

Robust Emotion Recognition in Context Debiasing

Dingkang Yang^{1,2,8} Kun Yang¹ Mingcheng Li^{1,2} Shunli Wang^{1,2}
Shuaibing Wang^{1,2} Lihua Zhang^{1,2,3,4,§}

¹Academy for Engineering and Technology, Fudan University ²Cognition and Intelligent Technology Laboratory (CIT Lab)

³Jilin Provincial Key Laboratory of Intelligence Science and Engineering, Changchun, China

⁴Engineering Research Center of AI and Robotics, Ministry of Education, Shanghai, China

dkyang20@fudan.edu.cn

Abstract

Context-aware emotion recognition (CAER) has recently boosted the practical applications of affective computing techniques in unconstrained environments. Mainstream CAER methods invariably extract ensemble representations from diverse contexts and subject-centred characteristics to perceive the target person’s emotional state. Despite advancements, the biggest challenge remains due to context bias interference. The harmful bias forces the models to rely on spurious correlations between background contexts and emotion labels in likelihood estimation, causing severe performance bottlenecks and confounding valuable context priors. In this paper, we propose a counterfactual emotion inference (CLEF) framework to address the above issue. Specifically, we first formulate a generalized causal graph to decouple the causal relationships among the variables in CAER. Following the causal graph, CLEF introduces a non-invasive context branch to capture the adverse direct effect caused by the context bias. During the inference, we eliminate the direct context effect from the total causal effect by comparing factual and counterfactual outcomes, resulting in bias mitigation and robust prediction. As a model-agnostic framework, CLEF can be readily integrated into existing methods, bringing consistent performance gains.

1. Introduction

“Context is the key to understanding, but it can also be the key to misunderstanding.”

—Jonathan Lockwood Huie

As the spiritual grammar of human life, emotions play an essential role in social communication and intelligent automation [21]. Accurately recognizing subjects’ emotional

[§]Corresponding author. ⁸Project lead.

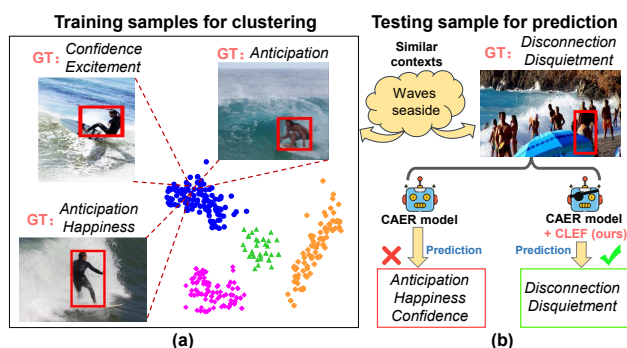


Figure 1. Illustration of the context bias in the CAER task. GT stands for the Ground Truth. Context-specific semantics easily yield spurious shortcuts with emotion labels during training to confound the model [32], giving erroneous results. Conversely, our CLEF effectively corrects biased predictions.

states from resource-efficient visual content has been extensively explored in various fields, including online education [15], driving monitoring [59], and human-computer interaction [1]. Conventional works have focused on extracting emotion-related information from subject attributes, such as facial expressions [9], body postures [3], acoustic behaviors [26], or multimodal combinations [22, 31, 54, 55, 57, 60]. Despite considerable advances in subject-oriented efforts, their performance suffers from severe bottlenecks in uncontrolled environments. As shown in Figure 1a, physical representations of subjects in wild-collected images are usually indistinguishable (e.g., ambiguous faces) due to natural occlusions that fail to provide usable emotional signals.

Inspired by psychological research [2], context-aware emotion recognition (CAER) [18] has been proposed to seek additional affective semantics from situational contexts. The contexts [19] are typically considered to include out-of-subject factors, such as background objects, place attributes, scene elements, and dynamic interactions of surrounding agents. These rich contextual stimuli promisingly

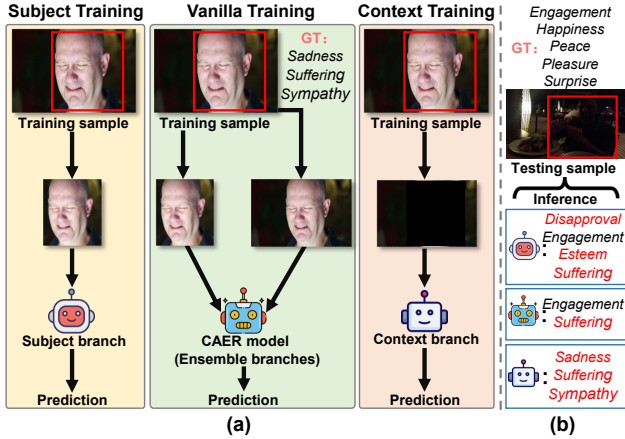


Figure 2. We conduct toy experiments to show the effects of context semantics. The indirect effect of the good context prior follows ensemble branches, narrowing the emotion candidate space. The bad direct effect follows the context branch, causing pure bias.

provide complementary emotion clues for accurate recognition. Most existing methods perform emotion inference by extracting ensemble representations from subjects and contexts using sophisticated structures [24, 32, 33, 53, 56, 65] or customized mechanisms [5, 7, 11, 14, 25, 42]. Nevertheless, a recent study [58] found that CAER models tend to rely on spurious correlations caused by a context bias rather than beneficial ensemble representations. An intuitive illustration is displayed in Figure 1. We first randomly choose some training samples on the EMOTIC dataset [19] and perform unsupervised clustering. From Figure 1a, samples containing seaside-related contexts form compact feature clusters, confirming the semantic similarity in the feature space. These samples have positive emotion categories, while negative emotions are nonexistent in similar contexts. In this case, the model [32] is easily misled to capture spurious dependencies between context-specific semantics and emotion labels. In the testing phase from Figure 1b, oriented to the sample with similar contexts but negative emotion categories, the model is confounded by the harmful context bias to infer completely wrong emotional states.

A straightforward solution is to conduct a randomized controlled trial by collecting images with all emotion annotations in all contexts. This manner is viewed as an approximate intervention for biased training. However, the current CAER debiasing effort [58] is sub-optimal since the predefined intervention fails to decouple good and bad context semantics. We argue that context semantics consists of the good prior and the bad bias. The toy experiments are performed to verify this insight. Specifically, we train on the EMOTIC dataset separately using the subject branch, the ensemble branches, and the context branch of a CAER baseline [18] in Figure 2a. Recognized subjects in sam-

ples during context training are masked to capture the direct context effect. Observing the testing results in Figure 2b, the context prior in ensemble learning as the valuable indirect effect helps the model filter out unnecessary candidates (*i.e.*, removing the “Disapproval” and “Esteem” categories) compared to the subject branch. Conversely, the harmful bias as the direct context effect in the context branch builds a misleading mapping between dim contexts and negative emotions during training, causing biased predictions.

