

# Domain Generalisation via Risk Distribution Matching

Toan Nguyen, Kien Do, Bao Duong, Thin Nguyen  
 Applied Artificial Intelligence Institute, Deakin University, Australia  
 {s222165627, k.do, duongng, thin.nguyen}@deakin.edu.au

## Abstract

We propose a novel approach for domain generalisation (DG) leveraging risk distributions to characterise domains, thereby achieving domain invariance. In our findings, risk distributions effectively highlight differences between training domains and reveal their inherent complexities. In testing, we may observe similar, or potentially intensifying in magnitude, divergences between risk distributions. Hence, we propose a compelling proposition: Minimising the divergences between risk distributions across training domains leads to robust invariance for DG. The key rationale behind this concept is that a model, trained on domain-invariant or stable features, may consistently produce similar risk distributions across various domains. Building upon this idea, we propose **Risk Distribution Matching (RDM)**. Using the maximum mean discrepancy (MMD) distance, RDM aims to minimise the variance of risk distributions across training domains. However, when the number of domains increases, the direct optimisation of variance leads to linear growth in MMD computations, resulting in inefficiency. Instead, we propose an approximation that requires only one MMD computation, by aligning just two distributions: that of the worst-case domain and the aggregated distribution from all domains. Notably, this method empirically outperforms optimising distributional variance while being computationally more efficient. Unlike conventional DG matching algorithms, RDM stands out for its enhanced efficacy by concentrating on scalar risk distributions, sidestepping the pitfalls of high-dimensional challenges seen in feature or gradient matching. Our extensive experiments on standard benchmark datasets demonstrate that RDM shows superior generalisation capability over state-of-the-art DG methods.

## 1. Introduction

In recent years, deep learning (DL) models have witnessed remarkable achievements and demonstrated super-human performance on training distributions [27]. Nonetheless, this success is accompanied by a caveat - deep

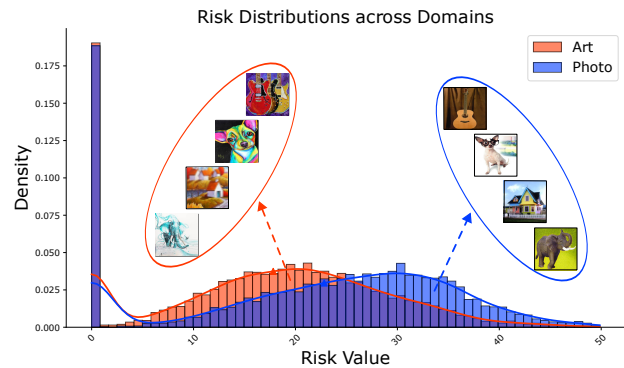


Figure 1. Risk distributions derived from training with ERM for the “Art” and “Photo” domains on the validation set of PACS dataset. Beyond low-risk samples, which may resemble training data, the “Photo” domain generally exhibits a larger distribution of risk values compared to “Art”, hinting at an inherent complexity in learning “Photo” samples. The figure indicates our motivation that *risk distributions* can effectively highlight differences between domains.

models are vulnerable to distributional shifts and exhibit catastrophic failures to unseen *out-of-domain* data [12, 34]. Such limitations hinder the widespread deployment of DL systems in real-world applications, where *domain difference* can be induced by several factors, such as spurious correlations [2] or variations in location or time [49].

In light of these challenges, domain generalisation (DG) aims to produce models capable of generalising to *unseen* target domains by leveraging data from diverse sets of training domains or environments [38]. An effective approach involves exploring and establishing *domain invariance* [32], with the expectation that these invariances will similarly apply to related, yet distinct, test domains. To this end, prevailing research focuses on characterising domains through sample representation [31, 38]. The objective is to seek for domain-invariant features by aligning the distributions of hidden representations across various domains. CORAL [53] trains a non-linear transformation that can align the second-order statistics of representations across different layers within deep networks. More, CausIRL [11]

aims to match representation distributions that have been intervened upon the spurious factors. While these methods show promise, they can face multiple challenges with the curse of dimensionality [6, 22]. The sparsity of high-dimensional representation spaces can lead to unreliable estimates of statistical properties, which in turn affects the quality of distribution matching techniques. Also, high-dimensional representations may contain many irrelevant or redundant dimensions, which can introduce noise to the true underlying similarities or differences between distributions. As dimensionality rises, computational complexity intensifies, reducing the efficacy of these methods [39]. Such challenges similarly present in DG methods that utilise gradients for domain alignment [46, 50].

In this paper, we propose to utilise *scalar risk distributions* as a means to characterise domains, leading to successfully exploring and enforcing domain invariance. Our research reveals that risk distributions can be a reliable indicator of *domain variation* as it effectively highlights differences between training domains. In Figure 1, we present a visual evidence through histograms, contrasting the risk distributions between the “Art” and “Photo” domains on the validation set of PACS dataset [30], derived from training with Empirical Risk Minimisation (ERM) [56]. The “Photo” domain generally exhibits a larger distribution of scalar risks than that of “Art”. This suggests an inherent complexity in learning “Photo” samples, or possibly due to a more limited training dataset compared to “Art”. During the testing phase, similar divergences between risk distributions may emerge, potentially intensifying in magnitude. Hence, we propose a compelling proposition: by *minimising the divergences between risk distributions across training domains*, we can achieve robust invariance for DG. The underlying rationale for this concept is that a model, when learning domain-invariant and stable features, tends to produce consistent risk distributions across domains.

