



Domain generalization through meta-learning: a survey

Downloaded from: <https://research.chalmers.se>, 2024-12-19 11:33 UTC

Citation for the original published paper (version of record):

Gholamzadeh Khoei, A., Yu, Y., Feldt, R. (2024). Domain generalization through meta-learning: a survey. *Artificial Intelligence Review*, 57(10). <http://dx.doi.org/10.1007/s10462-024-10922-z>

N.B. When citing this work, cite the original published paper.



Domain generalization through meta-learning: a survey

Arsham Gholamzadeh Khoee¹ · Yinan Yu¹ · Robert Feldt¹

Accepted: 22 August 2024
© The Author(s) 2024

Abstract

Deep neural networks (DNNs) have revolutionized artificial intelligence but often lack performance when faced with out-of-distribution data, a common scenario due to the inevitable domain shifts in real-world applications. This limitation stems from the common assumption that training and testing data share the same distribution—an assumption frequently violated in practice. Despite their effectiveness with large amounts of data and computational power, DNNs struggle with distributional shifts and limited labeled data, leading to overfitting and poor generalization across various tasks and domains. Meta-learning presents a promising approach by employing algorithms that acquire transferable knowledge across various tasks for fast adaptation, eliminating the need to learn each task from scratch. This survey paper delves into the realm of meta-learning with a focus on its contribution to domain generalization. We first clarify the concept of meta-learning for domain generalization and introduce a novel taxonomy based on the feature extraction strategy and the classifier learning methodology, offering a granular view of methodologies. Additionally, we present a decision graph to assist readers in navigating the taxonomy based on data availability and domain shifts, enabling them to select and develop a proper model tailored to their specific problem requirements. Through an exhaustive review of existing methods and underlying theories, we map out the fundamentals of the field. Our survey provides practical insights and an informed discussion on promising research directions.

Keywords Meta-learning · Domain generalization · Learning to generalize · Out-of-distribution data · Distributional shifts

✉ Arsham Gholamzadeh Khoee
khoee@chalmers.se

Yinan Yu
yinan@chalmers.se

Robert Feldt
robert.feldt@chalmers.se

¹ Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden

1 Introduction

Traditional machine learning approaches assume training and testing data are independent and identically distributed (i.i.d.). However, this assumption often fails in practice due to variations in data acquisition conditions, such as changes in sensor types, lighting conditions, or environmental factors, leading to substantial performance degradation when models encounter new domains (Khosla et al. 2012); thus, they cannot cope with the out-of-distribution (OOD) challenge (Ganin et al. 2016). A domain can be defined as a specific distribution of data characterized by its own set of features, statistical properties, and underlying concepts (Sugiyama and Storkey 2006). Domain generalization (DG) and domain adaptation (DA) techniques evolved to deal with domain shifts to address out-of-distribution problems (Ghifary et al. 2015; Tzeng et al. 2014). When we have two distinct domains, namely the source and target domains, we can apply domain adaptation techniques, which focus on adapting a model trained on a source domain to perform well on a target domain. Meanwhile, domain generalization aims to generalize over any domain, allowing models trained on one set of domains to perform well on unseen domains. It is important to note that DA techniques assume access to target domain data (Patel et al. 2015; Long et al. 2016), whereas DG techniques assume no access to the target domain data, which makes it a much harder problem to solve. Compared to DA, DG is more applicable in practice. First, machine learning models are trained on a specific large dataset to perform well on that exact task. However, in real-world scenarios, we often face new tasks with limited labeled data available. Essentially, collecting labeled data and retraining the model for such "unseen" domains can be extremely costly and time-consuming. Also, it is often impossible to enumerate all the "unseen" domains in advance. Accordingly, enhancing the generalization capability of machine learning models is crucial.

Meta-learning has emerged to enable a model to quickly adapt and learn from a small amount of data or even generalize to unseen tasks (Huisman et al. 2021). Meta-learning algorithms leverage prior knowledge, patterns, and experiences acquired from similar or related tasks to make learning transferable (Ravi and Larochelle 2016). Recently, some meta-learning algorithms have been used for domain generalization, allowing models to generalize across various domains with limited or no access to all domains during training. These algorithms learn a more generalized representation by identifying patterns that generalize across unseen domains. For instance, consider the PACS dataset (Li et al. 2017), which comprises images from four distinct domains: photo, art, cartoon, and sketch. The goal is to train a model on a subset of these domains, such as photos, art, and cartoons, in such a way that it can also accurately classify sketches despite never having "seen" sketches during training. This challenges the model to extract domain-invariant patterns essential for classifying images across all four domains. This "learning-to-generalize" paradigm aligns with the human learning process, where knowledge acquired from multiple situations enables individuals to adapt quickly to new situations (Finn et al. 2017).

