




Meta-Causal Feature Learning for Out-of-Distribution Generalization

Yuqing Wang, Xiangxian Li, Zhuang Qi, Jingyu Li, Xuelong Li,
Xiangxu Meng, and Lei Meng[✉] 

Shandong University, Jinan, Shandong, China
{wang-yuqing, xiangxian_lee, z_qi, jingyu_lee}@mail.sdu.edu.cn,
{lixuelong, mxm, lmeng}@sdu.edu.cn

Abstract. Causal inference has become a powerful tool to handle the out-of-distribution (OOD) generalization problem, which aims to extract the invariant features. However, conventional methods apply causal learners from multiple data splits, which may incur biased representation learning from imbalanced data distributions and difficulty in invariant feature learning from heterogeneous sources. To address these issues, this paper presents a balanced meta-causal learner (BMCL), which includes a balanced task generation module (BTG) and a meta-causal feature learning module (MCFL). Specifically, the BTG module learns to generate balanced subsets by a self-learned partitioning algorithm with constraints on the proportions of sample classes and contexts. The MCFL module trains a meta-learner adapted to different distributions. Experiments conducted on NICO++ dataset verified that BMCL effectively identifies the class-invariant visual regions for classification and may serve as a general framework to improve the performance of the state-of-the-art methods.

Keywords: Out-of-distribution · Causal inference · Meta-learning · Invariant feature · Balanced subsets

1 Introduction

Deep learning approaches have achieved impressive performance based on the independent and identically distributed (i.i.d.) hypothesis that testing and training data share similar distribution. However, real-world cases may violate the hypothesis due to the complex real data collection or generation mechanism (such as environmental differences [25], and selection bias [6]). Many studies have revealed the performance of classic machine learning methods has a sharp drop under distributional shifts [3, 15, 33], which indicates the necessity of learning an excellent model on out-of-distribution (OOD) data. To further promote the development of out-of-distribution generalization research, the NICO Challenge is held. The NICO Challenge is divided into two tracks: track 1 - public context generalization and track 2 - hybrid context generalization. The context of NICO++ includes two types: common context and unique context. The common context appears in all classes, so it supports the task of track 1, and the

unique context only appears in the specific class, which supports the task of track 2.

Commonly used neural networks, such as resnet18 [12] and resnet50 [12], are difficult to perform well in the OOD data scenarios. To solve this problem, many OOD generalization methods based on causal inference have been proposed [32, 33, 38, 39]. These algorithms typically aim to capture the invariant causal mechanism and focus on key regions of the image to reduce the impact of contextual diversity factors. CRLR [32] uses feature weighting to make images of different distributions have the same feature distribution. KeepingGood [35] adopts the idea of staged training to eliminate the adverse effect of data distribution on SGD momentum. CaaM [39] decouples causal features using CBAM [40] attention mechanisms, implements causal interventions based on IRM loss [4] to obtain cross-domain invariant causal features, and iteratively optimizing data partitioning to prevent excessive intervention. In general, these methods reduce the negative effects of confounding factors and achieve good results. However, we find facts that the subset of training data contains imbalance, which leads to ill-posed learning for causal learners and heterogeneous data leads to weight divergence of multiple classifiers, which may hinder model convergence. These reasons degrade the performance of existing methods.

To address these problems, this paper presents a balanced meta-causal learner (BMCL) that improve the conventional causal learning with balanced task generation and meta learning. The proposed framework follows the CaaM pipelines and it contains two novel modules: the balanced task generation and meta-causal feature learning modules. The balanced task generation module follows CaaM to generate raw data splits and samples meta-tasks therein with balanced three balancing strategies to alleviate the data imbalance problem, including manual balancing, loss balancing, and aggregation balancing with the final partitioning obtained via an aggregation balance algorithm. This balances the class and context of the training images in the tasks to allow the meta-learner to learn the invariant causal features under different visual contexts, and therefore alleviates the ill-posed learning with imbalanced data. The meta-causal feature learning module employs the meta-learning framework to enable the causal learner to retain knowledge learned from all tasks, i.e. data splits, and make it learn quickly to adapt to new tasks and learn unified features. This degrades the model complexity of CaaM by using a single meta-learner, and it fosters the model convergence caused by the weight divergence of its multiple causal learners.

Experiments have been conducted on the subset of the NICO++ dataset and the large datasets of NICO Challenge track1 and track2, including performance comparison, ablation study, and case study. The comparison results verify the generalization ability of the meta causal learner in the OOD case outperforms existing methods. The ablation further illustrates the effectiveness of each component. And case study shows the proposed BMCL can focus on class-invariant visual regions of images in different contexts. To summarize, this paper includes two main contributions:

- A meta causal learner is proposed, which can capture common causal features from different tasks. This can reduce the negative impacts of confounding factors and improve the generalization ability of the model in OOD case.
- It presents a balanced subset partitioning strategy and proves that balanced subsets can enhance the decision-making ability of the model.

