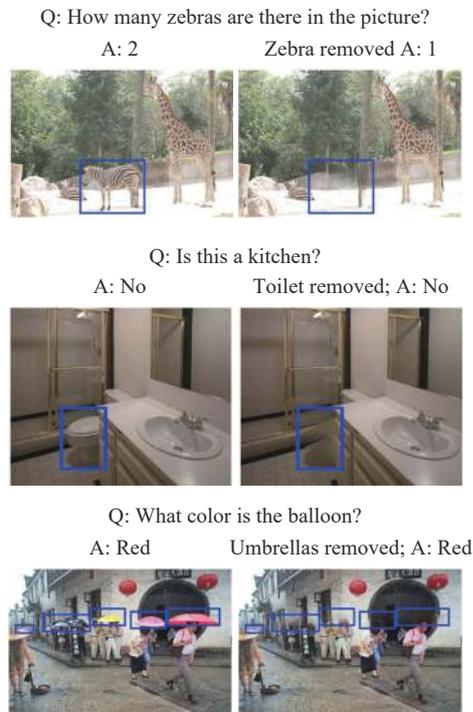


Fig. 6 Examples in IV-VQA and CV-VQA datasets<sup>[191]</sup>

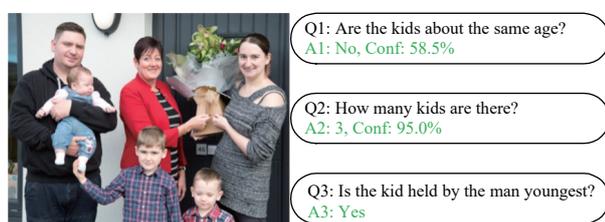


Fig. 7 The example QA pairs in AVQA<sup>[208]</sup>

Current image question answering uses various approaches to overcome the bias introduced by unbalanced data. However, there is still a lack of large-scale benchmark datasets that support fair and transparent evaluations of the causality behind the data and the reasoning ability of the method. Introducing the causal concept and causal methods like confounders and causal interventions when building benchmark datasets may help resolve the problem of the lack of specific causal reasoning benchmark datasets.

## 6.2 Video question answering

The video question answering task is more complex than the image question answering task due to the ubiquitous correlation between spatial and temporal information, i.e., the introduction of complex temporal relations. Thus, improving the spatial-temporal causal reasoning ability of models can improve the performance on this task, but simply approximating data distributions usually does not work. Thus, some recently released benchmark datasets are proposed to evaluate whether the model has the reasoning ability to understand the causal relation knowledge within the visual and linguistic content, as shown in Table 3.

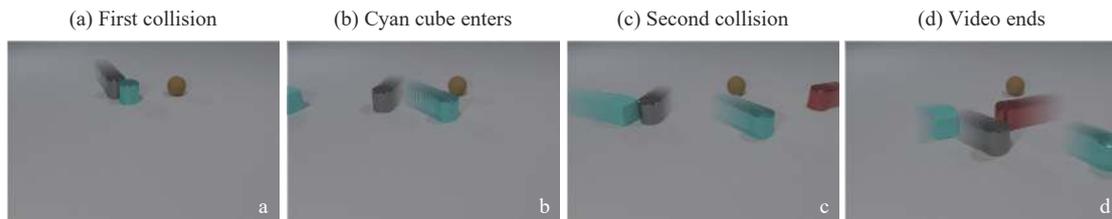
CLEVRER<sup>[210]</sup> contains synthesized videos and automatically generated questions describing the collision of geometric objects. A typical video and question types from CLEVRER are shown in Fig. 8. It is a balanced and synthetic dataset that contains diagnostic annotations and counterfactuals. VQuAD<sup>[211]</sup> is also a diagnostic synthesized dataset. It is constructed from a balanced dataset by separating objects into attributes like texture and color, and balancing the data distribution based on these attributes. A brief overview of the VQuAD objects is shown in Fig. 9. The VQuAD is a diagnostic dataset that can be used to evaluate the extent of reasoning abilities of various video QA methods. ComPhy<sup>[212]</sup> is a video QA dataset that focuses on understanding object-centric and relational physics properties hidden from visual appearances. As shown in Fig. 10, the ComPhy dataset studies object's intrinsic physical properties from object's interactions and how these properties affect their motions in future and counterfactual scenes to answer the corresponding questions. Action genome question answering (AGQA)<sup>[213]</sup> includes numerous QA pairs, which are automatically generated by the process. An overview of AGQA is shown in Fig. 11. The QA pairs in AGQA are generated by parsing videos to scene graphs and using the language composition inference by scene graph to generate QA pairs. SUTD-TrafficQA<sup>[214]</sup> is a traffic video question answering dataset with six challenging reasoning tasks, including basic understanding, event forecasting, reverse reasoning, counterfactual inference, introspection, and attribution analysis, to analyze the models' reasoning ability. Fig. 12 shows an example of a counterfactual traffic video question answering process from SUTD-TrafficQA. To be noticed, the counterfactual traffic video question answering task in Fig. 12 requires the outcome of a certain hypothesis that does not occur in the video. To accurately reason about the imagined events under the designated condition, the model is required to not only conduct relational reasoning in a hierarchical way but also fully explore the causal, logic, and spatial-temporal structures of the visual and linguistic content. NExT-QA<sup>[215]</sup> is a video question answering benchmark targeting the explanation of the content of the video, which requires a deeper understanding of videos and reasoning about causal and temporal actions from rich object inter-

Table 2 Current causal image question answering datasets

Datasets	Image source	Split (train/val/test)	Is collected?	Rebalanced?
VQA v1	COCO <sup>[209]</sup>	614K/--	Yes	No
VQA v2 <sup>[28]</sup>	COCO <sup>[209]</sup>	443K/214K/453K	No	Yes
VQA-CP v1 <sup>[190]</sup>	COCO <sup>[209]</sup>	245K/--/125K	Yes	No
VQA-CP v2 <sup>[190]</sup>	COCO <sup>[209]</sup>	438K/--/220K	No	Yes
IV-VQA <sup>[191]</sup>	COCO <sup>[209]</sup>	257K/11.6K/108K	No	Yes
CV-VQA <sup>[191]</sup>	COCO <sup>[209]</sup>	8.5K/0.4K/3.7K	No	Yes
AVQA <sup>[208]</sup>	Various	142.1K/8.7K/26.4K	Yes	Yes

Table 3 Video question answering (VideoQA) benchmarks related with reasoning. OC and OG denote open-ended question answering as problem of classification and generation, respectively. MC stands for multi-choice QA.

Datasets	Topic	QA pairs	QA task	Annotation	Real world	Diagnostic annotations	Counterfactual	Balanced
CLEVRER <sup>[210]</sup>	Object collision	282K	MC	Auto		✓	✓	✓
VQuAD <sup>[211]</sup>	Object movement	1.35M	MC	Auto		✓		✓
ComPhy <sup>[212]</sup>	Hidden physical properties	100K	OC&MC	Auto		✓	✓	✓
AGQA <sup>[213]</sup>	Natural video scenes	192M	OC&Compositional	Auto	✓			
SUTD-TrafficQA <sup>[214]</sup>	Traffic events	62K	MC&OC	Human	✓		✓	✓
NExT-QA <sup>[215]</sup>	Causal and temporal interactions	52K	MC&OG	Human	✓		✓	✓



**I. Descriptive**

**Q:** What shape is the object that collides with the cyan cylinder?  
**A:** Cylinder  
**Q:** How many metal objects are moving when the video ends?  
**A:** 3

**II. Explanatory**

**Q:** Which of the following is responsible for the gray cylinder-s colliding with the cube?  
 a) The presence of the sphere  
 b) The collision between the gray cylinder and the cyan cylinder  
**A:** b)

**III. Predictive**

**Q:** Which event will happen next?  
 a) The cube collides with the red object  
 b) The cyan cylinder collides with the red object  
**A:** a)

**IV. Counterfactual**

**Q:** Without the gray object, which event will not happen?  
 a) The cyan cylinder collides with the sphere  
 b) The red object and the sphere collide  
**A:** a), b)

Fig. 8 A sample in CLEVRER, including four question types: descriptive, explanatory, predictive, and counterfactual<sup>[210]</sup>

actions in daily activities. As shown in Fig. 13, the NExT-QA dataset contains rich object interactions and requires causal and temporal action reasoning in realistic videos. The NExT-QA dataset challenges QA models to reason about causal and temporal actions and understands rich object interactions in daily activities.

## 7 Extensive applications

Causal reasoning with visual representation learning

has a variety of applications. Modeling causal reasoning for a variety of tasks can achieve a better perception of the real world. In this section, we introduce the applications from five aspects: image/video analysis, explainable artificial intelligence, recommendation system, human-computer dialog and interaction, and crowd intelligence analysis. We also discuss how causal reasoning benefits various real-world applications, as shown in Fig. 14.

In image/video analysis, most of the existing work re-

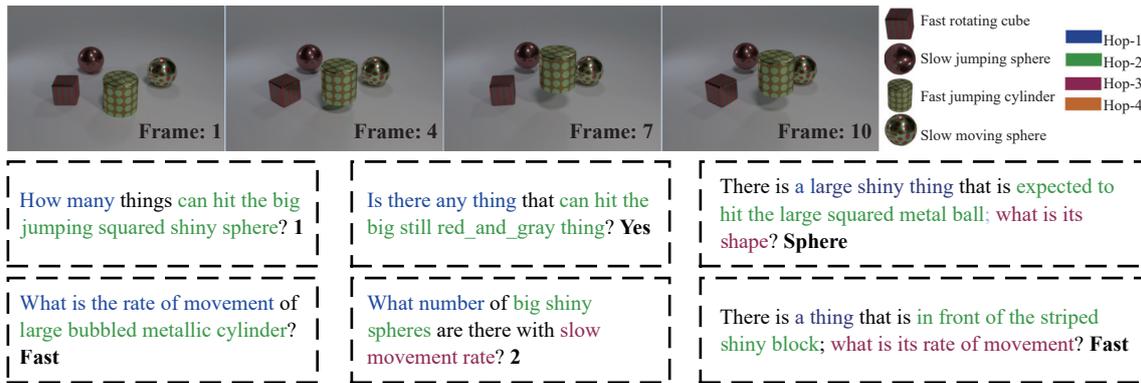


Fig. 9 Illustration of an instance of VQuAD dataset<sup>[211]</sup>, which shows various questions that are generated concerning the video created and the difference in complexity in terms of hops for the questions

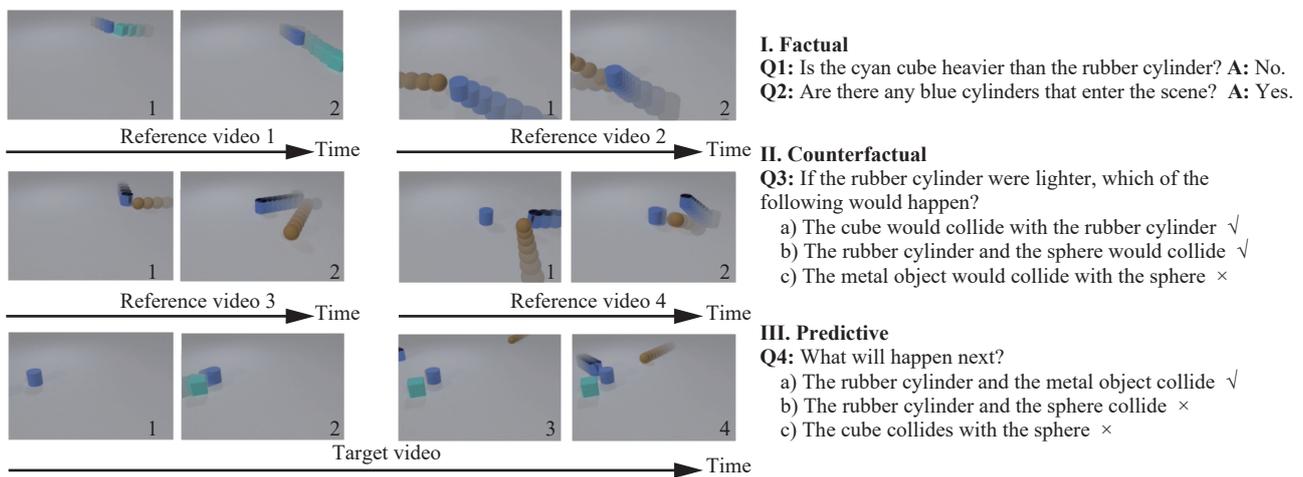


Fig. 10 Sample reference videos, target video, and question-answer pairs from ComPhy dataset<sup>[212]</sup>