Building upon this idea, we propose a novel matching approach for DG, namely *Risk Distribution Matching* (RDM). RDM’s objective is to minimise the *variance of risk distributions* across all training domains. Inspired by [38], we redefine the distributional variance metric to focus specifically on risk distributions and propose to compute it via the maximum mean discrepancy (MMD) distance [19]. However, when the number of training domains increases, directly optimising the variance induces a linear growth in MMD computations, reducing efficiency. Instead, we propose an approximation that requires only *one MMD computation* via aligning just two distributions: that of the *worst-case* (or worst-performing) domain and the aggregated distribution from all domains. Empirically, this approach outperforms optimising distributional variance while significantly reducing computational complexity. Unlike prevailing matching algorithms, RDM can ad-

dress the high-dimensional challenges and further improve efficacy by exclusively focusing on scalar risk distributions. Notably, our empirical studies show that RDM even exhibits enhanced generalisation while being more convenient to optimise. We summarise our contributions below:

- We propose RDM, a novel and efficient matching method for DG, based on our two hypotheses: i) risk distribution disparities offer insightful cues into domain variation; ii) reducing these divergences fosters a generalisable and invariant feature-learning predictor.
- We re-conceptualise the distributional variance metric to exclusively focus on risk distributions, with an objective to minimise it. We further provide an approximate version that aligns only the risk distribution of the worst-case domain with the aggregate from all domains, improving both performance and efficiency.
- Through extensive experiments on standard benchmark datasets, we empirically show that RDM consistently outperforms state-of-the-art DG methods, showcasing its remarkable generalisation capability.

## 2. Related Work

**Domain Generalisation (DG)** DG aims to develop models that can generalise well on unseen target domains by leveraging knowledge from multiple source domains. Typical DG methods include domain alignment [7, 32, 38], meta learning [4, 29], data augmentation [60, 62], disentangled representation learning [44, 54], robust optimisation [8, 48] and causality-based methods [13, 25, 41]. Our proposed method RDM is related to domain alignment, striving for *domain invariance* to enhance OOD generalisation. Existing research focuses on characterising domains through sample representations and aligning their distributions across domains to achieve domain-invariant features [1, 38]. CORAL [53] matches mean and variance of representation distributions, while MMD-AAE [31] and FedKA [55] consider matching all moments via the maximum mean discrepancy (MMD) distance [19]. Other methods promote domain invariance by minimising contrastive loss [10] between representations sharing the same labels [35, 37]. Many studies bypass the representation focus, instead characterising domains via gradients and achieving invariance by reducing inter-domain gradient variance [46, 50, 59].

Despite their potential, aligning these high-dimensional distributions may be affected by data sparsity, diversity, and high computational demands [6, 22]. Unlike these methods, RDM offers enhanced efficacy by focusing on *scalar risk distributions*, overcoming the high-dimensional challenges. Further, RDM adopts a novel strategy by efficiently aligning only two distributions: that of the worst-case domain

with the aggregate from all domains. From our experiments, RDM generally exhibits better generalisation performance while being more convenient to optimise compared to competing matching techniques. To the best of our knowledge, the incarnation of risk distributions for domain matching in RDM is novel and sensible.

**Distribution matching** Distribution matching has been an important topic with a wide range of applications in machine learning such as DG [31, 53], domain adaptation [9, 58], generative modelling [28, 33]. Early methods, like the MMD distance [19], leverage kernel-based approaches to quantify the distance between distributions, laying the foundation for many subsequent DG techniques [31, 38]. Further advancements have explored optimal transport methods, like the Wasserstein distance [3, 36], which provides a geometrically intuitive means to compare distributions. Other metrics, such as the Kullback-Leibler [26] or Jensen-Shannon [14, 16] divergences, can serve to measure the divergence between distributions and may require additional parameters for estimating the density ratio between the two distributions [52]. In this paper, we utilise the MMD distance to align risk distributions. Its inherent advantages include an analytical measure of the divergence between distributions without relying on distribution densities, and its non-parametric nature [19]. Alternative DG methods augment data by utilising distribution matching and style transfer to generate semantic-preserving samples [61, 62]. Our method differs as we emphasise *domain invariance via aligning risk distributions*, rather than augmenting representation distributions.

**Invariance and Causality in DG** Causal methods in DG assume that the causal mechanism of the target given causal input features is invariant while non-causal features may change across domains [2, 25, 45]. Based on this assumption, methods establish domain invariance to recover the causal mechanism, thereby improving generalisation. ICP [45] has shown that the causal predictor has an invariant distribution of residuals in regression models, however, is not suitable for deep learning. EQRM [13] and REx [25] leverage the invariance in the *average* risks over samples across domains. In contrast to above methods, we consider matching *entire risk distributions* over samples across domains, which, as our experiments demonstrate, is more powerful and enhances generalisation capability.

### 3. Preliminaries

Domain generalisation (DG) involves training a classifier  $f$  on data composed of multiple training domains (also called environments) so that  $f$  can perform well on unseen domains at test time. Mathematically, let  $\mathcal{D} =$

$\{\mathcal{D}_1, \dots, \mathcal{D}_m\}$  denote the training set consisting of  $m$  different domains/environments, and let  $\mathcal{D}_e := \{(x_e^i, y_e^i)\}_{i=1}^{n_e}$  denote the training data belonging to domain  $e$  ( $1 \leq e \leq m$ ). Given a loss function  $\ell$ , the risk of a particular domain sample  $(x_e^i, y_e^i)$  is denoted by  $R_e^i := \ell(f(x_e^i), y_e^i)$ , and the expected risk  $\bar{R}_e$  of domain  $e$  is defined as:

$$\bar{R}_e := \mathbb{E}_{(x_e, y_e) \sim \mathcal{D}_e} [\ell(f(x_e), y_e)] = \mathbb{E}_{\mathcal{D}_e} [R_e] \quad (1)$$

A common approach to train  $f$  is Empirical Risk Minimisation (ERM) [56] which minimises the expected risks across all training domains. Its loss function, denoted by  $\mathcal{L}_{\text{ERM}}$ , is computed as follows:

$$\mathcal{L}_{\text{ERM}} = \mathbb{E}_{e \sim \mathcal{E}} \mathbb{E}_{(x_e, y_e) \sim \mathcal{D}_e} [\ell(f(x_e), y_e)] \quad (2)$$

$$= \mathbb{E}_{e \sim \mathcal{E}} [\bar{R}_e] \quad (3)$$

where  $\mathcal{E} := \{1, \dots, m\}$  denotes the set of all domains.