2 Related Works

2.1 OOD Generalization

To further promote the research on the problem of agnostic distribution shifts between the training and testing sets, the out-of-distribution(OOD) generalization is proposed for learning a model that performs well under distribution shifts settings [1, 3, 16, 23, 33].

There are OOD generalization problems in many fields, specially in the field of image classification, where the diversity of OOD settings presents challenges: domain adaption [19, 37, 42], debiasing [7, 9, 21], long-tailed recognition [20, 24, 31, 35, 41]. To deal with the OOD problem, [13] proposed a real-world OOD dataset NICO, and an extended version of NICO called NICO++ [44] is released for NICO Challenge2022, which has a larger scale with images, contexts and classes. By adjusting the scale of the context, a variety of OOD situations can be simulated, allowing for an in-depth study. Domain adaptation tasks [19, 42] find a cross-domain invariant representation by separating task-related and task-independent features. Debiasing tasks [7, 21] improve generalization by training separate models based on data biases to remove biased information. Long-tail classification tasks facilitate classification by building networks to learn feature representations for the majority and minority classes [20], respectively, or by building a balance of head and tail classes [41].

2.2 Causal Inference

Causal inference is an effective method to solve the OOD problem. Pearl [26] believes that causality is divided into three levels: association, intervention and counterfactuals. Association is the correlation between data obtained by observing data. Intervention is using artificially controlling conditions to reduce the influence of confounding factors, usually through the front door criterion or back door criterion. Counterfactuals speculate on possible outcomes from conditions that did not occur. At present, there are outstanding causal intervention methods, such as re-weighting the samples so that samples in different environments have the same feature space [14, 27, 28, 32]. Furthermore, methods using invariant loss can obtain causal features that are invariant cross environments [2, 4, 18, 39]. The counterfactual method [11, 22, 34] is used to improve the generalization ability and robustness of the model by generating counterfactual samples. In the field of image classification, causal inference aims to make the model focus on the main regions of the image and ignore the impact of the context, that is, to obtain the invariant causal features.

2.3 Meta Learning

Meta-learning [17] is a branch of deep learning that aims to teach models to learn. One of the famous method is MAML [8], which is more like a learning strategy than a model and it can be used on other deep network models. Meta-learning builds tasks as the basic unit of training, learns from each task and quickly adapting to the next task. Meta-learning has a wide range of applications: Domain Adaptation and Domain Generalization [5, 36], in which meta-learning is used to learn regularizer or learn hyperparameters for feature transformation layers, so that the model can adapt to new domains. Another application is Few-Shot Learning [10, 29], which uses meta-learning and attention mechanisms to quickly extend to new classes by learning base classes without forgetting knowledge of old classes.

3 Relation Between BMCL, CBAM and CaaM

3.1 From Attentive Feature Learning to Causal Learning

OOD-distributed images amplify the impact of context on classification by creating pseudo-associations between context and categories, which requires the model to better focus on the subjects in images. To this goal, the work of CBAM [40] designed a lightweight attention module to focus on the class-predictive features \mathbf{F}_{att} , which is learned on the basis of raw features extracted by visual modal $V(\cdot)$, i.e., $\mathbf{F}_{att} = CBAM(V(x))$, where x is the images.

The effectiveness of the CBAM is limited because of the attentive bias caused by wrong associations between contextual information and classes. To addressed this, CaaM [39] improves the networks and the training pipeline to learn the context-irrelevant visual representation, as shown in Fig. 1(a). CaaM first learns to decouple causal features \mathbf{F}_c , confounding features \mathbf{F}_s , and mix features \mathbf{F}_x from images, that is, $\mathbf{F}_c, \mathbf{F}_s, \mathbf{F}_x = CaaM(x)$. Then CaaM generates training partitions $\mathcal{T} = \{t_1, t_2, \dots, t_m\}$ based on contextual information, and applies multiple causal learner for aggregating the knowledge of models.

3.2 The Enhanced Causal Learning by BMCL

As mentioned above, CaaM proposes to aggregate the knowledge learned in different partitions to obtain a better predictor, but the imbalance of partitions brings about mutual interference during aggregation, and the learning in different partitions needs to be more autonomous. The proposed BMCL addresses these problems and the training pipeline is shown in Fig. 1 (b).

BMCL follows the feature decoupling in CaaM, obtains $\mathbf{F}_c, \mathbf{F}_s, \mathbf{F}_x$, and designs the BTG module to get the balanced partition $\mathcal{T}' = \{t'_1, t'_2, \dots, t'_m\}$. In each partition, the learning is improved by means of setting corresponding task, and a meta-learner is used to learn unified class-level features, thereby enhancing the effects of causal learning.