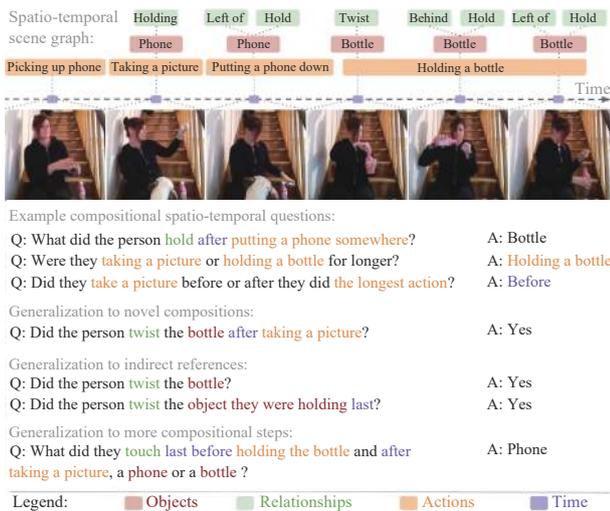


Fig. 11 An overview of AGQA dataset<sup>[213]</sup>

lies on learning data correlations rather than causal structures, and the superficial correlation within the image and video data makes the model vulnerable to visual changes in the dataset. Therefore, a causality-aware feature learning strategy is required to make the model learn essential

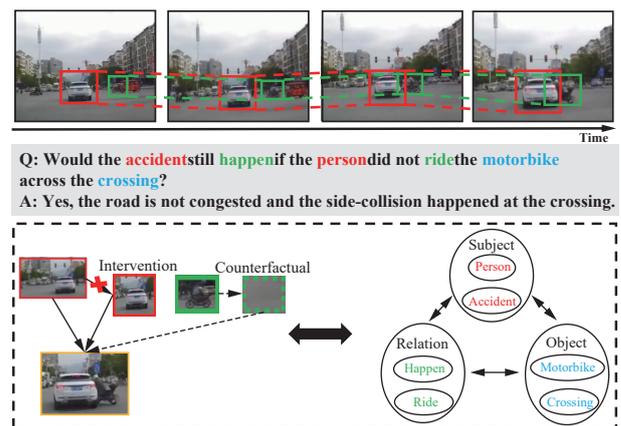


Fig. 12 An example of counterfactual question-answer pair in SUTD-TrafficQA dataset<sup>[214]</sup>

causal structures behind the data and robust to different data distributions. One of the main methods of dealing with superficial data correlations is using the causal intervention. Assume that commonsense knowledge exists in visual features, but commonsense might be confused by false observation bias. For example, the words “cup”, “table”, and “stool” have high co-occurrence frequencies be-

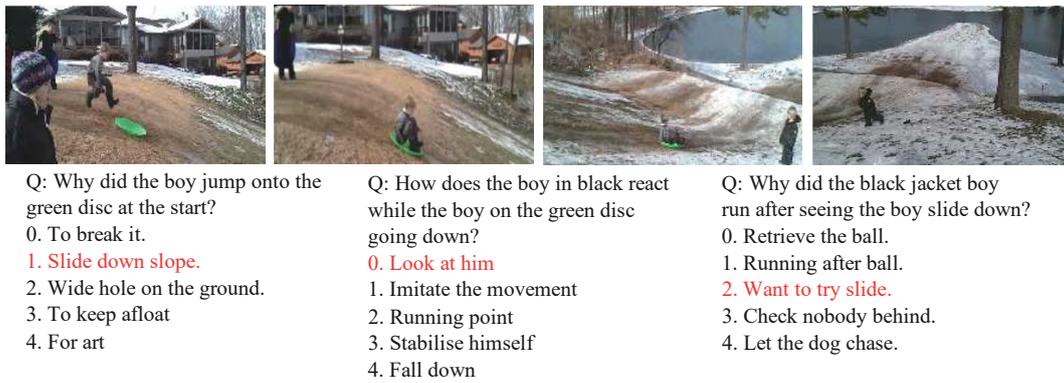


Fig. 13 Examples of multi-choice QA in NExT-QA dataset<sup>[215]</sup>

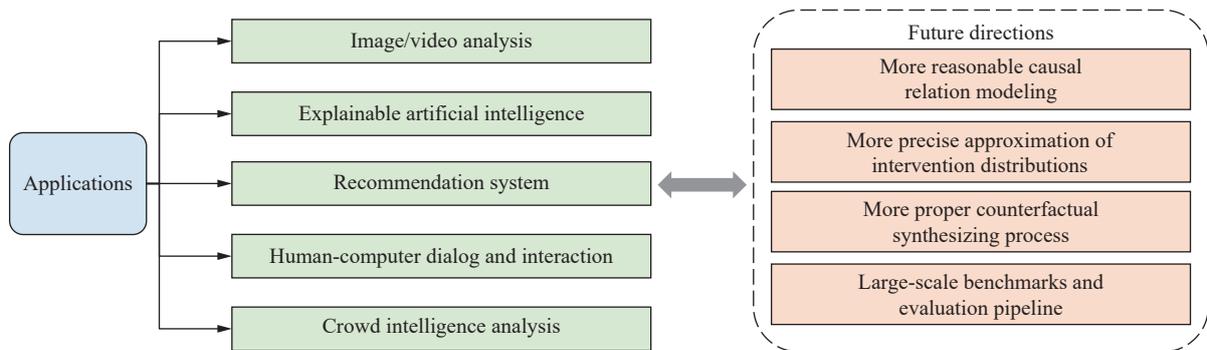


Fig. 14 Extensive applications and future directions

cause they commonly appear in daily life, but the commonsense knowledge usually wrongly predicts the class as table due to the observation bias. To reduce the observation bias, the causality-aware visual commonsense model is required, which regards the object category as a confounding factor and directly maximizes the likelihood after the intervention to learn the visual feature representation. By eliminating observation bias, the learned visual features are robust to image and video analysis tasks. A representative task is weakly supervised object localization, detection, and grounding<sup>[216–218]</sup>, which aims to localize objects described in the sentence to visual regions in the image/video. Despite recent progress, existing methods may suffer from the severe problem of spurious association such as: 1) the association is not object-relevant but extremely ambiguous due to weak supervision, and 2) the association is unavoidably confounded by the observational bias when using statistics-based methods. Therefore, a unified causal framework is required to learn the deconfounded object-relevant association for accurate and robust video object localization, detection and grounding.

With the development of deep learning across industries and disciplines, the applications of deep learning models in real-world scenes require a high degree of robustness, interpretability, and transparency. Unfortunately, the black-box properties of deep neural networks are still not fully explainable, and many machine decisions are still poorly understood<sup>[219]</sup>. In recent years,

causal interpretability has received increasing attention. These works<sup>[220–226]</sup> have made progress in explainable artificial intelligence based on causal interpretability. For example, in the current COVID-19 pandemic, causal mediation analysis helps disentangle different effects contributing to case fatality rates when an example of Simpson’s paradox was observed<sup>[227]</sup>. Learning the best treatment rules for each patient is one of the promising goals of applying explainable treatment effect estimation methods in the medical field. Since the effects of different available drugs can be estimated and explainable, doctors can prescribe better drugs accordingly.

At present, some causal reasoning works<sup>[228–235]</sup> have been applied to the recommendation system. The recommendation system is actually a problem of causal reasoning<sup>[228]</sup>. User embedding represents what type of person the user is and infers the user’s preferences based on the user’s attributes. The causal effect of a recommendation system is whether the user is satisfied with the recommendation. Superficial bias exists because the recommendation system is trained on biased samples (both users and items). An example is a personalized recommendation, where we wish to build a model of a customer’s shopping interest through various data sources, such as Web Browser records and shopping history. However, if we train a recommendation system on customers’ records in controlled settings, the system may provide little additional insight compared to the customers’ mental states and emotions, thus may fail when deployed. While

it may be useful to automate certain decisions, understanding causality may be necessary to recommend commodities that are personalized and reliable. A general approach to removing survival bias is to construct counterfactual mirror users, construct similarity measures using unbiased information, and construct matches from low-active to high-active users. In this way, we can alleviate the user's dissatisfaction with the previously recommended content and the low user activity.

For human-computer dialog and interaction, some emerging tasks contain the interaction between vision and language. Additionally, there exist multi-modal spatial-temporal information and complex relations captured by various devices. Most of the existing work relies on data correlation rather than causal relevant evidence, and the false correlation in the data makes the model vulnerable to language biases in the problems. Take a VQA task as an example, where we aim to remove visual objects that are unrelated to answering the question, and the prediction of the model is not expected to change. This can prevent the model from relying on superficial data correlations. When changing objects that are related to a question, the model is expected to change the answer accordingly. Adjusting question-related objects encourages the model to predict based on causality-aware objects. For a better user experience, the human-computer dialog and interaction system is required to understand people's purposes and make reliable decisions. Causal reasoning is beneficial to the pursuit of reliable human-computer interaction by uncovering and modeling heterogeneous spatial-temporal information in a reliable and explainable way. Especially for robot interaction<sup>[91, 233–235]</sup>, where the relevant environmental features are not known in advance, prior knowledge can be utilized as a good candidate for causal structures. The strong relation between causal reasoning and its ability to intervene in the world suggests that causal reasoning can greatly address this challenge for robotics, which benefits the application of robotics significantly.

The applications mentioned above usually focus on a single subject, whereas crowd intelligence analysis<sup>[236]</sup> aims to address related sensing and cognitive tasks for multiple subjects and their interaction. In recent years, we have been witnessing the explosive growth of multi-modal heterogeneous spatial/temporal/spatial-temporal data from different kinds of data sensors. Urban computing<sup>[237]</sup> is an example of crowd intelligence analysis, which aims to tackle traffic congestion, energy consumption, and pollution by using the data that has been generated by a large number of traffic vehicles in cities (e.g., traffic flow, human mobility, and geographical data). For example, huge amounts of heterogeneous traffic data come from various sources, including both static and dynamic data, such as traffic road networks, geographic information system (GIS) data, traffic flow, traffic mobility, traffic energy consumption, etc. Moreover, the heterogen-

eous spatial-temporal traffic data contains a large number of useful traffic rules with strong causal relations. Therefore, how to utilize different heterogeneous spatial-temporal data and discover their complex and entangled causal relations is beneficial to urban computing and crowd intelligence analysis.

## 8 More detailed discussions

Some researchers have successfully implemented causal reasoning for visual representation learning to discover causality and visual relations. However, causal reasoning for visual representation learning is still in its infancy stage, and many issues remain unsolved. Therefore, this section highlights several possible research directions and open problems to inspire further extensive and in-depth research on this topic. Potential research directions for causal visual representation learning can be summarized as: 1) more reasonable causal relation modeling; 2) more precise approximation of intervention distributions; 3) more proper counterfactual synthesizing process; 4) large-scale benchmarks and evaluation pipeline.

### 8.1 More reasonable causal relation modeling

Reasonable causality modeling is the basis for causal inference. Real-world data like visual information is usually unstructured, and the effect of causal relation may be unobserved. For example, momentum is likely to be detrimental under long-tailed distribution data<sup>[67]</sup>, and there is no consensus on how to properly model causality on many tasks because the real causality may be more complicated than expected. For the VQA task, Yang et al.<sup>[82]</sup> treated visual and language features as one vertex in the causal graph, and Niu et al.<sup>[81]</sup> consider the visual and linguistic features separately. However, these methods focus on causality discovery on either visual or linguistic modality without considering both of them. Therefore, future work should consider: 1) in-depth analysis of the relations between causal reasoning and visual representation learning; 2) model comprehensive and reasonable causal relation.

### 8.2 More precise approximation of intervention distributions

A precise estimation of the intervention distribution helps the implementation of a certain causal model. Most of the current intervention distribution approximation methods focus on identifying all confounders for a certain task, while these confounders are usually defined as the average of object features in visual tasks<sup>[61, 68, 193]</sup>. Actually, the average features may not properly describe a certain confounder, especially for complex heterogeneous visual data. Thus, how to approximate the confounders more accurately is a key future work that needs to be fur-

ther considered for causal intervention methods.

### 8.3 More proper counterfactual synthesizing process

Counterfactual inference-based methods usually focus on refining the training procedure, i.e., embedding the counterfactual inference process into the training procedure. Counterfactual synthesizing methods<sup>[62, 77, 84, 191]</sup> have proved their effectiveness in many tasks. Embedding counterfactual inference into models can effectively eliminate data bias within the data. A novel counterfactual framework<sup>[81]</sup> gives us insight into this potential. However, visual data is often entangled and heterogeneous, which makes the data bias hard to understand and model. Therefore, how to model a proper counterfactual synthesizing process is a potential direction of data debiasing in visual representation.