### 4. Risk Distribution Matching

A model  $f$  trained via ERM often struggles with generalisation to new test domains. This is because it tends to capture domain-specific features [2, 41], such as domain styles, to achieve low risks in training domains, rather than focusing on domain-invariant or semantic features. To overcome this issue, we present a novel training objective that bolsters generalisation through *domain invariance*. Our goal requires utilising a unique domain representative that both characterises each domain and provides valuable insights into domain variation. Specifically, we propose to leverage the *distribution of risks over all samples within a domain* (or shortly *risk distribution*) as this representative. Unlike other domain representatives, like latent representation or gradient distributions [31, 50], the risk distribution sidesteps high-dimensional challenges like data sparsity and high computational demands [6, 39]. In essence, a model capturing stable, domain-invariant features may consistently yield similar risk distributions across all domains. In pursuit of invariant models, we propose **Risk Distribution Matching** (RDM), a novel approach for DG that reduces the divergences between training risk distributions via *minimising the distributional variance across them*.

Let  $\mathcal{T}_e$  be the probability distribution over the risks of all samples in domain  $e$  (i.e.,  $\{R_e^i\}_{i=1}^{n_e}$ ). We refer to  $\mathcal{T}_e$  as the *risk distribution* of domain  $e$ , the representative that effectively captures the core characteristics of the domain. We denote  $\mathbb{V}_{\mathbb{R}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$  the distributional variance across the risk distributions  $\{\mathcal{T}_1, \dots, \mathcal{T}_m\}$  in the real number space. We achieve our objective by minimising the following loss function:

$$\mathcal{L}_{\text{final}} := \mathcal{L}_{\text{ERM}} + \lambda \mathbb{V}_{\mathbb{R}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) \quad (4)$$

where  $(\lambda \geq 0)$  is a coefficient balancing between reducing the total training risks with enforcing invariance across domains.  $\lambda$  is set to 1 unless specified otherwise.

To compute  $\mathbb{V}_{\mathbb{R}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$ , we require a suitable representation for the *implicit* risk distribution  $\mathcal{T}_e$  of domain  $e$ . Leveraging kernel mean embedding [51], we express  $\mathcal{T}_e$  as its embedding,  $\mu_{\mathcal{T}_e}$ , within a reproducing kernel Hilbert space (RKHS)  $\mathcal{H}$  using a feature map  $\phi : \mathbb{R} \rightarrow \mathcal{H}$  below:

$$\mu_{\mathcal{T}_e} := \mathbb{E}_{R_e \sim \mathcal{T}_e} [\phi(R_e)] \quad (5)$$

$$= \mathbb{E}_{R_e \sim \mathcal{T}_e} [k(R_e, \cdot)] \quad (6)$$

where a kernel function  $k(\cdot, \cdot) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  is introduced to bypass the explicit specification of  $\phi$ . Assuming the condition  $(\mathbb{E}_{R_e \sim \mathcal{T}_e} (k(R_e, R_e)) < \infty)$  holds, the mean map  $\mu_{\mathcal{T}_e}$  remains an element of  $\mathcal{H}$  [19, 31]. It is noteworthy that for a *characteristic* kernel  $k$ , the representation  $\mu_{\mathcal{T}_e}$  within  $\mathcal{H}$  is unique [19, 38]. Consequently, two distinct risk distributions  $\mathcal{T}_u$  and  $\mathcal{T}_v$  for any domains  $u, v$  respectively have different kernel mean embeddings in  $\mathcal{H}$ . In this work, we use the RBF kernel, a well-known characteristic kernel defined as  $k(x, x') := \exp\left(-\frac{1}{2\sigma} \|x - x'\|^2\right)$ , where  $\sigma > 0$  is the bandwidth parameter.

With the unique representation of  $\mathcal{T}_e$  established, our objective becomes computing the distributional variance between risk distributions within  $\mathcal{H}$ , represented by  $\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$ . Inspired by [38], we redefine the variance metric to focus specifically on risk distributions across multiple domains below:

$$\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) := \frac{1}{m} \sum_{e=1}^m \|\mu_{\mathcal{T}_e} - \mu_{\mathcal{T}}\|_{\mathcal{H}}^2 \quad (7)$$

where  $\mathcal{T} = \frac{1}{m} \sum_{e=1}^m \mathcal{T}_e$  denotes the probability distribution over the risks of all samples in the entire training set, or equivalently, the set of all  $m$  domains. Meanwhile,  $\mu_{\mathcal{T}_e}$  and  $\mu_{\mathcal{T}}$  represent the mean embeddings of  $\mathcal{T}_e$  and  $\mathcal{T}$ , respectively, and are computed as in Eq. 5. Incorporating  $\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$  into our loss function from Eq. 4, we get:

$$\mathcal{L}_{\text{final}} := \mathcal{L}_{\text{ERM}} + \lambda \mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) \quad (8)$$

Minimising  $\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$  in Eq. 8 facilitates our objective of equalising risk distributions across all domains, as proven by the theorem below.

**Theorem 1.** [38] *Given the distributional variance  $\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\})$  is calculated with a characteristic kernel  $k$ ,  $\mathbb{V}_{\mathcal{H}}(\{\mathcal{T}_1, \dots, \mathcal{T}_m\}) = 0$  if and only if  $\mathcal{T}_1 = \dots = \mathcal{T}_m (= \mathcal{T})$ .*

*Proof.* Please refer to our appendix for the proof.  $\square$

In the next part, we present how to compute the distributional variance using the Maximum Mean Discrepancy (MMD) distance [19], relying only on risk samples. Then, we propose an efficient approximation of optimising the distributional variance, yielding improved empirical performance.