Recently, there has been a growing interest in developing meta-learning methods for domain generalization. These methods aim to address the limitations of conventional machine learning approaches by incorporating meta-level learning procedures that enhance a model's ability to generalize across domains. More specifically, these models achieve cross-domain generalization by learning from a variety of tasks and domains during training, which prepares them to handle distributional shifts and

enables more effective generalization. It is important to note that in domain generalization via meta-learning, the assumption is that there are zero target domain examples during training, contrasting with typical meta-learning that often focuses on few-shot learning for task adaptation.

Wang et al. (2022) and Zhou et al. (2022) discussed recent advances in domain generalization, covering the formal definition of the problem, related fields, and theoretical foundations, and they highlight commonly used datasets and applications. Additionally, Liu et al. (2021) recently explored the OOD generalization problem, which addresses scenarios where the test data distribution differs from the training data. Their review encompasses methodologies categorized into unsupervised representation learning, supervised model learning, and optimization. Vilalta and Drissi (2002) published one of the first surveys on meta-learning to formalize self-adaptive learner models. Huisman et al. (2021) later provided a comprehensive survey on meta-learning to explore the common challenge of how to leverage meta-knowledge to improve the performance of learning algorithms. Hospedales et al. (2021) also discussed various perspectives on meta-learning by providing a novel taxonomy and investigating the utility of the most commonly used meta-learning algorithms. Vettoruzzo et al. (2024) have recently examined meta-learning algorithms, discussing their benefits and challenges, and categorizing methods into metric-, model-, and optimization-based techniques, with a focus on achieving consistent evaluation metrics and computational efficiency. In contrast to previous works, to our best knowledge, this paper is the first survey collecting literature and highlighting the overall potential of meta-learning for domain generalization. In this paper, we offer a fresh perspective and a new taxonomy for classifying meta-learning approaches designed for generalization to unseen domains.

As a more detailed explanation, this survey paper *provides a systematic overview of existing meta-learning methods tailored for domain generalization*. We introduce the concept of domain generalization and clarify its significance in real-world machine learning applications. Subsequently, we delve into the theoretical foundations of meta-learning and explain its applicability to the domain generalization setting. The paper reviews various frameworks and methodologies employed in existing meta-learning approaches for domain generalization, highlighting their strengths and limitations. Additionally, we present a taxonomy to help with the effective categorization of these algorithms. Furthermore, we survey the datasets and evaluation protocols prevalent in this research area and point out the principal challenges and open questions for future study. Our goal is to present an extensive reference for researchers and practitioners interested in understanding and advancing the field of meta-learning for domain generalization.

The rest of this survey is structured as follows. In Sect. 2, we provide background information on the topic and introduce the fundamental concepts of meta-learning and its applicability to the domain generalization setting. Section 3 presents a taxonomy that can be used to categorize techniques that will be discussed in this survey. We then proceed to explore the different frameworks of meta-learning methods for domain generalization in Sect. 4. Section 5 discusses the widely adopted datasets and evaluation protocols employed in the field. In Sect. 6, we cover the crucial applications of meta-learning for domain generalization and its significance in practical machine learning scenarios. Section 7 discusses the significance of the findings, key challenges, and promising research directions. Finally, Sect. 8 summarizes the main points of the survey and their implications.

2 Background

While deep learning models have achieved remarkable results in specific tasks, developing models that can generalize across tasks and domains remains a challenge. Initially, we will provide a concise overview of related research areas, highlighting their unique characteristics, with a comparison of the discussed learning paradigms available in Table 1. Subsequently, we will delve into more details of meta-learning and domain generalization, along with providing a formalization of meta-learning for domain generalization, as it represents a research area specifically aimed at addressing the challenges of generalization across various tasks and domains.

2.1 Different learning paradigms

In real-world applications, data availability is often limited and conditions change over time. A practical machine learning model must adapt to these changes without retraining from scratch, as traditional retraining is resource-intensive and inefficient. To address this challenge, several learning paradigms have been introduced. This section briefly introduces the key learning paradigms to build intuition for designing efficient machine learning models.