### 8.4 Large-scale benchmarks and evaluation pipeline

Although causal reasoning has been burgeoned for many visual learning tasks, most of the existing evaluation datasets are still traditional datasets for correlation learning without proper large-scale benchmarking datasets and pipelines to support fair and transparent evaluations of emerging research contributions. The only existing causal datasets discussed in the above sections have limited scale and lack comprehensive evaluation standards for causal reasoning. Therefore, more large-scale benchmark datasets and pipelines for specific visual representation learning tasks should be considered in future research.

Generally, causal visual representation learning is still an emerging and challenging research topic. Causal modeling, intervened distribution approximation, counterfactual inference, large-scale benchmarks, and evaluation pipelines have great potential for further exploration.

## 9 Conclusions

This paper has provided a comprehensive survey on causal reasoning for visual representation learning. In this paper, we focus on the prospective survey of related works, datasets, insights, future challenges and opportunities for causal reasoning, visual representation learning, and their integration. We mathematically present the basic concepts of causality, the SCM, the ICM principle, causal inference, and causal intervention. Then, based on the analysis, we further give some directions for conducting causal reasoning on visual representation learning tasks. We also review some recent popular visual learning tasks, including visual understanding, action detection and recognition, and visual question answering, including the discussions about the existing challenges of these visual learning methods. In addition, the related

causality-based visual representation learning works and datasets are also discussed systematically. Finally, extensive applications and some potential future research directions are provided for further exploration. We hope that this survey can help attract attention, encourage discussions, and bring to the forefront the urgency of developing novel causal reasoning methods, publicly available benchmarks, and consensus-building standards for reliable visual representation learning and related real-world applications more efficiently.

## Acknowledgements

This work was supported in part by National Natural Science Foundation of China (Nos. 62002395, 61976250 and U1811463), the National Key R&D Program of China (No. 2021ZD011601), the Guangdong Basic and Applied Basic Research Foundation, China (Nos. 2021A15150123 and 2020B1515020048).

## Open Access

This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made.

The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

- [1] K. M. He, X. Y. Zhang, S. Q. Ren, J. Sun. Deep residual learning for image recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, pp. 770–778, 2016. DOI: [10.1109/CVPR.2016.90](https://doi.org/10.1109/CVPR.2016.90).
- [2] T. S. Chen, L. Lin, R. Q. Chen, X. L. Hui, H. F. Wu. Knowledge-guided multi-label few-shot learning for general image recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 3, pp. 1371–1384, 2022. DOI: [10.1109/TPAMI.2020.3025814](https://doi.org/10.1109/TPAMI.2020.3025814).
- [3] A. R. Akula, K. Z. Wang, C. S. Liu, S. Saba-Sadiya, H. J. Lu, S. Todorovic, J. Chai, S. C. Zhu. CX-ToM: Counterfactual explanations with theory-of-mind for enhancing human trust in image recognition models. *iScience*, vol. 25, no. 1, Article number 103581, 2022. DOI: [10.1016/j.isci.2021.103581](https://doi.org/10.1016/j.isci.2021.103581).
- [4] L. M. Wang, Y. J. Xiong, Z. Wang, Y. Qiao, D. H. Lin, X.

- O. Tang, L. van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Proceedings of the 14th European Conference on Computer Vision*, Springer, Amsterdam, The Netherlands, pp.20–36, 2016. DOI: [10.1007/978-3-319-46484-8\\_2](https://doi.org/10.1007/978-3-319-46484-8_2).
- [5] B. L. Zhou, A. Andonian, A. Oliva, A. Torralba. Temporal relational reasoning in videos. In *Proceedings of the 15th European Conference on Computer Vision*, Springer, Munich, Germany, pp.831–846, 2018. DOI: [10.1007/978-3-030-01246-5\\_49](https://doi.org/10.1007/978-3-030-01246-5_49).
- [6] J. Lin, C. Gan, S. Han. TSM: Temporal shift module for efficient video understanding. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp.7082–7092, 2019. DOI: [10.1109/ICCV.2019.00718](https://doi.org/10.1109/ICCV.2019.00718).
- [7] Y. Liu, K. Z. Wang, L. B. Liu, H. Y. Lan, L. Lin. TCGL: Temporal contrastive graph for self-supervised video representation learning. *IEEE Transactions on Image Processing*, vol. 31, pp.1978–1993, 2022. DOI: [10.1109/TIP.2022.3147032](https://doi.org/10.1109/TIP.2022.3147032).
- [8] M. Buřta, L. Neumann, J. Matas. Deep TextSpotter: An end-to-end trainable scene text localization and recognition framework. In *Proceedings of IEEE International Conference on Computer Vision*, Venice, Italy, pp.2223–2231, 2017. DOI: [10.1109/ICCV.2017.242](https://doi.org/10.1109/ICCV.2017.242).
- [9] X. X. Chen, L. W. Jin, Y. Z. Zhu, C. J. Luo, T. W. Wang. Text recognition in the wild: A survey. *ACM Computing Surveys*, vol. 54, no. 2, Article number 42, 2022. DOI: [10.1145/3440756](https://doi.org/10.1145/3440756).
- [10] R. Rastgoo, K. Kiani, S. Escalera. Sign language recognition: A deep survey. *Expert Systems with Applications*, vol. 164, Article number 113794, 2021. DOI: [10.1016/j.eswa.2020.113794](https://doi.org/10.1016/j.eswa.2020.113794).
- [11] R. H. Gao, T. H. Oh, K. Grauman, L. Torresani. Listen to look: Action recognition by previewing audio. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.10454–10464, 2020. DOI: [10.1109/CVPR42600.2020.01047](https://doi.org/10.1109/CVPR42600.2020.01047).
- [12] Y. Cheng, R. Z. Wang, Z. H. Pan, R. Feng, Y. J. Zhang. Look, listen, and attend: Co-attention network for self-supervised audio-visual representation learning. In *Proceedings of the 28th ACM International Conference on Multimedia*, Seattle, USA, pp.3884–3892, 2020. DOI: [10.1145/3394171.3413869](https://doi.org/10.1145/3394171.3413869).
- [13] Y. B. Chen, Y. Q. Xian, A. S. Koepke, Y. Shan, Z. Akata. Distilling audio-visual knowledge by compositional contrastive learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp.7012–7021, 2021. DOI: [10.1109/CVPR46437.2021.00694](https://doi.org/10.1109/CVPR46437.2021.00694).
- [14] H. Y. Lan, Y. Liu, L. Lin. Audio-visual contrastive learning for self-supervised action recognition. [Online], Available: <https://arxiv.org/abs/2204.13386>, 2022.
- [15] Y. Liu, Z. Y. Lu, J. Li, T. Yang. Hierarchically learned view-invariant representations for cross-view action recognition. *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 29, no.8, pp.2416–2430, 2019. DOI: [10.1109/TCSVT.2018.2868123](https://doi.org/10.1109/TCSVT.2018.2868123).
- [16] Y. Liu, Z. Y. Lu, J. Li, C. Yao, Y. Z. Deng. Transferable feature representation for visible-to-infrared cross-dataset human action recognition. *Complexity*, vol.2018, Article number 5345241, 2018. DOI: [10.1155/2018/5345241](https://doi.org/10.1155/2018/5345241).
- [17] Y. Liu, Z. Y. Lu, J. Li, T. Yang, C. Yao. Deep image-to-video adaptation and fusion networks for action recognition. *IEEE Transactions on Image Processing*, vol.29, pp.3168–3182, 2020. DOI: [10.1109/TIP.2019.2957930](https://doi.org/10.1109/TIP.2019.2957930).
- [18] Y. Liu, K. Z. Wang, G. B. Li, L. Lin. Semantics-aware adaptive knowledge distillation for sensor-to-vision action recognition. *IEEE Transactions on Image Processing*, vol.30, pp.5573–5588, 2021. DOI: [10.1109/TIP.2021.3086590](https://doi.org/10.1109/TIP.2021.3086590).
- [19] Y. Y. Zhu, Y. Zhang, L. B. Liu, Y. Liu, G. B. Li, M. Z. Mao, L. Lin. Hybrid-order representation learning for electricity theft detection. *IEEE Transactions on Industrial Informatics*, to be published. DOI: [10.1109/THI.2022.3179243](https://doi.org/10.1109/THI.2022.3179243).
- [20] G. B. Li, Y. Xie, L. Lin, Y. Z. Yu. Instance-level salient object segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp.247–256, 2017. DOI: [10.1109/CVPR.2017.34](https://doi.org/10.1109/CVPR.2017.34).
- [21] X. D. Liang, K. Gong, X. H. Shen, L. Lin. Look into person: Joint body parsing & pose estimation network and a new benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 41, no. 4, pp.871–885, 2019. DOI: [10.1109/TPAMI.2018.2820063](https://doi.org/10.1109/TPAMI.2018.2820063).
- [22] S. B. Yang, G. B. Li, Y. Z. Yu. Relationship-embedded representation learning for grounding referring expressions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.43, no.8, pp.2765–2779, 2021. DOI: [10.1109/TPAMI.2020.2973983](https://doi.org/10.1109/TPAMI.2020.2973983).
- [23] X. Q. Zhang, R. H. Jiang, C. X. Fan, T. Y. Tong, T. Wang, P. C. Huang. Advances in deep learning methods for visual tracking: Literature review and fundamentals. *International Journal of Automation and Computing*, vol. 18, no. 3, pp.311–333, 2021. DOI: [10.1007/s11633-020-1274-8](https://doi.org/10.1007/s11633-020-1274-8).
- [24] Z. W. Wang, Q. She, A. Smolic. ACTION-Net: Multipath excitation for action recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp.13209–13218, 2021. DOI: [10.1109/CVPR46437.2021.01301](https://doi.org/10.1109/CVPR46437.2021.01301).
- [25] G. S. Pang, C. Yan, C. H. Shen, A. van den Hengel, X. Bai. Self-trained deep ordinal regression for end-to-end video anomaly detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.12170–12179, 2020. DOI: [10.1109/CVPR42600.2020.01219](https://doi.org/10.1109/CVPR42600.2020.01219).
- [26] Y. Liu, Z. Y. Lu, J. Li, T. Yang, C. Yao. Global temporal representation based cnns for infrared action recognition. *IEEE Signal Processing Letters*, vol. 25, no. 6, pp.848–852, 2018. DOI: [10.1109/LSP.2018.2823910](https://doi.org/10.1109/LSP.2018.2823910).
- [27] L. F. Wu, Q. Wang, M. Jian, Y. Qiao, B. X. Zhao. A comprehensive review of group activity recognition in videos. *International Journal of Automation and Computing*, vol. 18, no. 3, pp.334–350, 2021. DOI: [10.1007/s11633-020-1258-8](https://doi.org/10.1007/s11633-020-1258-8).
- [28] Y. Goyal, T. Khot, D. Summers-Stay, D. Batra, D. Parikh. Making the v in VQA matter: Elevating the role of image understanding in visual question answering. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp.6325–6334, 2017. DOI: [10.1109/CVPR.2017.670](https://doi.org/10.1109/CVPR.2017.670).
- [29] Q. X. Cao, B. L. Li, X. D. Liang, K. Z. Wang, L. Lin. Knowledge-routed visual question reasoning: Challenges for deep representation embedding. *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 7, pp.2758–2767, 2022. DOI: [10.1109/TNNLS.2020.3045034](https://doi.org/10.1109/TNNLS.2020.3045034).