#### 4.1. Maximum Mean Discrepancy

For domain  $e$ , the squared norm,  $\|\mu_{\mathcal{T}_e} - \mu_{\mathcal{T}}\|_{\mathcal{H}}^2$ , defined in Eq. 7, is identified as the squared MMD distance [18] between distributions  $\mathcal{T}_e$  and  $\mathcal{T}$ . It is expressed as follows:

$$\text{MMD}^2(\mathcal{T}_e, \mathcal{T}) = \|\mu_{\mathcal{T}_e} - \mu_{\mathcal{T}}\|_{\mathcal{H}}^2 \quad (9)$$

$$= \|\mathbb{E}_{R_e \sim \mathcal{T}_e} [\phi(R_e)] - \mathbb{E}_{R_f \sim \mathcal{T}} [\phi(R_f)]\|_{\mathcal{H}}^2 \quad (10)$$

$$= \mathbb{E}_{R_e, R'_e \sim \mathcal{T}_e} \langle \phi(R_e), \phi(R'_e) \rangle - 2\mathbb{E}_{R_e \sim \mathcal{T}_e; R_f \sim \mathcal{T}} \langle \phi(R_e), \phi(R_f) \rangle \quad (11)$$

$$+ \mathbb{E}_{R_f, R'_f \sim \mathcal{T}} \langle \phi(R_f), \phi(R'_f) \rangle$$

where  $\langle \cdot, \cdot \rangle$  denote the inner product operation in  $\mathcal{H}$ . Through the kernel trick, we can compute these inner products via the kernel function  $k$  without an explicit form of  $\phi$  below:

$$\text{MMD}^2(\mathcal{T}_e, \mathcal{T}) = \mathbb{E}_{R_e, R'_e \sim \mathcal{T}_e} k(R_e, R'_e) - 2\mathbb{E}_{R_e \sim \mathcal{T}_e; R_f \sim \mathcal{T}} k(R_e, R_f) \quad (12)$$

$$+ \mathbb{E}_{R_f, R'_f \sim \mathcal{T}} k(R_f, R'_f)$$

We reformulate our loss function in Eq. 8 to incorporate MMD as follows:

$$\mathcal{L}_{\text{final}} := \mathcal{L}_{\text{ERM}} + \frac{\lambda}{m} \sum_{e=1}^m \text{MMD}^2(\mathcal{T}_e, \mathcal{T}) \quad (13)$$

$$= \mathcal{L}_{\text{ERM}} + \lambda \mathcal{L}_{\text{RDM}} \quad (14)$$

The loss function  $\mathcal{L}_{\text{RDM}}$  involves minimising  $\text{MMD}^2(\mathcal{T}_e, \mathcal{T})$  for every domain  $e$ . Ideally, the distributional variance reaches its lowest value at 0 if  $\text{MMD}(\mathcal{T}_e, \mathcal{T}) = 0$ , equivalent to  $(\mathcal{T}_e = \mathcal{T})$  [18, 19], across  $e$  domains. The objective also entails aligning each individual risk distribution,  $\mathcal{T}_e$ , with the aggregated distribution spanning all domains,  $\mathcal{T}$ . With the characteristic RBF kernel, it can be viewed as *matching an infinite number of moments* across all risk distributions.

We emphasise our choice of MMD owing to its benefits for effective risk distribution matching: i) MMD is an



important member of the Integral Probability Metric family [40] that offers an analytical solution facilitated through RKHS, and ii) MMD enjoys the property of quantifying the dissimilarity between two implicit distributions via their finite samples in a non-parametric manner.

## 4.2. Further improvement of RDM

We find that effective alignment of risk distributions across  $m$  domains can be achieved by matching the risk distribution of the *worst-case* domain, denoted as  $w$ , with the combined risk distribution of all domains, offering an approximation to the optimisation of risk distributional variance seen in Eq. 13. This approximate version significantly reduces the MMD distances computation in  $\mathcal{L}_{\text{RDM}}$  from  $O(m)$  to  $O(1)$ , and further improves generalisation, as we demonstrate with empirical evidence in Section 5.

Denote by  $\left(w = \underset{e \in \mathcal{E}}{\operatorname{argmax}} \bar{R}_e\right)$  the worst-case domain, i.e., the domain that has the largest expected risk in  $\mathcal{E}$ . The approximate RDM’s loss,  $\hat{\mathcal{L}}_{\text{RDM}}$ , is computed as follows:

$$\hat{\mathcal{L}}_{\text{RDM}} = \text{MMD}^2(\mathcal{T}_w, \mathcal{T}) \quad (15)$$

$$\approx \mathcal{L}_{\text{RDM}} \quad (16)$$

In our experiments, we observed only a small gap between  $\hat{\mathcal{L}}_{\text{RDM}}$  and  $\mathcal{L}_{\text{RDM}}$ , while optimising  $\hat{\mathcal{L}}_{\text{RDM}}$  proving to be more computationally efficient. The key insight emerges from  $\bar{R}_e$ , the first moment (or mean) of  $\mathcal{T}_e$ . Often, the average risk can serve as a measure of domain uniqueness or divergence [25, 46]. Specifically, a domain with notably distinct mean risk is more likely to diverge greatly from other risk distributions. Under such circumstances,  $\hat{\mathcal{L}}_{\text{RDM}}$  will be an upper-bound of  $\mathcal{L}_{\text{RDM}}$ , as shown by:  $\mathcal{L}_{\text{RDM}} = \frac{1}{m} \sum_{e=1}^m \text{MMD}^2(\mathcal{T}_e, \mathcal{T}) \leq \frac{1}{m} \sum_{e=1}^m \text{MMD}^2(\mathcal{T}_w, \mathcal{T}) = \text{MMD}^2(\mathcal{T}_w, \mathcal{T}) = \hat{\mathcal{L}}_{\text{RDM}}$ . By optimising  $\hat{\mathcal{L}}_{\text{RDM}}$ , we can also potentially decrease  $\mathcal{L}_{\text{RDM}}$ , thus aligning risk distributions across domains effectively. More,  $\hat{\mathcal{L}}_{\text{RDM}}$  drives the model to prioritise the worst-case domain’s optimisation. This approach enhances the model’s robustness to extreme training scenarios, which further improves generalisation as proven in [25, 48]. These insights shed light on the superior performance of optimising  $\hat{\mathcal{L}}_{\text{RDM}}$  over  $\mathcal{L}_{\text{RDM}}$ . Therefore, we opted to use  $\hat{\mathcal{L}}_{\text{RDM}}$ , simplifying our model’s training and further bolstering its OOD performance.