2.1.1 Incremental learning

Incremental learning involves gradually training a model on new data over time. The model is updated incrementally as new data arrives, often in batches. However, previous data is often not retained, so incremental learning can suffer from catastrophic forgetting of earlier concepts. It is also necessary to carefully design training data to avoid bias and ensure that the AI system can generalize from the new data (Wu et al. 2019; He et al. 2011; Castro et al. 2018). The other important challenge is to deal with non-stationarity in the data, which can prevent the AI system from converging on a solution. Incremental learning can be done online or offline (He et al. 2020).

2.1.2 Online learning

Online learning involves training a model on data that becomes available in a sequential order, sample-by-sample or mini-batch by mini-batch (Hoi et al. 2021). The model is updated each time a new sample/mini-batch arrives. Online learning algorithms aim to perform updates efficiently to accommodate new data but do not explicitly retain old knowledge.

2.1.3 Continual learning

Continual learning refers to the ability of a model to learn sequentially over time from a stream of data. The goal is to learn new tasks without forgetting previous knowledge,

Table 1 Comparison of different learning paradigms highlighting similarities and dissimilarities among meta-learning, domain generalization, and related learning approaches

Learning Paradigm	$p(X)$ Consistency	Y Consistency	Test Access	$\hat{p}(Y X; \theta^1)$
Incremental/Online Learning (Hoi et al. 2021)	✓	✓	✓	$\min_{\theta} \mathbb{E}_{p(\mathcal{D})} [\mathcal{L}(\mathcal{D}; \theta)]$
Continual Learning (De Lange et al. 2021)	✗	✓	✓	$\min_{\theta} \mathbb{E}_{p(\mathcal{T})} \mathbb{E}_{p(\mathcal{D}_s)} [\mathcal{L}(\mathcal{D}_s; \theta)]$
Transfer Learning (Tripuraneni et al. 2020)	✗	✗	✓	$\min_{\theta} \mathbb{E}_{p(\mathcal{D}^T)} [\mathcal{L}(\mathcal{D}^T; \theta)]^2$
Multi-task Learning (Sener and Koltun 2018)	✓	✓	✓	$\min_{\theta^{\text{sh}}, \theta^T} \mathbb{E}_{p(\mathcal{T})} \mathbb{E}_{p(\mathcal{D}_s)} [\mathcal{L}(\mathcal{D}_s; \theta^{\text{sh}}, \theta^T)]^3$
Meta-Learning (Wang et al. 2024)	✓	✗	✓	$\min_{\theta} \mathbb{E}_{p(\mathcal{T})} \mathbb{E}_{p(\mathcal{D}_s)} [\mathcal{L}(\mathcal{D}_s^r; \mathcal{D}_s^e; \theta)]$, $\mathcal{D}_s = \mathcal{D}_s^r \cup \mathcal{D}_s^e$
Domain Adaptation (Farahani et al. 2021)	✗	✓	✓	$\min_{\theta} \mathbb{E}_{p(\mathcal{D}^T)} [\mathcal{L}(\mathcal{D}^T; \theta)]$
Homogeneous DG (Zhou et al. 2022)	✗	✓	✗	$\min_{\theta} \mathbb{E}_{p(S)} \mathbb{E}_{p(\mathcal{D}_s)} [\mathcal{L}(\mathcal{D}_s; \theta)]^5$
Heterogeneous DG (Wang et al. 2020; Shu et al. 2021)	✗	✗	✗	$\min_{\theta} \mathbb{E}_{p(S)} \mathbb{E}_{p(\mathcal{D}_s)} [\mathcal{L}(\mathcal{D}_s; \theta)]$

¹ In this table, we use θ to denote the set of all trainable parameters in the model

² \mathcal{D}^T represents the target domain

³ θ^{sh} and θ^T denote the shared and the task-specific parameters, respectively

⁴ \mathcal{D}_s^r and \mathcal{D}_s^e indicate the meta-train and meta-test splits of task s 's data, respectively

⁵ \mathcal{D}_s and S represent a dataset from source domain s and a set of source domains, respectively

handling different distributions in the data over time (De Lange et al. 2019). It involves retaining previously learned knowledge and effectively integrating new knowledge.