To disentangle the two effects in context semantics and achieve more appropriate context debiasing, we propose a unified counterfactual emotion inference (CLEF) framework from a causality perspective. CLEF focuses on assisting existing CAER methods to mitigate the context bias and breakthrough performance bottlenecks in a model-agnostic manner rather than beating them. Specifically, we first formulate a generalized causal graph to investigate causal relationships among variables in the CAER task. Along the causal graph, CLEF estimates the direct context effect caused by the harmful bias through a non-intrusive context branch during the training phase. Meanwhile, the valuable indirect effect of the context prior in ensemble representations of subjects and contexts is calculated following the vanilla CAER model. In the inference phase, we subtract the direct context effect from the total causal effect by depicting a counterfactual scenario to exclude bias interference. This scenario is described as follows:

Counterfactual CAER: *What would the prediction be, if the model only sees the confounded context and does not perform inference via vanilla ensemble representations?* Intuitively, ensemble representations in the counterfactual outcome are blocked in the no-treatment condition. As such, the model performs biased emotion estimation relying only on spurious correlations caused by the pure context bias, which results similarly to the predictions of the context branch in Figure 2b. By comparing factual and counterfactual outcomes, CLEF empowers the model to make unbiased predictions using the debiased causal effect. The main contributions are summarized as follows:

- We are the first to embrace counterfactual thinking to investigate causal effects in the CAER task and reveal that the context bias as the adverse direct causal effect misleads the models to produce spurious prediction shortcuts.
- We devise CLEF, a model-agnostic CAER debiasing framework that facilitates existing methods to capture valuable causal relationships and mitigate the harmful bias in context semantics through counterfactual inference. CLEF can be readily adapted to state-of-the-art (SOTA) methods with different structures, bringing consistent and significant performance gains.
- Extensive experiments are conducted on several large-scale CAER datasets. Comprehensive analyses show the broad applicability and effectiveness of our framework.

2. Related Work

Context-Aware Emotion Recognition. Benefiting from advances in deep learning algorithms [6, 27–29, 46–48, 50–52, 61–63], traditional emotion recognition typically infers emotional states from subject-oriented attributes, such as facial expressions [9, 23], body postures [3, 59], and acoustic behaviours [26, 31]. However, these efforts are potentially vulnerable in practical applications since subject characteristics in uncontrolled environments are usually indistinguishable, leading to severe performance deterioration. Recently, a pioneering work [18] inspired by psychological research [2] has advocated extracting complementary emotional clues from rich contexts, called context-aware emotion recognition (CAER). Kosti *et al.* [19] begin by utilizing a two-stream convolutional neural network (CNN) to capture effective semantics from subject-related regions and global contexts of complete images. The implementation is similar to the ensemble branch training in Figure 2a. After that, most CAER methods [5, 7, 11, 14, 20, 24, 25, 32, 33, 53, 56, 64, 65] follow an ensemble learning pattern: **i)** extracting unimodal/multimodal features from subject attributes; **ii)** learning emotionally relevant features from created contexts based on different definitions; and **iii)** producing ensemble representations for emotion predictions via fusion mechanisms. For instance, Yang *et al.* [56] discretize the context into scenes, agent dynamics, and agent-object interactions, using customized components to learn complementary contextual information. Despite achievements, they invariably suffer from performance bottlenecks due to spurious correlations caused by the context bias.

Causal Inference. Causal inference [12] is first extensively used in economics [45] and psychology [10] as a scientific theory that seeks causal relationships among variables. The investigation of event causality generally follows two directions: intervention and counterfactuals. Intervention [36] aims to actively manipulate the probability distributions of variables to obtain unbiased estimations or discover confounder effects. Counterfactuals [37] typically utilize distinct treatment conditions to imagine outcomes that are contrary to factual determinations, empowering systems to reason and think like humans. In recent years, several learning-based approaches have attempted to introduce causal inference in diverse fields to pursue desired model effects and exclude the toxicity of spurious shortcuts, including scene graph generation [44], visual dialogue [34, 40], image recognition [4, 30, 49], and adversarial learning [16, 17]. The CAER debiasing effort [58] most relevant to our work utilizes a predefined dictionary to approximate interventions and adopts memory-query operations to mitigate the bias dilemma. Nevertheless, the predefined-level intervention fails to capture pure bias effects in the context semantics, causing a sub-optimal solution. Inspired by [34], we remove the adverse context effect by empowering models

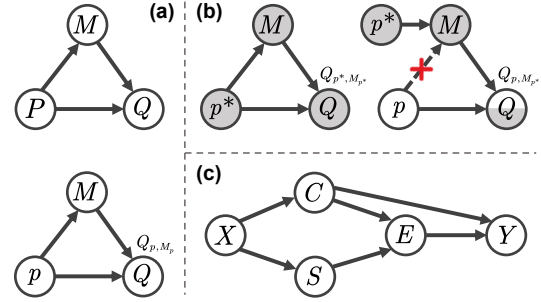


Figure 3. (a) Examples of a causal graph where nodes represent variables and arrows represent causal effects. (b) Examples of counterfactual notations. (c) The proposed CAER causal graph.

with the debiasing ability of twice-thinking through counterfactual causality, which is fundamentally different in design philosophy and methodology.

3. Preliminaries

Before starting, we first introduce the concepts and notations related to causal inference to facilitate a better understanding of our framework and philosophy.

Causal graph is a highly generalized analytical tool to reveal causal dependencies among variables. It usually follows the structured causal mode [39] defined as a directed acyclic graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where \mathcal{V} stands for a set of variables and \mathcal{E} implies the corresponding causal effects. A causal graph example with three variables is intuitively displayed in Figure 3a. Here, we represent a random variable as a capital letter (e.g., P), and denote its observed value as a lowercase letter (e.g., p). The causality from cause P to effect Q is reflected in two parts: the direct effect follows the causal link $P \rightarrow Q$, and the indirect effect follows the link $P \rightarrow M \rightarrow Q$ through the mediator variable M .

Counterfactual inference endows the models with the ability to depict counterfactual outcomes in factual observations through different treatment conditions [37]. In the factual outcome, the value of Q would be formalized under the conditions that P is set to p and M is set to m :

$$\begin{aligned} Q_{p,m} &= Q(P = p, M = m), \\ m &= M_p = M(P = p). \end{aligned} \quad (1)$$

Counterfactual outcomes can be obtained by exerting distinct treatments on the value of P . As shown in Figure 3b, when P is set to p^* , and the descendant M is changed, we have $Q_{p^*,M_{p^*}} = Q(P = p^*, M_{p^*} = M(P = p^*))$. Similarly, $Q_{p,M_{p^*}}$ reflects the counterfactual situation where $P = p$ and M is set to the value when $P = p^*$.

Causal effects reveal the difference between two corresponding outcomes when the value of the reference variable changes. Let $P = p$ denote the treated condition and $P = p^*$ represent the invisible counterfactual condition.