## 5. Experiments

We evaluate and analyse RDM using a synthetic ColoredMNIST dataset [2] and multiple benchmarks from the DomainBed suite [20]. Each of our claims is backed by empirical evidence in this section. Our source code to reproduce results is available at: <https://github.com/nktoan/risk-distribution-matching>

Algorithm	Initialisation	
	Rand.	ERM
ERM	27.9±1.5	27.9±1.5
GroupDRO	27.3±0.9	29.0±1.1
IGA	50.7±1.4	57.7±3.3
IRM	52.5±2.4	69.7±0.9
VREx	55.2±4.0	71.6±0.5
EQRM	53.4±1.7	71.4±0.4
CORAL	55.3±2.8	65.6±1.1
MMD	54.6±3.2	66.4±1.7
<b>RDM (ours)</b>	<b>56.3±1.5</b>	<b>72.4±1.0</b>
Oracle	72.1±0.7	
Optimum	75.0	

Table 1. ColoredMNIST test accuracy where the best results are marked as bold. Results of other methods are referenced from [13].

### 5.1. Synthetic Dataset: ColoredMNIST

We evaluate all baselines on a synthetic binary classification task, namely ColoredMNIST [2]. This dataset involves categorising digits (0-9) into two labels: “zero” for 0 to 4 range and “one” for 5 to 9 range, with each digit colored either red or green. The dataset is designed to assess the generalisation and robustness of baseline models against the influence of spurious color features. The dataset contains two training domains, where the chance of red digits being classified as “zero” is 80% and 90%, respectively, while this probability decreases to only 10% during testing. The goal is to train a predictor invariant to “digit color” features, capturing only “digit shape” features.

Following [13], we employ a two-hidden-layer MLP with 390 hidden units for all baselines. Optimised through the Adam optimiser [23] at a learning rate of 0.0001, with a dropout rate of 0.2, we train each algorithm for 600 iterations with a batch size of 25,000. We repeat the experiment ten times over different values of the penalty weight  $\lambda$ . We find our matching penalty quite small, yielding optimal RDM’s performance within the  $\lambda$  range of [1000, 10000]. We provide more details about experimental settings in the supplementary material.

We compare RDM with ERM and three different types of algorithms: robust optimisation (GroupDRO [48], IGA [24]), causal methods learning invariance (IRM [2], VREx [25], EQRM [13]) and representation distribution matching (MMD [31], CORAL [53]). All algorithms are run using two distinct network configurations: (i) initialising the network randomly via Xavier method [17]; (ii) pre-training the network with ERM for 400 iterations prior to performing the algorithms. Table 1 shows that our pro-




Figure 2. Histograms with their KDE curves depicting the risk distributions of ERM and RDM across four domains on ColoredMNIST. Vertical ticks denote the mean values of all distributions.

posed method RDM surpasses all algorithms, irrespective of the network configuration. RDM exhibits improvements of 1.0% and 6.8% over CORAL, both without and with pre-trained ERM, respectively, underlining the effectiveness of aligning risk distributions instead of high-dimensional representations. VREx and EQRm, which pursue invariant predictors by equalising average training risks across domains, demonstrate suboptimal performance compared to our approach. This improvement arises from our consideration of the entire risk distributions and the matching of all moments across them, which inherently foster stronger invariance for DG. Notably, all methods experience enhanced performance with ERM initialisation. RDM even excels beyond oracle performance (ERM trained on grayscale digits with 50% red and 50% green) and converges towards optimality.

Figure 2 demonstrates histograms with their KDE curves [42] depicting the risk distributions of ERM and RDM across four domains. The figure confirms our hypothesis that the disparities among risk distributions could serve as a valuable signal of *domain variation*. ERM’s histogram shows a clear difference between environments with 90% and 80% chance of red digits labelled “zero” and those with only 50% or 10%. More, ERM tends to overfit to training domains, which negatively impacts its generalisation to test domains. Remarkably, RDM effectively minimises the divergences between risk distributions across all domains, including *test domains with lower risks*. This also aligns with our motivation: an invariant or stable feature-learning predictor, by displaying similar risk distributions across domains, inherently boosts generalisation.

## 5.2. DomainBed

**Dataset and Protocol** Following previous works [13,20], we extensively evaluate all methods on five well-known DG benchmarks: VLCS [15], PACS [30], OfficeHome [57],

TerraIncognita [5], and DomainNet [43]. For a fair comparison, we reuse the training and evaluation protocol in DomainBed [20], including the dataset splits, training iterations, and model selection criteria. Our evaluation employs the leave-one-domain-out approach: each model is trained on all domains except one and then tested on the excluded domain. The final model is chosen based on its combined accuracy across all training-domain validation sets.

**Implementation Details** We use ResNet-50 [21] pre-trained on ImageNet [47] as the default backbone. The model is optimised via the Adam optimiser for 5,000 iterations on every dataset. We follow [13,25] to pre-train baselines with ERM for certain iterations before performing the algorithms. Importantly, we find that achieving accurate risk distribution matching using distribution samples requires larger batch sizes - details of which are examined in our ablation studies. For most datasets, the optimal batch size lies between [70, 100]. However, for huge datasets like TerraIncognita and DomainNet, it is between [30, 60]. Although computational resources limit us from testing larger batch sizes, these ranges consistently achieve strong performance on benchmarks. The matching coefficient  $\lambda$  in our method is set in [0.1, 10.0]. Additional hyper-parameters like learning rate, dropout rate, or weight decay, adhere to the preset ranges as detailed in [13]. We provide more implementation details in the supplementary material. We repeat our experiments ten times with varied seed values and hyper-parameters and report the average results.