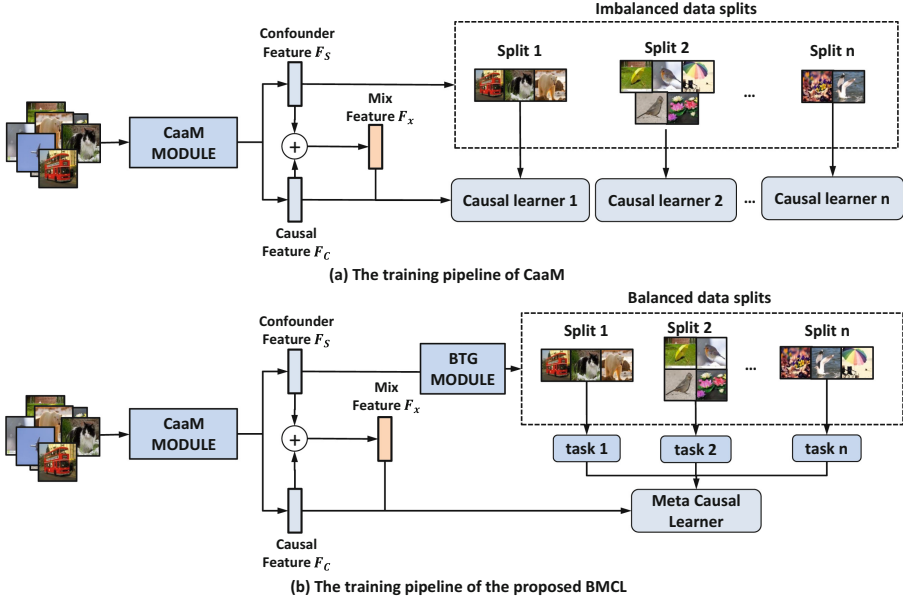


Fig. 1. Illustration of the training pipeline of CaaM and the proposed BMCL. BMCL makes the data partitioning more balanced, and adapts meta learner to enhance the effects of learning invariant features from different source data.

4 Our Approach

4.1 Overview

As shown in Fig. 2, the proposed BMCL has two main modules: balanced task generation (BTG) module and meta-causal feature learning (MCFL) module. The BTG module uses the aggregation balance method to train multiple partition matrices. This can divide the dataset into multiple subsets with different contexts in a balanced manner. The MCFL module aims to learn invariant causal features to reduce the negative effects of confounding features. This enables the model to focus on the main regions of the image with different contexts.

4.2 Balanced Task Generation (BTG) Module

The BTG module uses the confounder features \mathbf{F}_s extracted by the CaaM to train the partition matrix θ and update the data partition $\mathcal{T} = \{t_1, t_2, \dots, t_m\}$. BTG first uses a random partition matrix. To get the updated partition matrix, a bias classifier $h(\cdot)$ is trained for each matrix and the prediction $c = h(\mathbf{F}_s)$ can be obtained. For this, the BTG module can distinguish the contexts and it is optimized by minimizing the ERM Loss, i.e.

$$\mathcal{L}_{bias}^{erm} = E_{(x,y) \in \mathcal{D}} \ell(h(\mathbf{F}_s), y) \quad (1)$$

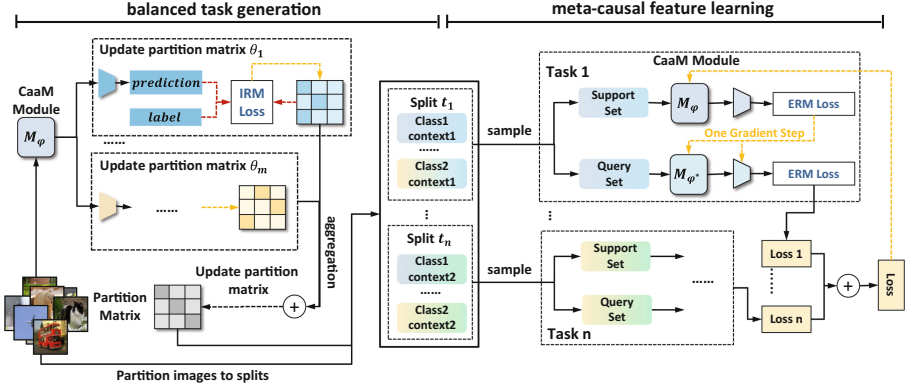


Fig. 2. Schematic diagram of the BMCL algorithm: the balanced task generation module divide the images into balanced splits with different contexts. The meta-causal feature learning module samples images from different partition to generate tasks, and learn invariant causal features from different tasks.

where \mathcal{D} is all training data, y is the label. Then using this classifier, under the constraint of an IRMLoss [4], a partition matrix θ is trained gradually. We maximize the loss to make sure the difference among all splits. After training we update the partition of the dataset in a fine-grained way:

$$\mathcal{L}_{split}^{irm} = \sum_{t \in \mathcal{T}_i(\theta)} R^t(h) + \lambda \cdot \|\nabla_{w|w=1.0} R^t(w \cdot h)\|^2 \quad (2)$$

where $\mathcal{T} = \{t_1, t_2, \dots, t_m\}$ is current data partition, $\mathcal{T}_i(\theta)$ denotes partition \mathcal{T}_i is decided by $\theta \in \mathcal{R}^{K \times m}$, K is the total number of training samples and m is the number of splits in a partition, $R^t(h) := E_{(x,y) \in t_i} \ell(h(F_s), y)$ is the loss under subset t_i , $h(\cdot)$ is the bias classifier trained in the previous step, $w = 1.0$ is a scalar and fixed “dummy classifier, the gradient norm penalty is used to measure the optimality of the dummy classifier at each subset t , and $\lambda \in [0, \infty]$ is a regularizer balancing weight between the ERM term and the invariance of the predictor $1 \cdot h$.