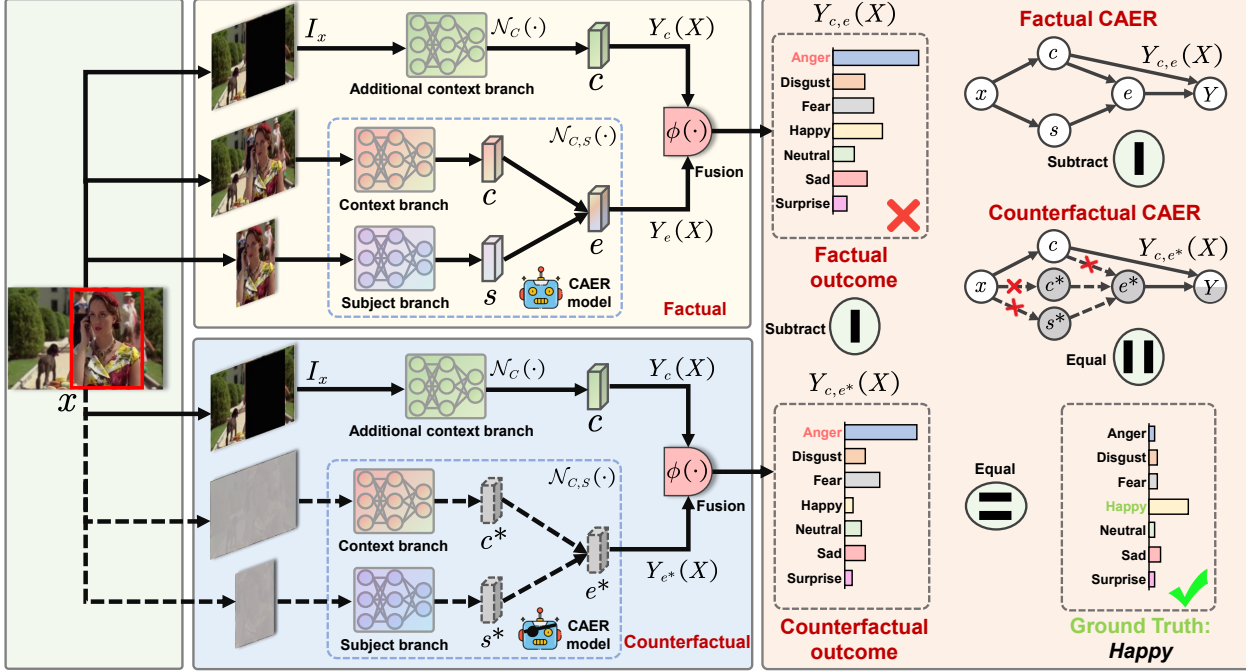


Figure 4. High-level overview of the proposed CLEF framework implementation. In addition to the vanilla CAER model, we introduce an additional context branch in a non-intrusive manner to capture the pure context bias as the direct context effect. By comparing factual and counterfactual outcomes, our framework effectively mitigates the interference of the harmful bias and achieves debiased emotion inference.

According to the causal theory [38], The Total Effect (TE) of treatment $P = p$ on Q by comparing the two hypothetical outcomes is formulated as:

$$\text{TE} = Q_{p, M_p} - Q_{p^*, M_{p^*}}. \quad (2)$$

TE can be disentangled into the Natural Direct Effect (NDE) and the Total Indirect Effect (TIE) [12]. NDE reflects the effect of $P = p$ on Q following the direct link $P \rightarrow Q$, and excluding the indirect effect along link $P \rightarrow M \rightarrow Q$ due to M is set to the value when P had been p^* . It reveals the response of Q when P converts from p to p^* :

$$\text{NDE} = Q_{p, M_{p^*}} - Q_{p^*, M_{p^*}}. \quad (3)$$

In this case, TIE is calculated by directly subtracting NDE from TE, which is employed to measure the unbiased prediction results in our framework:

$$\text{TIE} = \text{TE} - \text{NDE} = Q_{p, M_p} - Q_{p, M_{p^*}}. \quad (4)$$

4. The proposed CLEF Framework

4.1. Cause-Effect Look at CAER

As shown in Figure 3c, there are five variables in the proposed CAER causal graph, including input images X , subject features S , context features C , ensemble representations E , and emotion predictions Y . Note that our causal

graph has broad applicability and generality since it follows most CAER modelling paradigms.

Link $X \rightarrow C \rightarrow Y$ reflects the shortcut between the original inputs X and the model predictions Y through the harmful bias in the context features C . The adverse direct effect of the mediator C is obtained via a non-invasive branch of context modelling, which captures spurious correlations between context-specific semantics and emotion labels. Taking Figure 2b as an example, the context branch learns the undesired mapping between dim contexts and negative emotions during training.

Link $C \leftarrow X \rightarrow S$ portrays the total context and subject representations extracted from X via the corresponding encoders in vanilla CAER models. Based on design differences in distinct methods, C and S may come from a single feature or an aggregation of multiple sub-features. For instance, S is obtained from global body attributes and joint face-pose information in models [18] and [32], respectively.

Link $C/S \rightarrow E \rightarrow Y$ captures the indirect causal effect of C and S on the model predictions Y through the ensemble representations E . The mediator E is obtained depending on the feature integration mechanisms of different vanilla methods, such as feature concatenation [18] or attention fusion [20]. In particular, C provides the valuable context prior along the good causal link $C \rightarrow E \rightarrow Y$, which gives favorable estimations of potential emotional states when the subjects' characteristics are indistinguishable.

4.2. Counterfactual Inference

Our design philosophy is to mitigate the interference of the harmful context bias on model predictions by excluding the biased direct effect along the link $X \rightarrow C \rightarrow Y$. Following the notations on causal effects in Section 3, the causality in the factual scenarios is formulated as follows:

$$Y_{c,e}(X) = Y(C = c, E_{c,s} = E(C = c, S = s)|X). \quad (5)$$

$Y_{c,e}(X)$ reflects confounded emotion predictions because it suffers from the detrimental direct effect of C , *i.e.*, the pure context bias. To disentangle distinct causal effects in the context semantics, we calculate the Total Effect (TE) of $C = c$ and $S = s$, which is expressed as follows:

$$\text{TE} = Y_{c,e}(X) - Y_{c^*,e^*}(X). \quad (6)$$

Here, c^* and e^* represent the non-treatment conditions for observed values of C and E , where c and s leading to e are not given. Immediately, we estimate the Natural Direct Effect (NDE) for the harmful bias in context semantics:

$$\text{NDE} = Y_{c,e^*}(X) - Y_{c^*,e^*}(X). \quad (7)$$

$Y_{c,e^*}(X)$ describes a counterfactual outcome where C is set to c and E would be imagined to be e^* when C had been c^* and S had been s^* . The causal notation is expressed as:

$$Y_{c,e^*}(X) = Y(C = c, E_{c^*,s^*} = E(C = c^*, S = s^*)|X). \quad (8)$$

Since the indirect causal effect of ensemble representations E on the link $X \rightarrow C/S \rightarrow E \rightarrow Y$ is blocked, the model can only perform biased predictions by relying on the direct context effect in the link $X \rightarrow C \rightarrow Y$ that causes spurious correlations. To exclude the explicitly captured context bias in NDE, we subtract NDE from TE to estimate Total Indirect Effect (TIE):

$$\text{TIE} = Y_{c,e}(X) - Y_{c,e^*}(X). \quad (9)$$

We employ the reliable TIE as the unbiased prediction in the inference phase.