**Experimental Results** In Table 2, we show the average out-of-domain (OOD) accuracies of state-of-the-art DG methods on five benchmarks. Due to space constraints, domain-specific accuracies are detailed in the supplementary material. We compare RDM with ERM and various types of algorithms: distributional robustness (GroupDRO),

Algorithm	VLCS	PACS	OfficeHome	TerraIncognita	DomainNet	Avg
ERM	77.5±0.4	85.5±0.2	66.5±0.3	46.1±1.8	40.9±0.1	63.3
Mixup	77.4±0.6	84.6±0.6	68.1±0.3	<b>47.9±0.8</b>	39.2±0.1	63.4
MLDG	77.2±0.4	84.9±1.0	66.8±0.6	47.7±0.9	41.2±0.1	63.6
GroupDRO	76.7±0.6	84.4±0.8	66.0±0.7	43.2±1.1	33.3±0.2	60.9
IRM	78.5±0.5	83.5±0.8	64.3±2.2	47.6±0.8	33.9±2.8	61.6
VREx	78.3±0.2	84.9±0.6	66.4±0.6	46.4±0.6	33.6±2.9	61.9
EQRM	77.8±0.6	86.5±0.2	67.5±0.1	47.8±0.6	41.0±0.3	64.1
Fish	77.8±0.3	85.5±0.3	68.6±0.4	45.1±1.3	42.7±0.2	64.0
Fishr	77.8±0.1	85.5±0.4	67.8±0.1	47.4±1.6	41.7±0.0	64.0
CORAL	<b>78.8±0.6</b>	86.2±0.3	<b>68.7±0.3</b>	47.6±1.0	41.5±0.1	64.6
MMD	77.5±0.9	84.6±0.5	66.3±0.1	42.2±1.6	23.4±9.5	63.3
RDM (ours)	78.4±0.4	<b>87.2±0.7</b>	67.3±0.4	47.5±1.0	<b>43.4±0.3</b>	<b>64.8</b>

Table 2. DomainBed test accuracy where the best results are marked as bold. Results of other methods are referenced from [13, 50]. Model selection: training-domain validation set.

causal methods learning invariance (IRM, VREx, EQRM), gradient matching (Fish [50], Fishr [46]), representation distribution matching (MMD, CORAL) and other variants (Mixup [60], MLDG [29]). To ensure fairness in our evaluations, we have used the same training data volume across all baselines, although further employing augmentations can enhance models’ performance.

On average, RDM surpasses other baselines across all benchmarks, notably achieving a 1.5% average improvement over ERM. The significant improvement of RDM on DomainNet, a large-scale dataset with 586,575 images across 6 domains, is worth mentioning. This suggests that characterising domains with risk distributions to achieve invariance effectively enhances OOD performance. Compared to distributional robustness methods, RDM notably outperforms GroupDRO with improvement of 2.8% on PACS and a substantial 10.1% on DomainNet. RDM consistently improves over causality-based methods that rely on the average risk for domain invariance. This superiority attributes to our novel adoption of risk distributions, achieving enhanced invariance for DG. Our remarkable improvement over MMD suggests that aligning *risk distributions* via the MMD distance is more effective, easier to optimise than aligning representation distributions. While RDM typically outperforms CORAL and Fish in OOD scenarios, it only remains competitive or sometimes underperforms on certain datasets like OfficeHome. This decrease in performance may stem from the dataset’s inherent tendency to overfit within our risk distribution alignment objective. OfficeHome has only average about 240 samples per class, significantly fewer than other datasets with at least 1,400. This reduced sample size may not provide sufficiently diverse risk distributions to capture stable class features, resulting in overfitting on the training set. Despite these limitations,



(a) A small gap between two variants of RDM. (b) Learning curves of RDM and other baselines.

Figure 3a supports our claim about the approximation of the distributional variance, while Figure 3b compares the OOD performance learning curves of RDM with other methods. These insights are visualised every 15 iterations during PACS dataset training, excluding the OOD Sketch domain. After RDM is pre-trained with ERM for 100(×15) iterations, our visual analysis commences, ensuring a fair comparison.

our OfficeHome results still outperform several well-known baselines such as MLDG, VREx, or ERM. For a detailed discussion on this challenge, please refer to our supplementary material.

### 5.3. Analysis

In this section, we provide empirical evidence backing our claims in Section 4. In Figure 3a, we highlight a small gap when aligning the risk distribution of the worst-case domain with that of all domains combined (RDM with  $\hat{\mathcal{L}}_{\text{RDM}}$ ), compared to directly optimising the distributional variance (RDM with  $\mathcal{L}_{\text{RDM}}$ ). Notably,  $\hat{\mathcal{L}}_{\text{RDM}}$  consistently represents an upper bound of  $\mathcal{L}_{\text{RDM}}$ , which is sensible since the worst-case domain often exhibits the most distinct risk distribution. This suggests that optimising  $\hat{\mathcal{L}}_{\text{RDM}}$  also helps reduce the distributional variance  $\mathcal{L}_{\text{RDM}}$ , bringing the risk distributions across domains closer.

Algorithm	Training (s)	Mem (GiB)	Acc (%)
Fish	11,502	<b>5.26</b>	42.7
CORAL	11,504	17.00	41.5
RDM with $\mathcal{L}_{\text{RDM}}$	9,854	16.94	43.1
RDM with $\hat{\mathcal{L}}_{\text{RDM}}$	<b>7,749</b>	16.23	<b>43.4</b>

Table 3. Comparison between Fish, CORAL, and two variants of our method in terms of the training time (seconds), memory usage per iteration (GiB) and accuracy (%) on DomainNet.

When the number of training domains grows, especially with large-scale datasets like DomainNet, emphasising the risk distribution of the worst-case domain not only proves to be a more efficient approach but also significantly enhances OOD performance. In our exploration of training resources for DomainNet, we study three matching methods: Fish, CORAL and two variants of our RDM method. For a fair evaluation, all experiments were conducted with identical GPU resources, settings, and hyper-parameters, such as batch size or training iterations. Results can be seen in Table 3. Full details on training resources for these methods on other datasets are available in the supplementary material due to space constraints.