During this process, we found training only one partition matrix θ may lead to a large imbalance in subsets, which may lead to a poor performance of the model. In order to alleviate the imbalance, BMCL employs three balancing strategies and discusses this in the ablation study. Eventually, BMCL uses multiple matrices to decide the final partition in BTG module. It combines the probability distributions of multiple training to alleviate the problem of imbalance. And

then the aggregation method is used to obtain the final partition, i.e.

$$\theta_{final} = \sum_{i=0}^n \begin{pmatrix} p(1, 1) \cdots p(1, m) \\ \vdots \quad \ddots \quad \vdots \\ p(k, 1) \cdots p(k, m) \end{pmatrix}_i \quad (3)$$

where $\theta_{final} \in R^{K \times m}$, $p(k, m)$ denotes the probability that the k^{th} image is divided into the m^{th} partition. For θ_{final} , the index of the split to be divided into is:

$$Idx = \underset{\theta}{argmax}(Softmax(\theta_{final})) \quad (4)$$

Then, dividing the K images into corresponding data subsets based on Eq. 4. The BTG module divides the dataset into fine-grained subsets with different contexts, which helps the model adapt to other tasks and extract cross-domain invariant causal features.

4.3 Meta-Causal Feature Learning (MCFL) Module

Commonly used a way to solve the OOD problem of image classification is to extract causal features of image subjects from different contexts. There is a feasible solution that is to use the properties of meta-learning cross-task learning to extract invariant features from different tasks. Based on the idea of meta-learning, the MCFL module first organizes the training set into the form of a task, which consists of a support set and a query set for the training process. For m splits, we organize m tasks as an update. Each task is from one split. The total number of meta-tasks is s , which can be divided by m . The composition of the task is to randomly sample w classes, and then randomly sample $i + j$ images from each class, of which i images for the support set and j images for the query set.

During the training phase, each meta-task has its own independent training and testing sets, called support and query sets. The MCFL first promotes fast convergence of the model by simply updating the support set training in the meta-training process. The model first learns a parameter from the support set and then updates the model. Note, we define the feature extractor as M_φ with parameters φ and the classifier as f_μ with parameters μ . We use the feature extractor M_φ and classifier f_μ to get the causal feature F_c and then calculate an ERM Loss for the first model updating:

$$F_c = M_\varphi(x_{spt}) \quad (5)$$

$$\mathcal{L}_{\mathcal{K}_s} = \ell(f_\mu(F_c), y_{spt}) \quad (6)$$

where the x_{spt} and y_{spt} is the images and labels in support set of task \mathcal{K}_s , and the ℓ is the cross-entropy loss. The updated parameter vector φ^* and μ^* are updated using one or multiple gradient descent on the current support set calculated by $\mathcal{L}_{\mathcal{K}_s}$. For example, when using one gradient update:

$$\varphi^* = \varphi - \alpha \nabla_\varphi \mathcal{L}_{\mathcal{K}_s} \quad (7)$$

$$\mu^* = \mu - \alpha \nabla_{\mu} \mathcal{L}_{\mathcal{K}_s} \quad (8)$$

where the step size α may be fixed as a hyperparameter or meta-learned.

For meta-learning, a good meta learner should perform well on all meta tasks. Therefore, BMCL has only one meta learner and uses ERM loss to update this single learner. We use the updated model to train on the query set and get the training loss. Since the support set and the query set have the same classes, after learning on the support set to get φ^* , the learning of causal features can be promoted on the query set, and show a good performance. For building meta-task from splits, the updated model can be obtained after m tasks:

$$F_c = M_{\varphi^*}(x_{qry}) \quad (9)$$

$$\mathcal{L}_{train}^{meta} = \frac{1}{m} \sum_{i=0}^m \mathcal{L}_{\mathcal{K}_i} = \frac{1}{m} \sum_{i=0}^m \ell(f_{\mu^*}(F_c), y_{qry}) \quad (10)$$

It is worth noting that the meta-optimization is performed over the model parameters F_{φ} , whereas the objective is computed using the updated model parameters φ^* .

$$\varphi \leftarrow \varphi - \beta \nabla_{\varphi} \sum_{i=0}^m \mathcal{L}_{\mathcal{K}_i} \quad (11)$$

$$\mu \leftarrow \mu - \beta \nabla_{\mu} \sum_{i=0}^m \mathcal{L}_{\mathcal{K}_i} \quad (12)$$

where β is the meta step size. By training on multiple tasks and continuously minimizing the sum of losses on all tasks, the model can accurately extract causal features. This provides a guarantee for the model to achieve satisfactory performance on new tasks. In addition, a ERM loss is used to enhance the invariance of features, in this stage, we use mixup [43] to strengthen the images:

$$\begin{aligned} \tilde{x} &= \lambda x_i + (1 - \lambda) x_j \\ \tilde{y} &= \lambda y_i + (1 - \lambda) y_j \end{aligned} \quad (13)$$

And the loss was calculated as:

$$F_c = M_{\varphi}(\tilde{x}) \quad (14)$$

$$\mathcal{L}_{train}^{erm} = \ell(g(F_c), \tilde{y}) \quad (15)$$

where the $g(\cdot)$ is a classifier for this process alone.