- [30] Q. X. Cao, W. T. Wan, K. Z. Wang, X. D. Liang, L. Lin. Linguistically routing capsule network for out-of-distribution visual question answering. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.1594–1603, 2021. DOI: [10.1109/ICCV48922.2021.00164](https://doi.org/10.1109/ICCV48922.2021.00164).
- [31] Y. Liu, J. Li, Z. Y. Lu, T. Yang, Z. J. Liu. Combining multiple features for cross-domain face sketch recognition. In *Proceedings of the 11th Chinese Conference on Biometric Recognition*, Springer, Chengdu, China, pp.139–146, 2016. DOI: [10.1007/978-3-319-46654-5\\_16](https://doi.org/10.1007/978-3-319-46654-5_16).
- [32] L. C. Wang, Z. M. Ding, Z. Q. Tao, Y. Y. Liu, Y. Fu. Generative multi-view human action recognition. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp.6211–6220, 2019. DOI: [10.1109/ICCV.2019.00631](https://doi.org/10.1109/ICCV.2019.00631).
- [33] J. Y. Ni, R. Sarbajna, Y. Liu, A. H. H. Ngu, Y. Yan. Cross-modal knowledge distillation for vision-to-sensor action recognition. In *Proceedings of IEEE International Conference on Acoustics, Speech and Signal Processing*, Singapore, pp.4448–4452, 2022. DOI: [10.1109/ICASSP43922.2022.9746752](https://doi.org/10.1109/ICASSP43922.2022.9746752).
- [34] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, W. Brendel. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. In *Proceedings of the 7th International Conference on Learning Representations*, New Orleans, USA, 2019.
- [35] R. Shetty, B. Schiele, M. Fritz. Not using the car to see the sidewalk—Quantifying and controlling the effects of context in classification and segmentation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp.8210–8218, 2019. DOI: [10.1109/CVPR.2019.00841](https://doi.org/10.1109/CVPR.2019.00841).
- [36] D. Hendrycks, T. Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *Proceedings of the 7th International Conference on Learning Representations*, New Orleans, USA, 2019.
- [37] A. Azulay, Y. Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal of Machine Learning Research*, vol.20, no.184, pp.1–25, 2019.
- [38] J. Peters, D. Janzing, B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*, Cambridge, USA: The MIT Press, 2017.
- [39] J. Pearl. *Causality*, 2nd ed., New York, USA: Cambridge University Press, 2009.
- [40] B. Schölkopf, F. Locatello, S. Bauer, N. R. Ke, N. Kalchbrenner, A. Goyal, Y. Bengio. Toward causal representation learning. *Proceedings of the IEEE*, vol.109, no.5, pp.612–634, 2021. DOI: [10.1109/JPROC.2021.3058954](https://doi.org/10.1109/JPROC.2021.3058954).
- [41] L. Cheng, R. C. Guo, R. Moraffah, P. Sheth, K. S. Candan, H. Liu. Evaluation methods and measures for causal learning algorithms. *IEEE Transactions on Artificial Intelligence*, to be published. DOI: [10.1109/TAI.2022.3150264](https://doi.org/10.1109/TAI.2022.3150264).
- [42] Q. S. Zhang, S. C. Zhu. Visual interpretability for deep learning: A survey. *Frontiers of Information Technology & Electronic Engineering*, vol.19, no.1, pp.27–39, 2018. DOI: [10.1631/FITEE.1700808](https://doi.org/10.1631/FITEE.1700808).
- [43] Q. S. Zhang, Y. N. Wu, S. C. Zhu. Interpretable convolutional neural networks. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp.8827–8836, 2018. DOI: [10.1109/CVPR.2018.00920](https://doi.org/10.1109/CVPR.2018.00920).
- [44] Q. S. Zhang, Y. Yang, H. T. Ma, Y. N. Wu. Interpreting CNNs via decision trees. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp.6254–6263, 2019. DOI: [10.1109/CVPR.2019.00642](https://doi.org/10.1109/CVPR.2019.00642).
- [45] Q. S. Zhang, X. Wang, R. M. Cao, Y. N. Wu, F. Shi, S. C. Zhu. Extraction of an explanatory graph to interpret a CNN. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.43, no.11, pp.3863–3877, 2021. DOI: [10.1109/TPAMI.2020.2992207](https://doi.org/10.1109/TPAMI.2020.2992207).
- [46] Q. S. Zhang, J. Ren, G. Huang, R. M. Cao, Y. N. Wu, S. C. Zhu. Mining interpretable AOG representations from convolutional networks via active question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.43, no.11, pp.3949–3963, 2021. DOI: [10.1109/TPAMI.2020.2993147](https://doi.org/10.1109/TPAMI.2020.2993147).
- [47] Q. S. Zhang, X. Wang, Y. N. Wu, H. L. Zhou, S. C. Zhu. Interpretable CNNs for object classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.43, no.10, pp.3416–3431, 2021. DOI: [10.1109/TPAMI.2020.2982882](https://doi.org/10.1109/TPAMI.2020.2982882).
- [48] K. Yu, X. J. Guo, L. Liu, J. Y. Li, H. Wang, Z. L. Ling, X. D. Wu. Causality-based feature selection: Methods and evaluations. *ACM Computing Surveys*, vol.53, no.5, Article number 111, 2021. DOI: [10.1145/3409382](https://doi.org/10.1145/3409382).
- [49] K. Yu, L. Liu, J. Y. Li. A unified view of causal and non-causal feature selection. *ACM Transactions on Knowledge Discovery from Data*, vol.15, no.4, Article number 63, 2021. DOI: [10.1145/3436891](https://doi.org/10.1145/3436891).
- [50] K. Yu, Y. J. Yang, W. Ding. Causal feature selection with missing data. *ACM Transactions on Knowledge Discovery from Data*, vol.16, no.4, Article number 66, 2022. DOI: [10.1145/3488055](https://doi.org/10.1145/3488055).
- [51] X. J. Guo, K. Yu, F. Y. Cao, P. P. Li, H. Wang. Error-aware Markov blanket learning for causal feature selection. *Information Sciences*, vol.589, pp.849–877, 2022. DOI: [10.1016/j.ins.2021.12.118](https://doi.org/10.1016/j.ins.2021.12.118).
- [52] X. Li, Z. Z. Zhang, G. Q. Wei, C. L. Lan, W. J. Zeng, X. Jin, Z. B. Chen. Confounder identification-free causal visual feature learning. [Online], Available: <https://arxiv.org/abs/2111.13420>, 2021.
- [53] K. Yu, M. Z. Cai, X. Y. Wu, L. Liu, J. Y. Li. Multilabel feature selection: A local causal structure learning approach. *IEEE Transactions on Neural Networks and Learning Systems*, to be published. DOI: [10.1109/TNNLS.2021.3111288](https://doi.org/10.1109/TNNLS.2021.3111288).
- [54] M. Y. Yang, F. R. Liu, Z. T. Chen, X. W. Shen, J. Y. Hao, J. Wang. CausalVAE: Disentangled representation learning via neural structural causal models. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, 2021, pp.9588–9597. DOI: [10.1109/CVPR46437.2021.00947](https://doi.org/10.1109/CVPR46437.2021.00947).
- [55] S. Yang, H. Wang, K. Yu, F. Y. Cao, X. D. Wu. Towards efficient local causal structure learning. *IEEE Transactions on Big Data*, to be published. DOI: [10.1109/TB-DATA.2021.3062937](https://doi.org/10.1109/TB-DATA.2021.3062937).
- [56] L. Z. Li, Y. J. Lin, H. Zhao, J. K. Chen, S. Z. Li. Causality-based online streaming feature selection. *Concurrency and Computation: Practice and Experience*, vol.33, no.20, Article number e6347, 2021. DOI: [10.1002/cpe.6347](https://doi.org/10.1002/cpe.6347).
- [57] Z. L. Ling, K. Yu, H. Wang, L. Li, X. D. Wu. Using fea-

- ture selection for local causal structure learning. *IEEE Transactions on Emerging Topics in Computational Intelligence*, vol. 5, no. 4, pp. 530–540, 2021. DOI: [10.1109/TETCI.2020.2978238](https://doi.org/10.1109/TETCI.2020.2978238).
- [58] K. Yu, L. Liu, J. Y. Li, W. Ding, T. D. Le. Multi-source causal feature selection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 9, pp. 2240–2256, 2020. DOI: [10.1109/TPAMI.2019.2908373](https://doi.org/10.1109/TPAMI.2019.2908373).
- [59] X. Y. Wu, B. B. Jiang, K. Yu, H. Y. Chen, H. H. Chen. Accurate Markov boundary discovery for causal feature selection. *IEEE Transactions on Cybernetics*, vol. 50, no. 12, pp. 4983–4996, 2020. DOI: [10.1109/TCYB.2019.2940509](https://doi.org/10.1109/TCYB.2019.2940509).
- [60] K. Yu, L. Liu, J. Y. Li. Learning Markov blankets from multiple interventional data sets. *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 6, pp. 2005–2019, 2020. DOI: [10.1109/TNNLS.2019.2927636](https://doi.org/10.1109/TNNLS.2019.2927636).
- [61] T. Wang, C. Zhou, Q. R. Sun, H. W. Zhang. Causal attention for unbiased visual recognition. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp. 3071–3080, 2021. DOI: [10.1109/ICCV48922.2021.00308](https://doi.org/10.1109/ICCV48922.2021.00308).
- [62] Z. Q. Yue, T. Wang, Q. R. Sun, X. S. Hua, H. W. Zhang. Counterfactual zero-shot and open-set visual recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 15399–15409, 2021. DOI: [10.1109/CVPR46437.2021.01515](https://doi.org/10.1109/CVPR46437.2021.01515).
- [63] J. Q. Huang, Y. Qin, J. X. Qi, Q. R. Sun, H. W. Zhang. Deconfounded visual grounding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 1, pp. 998–1006, 2022. DOI: [10.1609/aaai.v36i1.19983](https://doi.org/10.1609/aaai.v36i1.19983).
- [64] C. Zhang, B. X. Jia, M. Edmonds, S. C. Zhu, Y. X. Zhu. ACRE: Abstract causal reasoning beyond covariation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 10638–10648, 2021. DOI: [10.1109/CVPR46437.2021.01050](https://doi.org/10.1109/CVPR46437.2021.01050).
- [65] D. Zhang, H. W. Zhang, J. H. Tang, X. S. Hua, Q. R. Sun. Causal intervention for weakly-supervised semantic segmentation. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Vancouver, Canada, Article number 56, 2020. DOI: [10.5555/3495724.3495780](https://doi.org/10.5555/3495724.3495780).
- [66] K. H. Tang, Y. L. Niu, J. Q. Huang, J. X. Shi, H. W. Zhang. Unbiased scene graph generation from biased training. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp. 3713–3722, 2020. DOI: [10.1109/CVPR42600.2020.00377](https://doi.org/10.1109/CVPR42600.2020.00377).
- [67] K. H. Tang, J. Q. Huang, H. W. Zhang. Long-tailed classification by keeping the good and removing the bad momentum causal effect. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Vancouver, Canada, Article number 128, 2020. DOI: [10.5555/3495724.3495852](https://doi.org/10.5555/3495724.3495852).
- [68] T. Wang, J. Q. Huang, H. W. Zhang, Q. R. Sun. Visual commonsense R-CNN. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp. 10757–10767, 2020. DOI: [10.1109/CVPR42600.2020.01077](https://doi.org/10.1109/CVPR42600.2020.01077).
- [69] L. Chen, H. W. Zhang, J. Xiao, X. N. He, S. L. Pu, S. F. Chang. Counterfactual critic multi-agent training for scene graph generation. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 4612–4622, 2019. DOI: [10.1109/ICCV.2019.00471](https://doi.org/10.1109/ICCV.2019.00471).
- [70] J. X. Shi, H. W. Zhang, J. Z. Li. Explainable and explicit visual reasoning over scene graphs. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 8368–8376, 2019. DOI: [10.1109/CVPR.2019.00857](https://doi.org/10.1109/CVPR.2019.00857).
- [71] K. H. Tang, M. Y. Tao, H. W. Zhang. Adversarial visual robustness by causal intervention. [Online], Available: <https://arxiv.org/abs/2106.09534>, 2021.
- [72] X. T. Hu, K. H. Tang, C. Y. Miao, X. S. Hua, H. W. Zhang. Distilling causal effect of data in class-incremental learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 3956–3965, 2021. DOI: [10.1109/CVPR46437.2021.00395](https://doi.org/10.1109/CVPR46437.2021.00395).
- [73] Z. Q. Yue, Q. R. Sun, X. S. Hua, H. W. Zhang. Transporting causal mechanisms for unsupervised domain adaptation. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp. 8579–8588, 2021. DOI: [10.1109/ICCV48922.2021.00848](https://doi.org/10.1109/ICCV48922.2021.00848).
- [74] Z. Q. Yue, H. W. Zhang, Q. R. Sun, X. S. Hua. Interventional few-shot learning. In *Proceedings of the 34th Conference on Neural Information Processing Systems*, Vancouver, Canada, pp. 2734–2746, 2020.
- [75] S. Yang, K. Yu, F. Y. Cao, L. Liu, H. Wang, J. Y. Li. Learning causal representations for robust domain adaptation. *IEEE Transactions on Knowledge and Data Engineering*, to be published. DOI: [10.1109/TKDE.2021.3119185](https://doi.org/10.1109/TKDE.2021.3119185).
- [76] R. Christiansen, N. Pfister, M. E. Jakobsen, N. Gnecco, J. Peters. A causal framework for distribution generalization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to be published. DOI: [10.1109/TPAMI.2021.3094760](https://doi.org/10.1109/TPAMI.2021.3094760).
- [77] C. Z. Mao, A. Cha, A. Gupta, H. Wang, J. F. Yang, C. Vondrick. Generative interventions for causal learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 3946–3955, 2021. DOI: [10.1109/CVPR46437.2021.00394](https://doi.org/10.1109/CVPR46437.2021.00394).
- [78] T. Kyono, M. van der Schaar. Exploiting causal structure for robust model selection in unsupervised domain adaptation. *IEEE Transactions on Artificial Intelligence*, vol. 2, no. 6, pp. 494–507, 2021. DOI: [10.1109/TAI.2021.3101185](https://doi.org/10.1109/TAI.2021.3101185).
- [79] F. Wu, X. Y. Duan, J. Xiao, Z. Zhao, S. L. Tang, Y. Zhang, Y. T. Zhuang. Temporal interaction and causal influence in community-based question answering. *IEEE Transactions on Knowledge and Data Engineering*, vol. 29, no. 10, pp. 2304–2317, 2017. DOI: [10.1109/TKDE.2017.2720737](https://doi.org/10.1109/TKDE.2017.2720737).
- [80] Y. L. Niu, H. W. Zhang. Introspective distillation for robust question answering. In *Proceedings of the 35th Conference on Neural Information Processing Systems*, pp. 16292–16304, 2021.
- [81] Y. L. Niu, K. H. Tang, H. W. Zhang, Z. W. Lu, X. S. Hua, J. R. Wen. Counterfactual VQA: A cause-effect look at language bias. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 12695–12705, 2021. DOI: [10.1109/](https://doi.org/10.1109/)