Our RDM with the  $\hat{\mathcal{L}}_{\text{RDM}}$  objective proves fastest in training and achieves the notably highest 43.4% accuracy on DomainNet. While RDM demands more memory than Fish, due to the storage of MMD distance values, it can be trained in less time - under an hour - and still delivers a 0.7% performance boost. This gain over Fish, a leading gradient matching method on DomainNet, is significant. Among two variants of RDM, the one using  $\hat{\mathcal{L}}_{\text{RDM}}$  is both the fastest and most accurate, justifying our claims on the benefits of aligning the risk distribution of the worst-case domain.

More, to further highlight the efficacy of risk distribution alignment for DG, we compare the OOD performance learning curves of RDM with competing baselines using representation (CORAL) and gradient (Fish) alignments, as depicted in Figure 3b. Impressively, RDM consistently outperforms, demonstrating enhanced generalisation throughout the training process.

#### 5.4. Ablation studies

We explore the impact of the matching coefficient  $\lambda$  and training batch size on risk distribution matching, using primarily the PACS dataset for brevity. While other datasets exhibit similar trends, their detailed results are provided in the supplementary material.

**Matching coefficient  $\lambda$**  Figure 4a illustrates the performance of RDM on the PACS dataset for vary-

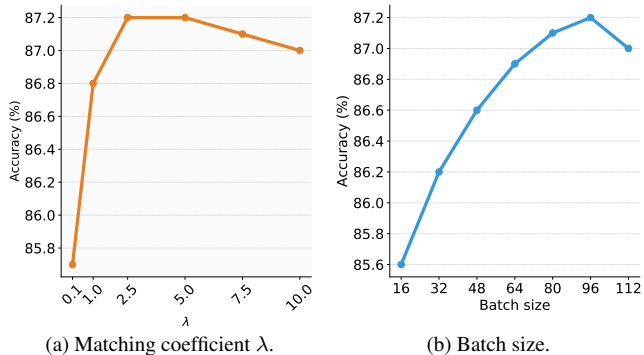


Figure 4. Ablation studies on the effects of the matching coefficient  $\lambda$  and the training batch size on the PACS dataset.

ing values of the matching coefficient  $\lambda$ , spanning  $\{0.1, 1.0, 2.5, 5.0, 7.5, 10.0\}$ . Notably, as  $\lambda$  increases, RDM’s accuracy consistently improves, justifying the significance of our risk distribution matching module in fostering generalisation. In particular, when  $\lambda = 5.0$ , RDM demonstrates a notable 1.6% average accuracy boost across all domains, in contrast to when using only  $\lambda = 0.1$ . Across most datasets, a  $\lambda$  value within  $[0.1, 10.0]$  appears sufficient to produce good results.

**Batch size** We study the impact of batch size on RDM’s performance. Our assumption is that achieving accurate risk distribution matching through data samples would require larger batch sizes. Figure 4b validates this, revealing enhanced generalisation results on PACS with increased batch sizes. For PACS, sizes between  $[70, 100]$  yield promising, potentially optimal outcomes, despite computational limitations restrict our exploration of larger sizes.

## 6. Conclusion

We have demonstrated that RDM, a novel matching method for domain generalisation (DG), provides enhanced generalisation capability by aligning risk distributions across domains. RDM efficiently overcomes high-dimensional challenges of conventional DG matching methods. RDM is built on our observation that risk distributions can effectively represent the differences between training domains. By minimising these divergences, we can achieve an invariant and generalisable predictor. We further improve RDM by matching only the risk distribution of the worst-case domain with the aggregate from all domains, bypassing the need to directly compute the distributional variance. This approximate version not only offers computational efficiency but also delivers improved out-of-domain results. Our extensive experiments on several benchmarks reveal that RDM surpasses leading DG techniques. We hope our work can inspire further investigations into the benefits of risk distributions for DG.