4.3. Implementation Instantiation

Framework Structure. From Figure 4, CLEF’s predictions consist of two parts: the prediction $Y_c(X) = \mathcal{N}_C(c|x)$ of the additional context branch (*i.e.*, $X \rightarrow C \rightarrow Y$) and $Y_e(X) = \mathcal{N}_{C,S}(c, s|x)$ of the vanilla CAER model (*i.e.*, $X \rightarrow C/S \rightarrow E \rightarrow Y$). The context branch is instantiated as a simple neural network $\mathcal{N}_C(\cdot)$ (*e.g.*, ResNet [13]) to receive context images with masked recognized subjects. The masking operation forces the network to focus on pure context semantics for estimating the direct effect. For a given input x , its corresponding context image I_x is expressed as:

$$I_x = \begin{cases} x(i, j) & \text{if } x(i, j) \notin \text{bbox}_{\text{subject}}, \\ 0 & \text{otherwise,} \end{cases} \quad (10)$$

where $\text{bbox}_{\text{subject}}$ means the bounding box of the subject. $\mathcal{N}_{C,S}(\cdot)$ denotes any CAER model based on their specific mechanisms to learn ensemble representations e from c and s for prediction. Subsequently, a pragmatic fusion strategy $\phi(\cdot)$ is introduced to obtain the final score $Y_{c,e}(X)$:

$$Y_{c,e}(X) = \phi(Y_c(X), Y_e(X)) = \log\sigma(Y_c(X) + Y_e(X)), \quad (11)$$

where σ is the sigmoid activation.

Training Procedure. As a universal framework, we take the multi-class classification task in Figure 4 as an example to adopt the cross-entropy loss $\mathcal{CE}(\cdot)$ as the optimization objective. The task-specific losses for $Y_{c,e}(X)$ and $Y_{c,e^*}(X)$ are as follows:

$$\mathcal{L}_{\text{task}} = \mathcal{CE}(Y_{c,e}(X), y) + \mathcal{CE}(Y_{c,e^*}(X), y), \quad (12)$$

where y means the ground truth. Since neural models cannot handle no-treatment conditions where the inputs are void, we devise a trainable parameter initialized by the uniform distribution in practice to represent the imagined $Y_{e^*}(X)$, which is shared by all samples. The design intuition is that the uniform distribution ensures a safe estimation for NDE, which is justified in subsequent ablation studies. To avoid inappropriate $Y_{e^*}(X)$ that potentially causes TIE to be dominated by TE or NDE, we employ the Kullback-Leibler divergence $\mathcal{KL}(\cdot)$ to regularize the difference between $Y_{c,e^*}(X)$ and $Y_{c,e}(X)$ to estimate $Y_{e^*}(X)$:

$$\mathcal{L}_{kl} = \mathcal{KL}(Y_{c,e^*}(X), Y_{c,e}(X)). \quad (13)$$

The final loss is expressed as:

$$\mathcal{L}_{\text{fin}} = \sum_{(c,s,y) \in \mathcal{D}} \mathcal{L}_{\text{task}} + \mathcal{L}_{kl}. \quad (14)$$

Inference Procedure. According to Eq. (9), the debiased prediction is performed as follows:

$$\text{TIE} = \phi(Y_c(X), Y_e(X)) - \phi(Y_c(X), Y_{e^*}(X)). \quad (15)$$

5. Experiments

5.1. Datasets and Evaluation Metrics

Experiments are conducted on two large-scale image-based CAER datasets, including EMOTIC [19] and CAER-S [20]. **EMOTIC** is the first benchmark to support emotion recognition in real-world contexts, which has 23,571 images of 34,320 annotated subjects. All samples are collected from non-controlled environments to provide rich context resources. Each recognized subject is annotated with 26 discrete emotion categories and body bounding box information. The dataset is partitioned into 70% samples for training, 10% samples for validation, and 20% samples for testing. **CAER-S** consists of 70k static images extracted

Table 1. Quantitative results of CLEF-based methods for each emotion category on the EMOTIC dataset. We report the average precision of each category to provide comprehensive comparison experiments. The improved results are marked in **bold**.

Category	EMOT-Net [19]	EMOT-Net + CLEF	CAER-Net [20]	CAER-Net + CLEF	GNN-CNN [65]	GNN-CNN + CLEF	CD-Net [53]	CD-Net + CLEF	EmotiCon [32]	EmotiCon + CLEF
Affection	26.47	35.28	22.36	28.62	47.52	61.84	28.44	35.51	38.55	43.72
Anger	11.24	11.76	12.88	14.01	11.27	16.37	12.12	14.6	14.69	17.09
Annoyance	15.26	17.46	14.42	12.85	12.33	11.08	19.71	16.94	24.68	25.40
Anticipation	57.31	94.29	52.85	82.27	63.20	93.25	57.65	89.05	60.73	92.24
Aversion	7.44	13.14	3.26	10.23	6.81	10.30	9.94	16.83	11.33	15.51
Confidence	80.33	74.48	72.68	73.18	74.83	69.02	69.26	73.11	68.12	65.90
Disapproval	16.14	19.73	15.37	17.04	12.64	15.16	22.78	27.45	18.55	21.47
Disconnection	20.64	30.66	22.01	24.76	23.17	28.35	27.55	31.70	28.73	33.31
Disquietment	19.57	19.73	10.84	13.47	17.66	20.11	21.04	23.37	22.14	24.56
Doubt/Confusion	31.88	19.81	26.07	22.15	19.67	16.57	24.23	19.55	38.43	32.87
Embarrassment	3.05	6.53	1.88	5.31	1.58	4.08	4.50	7.24	10.31	12.98
Engagement	86.69	97.39	73.71	90.46	87.31	92.88	85.32	94.38	86.23	92.75
Esteem	17.86	22.30	15.38	17.91	12.05	18.69	18.66	23.01	25.75	29.13
Excitement	78.05	73.36	70.42	63.01	72.68	65.21	70.07	60.42	80.75	72.64
Fatigue	8.87	10.34	6.29	8.66	12.93	17.67	11.56	14.67	19.35	22.34
Fear	15.70	8.46	7.47	10.12	6.15	10.34	10.38	11.23	16.99	18.71
Happiness	58.92	77.89	53.73	72.37	72.90	81.79	68.46	84.24	80.45	87.06
Pain	9.46	13.97	8.16	10.32	8.22	11.94	13.82	16.44	14.68	15.45
Peace	22.35	23.23	19.55	20.05	30.68	31.56	28.18	26.05	35.72	35.96
Pleasure	46.72	45.92	34.12	34.46	48.37	51.73	47.64	50.92	67.31	68.42
Sadness	18.69	27.19	17.75	23.06	23.90	33.28	32.99	37.43	40.26	45.25
Sensitivity	9.05	7.84	6.94	8.12	4.74	5.14	7.21	10.70	13.94	15.07
Suffering	17.67	18.05	14.85	15.63	23.71	25.60	35.19	30.85	48.05	43.16
Surprise	22.38	12.27	17.46	14.70	8.44	6.01	7.42	7.21	19.60	20.18
Sympathy	15.23	30.15	14.89	15.53	19.45	25.13	10.33	13.66	16.74	20.64
Yearning	9.22	12.13	4.84	5.16	9.86	13.64	6.24	8.63	15.08	17.39
mAP (%)	27.93	31.67	23.85	27.44	28.16	32.18	28.87	32.51	35.28	38.05