- CVPR46437.2021.01251.
- [82] X. Yang, H. W. Zhang, G. J. Qi, J. F. Cai. Causal attention for vision-language tasks. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp.9842–9852, 2021. DOI: [10.1109/CVPR46437.2021.00972](https://doi.org/10.1109/CVPR46437.2021.00972).
- [83] J. X. Qi, Y. L. Niu, J. Q. Huang, H. W. Zhang. Two causal principles for improving visual dialog. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.10857–10866, 2020. DOI: [10.1109/CVPR42600.2020.01087](https://doi.org/10.1109/CVPR42600.2020.01087).
- [84] L. Chen, X. Yan, J. Xiao, H. W. Zhang, S. L. Pu, Y. T. Zhuang. Counterfactual samples synthesizing for robust visual question answering. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.10797–10806, 2020. DOI: [10.1109/CVPR42600.2020.01081](https://doi.org/10.1109/CVPR42600.2020.01081).
- [85] P. Wu, J. Liu. Learning causal temporal relation and feature discrimination for anomaly detection. *IEEE Transactions on Image Processing*, vol.30, pp.3513–3527, 2021. DOI: [10.1109/TIP.2021.3062192](https://doi.org/10.1109/TIP.2021.3062192).
- [86] W. J. Shi, G. Huang, S. J. Song, C. Wu. Temporal-spatial causal interpretations for vision-based reinforcement learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to be published. DOI: [10.1109/TPAMI.2021.3133717](https://doi.org/10.1109/TPAMI.2021.3133717).
- [87] X. H. Zhang, Y. K. Wong, X. F. Wu, J. W. Lu, M. Kankanhalli, X. D. Li, W. D. Geng. Learning causal representation for training cross-domain pose estimator via generative interventions. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.11250–11260, 2021. DOI: [10.1109/ICCV48922.2021.01108](https://doi.org/10.1109/ICCV48922.2021.01108).
- [88] Z. W. Xu, X. D. Shen, Y. Wong, M. S. Kankanhalli. Unsupervised motion representation learning with capsule autoencoders. In *Proceedings of the 35th Conference on Neural Information Processing Systems*, pp.3205–3217, 2021.
- [89] A. Fire, S. C. Zhu. Inferring hidden statuses and actions in video by causal reasoning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition Workshops*, IEEE, Honolulu, USA, pp.48–56, 2017. DOI: [10.1109/CVPRW.2017.13](https://doi.org/10.1109/CVPRW.2017.13).
- [90] V. N. Gangapure, S. Nanda, A. S. Chowdhury. Superpixel-based causal multisensor video fusion. *IEEE Transactions on Circuits and Systems for Video Technology*, vol.28, no.6, pp.1263–1272, 2018. DOI: [10.1109/TCSVT.2017.2662743](https://doi.org/10.1109/TCSVT.2017.2662743).
- [91] C. M. Xiong, N. Shukla, W. L. Xiong, S. C. Zhu. Robot learning with a spatial, temporal, and causal and-or graph. In *Proceedings of IEEE International Conference on Robotics and Automation*, IEEE, Stockholm, Sweden, pp.2144–2151, 2016. DOI: [10.1109/ICRA.2016.7487364](https://doi.org/10.1109/ICRA.2016.7487364).
- [92] Y. Liu, K. Z. Wang, H. Y. Lan, L. Lin. Temporal contrastive graph learning for video action recognition and retrieval. [Online], Available: <https://arxiv.org/abs/2101.00820>, 2021.
- [93] X. Yang, H. W. Zhang, J. F. Cai. Deconfounded image captioning: A causal retrospect. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to be published. DOI: [10.1109/TPAMI.2021.3121705](https://doi.org/10.1109/TPAMI.2021.3121705).
- [94] Z. Y. Shen, J. S. Liu, Y. He, X. X. Zhang, R. Z. Xu, H. Yu, P. Cui. Towards out-of-distribution generalization: A survey. [Online], Available: <https://arxiv.org/abs/2108.13624>, 2021.
- [95] J. W. Chen, H. D. Dong, X. Wang, F. L. Feng, M. N. Wang, X. He. Bias and debias in recommender system: A survey and future directions. [Online], Available: <https://arxiv.org/abs/2010.03240>, 2020.
- [96] J. Mitrovic, B. McWilliams, J. Walker, L. Buesing, C. Blundell. Representation learning via invariant causal mechanisms. In *Proceedings of the 9th International Conference on Learning Representations*, 2021.
- [97] X. W. Shen, F. R. Liu, H. Z. Dong, Q. Lian, Z. T. Chen, T. Zhang. Disentangled generative causal representation learning. [Online], Available: <https://arxiv.org/abs/2010.02637>, 2020.
- [98] R. Girshick, J. Donahue, T. Darrell, J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, Columbus, USA, pp.580–587, 2014. DOI: [10.1109/CVPR.2014.81](https://doi.org/10.1109/CVPR.2014.81).
- [99] K. M. He, X. Y. Zhang, S. Q. Ren, J. Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol.37, no.9, pp.1904–1916, 2015. DOI: [10.1109/TPAMI.2015.2389824](https://doi.org/10.1109/TPAMI.2015.2389824).
- [100] R. Girshick. Fast R-CNN. In *Proceedings of IEEE International Conference on Computer Vision*, IEEE, Santiago, Chile, pp.1440–1448, 2015. DOI: [10.1109/ICCV.2015.169](https://doi.org/10.1109/ICCV.2015.169).
- [101] S. Q. Ren, K. M. He, R. Girshick, J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. In *Proceedings of the 28th International Conference on Neural Information Processing Systems*, Montreal, Canada, pp.91–99, 2015. DOI: [10.5555/2969239.2969250](https://doi.org/10.5555/2969239.2969250).
- [102] T. Y. Lin, P. Dollár, R. Girshick, K. M. He, B. Hariharan, S. Belongie. Feature pyramid networks for object detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, Honolulu, USA, pp.936–944, 2017. DOI: [10.1109/CVPR.2017.106](https://doi.org/10.1109/CVPR.2017.106).
- [103] J. F. Dai, Y. Li, K. M. He, J. Sun. R-FCN: Object detection via region-based fully convolutional networks. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Barcelona, Spain, pp.379–387, 2016. DOI: [10.5555/3157096.3157139](https://doi.org/10.5555/3157096.3157139).
- [104] K. M. He, G. Gkioxari, P. Dollár, R. Girshick. Mask R-CNN. In *Proceedings of IEEE International Conference on Computer Vision*, IEEE, Venice, Italy, pp.2980–2988, 2017. DOI: [10.1109/ICCV.2017.322](https://doi.org/10.1109/ICCV.2017.322).
- [105] D. Erhan, C. Szegedy, A. Toshev, D. Anguelov. Scalable object detection using deep neural networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, Columbus, USA, pp.2155–2162, 2014. DOI: [10.1109/CVPR.2014.276](https://doi.org/10.1109/CVPR.2014.276).
- [106] D. Yoo, S. Park, J. Y. Lee, A. S. Paek, I. S. Kweon. Attentionnet: Aggregating weak directions for accurate object detection. In *Proceedings of IEEE International Conference on Computer Vision*, Santiago, Chile, pp.2659–2667, 2015. DOI: [10.1109/ICCV.2015.305](https://doi.org/10.1109/ICCV.2015.305).
- [107] M. Najibi, M. Rastegari, L. S. Davis. G-CNN: AN iterative grid based object detector. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, pp.2369–2377, 2016. DOI: [10.1109/CVPR.2016.260](https://doi.org/10.1109/CVPR.2016.260).