2.1.4 Transfer learning

Transfer learning refers to the process of reusing a pre-trained model as the basis for a new model, where a model trained on one task is repurposed for a second related task to allow rapid progress when modeling the second task (Pan and Yang 2009). Essentially, transfer learning involves using the knowledge learned from one task to improve the performance of another related task.

2.1.5 Multi-task learning

Multi-task learning is a machine learning approach in which we try to learn multiple tasks simultaneously, optimizing multiple loss functions at once. The goal of multi-task learning is to improve learning efficiency and prediction accuracy for the task-specific models, when compared to training the models separately (Yang and Hospedales 2014). It leverages shared representations, which can help other tasks be learned better.

2.1.6 Meta-learning

Meta-learning refers to the process of learning how to learn (Finn et al. 2017). It involves training a model to learn new tasks quickly with minimal data by leveraging knowledge learned from previous tasks. The goal of meta-learning is to learn a good initialization of the model's parameters that can be adapted quickly to new tasks with only a few samples (also referred to as few-shot learning) (Snell et al. 2017). In other words, it transfers knowledge from many source tasks, so it can learn a new task more quickly, more proficiently, and more stably (Hospedales et al. 2021).

2.1.7 Domain adaptation

DA focuses on adapting a model from one source domain to perform well on a related target domain (Wang and Deng 2018; Csurka 2017; Bousmalis et al. 2016; Long et al. 2015; Ben-David et al. 2006). It assumes some similarity or overlap between the two domains. It leverages labeled data from the source domain and unlabeled/limited labeled data from the target domain. Also, it is closely related to transfer learning, which also aims to transfer knowledge from the source domain to the target domain.

2.1.8 Domain generalization

DG goes a step further in comparison to DA. It involves training a model on multiple source domains with the goal of making it perform well on any unseen or new target domains (Muandet et al. 2013; Li et al. 2017; Shankar et al. 2018). So, it does not assume access to target domain data during training. DG addresses the challenge of adapting to diverse and unknown domains by learning representations that capture the underlying invariant factors across various domains.

2.2 Meta-learning for domain generalization: promises and challenges

Meta-learning is considered a promising approach for DG, as it trains models to build transferable knowledge that can be efficiently applied to different tasks and domains. This aligns with the goal of DG to generalize to unseen domains. By training on a variety of learning tasks, meta-learning enables models to acquire knowledge that is transferable across different domains. In other words, meta-learning enables models to learn how to learn, which significantly enhances their ability to adapt to new and unseen domains by leveraging previously acquired knowledge. This adaptability is crucial for DG, where the goal is to generalize beyond the training domains (see Table 2 for a comparison between traditional machine learning and meta-learning). In meta-learning, models are trained episodically where each task or episode simulates a domain shift. This procedure effectively equips the model with the ability to handle domain shifts by encouraging the development of features that are robust to the variations between training and test domains. By repeatedly encountering a variety of tasks, the model learns to extract features that are crucial for performance across different domains, rather than overfitting to the idiosyncrasies of any single domain. These robust features are more likely to be relevant in unseen domains, thus improving the overall generalization of the model. Additionally, the intrinsic sample efficiency of meta-learning can lead to more effective utilization of limited data. Meta-learning is designed to achieve high performance with fewer examples by leveraging the knowledge gained from previous tasks. This sample efficiency is particularly beneficial in scenarios where data from new domains is scarce or expensive to obtain.

One of the main challenges of using meta-learning for DG is ensuring a sufficient diversity of learning tasks during the meta-training phase. Without enough diversity, the model may lack the capacity to generalize effectively to completely new domains. Additionally, the degree of the distributional shift between the training domains and the test domain can be substantial. If the unseen test domain differs significantly from the training domains, the model may encounter difficulties in generalizing. Furthermore, while meta-learning is designed to be sample-efficient, in reality, it often requires a large number of tasks to achieve high performance, which might not always be feasible or accessible. It is also important to note that effective domain generalization often requires understanding causal relationships within the data, which meta-learning algorithms might not capture without proper inductive biases or domain knowledge integration.

2.3 Formalization of meta-learning for domain generalization

In meta-learning, the goal is to solve new unseen tasks by leveraging the knowledge gained from solving N previous tasks. Each task \mathcal{T}_i corresponds to a dataset $\mathcal{D}_i = \{(x_j^i, y_j^i) \mid x_j \in \mathcal{X}, y_j \in \mathcal{Y}\}_{j=1}^n$, where \mathcal{D}_i is divided into a support set \mathcal{D}_i^S and a query set \mathcal{D}_i^Q . It is important to note that in meta-learning, all tasks are sampled from a unique task distribution $p(\mathcal{T})$.