## References

- [1] Isabela Albuquerque, João Monteiro, Mohammad Darvishi, Tiago H Falk, and Ioannis Mitliagkas. Generalizing to unseen domains via distribution matching. *arXiv preprint arXiv:1911.00804*, 2019. [2](#)
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. [1](#), [3](#), [5](#)
- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, pages 214–223, 2017. [3](#)
- [4] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *NIPS*, 31, 2018. [2](#)
- [5] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *ECCV*, pages 456–473, 2018. [6](#)
- [6] Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966. [2](#), [3](#)
- [7] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79:151–175, 2010. [2](#)
- [8] Peter Bühlmann. Invariance, causality and robustness. 2020. [2](#)
- [9] Chaoqi Chen, Weiping Xie, Wenbing Huang, Yu Rong, Xinghao Ding, Yue Huang, Tingyang Xu, and Junzhou Huang. Progressive feature alignment for unsupervised domain adaptation. In *CVPR*, pages 627–636, 2019. [3](#)
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *ICML*, pages 1597–1607, 2020. [2](#)
- [11] Mathieu Chevalley, Charlotte Bunne, Andreas Krause, and Stefan Bauer. Invariant causal mechanisms through distribution matching. *arXiv preprint arXiv:2206.11646*, 2022. [1](#)
- [12] Andrea Dittadi, Frederik Träuble, Francesco Locatello, Manuel Wüthrich, Vaibhav Agrawal, Ole Winther, Stefan Bauer, and Bernhard Schölkopf. On the transfer of disentangled representations in realistic settings. *ICLR*, 2021. [1](#)
- [13] Cian Eastwood, Alexander Robey, Shashank Singh, Julius Von Kügelgen, Hamed Hassani, George J Pappas, and Bernhard Schölkopf. Probable domain generalization via quantile risk minimization. *NeurIPS*, 35:17340–17358, 2022. [2](#), [3](#), [5](#), [6](#), [7](#)
- [14] Dominik Maria Endres and Johannes E Schindelin. A new metric for probability distributions. *IEEE Transactions on Information theory*, pages 1858–1860, 2003. [3](#)
- [15] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *ICCV*, pages 1657–1664, 2013. [6](#)
- [16] Bent Fuglede and Flemming Topsøe. Jensen-shannon divergence and hilbert space embedding. In *ISIT*, page 31. IEEE, 2004. [3](#)
- [17] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *AISTATS*, pages 249–256, 2010. [5](#)
- [18] Arthur Gretton, Karsten Borgwardt, Malte Rasch, Bernhard Schölkopf, and Alex Smola. A kernel method for the two-sample-problem. *NIPS*, 19, 2006. [4](#)
- [19] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, pages 723–773, 2012. [2](#), [3](#), [4](#)
- [20] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *ICLR*, 2021. [5](#), [6](#)
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [6](#)
- [22] Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proc. Annu. ACM Symp. Theory Comput.*, pages 604–613, 1998. [2](#)
- [23] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015. [5](#)
- [24] Masanori Koyama and Shoichiro Yamaguchi. Out-of-distribution generalization with maximal invariant predictor. 2020. [5](#)
- [25] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*, pages 5815–5826, 2021. [2](#), [3](#), [5](#), [6](#)
- [26] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86, 1951. [3](#)
- [27] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015. [1](#)
- [28] Chun-Liang Li, Wei-Cheng Chang, Yu Cheng, Yiming Yang, and Barnabás Póczos. Mmd-gan: Towards deeper understanding of moment matching network. *NIPS*, 30, 2017. [3](#)
- [29] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *AAAI*, volume 32, 2018. [2](#), [7](#)
- [30] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *ICCV*, pages 5542–5550, 2017. [2](#), [6](#)
- [31] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *CVPR*, pages 5400–5409, 2018. [1](#), [2](#), [3](#), [4](#), [5](#)
- [32] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representations. In *AAAI*, volume 32, 2018. [1](#), [2](#)
- [33] Yujia Li, Kevin Swersky, and Rich Zemel. Generative moment matching networks. In *ICML*, pages 1718–1727, 2015. [3](#)
- [34] Chaochao Lu, Yuhuai Wu, José Miguel Hernández-Lobato, and Bernhard Schölkopf. Invariant causal representation learning for out-of-distribution generalization. In *ICLR*, 2021. [1](#)
- [35] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. In *ICML*, pages 7313–7324, 2021. [2](#)
- [36] Facundo Mémoli. Gromov–wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011. [3](#)

- [37] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *ICCV*, pages 5715–5725, 2017. [2](#)
- [38] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *ICML*, pages 10–18, 2013. [1](#), [2](#), [3](#), [4](#)
- [39] Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *TPAMI*, 36(11):2227–2240, 2014. [2](#), [3](#)
- [40] Alfred Müller. Integral probability metrics and their generating classes of functions. *Advances in applied probability*, 29(2):429–443, 1997. [5](#)
- [41] Toan Nguyen, Kien Do, Duc Thanh Nguyen, Bao Duong, and Thin Nguyen. Causal inference via style transfer for out-of-distribution generalisation. In *KDD*, pages 1746–1757, 2023. [2](#), [3](#)
- [42] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3):1065–1076, 1962. [6](#)
- [43] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *ICCV*, pages 1406–1415, 2019. [6](#)
- [44] Xingchao Peng, Zijun Huang, Ximeng Sun, and Kate Saenko. Domain agnostic learning with disentangled representations. In *ICML*, pages 5102–5112, 2019. [2](#)
- [45] Jonas Peters, Peter Bühlmann, and Nicolai Meinshausen. Causal inference by using invariant prediction: identification and confidence intervals. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 78(5):947–1012, 2016. [3](#)
- [46] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *ICML*, pages 18347–18377, 2022. [2](#), [5](#), [7](#)
- [47] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015. [6](#)
- [48] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. In *ICLR*, 2020. [2](#), [5](#)
- [49] Vaishal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do image classifiers generalize across time? In *ICCV*, pages 9661–9669, 2021. [1](#)
- [50] Yuge Shi, Jeffrey Seely, Philip Torr, Siddharth N, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. In *ICLR*, 2022. [2](#), [3](#), [7](#)
- [51] Alex Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A hilbert space embedding for distributions. In *International conference on algorithmic learning theory*, pages 13–31, 2007. [4](#)
- [52] Masashi Sugiyama, Taiji Suzuki, and Takafumi Kanamori. *Density ratio estimation in machine learning*. Cambridge University Press, 2012. [3](#)
- [53] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV Workshops*, pages 443–450. Springer, 2016. [1](#), [2](#), [3](#), [5](#)
- [54] Xinwei Sun, Botong Wu, Xiangyu Zheng, Chang Liu, Wei Chen, Tao Qin, and Tie-Yan Liu. Recovering latent causal factor for generalization to distributional shifts. *NeurIPS*, 34:16846–16859, 2021. [2](#)
- [55] Yuwei Sun, Ng Chong, and Hideya Ochiai. Feature distribution matching for federated domain generalization. In *ACML*, pages 942–957, 2023. [2](#)
- [56] Vladimir N Vapnik. An overview of statistical learning theory. *IEEE transactions on neural networks*, 10(5):988–999, 1999. [2](#), [3](#)
- [57] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. [6](#)
- [58] Jindong Wang, Wenjie Feng, Yiqiang Chen, Han Yu, Meiyu Huang, and Philip S Yu. Visual domain adaptation with manifold embedded distribution alignment. In *ACMMM*, pages 402–410, 2018. [3](#)
- [59] Pengfei Wang, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. Sharpness-aware gradient matching for domain generalization. In *CVPR*, pages 3769–3778, 2023. [2](#)
- [60] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *ICLR*, 2018. [2](#), [7](#)
- [61] Yabin Zhang, Minghan Li, Ruihuang Li, Kui Jia, and Lei Zhang. Exact feature distribution matching for arbitrary style transfer and domain generalization. In *CVPR*, pages 8035–8045, 2022. [3](#)
- [62] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *ICLR*, 2021. [2](#), [3](#)