4.4 Training Strategy

BMCL incorporates two stages model training process:

- **Stage 1:** Training balanced subsets partition. The initial partition matrix is randomly generated. And then a balanced constrain is used to update the partition matrices. First, we train a biased classifier using the empirical risk loss \mathcal{L}_{bias}^{erm} and use the classifier to constrain the data partition by an invariant risk loss $\mathcal{L}_{split}^{irm}$. Significantly, we minimize the empirical risk loss \mathcal{L}_{bias}^{erm} , but maximize the invariant risk loss $\mathcal{L}_{split}^{irm}$, i.e.

$$\min \mathcal{L}_{bias}^{erm} - \lambda \mathcal{L}_{split}^{irm} \quad (16)$$

- **Stage 2:** The extraction of training features is jointly constrained by the empirical risk loss from different tasks and from the whole dataset. In order to obtain robust features, we minimize the loss:

$$\min \mathcal{L}_{train}^{meta} + \mathcal{L}_{train}^{erm} \quad (17)$$

5 Experiments

5.1 Datasets

NICO++ [44]. A real-world image classification dataset under OOD settings, that is, the contexts of images in the testing may be unseen during training. To achieve this, NICO++ decomposes images into subject concepts and visual contexts, so that the contextual distribution in training and testing can be easily adjusted. There are typically two types of contextual setting in NICO++, namely, the common context setting (as in track 1 of NICO++ Challenge) and the hybrid context setting (as in track 2 of NICO++ Challenge). The common contexts setting means that all classes share identical contexts both in training and testing; and the unique contexts setting means each class has unique contexts. In details, NICO++ currently includes 200,000 images in 80 categories, ranging from animals, plants, traffic to objects; images in each category are organized into 10 public contexts and 10 unique contexts.

Track1. A subset sampled from NICO++ under common context setting. The track1 dataset has 88,866 images for training, 13,907 for public testing, and 35,920 for private testing. All contexts in track1 is existing in all categories.

Track2. A subset sampled from NICO++ under hybrid context setting. The dataset for track2 has 57,425 images for training, 8715 for public testing, and 20,079 for private testing. in track2, each category includes some unique contexts, which makes the situation more complected.

NICO++(subset). To further investigate the ability of OOD generation, we used a self-made dataset termed as NICO++(subset) in experiments. In details, sampled images contain 6 contexts in 10 classes including animals, vehicles, and others from the NICO++. Furthermore, we follow the long-tailed and zero-shot settings in previous work [39], and additionally set 4 contexts for training and 2 contexts unseen during training for testing, the size of training, validating, and testing are 2,870, 1,754, and 1755 respectively.

5.2 Evaluation Protocol

We use the Top-1 Accuracy on the validation set and test set for evaluation as the previous work [39] of causal learning did. The formula is as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \quad (18)$$

where TP, TN, FP, FN are the number of true positive, false positive samples, true negative samples, and false negative samples.

5.3 Implementation Details

For experiments in all the datasets, we used SGD as the optimizer, and setting learning rate from 0.01 to 0.03. Models were trained for 220 epochs and the learning rate was decayed by 0.1 at 200 epoch. In settings of partition training in CaaM and BMCL, the number of partition is 4 and from the 40-th epoch, the data partitions were updated every 20 epochs. For the process of updating splits, we followed the previous paper and trained each for 100 epochs with early stopping, when epoch is over 40 and accuracy no longer increases more than 5 epoch. The optimizer for this process was set to SGD with learning rate as 0.1. And the λ for IRM Loss was set to $1e6$. For the meta-task, we choose 3 classes for one task, and 1 or 2 image(s) for the support sets, 11–15 images for the query set. The updated learning rate of the one step gradient of meta was set to 0.005–0.01. For the erm training process, we set the batch size to $32*4$.

5.4 Performance Comparison

This section shows the performance comparison of BMCL with other image classification methods. We typically divided them into visual learning based on convolutional networks (Conv. Method) including the traditional ResNet-18 [12], the widely-used data augmentation methods Mixup [43]; and methods using causal inference, including the CBAM [40] and the CaaM [39]. The drawn following observations from Table 1:

- Generally speaking, BMCL achieved better performance than other algorithms in all cases. It is reasonable since BMCL is able to generate balanced subsets to away from ill-posed learning. Benefiting from this property, BMCL typically outperformed the other methods on both datasets with a large margin up to 12%.

Table 1. Classification accuracies (%) of algorithms on track1, track2, and NICO++(subset) datasets using ResNet18 as backbone. “val” and “test” denote the accuracies on validating set and testing set. “pub test” denote the accuracies on the publish testing set of Nico++ Challenge.