- [108] J. Redmon, S. Divvala, R. Girshick, A. Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, pp. 779–788, 2016. DOI: [10.1109/CVPR.2016.91](https://doi.org/10.1109/CVPR.2016.91).
- [109] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C. Y. Fu, A. C. Berg. SSD: Single shot multibox detector. In *Proceedings of the 14th European Conference on Computer Vision*, Springer, Amsterdam, The Netherlands, pp. 21–37, 2016. DOI: [10.1007/978-3-319-46448-0\\_2](https://doi.org/10.1007/978-3-319-46448-0_2).
- [110] J. Redmon, A. Farhadi. YOLO9000: Better, faster, stronger. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp. 6517–6525, 2017. DOI: [10.1109/CVPR.2017.690](https://doi.org/10.1109/CVPR.2017.690).
- [111] Z. Q. Shen, Z. Liu, J. G. Li, Y. G. Jiang, Y. R. Chen, X. Y. Xue. DSOD: Learning deeply supervised object detectors from scratch. In *Proceedings of IEEE International Conference on Computer Vision*, Venice, Italy, pp. 1937–1945, 2017. DOI: [10.1109/ICCV.2017.212](https://doi.org/10.1109/ICCV.2017.212).
- [112] C. Y. Fu, W. Liu, A. Ranga, A. Tyagi, A. C. Berg. DSSD: Deconvolutional single shot detector. [Online], Available: <https://arxiv.org/abs/1701.06659>, 2017.
- [113] G. B. Li, Y. Xie, T. H. Wei, K. Z. Wang, L. Lin. Flow guided recurrent neural encoder for video salient object detection. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp. 3243–3252, 2018. DOI: [10.1109/CVPR.2018.00342](https://doi.org/10.1109/CVPR.2018.00342).
- [114] H. F. Li, G. Q. Chen, G. B. Li, Y. Z. Yu. Motion guided attention for video salient object detection. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 7273–7282, 2019. DOI: [10.1109/ICCV.2019.00737](https://doi.org/10.1109/ICCV.2019.00737).
- [115] P. X. Yan, G. B. Li, Y. Xie, Z. Li, C. Wang, T. S. Chen, L. Lin. Semi-supervised video salient object detection using pseudo-labels. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 7283–7292, 2019. DOI: [10.1109/ICCV.2019.00738](https://doi.org/10.1109/ICCV.2019.00738).
- [116] I. Armeni, Z. Y. He, A. Zamir, J. Gwak, J. Malik, M. Fischer, S. Savarese. 3D scene graph: A structure for unified semantics, 3D space, and camera. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 5663–5672, 2019. DOI: [10.1109/ICCV.2019.00576](https://doi.org/10.1109/ICCV.2019.00576).
- [117] J. Johnson, R. Krishna, M. Stark, L. J. Li, D. A. Shamma, M. S. Bernstein, F. F. Li. Image retrieval using scene graphs. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, Boston, USA, pp. 3668–3678, 2015. DOI: [10.1109/CVPR.2015.7298990](https://doi.org/10.1109/CVPR.2015.7298990).
- [118] R. Z. Wang, Z. Y. Wei, P. J. Li, Q. Zhang, X. J. Huang. Storytelling from an image stream using scene graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 5, 2020, pp. 9185–9192. DOI: [10.1609/aaai.v34i05.6455](https://doi.org/10.1609/aaai.v34i05.6455).
- [119] H. Qi, Y. L. Xu, T. Yuan, T. F. Wu, S. C. Zhu. Scene-centric joint parsing of cross-view videos. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence and 30th Innovative Applications of Artificial Intelligence Conference and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence*, New Orleans, USA, Article number 893, 2018.
- [120] B. Dai, Y. Q. Zhang, D. H. Lin. Detecting visual relationships with deep relational networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp. 3298–3308, 2017. DOI: [10.1109/CVPR.2017.352](https://doi.org/10.1109/CVPR.2017.352).
- [121] H. W. Zhang, Z. Kyaw, S. F. Chang, T. S. Chua. Visual translation embedding network for visual relation detection. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp. 3107–3115, 2017. DOI: [10.1109/CVPR.2017.331](https://doi.org/10.1109/CVPR.2017.331).
- [122] Z. S. Hung, A. Mallya, S. Lazebnik. Union visual translation embedding for visual relationship detection and scene graph generation. [Online], Available: <https://arxiv.org/abs/1905.11624v1>, 2019.
- [123] Y. N. Chen, Y. J. Wang, Y. Zhang, Y. W. Guo. PANet: A context based predicate association network for scene graph generation. In *Proceedings of IEEE International Conference on Multimedia and Expo*, Shanghai, China, pp. 508–513, 2019. DOI: [10.1109/ICME.2019.00094](https://doi.org/10.1109/ICME.2019.00094).
- [124] K. H. Tang, H. W. Zhang, B. Y. Wu, W. H. Luo, W. Liu. Learning to compose dynamic tree structures for visual contexts. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 6612–6621, 2019. DOI: [10.1109/CVPR.2019.00678](https://doi.org/10.1109/CVPR.2019.00678).
- [125] Y. K. Li, W. L. Ouyang, X. G. Wang, X. O. Tang. ViP-CNN: Visual phrase guided convolutional neural network. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp. 7244–7253, 2017. DOI: [10.1109/CVPR.2017.766](https://doi.org/10.1109/CVPR.2017.766).
- [126] Y. Z. Liang, Y. L. Bai, W. Zhang, X. M. Qian, L. Zhu, T. Mei. VrR-VG: Refocusing visually-relevant relationships. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 10402–10411, 2019. DOI: [10.1109/ICCV.2019.01050](https://doi.org/10.1109/ICCV.2019.01050).
- [127] Y. K. Li, W. L. Ouyang, B. L. Zhou, J. P. Shi, C. Zhang, X. G. Wang. Factorizable Net: An efficient subgraph-based framework for scene graph generation. In *Proceedings of the 15th European Conference on Computer Vision*, Springer, Munich, Germany, pp. 346–363, 2018. DOI: [10.1007/978-3-030-01246-5\\_21](https://doi.org/10.1007/978-3-030-01246-5_21).
- [128] M. S. Qi, W. J. Li, Z. Y. Yang, Y. H. Wang, J. B. Luo. Attentive relational networks for mapping images to scene graphs. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 3952–3961, 2019. DOI: [10.1109/CVPR.2019.00408](https://doi.org/10.1109/CVPR.2019.00408).
- [129] C. W. Lu, R. Krishna, M. Bernstein, F. F. Li. Visual relationship detection with language priors. In *Proceedings of the 14th European Conference on Computer Vision*, Springer, Amsterdam, The Netherlands, pp. 852–869, 2016. DOI: [10.1007/978-3-319-46448-0\\_51](https://doi.org/10.1007/978-3-319-46448-0_51).
- [130] T. S. Chen, W. H. Yu, R. Q. Chen, L. Lin. Knowledge-embedded routing network for scene graph generation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 6156–6164, 2019. DOI: [10.1109/CVPR.2019.00632](https://doi.org/10.1109/CVPR.2019.00632).
- [131] J. X. Gu, H. D. Zhao, Z. Lin, S. Li, J. F. Cai, M. Y. Ling. Scene graph generation with external knowledge and image reconstruction. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 1969–1978, 2019. DOI: [10.1109/CVPR.2019.00207](https://doi.org/10.1109/CVPR.2019.00207).
- [132] R. Zellers, M. Yatskar, S. Thomson, Y. Choi. Neural motifs: Scene graph parsing with global context. In *Proceed-*

- ings of *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp. 5831–5840, 2018. DOI: [10.1109/CVPR.2018.00611](https://doi.org/10.1109/CVPR.2018.00611).
- [133] B. A. Plummer, A. Mallya, C. M. Cervantes, J. Hockenmaier, S. Lazebnik. Phrase localization and visual relationship detection with comprehensive image-language cues. In *Proceedings of IEEE International Conference on Computer Vision*, Venice, Italy, pp. 1946–1955, 2017. DOI: [10.1109/ICCV.2017.213](https://doi.org/10.1109/ICCV.2017.213).
- [134] S. B. Yang, G. B. Li, Y. Z. Yu. Cross-modal relationship inference for grounding referring expressions. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 4140–4149, 2019. DOI: [10.1109/CVPR.2019.00427](https://doi.org/10.1109/CVPR.2019.00427).
- [135] X. R. Lin, G. B. Li, Y. Z. Yu. Scene-intuitive agent for remote embodied visual grounding. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern*, IEEE, Nashville, USA, pp. 7032–7041, 2021. DOI: [10.1109/CVPR46437.2021.00696](https://doi.org/10.1109/CVPR46437.2021.00696).
- [136] H. L. Liu, A. R. Lin, X. G. Han, L. Yang, Y. Z. Yu, S. G. Cui. Refer-it-in-RGBD: A bottom-up approach for 3D visual grounding in RGBD images. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 6028–6037, 2021. DOI: [10.1109/CVPR46437.2021.00597](https://doi.org/10.1109/CVPR46437.2021.00597).
- [137] M. J. Sun, J. M. Xiao, E. G. Lim. Iterative shrinking for referring expression grounding using deep reinforcement learning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 14055–14064, 2021. DOI: [10.1109/CVPR46437.2021.01384](https://doi.org/10.1109/CVPR46437.2021.01384).
- [138] A. Kamath, M. Singh, Y. LeCun, G. Synnaeve, I. Misra, N. Carion. MDETR-modulated detection for end-to-end multi-modal understanding. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp. 1760–1770, 2021. DOI: [10.1109/ICCV48922.2021.00180](https://doi.org/10.1109/ICCV48922.2021.00180).
- [139] J. J. Deng, Z. Y. Yang, T. L. Chen, W. G. Zhou, H. Q. Li. TransVG: End-to-end visual grounding with transformers. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp. 1749–1759, 2021. DOI: [10.1109/ICCV48922.2021.00179](https://doi.org/10.1109/ICCV48922.2021.00179).
- [140] J. Wu, G. B. Li, S. Liu, L. Lin. Tree-structured policy based progressive reinforcement learning for temporally language grounding in video. *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 7, pp. 12386–12393, 2020. DOI: [10.1609/aaai.v34i07.6924](https://doi.org/10.1609/aaai.v34i07.6924).
- [141] L. Chen, W. B. Ma, J. Xiao, H. W. Zhang, S. F. Chang. REF-NMS: Breaking proposal bottlenecks in two-stage referring expression grounding. *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 2, pp. 1036–1044, 2021.
- [142] J. Wu, G. B. Li, X. G. Han, L. Lin. Reinforcement learning for weakly supervised temporal grounding of natural language in untrimmed videos. In *Proceedings of the 28th ACM International Conference on Multimedia*, Seattle, USA, pp. 1283–1291, 2020. DOI: [10.1145/3394171.3413862](https://doi.org/10.1145/3394171.3413862).
- [143] R. A. Yeh, M. N. Do, A. G. Schwing. Unsupervised textual grounding: Linking words to image concepts. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp. 6125–6134, 2018. DOI: [10.1109/CVPR.2018.00641](https://doi.org/10.1109/CVPR.2018.00641).
- [144] C. L. Zitnick, P. Dollár. Edge boxes: Locating object proposals from edges. In *Proceedings of the 13th European Conference on Computer Vision*, Springer, Zurich, Switzerland, pp. 391–405, 2014. DOI: [10.1007/978-3-319-10602-1\\_26](https://doi.org/10.1007/978-3-319-10602-1_26).
- [145] J. R. R. Uijlings, K. E. A. van de Sande, T. Gevers, A. W. M. Smeulders. Selective search for object recognition. *International Journal of Computer Vision*, vol. 104, no. 2, pp. 154–171, 2013. DOI: [10.1007/s11263-013-0620-5](https://doi.org/10.1007/s11263-013-0620-5).
- [146] Y. F. Liu, B. Wan, L. Ma, X. M. He. Relation-aware instance refinement for weakly supervised visual grounding. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 5608–5617, 2021. DOI: [10.1109/CVPR46437.2021.00556](https://doi.org/10.1109/CVPR46437.2021.00556).
- [147] L. W. Wang, J. Huang, Y. Li, K. Xu, Z. Y. Yang, D. Yu. Improving weakly supervised visual grounding by contrastive knowledge distillation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 14085–14095, 2021. DOI: [10.1109/CVPR46437.2021.01387](https://doi.org/10.1109/CVPR46437.2021.01387).
- [148] J. Wang, L. Specia. Phrase localization without paired training examples. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 4662–4671, 2019. DOI: [10.1109/ICCV.2019.00476](https://doi.org/10.1109/ICCV.2019.00476).
- [149] S. B. Yang, G. B. Li, Y. Z. Yu. Dynamic graph attention for referring expression comprehension. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp. 4643–4652, 2019. DOI: [10.1109/ICCV.2019.00474](https://doi.org/10.1109/ICCV.2019.00474).
- [150] S. B. Yang, G. B. Li, Y. Z. Yu. Graph-structured referring expression reasoning in the wild. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp. 9949–9958, 2020. DOI: [10.1109/CVPR42600.2020.00997](https://doi.org/10.1109/CVPR42600.2020.00997).
- [151] R. Zellers, Y. Bisk, A. Farhadi, Y. Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp. 6713–6724, 2019. DOI: [10.1109/CVPR.2019.00688](https://doi.org/10.1109/CVPR.2019.00688).
- [152] A. M. Wu, L. C. Zhu, Y. H. Han, Y. Yang. Connective cognition network for directional visual commonsense reasoning. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, Vancouver, Canada, Article number 509, 2019. DOI: [10.5555/3454287.3454796](https://doi.org/10.5555/3454287.3454796).
- [153] W. J. Yu, J. W. Zhou, W. H. Yu, X. D. Liang, N. Xiao. Heterogeneous graph learning for visual commonsense reasoning. In *Proceedings of the 33rd Conference on Neural Information Processing Systems*, Vancouver, Canada, pp. 2765–2775, 2019.
- [154] J. X. Lin, U. Jain, A. G. Schwing. TAB-VCR: Tags and attributes based visual commonsense reasoning baselines. In *Proceedings of the 32nd Conference on Neural Information Processing Systems*, Vancouver, Canada, 2019.
- [155] X. Zhang, F. F. Zhang, C. S. Xu. Multi-level counterfactual contrast for visual commonsense reasoning. In *Proceedings of the 29th ACM International Conference on Multimedia*, pp. 1793–1802, 2021. DOI: [10.1145/3474085.3475328](https://doi.org/10.1145/3474085.3475328).
- [156] J. S. Lu, D. Batra, D. Parikh, S. Lee. ViLBERT: Pre-training task-agnostic visiolinguistic representations for vision-and-language tasks. In *Proceedings of the 33rd In-*