from video clips. These images record 7 emotional states of different subjects in various context scenarios from 79 TV shows. The data samples are randomly divided into training, validation, and testing sets in the ratio of 7:1:2. We utilize the standard mean Average Precision (mAP) and classification accuracy to evaluate the results on the EMOTIC and CAER-S datasets, respectively.

5.2. Model Zoo

We evaluate the effectiveness of the proposed CLEF using five representative methods, which have completely different network structures and contextual modelling paradigms. Concretely, **EMOT-Net** [18] is a two-stream classical CNN model where one stream extracts human features from body regions, and the other captures global context semantics. **CAER-Net** [20] extracts subject attributes from faces and uses the images after hiding faces as background contexts. **GNN-CNN** [65] utilizes the graph neural network (GNN) to integrate emotion-related objects in contexts and distills subject information with a VGG-16 [43]. **CD-Net** [53] designs a tube-transformer to perform fine-grained interactions from facial, bodily, and contextual features. **EmotiCon** [32] employs attention and depth maps to model context representations. Subject-relevant features are extracted from facial expressions and body postures.

5.3. Implementation Details

We use a ResNet-152 [13] pre-trained on the Places365 [66] dataset to parameterize the non-invasive context branch in

CLEF. The output of the last linear layer is replaced to produce task-specific numbers of neurons for predictions. Rich scene attributes in Places365 provide proper semantics for distilling the context bias. In addition to the annotated EMOTIC, we employ the Faster R-CNN [41] to detect bounding boxes of recognized subjects in CAER-S. Immediately, the context images I_x are obtained by masking the target subjects in samples based on the corresponding bounding boxes. For a fair comparison, the five selected CAER methods are reproduced via the PyTorch toolbox [35] following their reported training settings, including the optimizer, loss function, learning rate strategy, etc. All models are implemented on NVIDIA Tesla V100 GPUs.

5.4. Comparison with State-of-the-art Methods

We compare the five CLEF-based methods with existing SOTA models, including HLCR [7], TEKG [5], RRLA [24], VRD [14], SIB-Net [25], MCA [56], and GRERN [11].

Quantitative Results on the EMOTIC. Table 1 shows the Average Precision (AP) of the vanilla methods and their counterparts in the CLEF framework for each emotion category. We have the following critical observations. **i)** CLEF significantly improves the performance of all models in most categories. For instance, CLEF yields average gains of 8.33% and 6.52% on the AP scores for “Affection” and “Sadness”, reflecting positivity and negativity, respectively. **ii)** Our framework favorably improves several categories heavily confounded by the harmful context bias due to uneven distributions of emotional states across distinct




Figure 5. Emotion classification accuracy (%) for each category of different CLEF-based methods on the CAER-S dataset.

Table 2. Quantitative results of different models and CLEF-based methods on the EMOTIC dataset. \uparrow represents the improvement of the CLEF-based version over the vanilla method.

Methods	mAP (%)
HLCR [7]	30.02
TEKG [5]	31.36
RRLA [24]	32.41
VRD [14]	35.16
SIB-Net [25]	35.41
MCA [56]	37.73
EMOT-Net [19]	27.93
EMOT-Net + CLEF	31.67 (\uparrow 3.74)
CAER-Net [20]	23.85
CAER-Net + CLEF	27.44 (\uparrow 3.59)
GNN-CNN [65]	28.16
GNN-CNN + CLEF	32.18 (\uparrow 4.02)
CD-Net [53]	28.87
CD-Net + CLEF	32.51 (\uparrow 3.64)
EmotiCon [32]	35.28
EmotiCon + CLEF	38.05 (\uparrow 2.77)

contexts. For example, the CLEF-based models improve the AP scores for “Engagement” and “Happiness” categories to 90.46%~97.39% and 72.37%~87.06%, outperforming the results in the vanilla baselines by large margins.

Table 2 presents the comparison results with existing models regarding the mean AP (mAP) scores. **i)** Thanks to CLEF’s bias exclusion, the mAP scores of EMOT-Net, CAER-Net, GNN-CNN, CD-Net, and EmotiCon are consistently increased by 3.74%, 3.59%, 4.02%, 3.64%, and 2.77%, respectively. Among them, the most noticeable improvement in GNN-CNN is because the vanilla model more easily captures spurious *context-emotion* correlations based on fine-grained context element exploration [65], leading to the better debiasing effect with CLEF. **ii)** Compared to SIB-Net and MCA with complex module stacking [56] and massive parameters [25], the CLEF-based EmotiCon achieves the best performance with the mAP score of 38.05% through efficient counterfactual inference.

Quantitative Results on the CAER-S. Table 3 provides the evaluation results on the CAER-S dataset. **i)** Evidently, CLEF consistently improves different baselines by decoupling and excluding the prediction bias of emotional states in the TV show contexts. Concretely, the overall accuracies of EMOT-Net, CAER-Net, GNN-CNN, CD-Net, and EmotiCon are improved by 2.52%, 2.39%, 2.32%, 3.08%,

Table 3. Quantitative results of different models and CLEF-based methods on the CAER-S dataset.

Methods	Accuracy (%)
Fine-tuned VGGNet [43]	64.85
Fine-tuned ResNet [13]	68.46
SIB-Net [25]	74.56
MCA [56]	79.57
GRERN [11]	81.31
RRLA [24]	84.82
VRD [14]	90.49
EMOT-Net [19]	74.51
EMOT-Net + CLEF	77.03 (\uparrow 2.52)
CAER-Net [20]	73.47
CAER-Net + CLEF	75.86 (\uparrow 2.39)
GNN-CNN [65]	77.21
GNN-CNN + CLEF	79.53 (\uparrow 2.32)
CD-Net [53]	85.33
CD-Net + CLEF	88.41 (\uparrow 3.08)
EmotiCon [32]	88.65
EmotiCon + CLEF	90.62 (\uparrow 1.97)

and 1.97%, respectively. **ii)** The gains of our framework on the CAER-S are slightly weaker than those on the EMOTIC. A reasonable explanation is that the EMOTIC contains richer context semantics than the CAER-S, such as scene elements and agent dynamics [19]. As a result, CLEF more accurately estimates the adverse context effect and favorably removes its interference. **iii)** Also, we find in Figure 5 that the classification accuracies of most emotion categories across the five methods are improved appropriately.