Method	Model	NICO++(subset)		NICO++-track1	NICO++-track2
		val acc	test acc	pub test	pub test
Conv. method	ResNet-18	44.73	45.93	58.07	47.11
	Mixup	49.00	49.06	66.47	56.25
Causal method	CBAM	44.27	45.47	65.35	58.71
	CaaM	43.93	46.44	72.93	66.44
	BMCL	51.45	52.02	84.66	78.81

- Basic visual networks such as ResNet-18 are prone to bias the context and lead to poor predictions when faced with OOD images. The Mixup enhancement enriches the context type corresponding to the category subject to a certain extent through image interpolation during training, thereby implicitly enhancing the model’s attention to the subject part, and achieving an improvement of about 8% on the small data set NICO++ (subset). On the larger datasets Track 1 and Track 2, due to the richer contextual information, the improvement is also more obvious by 13% and 19%.
- Methods based on causal learning generally perform better than traditional vision networks on large datasets, mainly because they design explicit methods to pay attention to the predictive features of images, and the improvement is more stable in the presence of sufficient data.
- However, on the small data set NICO++ (subset), the method based on causal learning did not bring significant improvement to the results. For CBAM, the attention mechanism is easy to focus on the background on a small amount of data, which aggravates the impact of OOD on model prediction. However, the update of CaaM is prone to imbalance. In a small data set, some predictors only have few samples, which is likely to have a bad impact on the overall prediction.

5.5 Ablation Study

In this section, we investigate the effectiveness of the proposed algorithm. The experiment selected resnet18 [12] as the baseline. Our BMCL method is divided into two modules, a balancing module and a meta-learning module. As can be seen from Table 2, we tested two modules based on CaaM respectively, and both achieved good results. Then we use the two modules at the same time, and use the meta-model to learn causal features while ensuring the balance of the dataset partition, and achieve the current best performance.

Table 2. Results of the ablation study, showing the effectiveness of each module in BMCL on classification performance.

Model	NICO++(subset)	
	val acc	test acc
Baseline	44.73	45.93
+CBAM	44.27	45.47
+CBAM+Causal	43.93	46.44
+CBAM+Causal+BTG	46.38	47.46
+CBAM+Causal+meta	49.74	49.52
+CBAM+Causal+BTG+meta	51.45	52.02

Table 3. The results of different combination of feature learning and balancing strategy.

Feature learning	Balancing strategy	NICO++(subset)	
		val acc	test acc
Backbone	-	43.93	46.44
	LB	46.04	47.24
	MB	44.05	47.41
	GB	46.38	47.46
Meta learner	-	49.74	49.52
	LB	49.86	50.26
	MB	51.62	49.57
	GB	51.45	52.02

5.6 In Depth Analysis of Balancing Strategy and Meta Learner

Table 3 shows the performance of using different balancing methods, including Loss Balance (LB), Manual Balance (MB), Aggregation Balance (GB). Among them, LB is to add loss in the process of training the partition matrix, it can automatically balance the number of images in different subsets during the training process. MB is to manually select the images of each class when dividing the splits, and follow the principle of more deletion, less complement, and move the extra images from a split to other splits. GB is a smooth operation that reduces the chance and extremeness of the partition by training multiple partition matrices, so as to alleviate the imbalance.

By testing different balancing methods, we found that GB is the best, and the reason why LB and MB are not good may be because the forced full balancing leads to destroying the structure of data partitioning, which is disadvantageous to improve performance. While GB can alleviate the imbalance, it can retain the structure of data partition. As observed, GB balance method achieved the

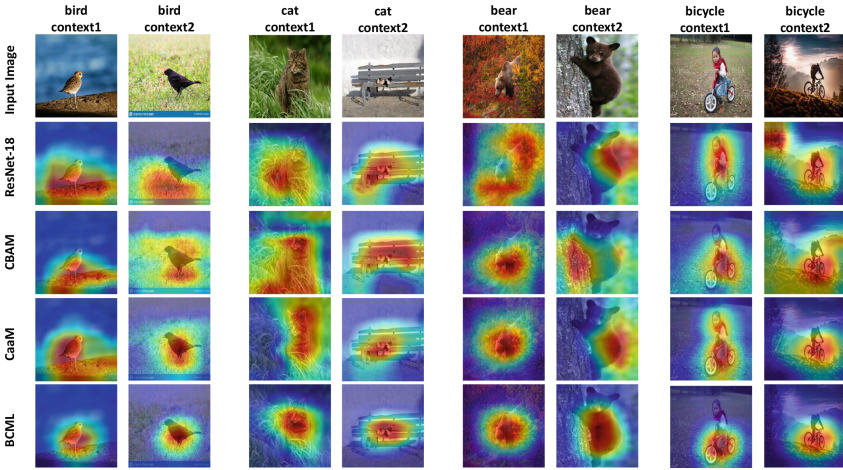


Fig. 3. Visualization of attention maps with base model, CBAM, CaaM, and BMCL

best results both under CaaM and under the meta learner and outperformed the other on NICO++(subset) with 2%.