- ternational Conference on Neural Information Processing Systems, Vancouver, Canada, Article number 2, 2019.
- [157] Y. C. Chen, L. J. Li, L. C. Yu, A. El Kholy, F. Ahmed, Z. Gan, Y. Cheng, J. J. Liu. UNITER: Universal image-text representation learning. In *Proceedings of the 16th European Conference on Computer Vision*, Springer, Glasgow, UK, pp.104–120, 2020. DOI: [10.1007/978-3-030-58577-8\\_7](https://doi.org/10.1007/978-3-030-58577-8_7).
- [158] W. J. Su, X. Z. Zhu, Y. Cao, B. Li, L. W. Lu, F. R. Wei, J. F. Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *the 8th International Conference on Learning Representations*, Addis Ababa, Ethiopia, 2020.
- [159] P. Sharma, N. Ding, S. Goodman, R. Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, Melbourne, Australia, pp.2556–2565, 2018. DOI: [10.18653/v1/P18-1238](https://doi.org/10.18653/v1/P18-1238).
- [160] J. Carreira, A. Zisserman. Quo vadis, action recognition? A new model and the kinetics dataset. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Honolulu, USA, pp.4724–4733, 2017. DOI: [10.1109/CVPR.2017.502](https://doi.org/10.1109/CVPR.2017.502).
- [161] C. Sun, A. Shrivastava, C. Vondrick, K. Murphy, R. Sukthankar, C. Schmid. Actor-centric relation network. In *Proceedings of the 15th European Conference on Computer Vision*, Springer, Munich, Germany, pp.335–351, 2018. DOI: [10.1007/978-3-030-01252-6\\_20](https://doi.org/10.1007/978-3-030-01252-6_20).
- [162] C. Y. Wu, C. Feichtenhofer, H. Q. Fan, K. M. He, P. Krähenbühl, R. Girshick. Long-term feature banks for detailed video understanding. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp.284–293, 2019. DOI: [10.1109/CVPR.2019.00037](https://doi.org/10.1109/CVPR.2019.00037).
- [163] C. Y. Yang, Y. H. Xu, J. P. Shi, B. Dai, B. L. Zhou. Temporal pyramid network for action recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.588–597, 2020. DOI: [10.1109/CVPR42600.2020.00067](https://doi.org/10.1109/CVPR42600.2020.00067).
- [164] C. Feichtenhofer. X3D: Expanding architectures for efficient video recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.200–210, 2020.
- [165] C. Feichtenhofer, H. Q. Fan, J. Malik, K. M. He. Slow-fast networks for video recognition. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Seoul, Korea, pp.6201–6210, 2019. DOI: [10.1109/ICCV.2019.00630](https://doi.org/10.1109/ICCV.2019.00630).
- [166] W. T. Bao, Q. Yu, Y. Kong. Evidential deep learning for open set action recognition. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.13329–13338, 2021. DOI: [10.1109/ICCV48922.2021.01310](https://doi.org/10.1109/ICCV48922.2021.01310).
- [167] A. Aich, M. Zheng, S. Karanam, T. Chen, A. K. Roy-Chowdhury, Z. Y. Wu. Spatio-temporal representation factorization for video-based person re-identification. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.152–162, 2021. DOI: [10.1109/ICCV48922.2021.00022](https://doi.org/10.1109/ICCV48922.2021.00022).
- [168] J. Tan, J. Q. Tang, L. M. Wang, G. S. Wu. Relaxed transformer decoders for direct action proposal generation. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.13506–13515, 2021. DOI: [10.1109/ICCV48922.2021.01327](https://doi.org/10.1109/ICCV48922.2021.01327).
- [169] G. Bertasius, H. Wang, L. Torresani. Is space-time attention all you need for video understanding? In *Proceedings of the 38th International Conference on Machine Learning*, pp.813–824, 2021.
- [170] X. Wang, S. W. Zhang, Z. W. Qing, Y. J. Shao, Z. R. Zuo, C. X. Gao, N. Sang. OadTR: Online action detection with transformers. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp.7545–7555, 2021. DOI: [10.1109/ICCV48922.2021.00747](https://doi.org/10.1109/ICCV48922.2021.00747).
- [171] C. H. Zhang, A. Gupta, A. Zisserman. Temporal query networks for fine-grained video understanding. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, 2021, pp.4484–4494. DOI: [10.1109/CVPR46437.2021.00446](https://doi.org/10.1109/CVPR46437.2021.00446).
- [172] S. J. Yan, Y. J. Xiong, D. H. Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence*, New Orleans, USA, Article number.912, 2018. DOI: [10.5555/3504035.3504947](https://doi.org/10.5555/3504035.3504947).
- [173] C. Y. Si, W. T. Chen, W. Wang, L. Wang, T. N. Tan. An attention enhanced graph convolutional LSTM network for skeleton-based action recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp.1227–1236, 2019. DOI: [10.1109/CVPR.2019.00132](https://doi.org/10.1109/CVPR.2019.00132).
- [174] L. Shi, Y. F. Zhang, J. Cheng, H. Q. Lu. Two-stream adaptive graph convolutional networks for skeleton-based action recognition. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Long Beach, USA, pp.12018–12027, 2019. DOI: [10.1109/CVPR.2019.01230](https://doi.org/10.1109/CVPR.2019.01230).
- [175] K. Lin, L. J. Wang, Z. C. Liu. End-to-end human pose and mesh reconstruction with transformers. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp.1954–1963, 2021. DOI: [10.1109/CVPR46437.2021.00199](https://doi.org/10.1109/CVPR46437.2021.00199).
- [176] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. H. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, N. Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *Proceedings of the 9th International Conference on Learning Representations*, 2021.
- [177] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučić, C. Schmid. ViViT: A video vision transformer. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, USA, pp.6816–6826, 2021. DOI: [10.1109/ICCV48922.2021.00676](https://doi.org/10.1109/ICCV48922.2021.00676).
- [178] P. Anderson, X. D. He, C. Buehler, D. Teney, M. Johnson, S. Gould, L. Zhang. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp.6077–6086, 2018. DOI: [10.1109/CVPR.2018.00636](https://doi.org/10.1109/CVPR.2018.00636).
- [179] S. Hochreiter, J. Schmidhuber. Long short-term memory. *Neural Computation*, vol.9, no.8, pp.1735–1780, 1997. DOI: [10.1162/neco.1997.9.8.1735](https://doi.org/10.1162/neco.1997.9.8.1735).
- [180] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Long Beach, USA, pp.6000–6010, 2017. DOI: [10.5555/](https://doi.org/10.5555/)

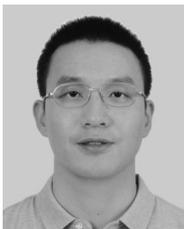
3295222.3295349.

- [181] J. Devlin, M. W. Chang, K. Lee, K. Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Minneapolis, USA, pp.4171–4186, 2019. DOI: [10.18653/v1/N19-1423](https://doi.org/10.18653/v1/N19-1423).
- [182] S. Antol, A. Agrawal, J. S. Lu, M. Mitchell, D. Batra, C. L. Zitnick, D. Parikh. VQA: Visual question answering. In *Proceedings of IEEE International Conference on Computer Vision*, Santiago, Chile, pp.2425–2433, 2015. DOI: [10.1109/ICCV.2015.279](https://doi.org/10.1109/ICCV.2015.279).
- [183] Z. C. Yang, X. D. He, J. F. Gao, L. Deng, A. Smola. Stacked attention networks for image question answering. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, pp.21–29, 2016. DOI: [10.1109/CVPR.2016.10](https://doi.org/10.1109/CVPR.2016.10).
- [184] D. J. Xu, Z. Zhao, J. Xiao, F. Wu, H. W. Zhang, X. N. He, Y. T. Zhuang. Video question answering via gradually refined attention over appearance and motion. In *Proceedings of the 25th ACM International Conference on Multimedia*, Mountain View, USA, pp.1645–1653, 2017. DOI: [10.1145/3123266.3123427](https://doi.org/10.1145/3123266.3123427).
- [185] T. M. Le, V. Le, S. Venkatesh, T. Tran. Hierarchical conditional relation networks for video question answering. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.9969–9978, 2020. DOI: [10.1109/CVPR42600.2020.00999](https://doi.org/10.1109/CVPR42600.2020.00999).
- [186] P. Jiang, Y. H. Han. Reasoning with heterogeneous graph alignment for video question answering. *Proceedings of the AAAI Conference on Artificial Intelligence*, vol.34, no.7, pp.11109–11116, 2020. DOI: [10.1609/aaai.v34i07.6767](https://doi.org/10.1609/aaai.v34i07.6767).
- [187] D. Huang, P. H. Chen, R. H. Zeng, Q. Du, M. K. Tan, C. Gan. Location-aware graph convolutional networks for video question answering. *Proceedings of the AAAI Conference on Artificial Intelligence*, vol.34, no.7, pp.11021–11028, 2020. DOI: [10.1609/aaai.v34i07.6737](https://doi.org/10.1609/aaai.v34i07.6737).
- [188] J. Lei, L. J. Li, L. W. Zhou, Z. Gan, T. L. Berg, M. Bansal, J. J. Liu. Less is more: CLIPBERT for video-and-language learning via sparse sampling. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp.7327–7337, 2021. DOI: [10.1109/CVPR46437.2021.00725](https://doi.org/10.1109/CVPR46437.2021.00725).
- [189] F. Liu, J. Liu, W. N. Wang, H. Q. Lu. HAIR: Hierarchical visual-semantic relational reasoning for video question answering. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, USA, pp.1678–1787, 2021. DOI: [10.1109/ICCV48922.2021.00172](https://doi.org/10.1109/ICCV48922.2021.00172).
- [190] A. Agrawal, D. Batra, D. Parikh, A. Kembhavi. Don't just assume; look and answer: Overcoming priors for visual question answering. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, USA, pp.4971–4980, 2018. DOI: [10.1109/CVPR.2018.00522](https://doi.org/10.1109/CVPR.2018.00522).
- [191] V. Agarwal, R. Shetty, M. Fritz. Towards causal VQA: Revealing and reducing spurious correlations by invariant and covariant semantic editing. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Seattle, USA, pp.9687–9695, 2020. DOI: [10.1109/CVPR42600.2020.00971](https://doi.org/10.1109/CVPR42600.2020.00971).
- [192] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio. Generative adversarial nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems*, MIT Press, Montreal, Canada, pp.2672–2680, 2014. DOI: [10.5555/2969033.2969125](https://doi.org/10.5555/2969033.2969125).
- [193] S. Y. Zhang, T. Jiang, T. Wang, K. Kuang, Z. Zhao, J. K. Zhu, J. Yu, H. X. Yang, F. Wu. DeVLBERT: Learning deconfounded visio-linguistic representations. In *Proceedings of the 28th ACM International Conference on Multimedia*, Seattle, USA, pp.4373–4382, 2020. DOI: [10.1145/3394171.3413518](https://doi.org/10.1145/3394171.3413518).
- [194] Y. C. Li, X. Wang, J. B. Xiao, W. Ji, T. S. Chua. Invariant grounding for video question answering. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.2928–2937, 2022.
- [195] R. Y. Liu, H. Liu, G. Li, H. D. Hou, T. H. Yu, T. Yang. Contextual debiasing for visual recognition with causal mechanisms. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.12755–12765, 2022.
- [196] Y. J. Liu, R. Cadei, J. Schweizer, S. Bahmani, A. Alahi. Towards robust and adaptive motion forecasting: A causal representation perspective. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.17081–17092, 2022.
- [197] F. R. Lv, J. Liang, S. Li, B. Zang, C. H. Liu, Z. T. Wang, D. Liu. Causality inspired representation learning for domain generalization. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.8046–8056, 2022.
- [198] X. R. Lin, Y. Y. Chen, G. B. Li, Y. Z. Yu. A causal inference look at unsupervised video anomaly detection. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pp.1620–1629, 2022. DOI: [10.1609/aaai.v36i2.20053](https://doi.org/10.1609/aaai.v36i2.20053).
- [199] X. R. Lin, Z. Y. Wu, G. Q. Chen, G. B. Li, Y. Z. Yu. A causal debiasing framework for unsupervised salient object detection. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence*, pp.1610–1619, 2022. DOI: [10.1609/aaai.v36i2.20052](https://doi.org/10.1609/aaai.v36i2.20052).
- [200] Y. Liu, G. B. Li, L. Lin. Cross-modal causal relational reasoning for event-level visual question answering. [Online], Available: <https://arxiv.org/abs/2207.12647>, 2022.
- [201] M. Ren, Y. L. Wang, Z. F. He. Towards interpretable defense against adversarial attacks via causal inference. *Machine Intelligence Research*, vol.19, no.3, pp.209–226, 2022. DOI: [10.1007/s11633-022-1330-7](https://doi.org/10.1007/s11633-022-1330-7).
- [202] R. J. Bowden, D. A. Turkington. *Instrumental Variables*, Cambridge, UK: Cambridge University Press, 1984.
- [203] J. Y. Zhu, T. Park, P. Isola, A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of IEEE International Conference on Computer Vision*, Venice, Italy, pp.2242–2251, 2017. DOI: [10.1109/ICCV.2017.244](https://doi.org/10.1109/ICCV.2017.244).
- [204] D. P. Kingma, M. Welling. Auto-encoding variational bayes. [Online], Available: <https://arxiv.org/abs/1312.6114>, 2013.
- [205] P. Zhang, Y. Goyal, D. Summers-Stay, D. Batra, D. Parikh. Yin and Yang: Balancing and answering binary visual questions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, pp.5014–5022, 2016. DOI: [10.1109/CVPR.2016.542](https://doi.org/10.1109/CVPR.2016.542).