5.5. Ablation Studies

In Table 4, we select the SOTA CD-Net and EmotiCon to perform thorough ablation studies on both datasets to evaluate the importance of all designs in CLEF.

Necessity of Framework Structure. **i)** When removing CAER models from CLEF, the significant performance deterioration suggests that the indirect causal effect in ensemble representations provides valuable emotion semantics. **ii)** When the additional context branch (ACB) is excluded, CLEF degrades to a debiased pattern that is not context-conditional, treated as TE. TE’s gains are inferior to TIE’s since it reduces the general bias over the whole dataset rather than the specific context bias. **iii)** Also, we find that the $\mathcal{KL}(\cdot)$ regularization is indispensable for estimating the proper $Y_{e^*}(X)$ and improving debiasing gains.

Rationality of Context Modelling. **i)** We observe that per-

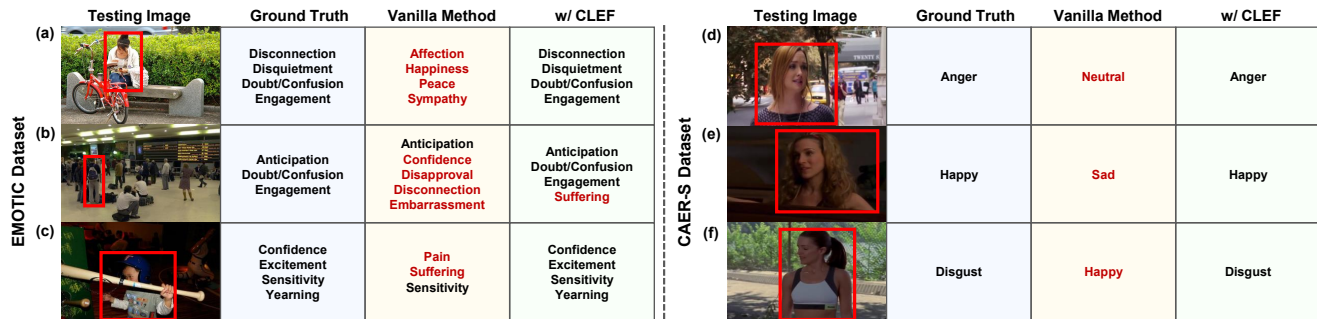


Figure 6. Qualitative results of the vanilla and CLEF-based CD-Net [53] on the EMOTIC and CAER-S datasets. Three testing sample images on each dataset are randomly selected. Incorrectly predicted categories are marked in red.

Table 4. Ablation study results on the EMOTIC and CAER-S datasets. “ACB” means the additional context branch. “w/” and “w/o” are short for the with and without, respectively.

Setting	EMOTIC [19]		CAER-S [20]	
	CD-Net	EmotiCon	CD-Net	EmotiCon
Vanilla Method	28.87	35.28	85.33	88.65
Necessity of Framework Structure				
+ CLEF	32.51	38.05	88.41	90.62
w/o CAER Model	19.64	19.64	62.87	62.87
w/o ACB	28.55	35.43	85.54	88.28
w/o $\mathcal{KL}(\cdot)$ Regularization	32.26	37.44	88.09	90.36
Rationality of Context Modelling				
w/o Masking Operation	31.38	36.95	87.68	89.85
w/ ImageNet Pre-training	30.74	36.62	87.35	89.27
w/ ResNet-50 [13]	31.45	37.54	87.83	90.04
w/ VGG-16 [43]	29.93	36.48	86.76	89.39
Effectiveness of No-treatment Assumption				
w/ Average Feature Embedding	27.85	33.18	83.06	85.67
w/ Random Feature Embedding	24.61	28.77	76.43	78.25

forming the masking operation on target subjects in input images of ACB is essential for ensuring reliable capture of the context-oriented adverse direct effect. **ii)** When the ResNet-152 pre-trained on Places365 [66] is replaced with the one pre-trained on ImageNet [8] in ACB, the gain drops prove that scene-level semantics are more expressive than object-level semantics in reflecting the context bias. This makes sense since scene attributes usually contain diverse object concepts. **iii)** Moreover, the improvements from CLEF gradually increase as more advanced pre-training backbones are used, which shows that our framework does not rely on a specific selection of instantiated networks.

Effectiveness of No-treatment Assumption. We provide two alternatives regarding the no-treatment condition assumption, where random and average feature embeddings are obtained by the random initialization and the prior distribution of the training set, respectively. The worse-than-baseline results imply that our uniform distribution assumption ensures a safe estimation of the biased context effect.

Debiasing Ability Comparison. A gain comparison between our CLEF and the previous CAER debiasing effort CCIM on both datasets is presented in Table 5. Intuitively,

Table 5. Debiasing comparison results of CCIM [58] and the proposed CLEF on the EMOTIC and CAER-S datasets.

Dataset	EMOT-Net [19]		CAER-Net [20]	
	Vanilla	w/ CCIM	Vanilla	w/ CCIM
EMOTIC	27.93	30.88	31.67	23.85
CAER-S	74.51	75.82	77.03	26.51

our framework consistently outperforms CCIM [58] in both methods. The reasonable reason is that CCIM fails to capture the pure context bias due to over-reliance on the predefined context confounders, causing sub-optimal solutions. In contrast, CLEF decouples the good context prior and the bad context effect, enabling robust debiased predictions.

5.6. Qualitative Evaluation

Figure 6 shows the performance of vanilla CD-Net before and after counterfactual debiasing via CLEF. Intuitively, our framework effectively corrects the misjudgments of the vanilla method for emotional states in diverse contexts. Taking Figure 6a as an example, CLEF eliminates spurious correlations between vegetation-related contexts and positive emotions, giving negative categories aligned with ground truths. Moreover, the CLEF-based CD-Net in Figure 6e excludes misleading clues about negative emotions provided by dim contexts and achieves an unbiased prediction.

6. Conclusion

This paper proposes CLEF, a causal debiasing framework based on counterfactual inference to address the context bias interference in CAER. CLEF reveals that the harmful bias confounds model performance along the direct causal effect via the tailored causal graph, and accomplishes bias mitigation by subtracting the direct context effect from the total causal effect. Extensive experiments prove that CLEF brings favorable improvements to existing models.

Acknowledgements

This work is supported in part by the National Key R&D Program of China (2021ZD0113503).