In domain generalization, we are given data from M different source domains and aim to perform well on a new unseen domain, known as the target domain. Specifically, each domain d_i is defined as follows:

$$d_i \triangleq \{p_i(x), p_i(y|x)\}, \quad (1)$$

Table 2 This table outlines the key differences between traditional machine learning and meta-learning across various aspects such as tasks, training datasets, purpose, loss functions, weights, and the role of these weights. The comparison highlights how meta-learning enhances adaptability and generalization to new and unseen domains by leveraging knowledge from a variety of tasks, in contrast to traditional machine learning which typically focuses on optimizing performance for a specific task

	Machine Learning	Meta-Learning
Task	Image recognition, semantic analysis of text	
Training Dataset	Typically a large dataset annotated for the current task	Variety of tasks, often organized into training episodes. Each episode contains a small amount of data from a particular task
Purpose	Learn to perform a specific task from a dataset	Learn to perform specific tasks and also adapt to new tasks using only a few examples or new datasets (domains) with unseen distributions
Loss Function	A mathematical function that measures the discrepancy between the predicted output and the ground truth output	Measures not only the prediction error on the current task (same as machine learning loss) but also optimizes adaptability to new tasks
Weights	Learned parameters that are adjusted during the training process to minimize the loss on that task	Base-level weights which are task-specific, and meta-level weights or hyperparameters that guide the learning process across tasks and datasets (different domains)
What are these weights good for	Transforming the raw features into a more representative space (for example, a more linearly separable space for classification tasks)	Should be able to ignore features only relevant to domain differences (robust against domain shift), focusing on the features relevant to the task

where $p_i(x)$ refers to the distribution of the input data for domain d_i , and this distribution may vary across different domains. The term $p_i(y|x)$ represents the conditional probability distribution of the target y given the input data x . In a homogeneous DG setting, $p_i(y|x)$ is assumed to be consistent across all domains, meaning the label space is the same. However, in a heterogeneous DG setting, $p_i(y|x)$ can vary, as there can be disjoint label spaces across different domains.

For instance, consider the PACS (Li et al. 2017) dataset for homogeneous DG; this dataset comprises images from 4 distinct domains: photo, art painting, cartoon, and sketch. Each domain shares the same set of 7 classes for labeling. Tasks are drawn from these varying domains, each with its unique statistical characteristics. Consequently, there is a significant domain shift between tasks, necessitating a model capable of generalizing across these different domains to effectively handle tasks from varied distributions. In contrast, the Visual Decathlon (VD) (Rebuffi et al. 2017) dataset includes 10 different domains, each encompassing a variety of categories, and is thus suitable for designing tasks for heterogeneous DG. While the aforementioned datasets focus on multi-source DG, where the model is trained across several source domains to provide representations that generalize to novel unseen domains, there are scenarios where only a single source domain is available. In these cases, the model is expected to generalize to target domains without prior exposure to multiple source domains. This challenge is known as single-domain generalization. Additionally, some problems require the model to handle an open-set of target domains

as opposed to a limited or closed-set of target domains. This requirement significantly increases the difficulty of domain generalization problems.

As it can be conceived, a domain can be viewed as a special case of a task from the perspective of meta-learning. Therefore, the meta-learning paradigm can also be applied to domain generalization problems, where each domain is treated as if it were a distinct task, defined by a specific source \mathcal{S}_i and associated with a dataset \mathcal{D}_i . It's also worth mentioning that we can leverage zero-shot learning capabilities for new unseen domains when integrating domain generalization techniques with meta-learning.

3 Taxonomy

To advance DG through meta-learning, it is important to get an overview of the existing approaches and their similarities and differences. Here, we introduce a taxonomy and its key dimensions (axes) with which different meta-learning approaches can be analyzed. The taxonomy provides two axes, the discriminability axis, and the generalizability axis, representing how a model learns to generalize to unseen domains. This taxonomy aims to highlight the interaction and knowledge transferability between data points from different domains and classes. The end goal is to meta-learn domain-agnostic features and discriminative classifiers that perform well on unseen domains.