5.7 Case Study

Figure 3 shows the attention maps generated by the proposed BMCL, the commonly-used backbone ResNet18, the attention mechanism CBAM, and the causal inference method CaaM. This is achieved by using the Grad-CAM [30]. In this experiment, four classes images from the NICO++ test dataset are selected and each class of images contains two contexts. As observed, BMCL can focus on the class meaningful regions of the image rather than the contexts. For the ResNet18, it can easily be influenced by the surrounding environment and focus on meaningless regions, such as a cyclist and a light in the background. For the CBAM, it always pays attention to a wide range, this contains the important regions and a large of noisy regions besides. The CaaM achieves the good performance compare with ResNet18 and CBAM, which indicates the reliability of causal inference. But CaaM may learn biased feature attention maps, such as a cyclist in bicycle image and a bench in cat image. This could be caused by unbalanced data learning. Significantly, BMCL achieves the best performance. It can focus precisely on the class-invariant visual regions, and exclude the interference of complex background. These observations verify the effectiveness of BMCL for invariant causal feature learning.

6 Conclusions

This paper presents a novel approach, termed BMCL, to cope with the challenge of agnostic distribution shifts in out-of-distribution (OOD) settings, which can

perform a self-learning balanced subset partition method to generate balanced subsets and learn the invariant causal features based on meta-learning. Notably, BMCL can keep the causal learner away from ill-posed learning and reduce the model complexity. Experimental results show that BMCL can alleviate the interference of confounder factors and enhance the learning and generalization ability of the model in the OOD case.

This study can be further explored in two directions. First, a cost-effective subsets partition method can be explored to reduce the time cost, such as combining curriculum learning to partially sample the data. Second, BMCL can be extended to more challenging settings, such as Federated Learning.

Acknowledgment. This work is supported in part by the Excellent Youth Scholars Program of Shandong Province (Grant no. 2022HWYQ-048) and the Oversea Innovation Team Project of the “20 Regulations for New Universities” funding program of Jinan (Grant no. 2021GXRC073)

References

1. Achille, A., Soatto, S.: Information dropout: Learning optimal representations through noisy computation. *IEEE Trans. Pattern Anal. Mach. Intell.* **40**(12), 2897–2905 (2018)
2. Ahuja, K., Shanmugam, K., Varshney, K., Dhurandhar, A.: Invariant risk minimization games. In: *International Conference on Machine Learning*. pp. 145–155. PMLR (2020)
3. Arjovsky, M.: Out of distribution generalization in machine learning. Ph.D. thesis, New York University (2020)
4. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. *arXiv preprint [arXiv:1907.02893](https://arxiv.org/abs/1907.02893)* (2019)
5. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. *Advances in neural information processing systems* 31 (2018)
6. Bengio, Y., Deleu, T., Rahaman, N., Ke, R., Lachapelle, S., Bilaniuk, O., Goyal, A., Pal, C.: A meta-transfer objective for learning to disentangle causal mechanisms. *arXiv preprint [arXiv:1901.10912](https://arxiv.org/abs/1901.10912)* (2019)
7. Clark, C., Yatskar, M., Zettlemoyer, L.: Don’t take the easy way out: Ensemble based methods for avoiding known dataset biases. *arXiv preprint [arXiv:1909.03683](https://arxiv.org/abs/1909.03683)* (2019)
8. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *International conference on machine learning*. pp. 1126–1135. PMLR (2017)
9. Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F.A., Brendel, W.: Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint [arXiv:1811.12231](https://arxiv.org/abs/1811.12231)* (2018)
10. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4367–4375 (2018)