References

- [1] Abeer Ali Alnuaim, Mohammed Zakariah, Aseel Alhadlaq, Chitra Shashidhar, Wesam Atef Hatamleh, Hussam Tarazi, Prashant Kumar Shukla, Rajnish Ratna, et al. Human-computer interaction with detection of speaker emotions using convolution neural networks. *Computational Intelligence and Neuroscience*, 2022, 2022. 1
- [2] Lisa Feldman Barrett, Batja Mesquita, and Maria Gendron. Context in emotion perception. *Current Directions in Psychological Science*, 20(5):286–290, 2011. 1, 3
- [3] Uttaran Bhattacharya, Trisha Mittal, Rohan Chandra, Tanmay Randhavane, Aniket Bera, and Dinesh Manocha. Step: Spatial temporal graph convolutional networks for emotion perception from gaits. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1342–1350, 2020. 1, 3
- [4] Krzysztof Chalupka, Pietro Perona, and Frederick Eberhardt. Visual causal feature learning. *arXiv preprint arXiv:1412.2309*, 2014. 3
- [5] Jing Chen, Tao Yang, Ziqiang Huang, Kejun Wang, Meichen Liu, and Chunyan Lyu. Incorporating structured emotion commonsense knowledge and interpersonal relation into context-aware emotion recognition. *Applied Intelligence*, 53(4):4201–4217, 2023. 2, 3, 6, 7
- [6] Jiawei Chen, Dingkan Yang, Yue Jiang, Yuxuan Lei, and Lihua Zhang. Miss: A generative pretraining and finetuning approach for med-vqa. *arXiv preprint arXiv:2401.05163*, 2024. 3
- [7] Willams de Lima Costa, Estefania Talavera, Lucas Silva Figueiredo, and Veronica Teichrieb. High-level context representation for emotion recognition in images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop (CVPRW)*, pages 326–334, 2023. 2, 3, 6, 7
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255, 2009. 8
- [9] Yangtao Du, Dingkan Yang, Peng Zhai, Mingchen Li, and Lihua Zhang. Learning associative representation for facial expression recognition. In *IEEE International Conference on Image Processing (ICIP)*, pages 889–893, 2021. 1, 3
- [10] E Michael Foster. Causal inference and developmental psychology. *Developmental Psychology*, 46(6):1454, 2010. 3
- [11] Qinquan Gao, Hanxin Zeng, Gen Li, and Tong Tong. Graph reasoning-based emotion recognition network. *IEEE Access*, 9:6488–6497, 2021. 2, 3, 6, 7
- [12] Madelyn Glymour, Judea Pearl, and Nicholas P Jewell. *Causal inference in statistics: A primer*. John Wiley & Sons, 2016. 3, 4
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 5, 6, 7, 8
- [14] Manh-Hung Hoang, Soo-Hyung Kim, Hyung-Jeong Yang, and Guee-Sang Lee. Context-aware emotion recognition based on visual relationship detection. *IEEE Access*, 9:90465–90474, 2021. 2, 3, 6, 7
- [15] Scott E Jacobs and James J Gross. Emotion regulation in education: Conceptual foundations, current applications, and future directions. *International Handbook of Emotions in Education*, pages 183–201, 2014. 1
- [16] Diviyani Kalainathan, Olivier Goudet, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Sam: Structural agnostic model, causal discovery and penalized adversarial learning. 2018. 3
- [17] Murat Kocaoglu, Christopher Snyder, Alexandros G Dimakis, and Sriram Vishwanath. Causalgan: Learning causal implicit generative models with adversarial training. *arXiv preprint arXiv:1709.02023*, 2017. 3
- [18] Ronak Kosti, Jose M Alvarez, Adria Recasens, and Agata Lapedriza. Emotion recognition in context. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1667–1675, 2017. 1, 2, 3, 4, 6
- [19] Ronak Kosti, Jose M Alvarez, Adria Recasens, and Agata Lapedriza. Context based emotion recognition using emotic dataset. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 42(11):2755–2766, 2019. 1, 2, 3, 5, 6, 7, 8
- [20] Jiyoung Lee, Seungryong Kim, Sunok Kim, Jungin Park, and Kwanghoon Sohn. Context-aware emotion recognition networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10143–10152, 2019. 3, 4, 5, 6, 7, 8
- [21] Yuxuan Lei, Dingkan Yang, Mingcheng Li, Shunli Wang, Jiawei Chen, and Lihua Zhang. Text-oriented modality reinforcement network for multimodal sentiment analysis from unaligned multimodal sequences. *arXiv preprint arXiv:2307.13205*, 2023. 1
- [22] Mingcheng Li, Dingkan Yang, and Lihua Zhang. Towards robust multimodal sentiment analysis under uncertain signal missing. *IEEE Signal Processing Letters*, 30:1497–1501, 2023. 1
- [23] Shan Li and Weihong Deng. Deep facial expression recognition: A survey. *IEEE Transactions on Affective Computing*, 13(3):1195–1215, 2020. 3
- [24] Weixin Li, Xuan Dong, and Yunhong Wang. Human emotion recognition with relational region-level analysis. *IEEE Transactions on Affective Computing*, 2021. 2, 3, 6, 7
- [25] Xinpeng Li, Xiaojiang Peng, and Changxing Ding. Sequential interactive biased network for context-aware emotion recognition. In *IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–6, 2021. 2, 3, 6, 7
- [26] Zheng Lian, Bin Liu, and Jianhua Tao. Ctnet: Conversational transformer network for emotion recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:985–1000, 2021. 1, 3
- [27] Yang Liu, Jing Liu, Mengyang Zhao, Dingkan Yang, Xi-aoguang Zhu, and Liang Song. Learning appearance-motion normality for video anomaly detection. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2022. 3