- [206] H. Tan, M. Bansal. LXMERT: Learning cross-modality encoder representations from transformers. In *Proceedings of Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, Hong Kong, China, pp. 5100–5111, 2019. DOI: [10.18653/v1/D19-1514](https://doi.org/10.18653/v1/D19-1514).
- [207] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 8748–8763, 2021.
- [208] L. J. Li, J. Lei, Z. Gan, J. J. Liu. Adversarial VQA: A new benchmark for evaluating the robustness of VQA models. In *Proceedings of IEEE/CVF International Conference on Computer Vision*, IEEE, Montreal, Canada, pp. 2022–2031, 2021. DOI: [10.1109/ICCV48922.2021.00205](https://doi.org/10.1109/ICCV48922.2021.00205).
- [209] X. L. Chen, H. Fang, T. Y. Lin, R. Vedantam, S. Gupta, P. Dollár, C. L. Zitnick. Microsoft COCO captions: Data collection and evaluation server. [Online], Available: <https://arxiv.org/abs/1504.00325>, 2015.
- [210] K. X. Yi, C. Gan, Y. Z. Li, P. Kohli, J. J. Wu, A. Torralba, J. B. Tenenbaum. Cleverer: Collision events for video representation and reasoning. In *Proceedings of the 8th International Conference on Learning Representations*, Addis Ababa, Ethiopia, 2020.
- [211] V. Gupta, B. N. Patro, H. Parihar, V. P. Namboodiri. VQuAD: Video question answering diagnostic dataset. In *Proceedings of IEEE/CVF Winter Conference on Applications of Computer Vision*, IEEE, Waikoloa, USA, pp. 282–291, 2022. DOI: [10.1109/WACVW54805.2022.00034](https://doi.org/10.1109/WACVW54805.2022.00034).
- [212] Z. F. Chen, K. X. Yi, Y. Z. Li, M. Y. Ding, A. Torralba, J. B. Tenenbaum, C. Gan. ComPhy: Compositional physical reasoning of objects and events from videos. In *Proceedings of the 10th International Conference on Learning Representations*, 2022.
- [213] M. Grunde-McLaughlin, R. Krishna, M. Agrawala. AGQA: A benchmark for compositional spatio-temporal reasoning. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 11282–11292, 2021. DOI: [10.1109/CVPR46437.2021.01113](https://doi.org/10.1109/CVPR46437.2021.01113).
- [214] L. Xu, H. Huang, J. Liu. SUTD-trafficQA: A question answering benchmark and an efficient network for video reasoning over traffic events. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, 2021, pp. 9873–9883. DOI: [10.1109/CVPR46437.2021.00975](https://doi.org/10.1109/CVPR46437.2021.00975).
- [215] J. B. Xiao, X. D. Shang, A. Yao, T. S. Chua. NEXT-QA: Next phase of question-answering to explaining temporal actions. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Nashville, USA, pp. 9772–9781, 2021. DOI: [10.1109/CVPR46437.2021.00965](https://doi.org/10.1109/CVPR46437.2021.00965).
- [216] D. W. Zhang, W. Y. Zeng, J. R. Yao, J. W. Han. Weakly supervised object detection using proposal-and semantic-level relationships. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 6, pp. 3349–3363, 2022. DOI: [10.1109/TPAMI.2020.3046647](https://doi.org/10.1109/TPAMI.2020.3046647).
- [217] D. W. Zhang, J. W. Han, G. Cheng, M. H. Yang. Weakly supervised object localization and detection: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 9, pp. 5866–5885, 2022. DOI: [10.1109/TPAMI.2021.3074313](https://doi.org/10.1109/TPAMI.2021.3074313).
- [218] W. Wang, J. Y. Gao, C. S. Xu. Weakly-supervised video object grounding via causal intervention. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, to be published. DOI: [10.1109/TPAMI.2022.3180025](https://doi.org/10.1109/TPAMI.2022.3180025).
- [219] E. Tjoa, C. T. Guan. A survey on explainable artificial intelligence (XAI): Toward medical XAI. *IEEE Transactions on Neural Networks and Learning Systems*, vol. 32, no. 11, pp. 4793–4813, 2021. DOI: [10.1109/TNNLS.2020.3027314](https://doi.org/10.1109/TNNLS.2020.3027314).
- [220] Á. Parafita, J. Vitriá. Explaining visual models by causal attribution. In *Proceedings of IEEE/CVF International Conference on Computer Vision Workshop*, IEEE, Seoul, Korea, pp. 4167–4175, 2019. DOI: [10.1109/ICCVW.2019.00512](https://doi.org/10.1109/ICCVW.2019.00512).
- [221] T. Narendra, A. Sankaran, D. Vijaykeerthy, S. Mani. Explaining deep learning models using causal inference. [Online], Available: <https://arxiv.org/abs/1811.04376>, 2018.
- [222] M. Harradon, J. Druce, B. Ruttenberg. Causal learning and explanation of deep neural networks via autoencoded activations. [Online], Available: <https://arxiv.org/abs/1802.00541>, 2018.
- [223] A. Chattopadhyay, P. Manupriya, A. Sarkar, V. N. Balasubramanian. Neural network attributions: A causal perspective. In *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, USA, 2019, pp. 981–990.
- [224] R. Moraffah, M. Karami, R. C. Guo, A. Raglin, H. Liu. Causal interpretability for machine learning-problems, methods and evaluation. *ACM SIGKDD Explorations Newsletter*, vol. 2, no. 1, pp. 18–33, 2020. DOI: [10.1145/3400051.3400058](https://doi.org/10.1145/3400051.3400058).
- [225] M. O’Shaughnessy, G. Canal, M. Connor, M. Davenport, C. Rozell. Generative causal explanations of black-box classifiers. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, Vancouver, Canada, pp. 5453–5467, 2020. DOI: [10.5555/3495724.3496182](https://doi.org/10.5555/3495724.3496182).
- [226] W. Y. Lin, H. Lan, B. C. Li. Generative causal explanations for graph neural networks. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 6666–6679, 2021.
- [227] J. von Kügelgen, L. Gresele, B. Schölkopf. Simpson’s paradox in COVID-19 case fatality rates: A mediation analysis of age-related causal effects. *IEEE Transactions on Artificial Intelligence*, vol. 2, no. 1, pp. 18–27, 2021. DOI: [10.1109/TAI.2021.3073088](https://doi.org/10.1109/TAI.2021.3073088).
- [228] Y. Zheng, C. Gao, X. Li, X. N. He, Y. Li, D. P. Jin. Disentangling user interest and conformity for recommendation with causal embedding. In *Proceedings of the Web Conference*, ACM, Ljubljana, Slovenia, pp. 2980–2991, 2021. DOI: [10.1145/3442381.3449788](https://doi.org/10.1145/3442381.3449788).
- [229] D. G. Liu, P. X. Cheng, H. Zhu, Z. H. Dong, X. Q. He, W. K. Pan, Z. Ming. Mitigating confounding bias in recommendation via information bottleneck. In *Proceedings of the 15th ACM Conference on Recommender Systems*, Amsterdam, The Netherlands, pp. 351–360, 2021. DOI: [10.1145/3460231.3474263](https://doi.org/10.1145/3460231.3474263).
- [230] T. X. Wei, F. L. Feng, J. W. Chen, Z. W. Wu, J. F. Yi, X. N. He. Model-agnostic counterfactual reasoning for eliminating popularity bias in recommender system. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, Singapore, pp. 1791–1800, 2021. DOI: [10.1145/3447548.3467289](https://doi.org/10.1145/3447548.3467289).
- [231] W. J. Wang, F. L. Feng, X. N. He, H. W. Zhang, T. S.

Chua. Clicks can be cheating: Counterfactual recommendation for mitigating clickbait issue. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1288–1297, 2021. DOI: [10.1145/3404835.3462962](https://doi.org/10.1145/3404835.3462962).

- [232] Y. Zhang, F. Feng, X. He, T. Wei, C. Song, G. Ling, Y. Zhang. Causal intervention for leveraging popularity bias in recommendation. In *Proceedings the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 11–20, 2021.
- [233] K. C. Stocking, A. Gopnik, and C. Tomlin. From robot learning to robot understanding: Leveraging causal graphical models for robotics. In *Proceedings of Conference on Robot Learning*. pp. 1776–1781, 2022.
- [234] T. E. Lee, J. A. Zhao, A. S. Sawhney, S. Girdhar, O. Kroemer. Causal reasoning in simulation for structure and transfer learning of robot manipulation policies. In *Proceedings of IEEE International Conference on Robotics and Automation*. IEEE, pp. 4776–4782, 2021.
- [235] S. C. Smith and S. Ramamoorthy. Counterfactual explanation and causal inference in service of robustness in robot control. In *Proceedings of the 10th IEEE Joint International Conference on Development and Learning and Epigenetic Robotics*. IEEE, 2020.
- [236] F. Hou, Y. Pei, and J. Sun. *Mobile Crowd Sensing: Incentive Mechanism Design*. Springer, 2019.
- [237] Y. Zheng, L. Capra, O. Wolfson, H. Yang. *Urban computing: concepts, methodologies, and applications*. ACM Transactions on Intelligent Systems and Technology, vol. 5, no. 3, pp. 1–55, 2014.



**Yang Liu** received the B.Sc. degree in telecommunications engineering from Chang'an University, China in 2014, and the Ph.D. degree in telecommunications and information systems from Xidian University, China in 2019. He is currently a research associate professor working at School of Computer Science and Engineering, Sun Yat-sen University, China. He has

authorized and co-authored more than 20 papers in top-tier academic journals and conferences. He has been serving as a reviewer for numerous academic journals and conferences such as IEEE TIP, TNNLS, TMM, TCSVT, TCyb, CVPR, ICCV, AAAI, and ECCV. He is a member of IEEE and CSIG.

His research interests include video understanding, causal reasoning, and computer vision.

E-mail: [liuy856@mail.sysu.edu.cn](mailto:liuy856@mail.sysu.edu.cn)

ORCID iD: 0000-0002-9423-9252



**Yu-Shen Wei** received the B.Sc. degree in computer science and technology from Sun Yat-sen University, China in 2020. He is currently a master student at School of Computer Science and Engineering, Sun Yat-sen University, China.

His current research interests include video understanding, computer vision and machine learning.

E-mail: [weiysh8@mail2.sysu.edu.cn](mailto:weiysh8@mail2.sysu.edu.cn)

ORCID iD: 0000-0002-0527-5463



**Hong Yan** received the B.Sc. degree in computer science and technology from Nanchang University, China in 2020. He is currently a master student at School of Computer Science and Engineering, Sun Yat-sen University, China.

His research interests include video understanding, computer vision and machine learning.

E-mail: [yanh36@mail2.sysu.edu.cn](mailto:yanh36@mail2.sysu.edu.cn)

ORCID iD: 0000-0003-4100-6751



**Guan-Bin Li** received the Ph.D. degree from the University of Hong Kong, China in 2016. He is currently an associate professor in School of Computer Science and Engineering, Sun Yat-Sen University, China. He is a recipient of ICCV 2019 Best Paper Nomination Award. He has authorized and co-authored on more than 70 papers in top-tier academic journals and

conferences. He serves as an area chair for the conference of VIS-APP. He has been serving as a reviewer for numerous academic journals and conferences such as TPAMI, IJCV, TIP, TMM, TCyb, CVPR, ICCV, ECCV and NeurIPS.

His research interests include computer vision, image processing, and machine learning.

E-mail: [ligranbin@mail.sysu.edu.cn](mailto:ligranbin@mail.sysu.edu.cn)

ORCID iD: 0000-0002-4805-0926



**Liang Lin** received the Ph. D. degree from Beijing Institute of Technology, China in 2008. He is a full professor of computer science at Sun Yat-sen University, China. He served as the executive director and distinguished scientist of SenseTime Group from 2016 to 2018, leading the R&D teams for cutting-edge technology transferring. He has authored or co-authored more than

200 papers in leading academic journals and conferences, and his papers have been cited by more than 21 000 times. He is an associate editor of *IEEE Transactions on Neural Networks and Learning Systems* and *IEEE Transactions on Human-Machine Systems*, and served as area Chairs for numerous conferences such as CVPR, ICCV, SIGKDD and AAAI. He is the recipient of numerous awards and honors including Wu Wen-Jun Artificial Intelligence Award, the First Prize of China Society of Image and Graphics, ICCV Best Paper Nomination in 2019, Annual Best Paper Award by Pattern Recognition (Elsevier) in 2018, Best Paper Dimond Award in IEEE ICME 2017, Google Faculty Award in 2012. His supervised Ph.D. students received ACM China Doctoral Dissertation Award, CCF Best Doctoral Dissertation and CAAI Best Doctoral Dissertation. He is a fellow of IET/IAPR.

His research interests include artificial intelligence, computer vision, machine learning, multimedia, and NLP/Dialogue.

E-mail: [linliang@ieee.org](mailto:linliang@ieee.org) (Corresponding author)

ORCID iD: 0000-0003-2248-3755