11. Goyal, Y., Wu, Z., Ernst, J., Batra, D., Parikh, D., Lee, S.: Counterfactual visual explanations. In: International Conference on Machine Learning. pp. 2376–2384. PMLR (2019)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
13. He, Y., Shen, Z., Cui, P.: Towards non-iid image classification: A dataset and baselines. *Pattern Recogn.* **110**, 107383 (2021)
14. Heinze-Deml, C., Peters, J., Meinshausen, N.: Invariant causal prediction for non-linear models. *Journal of Causal Inference* 6(2) (2018)
15. Hendrycks, D., Basart, S., Mu, N., Kadavath, S., Wang, F., Dorundo, E., Desai, R., Zhu, T., Parajuli, S., Guo, M., et al.: The many faces of robustness: A critical analysis of out-of-distribution generalization. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8340–8349 (2021)
16. Hendrycks, D., Gimpel, K.: A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint [arXiv:1610.02136](https://arxiv.org/abs/1610.02136) (2016)
17. Hospedales, T.M., Antoniou, A., Micaelli, P., Storkey, A.J.: Meta-learning in neural networks: A survey. *IEEE transactions on pattern analysis and machine intelligence* (2021)
18. Jin, W., Barzilay, R., Jaakkola, T.: Domain extrapolation via regret minimization. arXiv preprint [arXiv:2006.03908](https://arxiv.org/abs/2006.03908) (2020)
19. Jin, X., Lan, C., Zeng, W., Chen, Z.: Style normalization and restitution for domain generalization and adaptation. *IEEE Transactions on Multimedia* (2021)
20. Khan, S.H., Hayat, M., Bennamoun, M., Sohel, F.A., Togneri, R.: Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems* **29**(8), 3573–3587 (2017)
21. Kim, B., Kim, H., Kim, K., Kim, S., Kim, J.: Learning not to learn: Training deep neural networks with biased data. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9012–9020 (2019)
22. Kim, H., Shin, S., Jang, J., Song, K., Joo, W., Kang, W., Moon, I.C.: Counterfactual fairness with disentangled causal effect variational autoencoder. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 8128–8136 (2021)
23. Liang, S., Li, Y., Srikant, R.: Enhancing the reliability of out-of-distribution image detection in neural networks. arXiv preprint [arXiv:1706.02690](https://arxiv.org/abs/1706.02690) (2017)
24. Mahajan, Dhruv, Girshick, Ross, Ramanathan, Vignesh, He, Kaiming, Paluri, Manohar, Li, Yixuan, Bharambe, Ashwin, van der Maaten, Laurens: Exploring the limits of weakly supervised pretraining. In: Ferrari, Vittorio, Hebert, Martial, Sminchisescu, Cristian, Weiss, Yair (eds.) ECCV 2018. LNCS, vol. 11206, pp. 185–201. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-01216-8_12
25. Meinshausen, N., Bühlmann, P.: Maximin effects in inhomogeneous large-scale data. *Ann. Stat.* **43**(4), 1801–1830 (2015)
26. Pearl, J., Mackenzie, D.: The Book of Why: The New Science of Cause and Effect. Basic books (2018)
27. Peters, J., Bühlmann, P., Meinshausen, N.: Causal inference by using invariant prediction: identification and confidence intervals. *J. R. Stat. Soc. Ser B (Stat. Methodol.)* **78**(5), 947–1012 (2016)
28. Pfister, N., Bühlmann, P., Peters, J.: Invariant causal prediction for sequential data. *J. Am. Stat. Assoc.* **114**(527), 1264–1276 (2019)

29. Ren, M., Liao, R., Fetaya, E., Zemel, R.: Incremental few-shot learning with attention attractor networks. In: 32nd Advances in Neural Information Processing Systems (2019)
30. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 618–626 (2017)
31. Shen, L., Lin, Z., Huang, Q.: Relay backpropagation for effective learning of deep convolutional neural networks. In: European Conference on Computer Vision, pp. 467–482. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46478-7_29
32. Shen, Z., Cui, P., Kuang, K., Li, B., Chen, P.: Causally regularized learning with agnostic data selection bias. In: Proceedings of the 26th ACM International Conference on Multimedia, pp. 411–419 (2018)
33. Shen, Z., et al.: Towards out-of-distribution generalization: A survey. arXiv preprint [arXiv:2108.13624](https://arxiv.org/abs/2108.13624) (2021)
34. Subbaswamy, A., Saria, S.: Counterfactual normalization: proactively addressing dataset shift and improving reliability using causal mechanisms. arXiv preprint [arXiv:1808.03253](https://arxiv.org/abs/1808.03253) (2018)
35. Tang, K., Huang, J., Zhang, H.: Long-tailed classification by keeping the good and removing the bad momentum causal effect. Adv. Neural. Inf. Process. Syst. **33**, 1513–1524 (2020)
36. Tseng, H.Y., Lee, H.Y., Huang, J.B., Yang, M.H.: Cross-domain few-shot classification via learned feature-wise transformation. arXiv preprint [arXiv:2001.08735](https://arxiv.org/abs/2001.08735) (2020)
37. Wang, J., et al.: Generalizing to unseen domains: A survey on domain generalization. IEEE Transactions on Knowledge and Data Engineering (2022)
38. Wang, R., Yi, M., Chen, Z., Zhu, S.: Out-of-distribution generalization with causal invariant transformations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 375–385 (2022)
39. Wang, T., Zhou, C., Sun, Q., Zhang, H.: Causal attention for unbiased visual recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3091–3100 (2021)
40. Woo, S., Park, J., Lee, J.Y., Kweon, I.S.: CBAM: convolutional block attention module. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 3–19 (2018)
41. Xu, Y., Li, Y.L., Li, J., Lu, C.: Constructing balance from imbalance for long-tailed image recognition. arXiv preprint [arXiv:2208.02567](https://arxiv.org/abs/2208.02567) (2022)
42. Yue, Z., Sun, Q., Hua, X.S., Zhang, H.: Transporting causal mechanisms for unsupervised domain adaptation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 8599–8608 (2021)
43. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: beyond empirical risk minimization. arXiv preprint [arXiv:1710.09412](https://arxiv.org/abs/1710.09412) (2017)
44. Zhang, X., Zhou, L., Xu, R., Cui, P., Shen, Z., Liu, H.: Nico++: Towards better benchmarking for domain generalization. ArXiv abs/2204.08040 (2022)