- [28] Yang Liu, Jing Liu, Kun Yang, Bobo Ju, Siao Liu, Yuzheng Wang, Dingkang Yang, Peng Sun, and Liang Song. Ampnet: Appearance-motion prototype network assisted automatic video anomaly detection system. *IEEE Transactions on Industrial Informatics*, 2023.
- [29] Yang Liu, Dingkang Yang, Yan Wang, Jing Liu, and Liang Song. Generalized video anomaly event detection: Systematic taxonomy and comparison of deep models. *arXiv preprint arXiv:2302.05087*, 2023. 3
- [30] David Lopez-Paz, Robert Nishihara, Soumith Chintala, Bernhard Scholkopf, and Léon Bottou. Discovering causal signals in images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6979–6987, 2017. 3
- [31] Trisha Mittal, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. M3er: Multiplicative multimodal emotion recognition using facial, textual, and speech cues. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, pages 1359–1367, 2020. 1, 3
- [32] Trisha Mittal, Pooja Guhan, Uttaran Bhattacharya, Rohan Chandra, Aniket Bera, and Dinesh Manocha. Emoticon: Context-aware multimodal emotion recognition using frege’s principle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14234–14243, 2020. 1, 2, 3, 4, 6, 7
- [33] Trisha Mittal, Aniket Bera, and Dinesh Manocha. Multimodal and context-aware emotion perception model with multiplicative fusion. *IEEE MultiMedia*, 28(2):67–75, 2021. 2, 3
- [34] Yulei Niu, Kaihua Tang, Hanwang Zhang, Zhiwu Lu, Xian-Sheng Hua, and Ji-Rong Wen. Counterfactual vqa: A cause-effect look at language bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12700–12710, 2021. 3
- [35] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [36] Judea Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3:96–146, 2009. 3
- [37] Judea Pearl. *Causality*. Cambridge University Press, 2009. 3
- [38] Judea Pearl. Interpretation and identification of causal mediation. *Psychological Methods*, 19(4):459, 2014. 4
- [39] Judea Pearl et al. Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*, 19:2, 2000. 3
- [40] Jiaxin Qi, Yulei Niu, Jianqiang Huang, and Hanwang Zhang. Two causal principles for improving visual dialog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10860–10869, 2020. 3
- [41] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 28, 2015. 6
- [42] Shulan Ruan, Kun Zhang, Yijun Wang, Hanqing Tao, Weidong He, Guangyi Lv, and Enhong Chen. Context-aware generation-based net for multi-label visual emotion recognition. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE Computer Society, 2020. 2
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 6, 7, 8
- [44] Kaihua Tang, Yulei Niu, Jianqiang Huang, Jiaxin Shi, and Hanwang Zhang. Unbiased scene graph generation from biased training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3716–3725, 2020. 3
- [45] Hal R Varian. Causal inference in economics and marketing. *Proceedings of the National Academy of Sciences*, 113(27):7310–7315, 2016. 3
- [46] Shunli Wang, Dingkang Yang, Peng Zhai, Chixiao Chen, and Lihua Zhang. Tsa-net: Tube self-attention network for action quality assessment. In *Proceedings of the 29th ACM International Conference on Multimedia (ACM MM)*, pages 4902–4910, 2021. 3
- [47] Shunli Wang, Shuaiqing Wang, Bo Jiao, Dingkang Yang, Liuzhen Su, Peng Zhai, Chixiao Chen, and Lihua Zhang. Caspacenet: Counterfactual analysis for 6d pose estimation in space. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10627–10634, 2022.
- [48] Shunli Wang, Dingkang Yang, Peng Zhai, and Lihua Zhang. Cpr-clip: Multimodal pre-training for composite error recognition in cpr training. *IEEE Signal Processing Letters*, 2023. 3
- [49] Tan Wang, Jianqiang Huang, Hanwang Zhang, and Qianru Sun. Visual commonsense r-cnn. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10760–10770, 2020. 3
- [50] Yuzheng Wang, Zhaoyu Chen, Dingkang Yang, Pinxue Guo, Kaixun Jiang, Wenqiang Zhang, and Lizhe Qi. Model robustness meets data privacy: Adversarial robustness distillation without original data. *arXiv preprint arXiv:2303.11611*, 2023. 3
- [51] Yuzheng Wang, Zhaoyu Chen, Dingkang Yang, Yang Liu, Siao Liu, Wenqiang Zhang, and Lizhe Qi. Adversarial contrastive distillation with adaptive denoising. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023.
- [52] Yuzheng Wang, Zhaoyu Chen, Jie Zhang, Dingkang Yang, Zuhao Ge, Yang Liu, Siao Liu, Yunqian Sun, Wenqiang Zhang, and Lizhe Qi. Sampling to distill: Knowledge transfer from open-world data. *arXiv preprint arXiv:2307.16601*, 2023. 3
- [53] Zili Wang, Lingjie Lao, Xiaoya Zhang, Yong Li, Tong Zhang, and Zhen Cui. Context-dependent emotion recognition. *Journal of Visual Communication and Image Representation*, 89:103679, 2022. 2, 3, 6, 7, 8
- [54] Dingkang Yang, Shuai Huang, Haopeng Kuang, Yangtao Du, and Lihua Zhang. Disentangled representation learning for multimodal emotion recognition. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, pages 1642–1651, 2022. 1

- [55] Dingkang Yang, Shuai Huang, Yang Liu, and Lihua Zhang. Contextual and cross-modal interaction for multi-modal speech emotion recognition. *IEEE Signal Processing Letters*, 29:2093–2097, 2022. 1
- [56] Dingkang Yang, Shuai Huang, Shunli Wang, Yang Liu, Peng Zhai, Liuzhen Su, Mingcheng Li, and Lihua Zhang. Emotion recognition for multiple context awareness. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 144–162, 2022. 2, 3, 6, 7
- [57] Dingkang Yang, Haopeng Kuang, Shuai Huang, and Lihua Zhang. Learning modality-specific and -agnostic representations for asynchronous multimodal language sequences. In *Proceedings of the 30th ACM International Conference on Multimedia (ACM MM)*, pages 1708–1717, 2022. 1
- [58] Dingkang Yang, Zhaoyu Chen, Yuzheng Wang, Shunli Wang, Mingcheng Li, Siao Liu, Xiao Zhao, Shuai Huang, Zhiyan Dong, Peng Zhai, and Lihua Zhang. Context deconfounded emotion recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 19005–19015, 2023. 2, 3, 8
- [59] Dingkang Yang, Shuai Huang, Zhi Xu, Zhenpeng Li, Shunli Wang, Mingcheng Li, Yuzheng Wang, Yang Liu, Kun Yang, Zhaoyu Chen, et al. Aide: A vision-driven multi-view, multi-modal, multi-tasking dataset for assistive driving perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20459–20470, 2023. 1, 3
- [60] Dingkang Yang, Yang Liu, Can Huang, Mingcheng Li, Xiao Zhao, Yuzheng Wang, Kun Yang, Yan Wang, Peng Zhai, and Lihua Zhang. Target and source modality co-reinforcement for emotion understanding from asynchronous multimodal sequences. *Knowledge-Based Systems*, 265:110370, 2023. 1
- [61] Dingkang Yang, Kun Yang, Yuzheng Wang, Jing Liu, Zhi Xu, Rongbin Yin, Peng Zhai, and Lihua Zhang. How2comm: Communication-efficient and collaboration-pragmatic multi-agent perception. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023. 3
- [62] Kun Yang, Dingkang Yang, Jingyu Zhang, Mingcheng Li, Yang Liu, Jing Liu, Hanqi Wang, Peng Sun, and Liang Song. Spatio-temporal domain awareness for multi-agent collaborative perception. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 23383–23392, 2023.
- [63] Kun Yang, Dingkang Yang, Jingyu Zhang, Hanqi Wang, Peng Sun, and Liang Song. What2comm: Towards communication-efficient collaborative perception via feature decoupling. In *Proceedings of the 31th ACM International Conference on Multimedia (ACM MM)*, page 7686–7695, 2023. 3
- [64] Hanxin Zeng, Gen Li, Tong Tong, and Qinquan Gao. A graph convolutional network for emotion recognition in context. In *2020 Cross Strait Radio Science & Wireless Technology Conference (CSRSWTC)*, pages 1–3. IEEE, 2020. 3
- [65] Minghui Zhang, Yumeng Liang, and Huadong Ma. Context-aware affective graph reasoning for emotion recognition. In *IEEE International Conference on Multimedia and Expo (ICME)*, pages 151–156. IEEE, 2019. 2, 3, 6, 7
- [66] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(6):1452–1464, 2017. 6, 8