Compared to other meta-learning techniques, a unique characteristic of DG is its design principle that prevents the model from adapting to specific test domains. This is because, by definition, DG focuses on generalizing to completely unseen test domains, unlike domain adaptation, where the model has some form of access or exposure to information from the test domain during training. This distinct feature requires the meta-learning approach to *generalize* instead of *adapt*. Consequently, it is crucial to balance the trade-off between domain-invariant and domain-specific features (Chattopadhyay et al. 2020). Specifically, in classification tasks, the domain-specific features that are learned must be carefully curated to enhance class separability even for unseen domains. The taxonomy presented in this survey is designed to address these aspects. We divide the categorization into two principal axes: the first concerning the treatment of domain features by the feature extractor (generalizability) and the second relating to the training methodology of the classifier within the model (discriminability).

3.1 Generalizability axis

The feature extractor is a pivotal component in domain generalization, as it determines how the model perceives and processes data from diverse domains to learn generalized representations applicable across various settings. Approaches along this axis can be divided into two categories: (1) *Minimization of Inter-Domain Distances* and (2) *Maximization of Intra-Domain Distances*.

3.1.1 Minimization of inter-domain distances

Some approaches aim to extract domain-invariant features by minimizing the distance between feature representations across different domains. The primary objective is to

identify and harness the characteristics that are consistent across domains. In other words, these methods mainly focus on directly inducing invariant representations without diversifying inputs, concentrating on the consistent and discriminative aspects of the data. By doing so, they reduce the model's sensitivity to the idiosyncrasies of any given domain, thereby enhancing the robustness of the model against the shifts that occur in novel environments.

3.1.2 Maximization of intra-domain distances

Alternatively, certain approaches seek to construct a feature extractor capable of handling a more extensive variety of data by maximizing inter-domain distances. These methods utilize strategies like data augmentation, noise injection, and domain randomization to expand the training data spectrum and encourage the model to learn features that are broadly applicable rather than domain-specific. By exposing the model to a more diverse set of modified inputs, these methods strive to replicate the variability that the model will encounter in real-world scenarios, thus promoting the learning of features with better generalizability. Through the explicit diversification of inputs from a domain, these approaches develop feature extractors prepared to handle greater variance in the data.

3.2 Discriminability axis

The performance of DG models is highly dependent on the effectiveness of the classifier's training methodology. The strategies in this context are categorized into two groups based on how the classifier manages distances: (1) *Minimization of Intra-Class Distances* and (2) *Maximization of Inter-Class Distances*.

3.2.1 Minimization of intra-class distances

Most approaches simply rely on minimizing the distances between instances within the same class. The intention is to encourage the model to group similar examples together. This method ensures that the model's predictions are consistent for similar instances.

3.2.2 Maximization of inter-class distances

Some methodologies also explicitly maximize the distance between different classes by incorporating triplet loss to train the classifier. This not only minimizes inter-class distances but also explicitly maximizes intra-class distances. By doing so, the model is encouraged to learn robust representations that distinctly separate different classes. The use of triplet loss helps to develop a feature space where classes are well-delineated, resulting in clearly separated clusters. This enhances the classifier's discriminative power and its ability to generalize across domains.

Taken together, these two axes allow us to categorize domain generalization techniques based on how they train the feature extractor and classifier components to improve generalization. Methods can differ in whether they align or diversify representations across domains, as well as how they distinguish between classes during training, as illustrated in Fig. 1.

3.2.2.1 Application of the taxonomy In Fig. 2, we present a decision graph to guide readers through the taxonomy to select the most appropriate meta-learning method for DG based on data availability and domain shifts. It categorizes meta-learning approaches based on the strategies employed for the feature extractor (generalizability axis) and the classifier training process (discriminability axis), facilitating an understanding of the relationships between various methods and their underlying principles.

4 Methodologies

In recent years, there has been significant interest in utilizing the meta-learning procedure to create models for DG. In this section, we will explore individual meta-learning methods for DG by describing the methodologies and motivations behind their designs. We discuss the strengths and weaknesses of each method in detail.

Meta-learning focuses on building a general model that can be used to adapt rapidly to new situations. For the first time, Li et al. exploited the idea behind meta-learning for DG by dividing the data from multiple source domains into non-overlapping meta-train and meta-test sets to simulate domain shifts in the training procedure. Therefore, we will



Fig. 1 An illustrative diagram of the meta-learning taxonomy for domain generalization. The quadrant chart highlights two principal axes: the first axis (the generalizability axis) represents the strategy of the feature extractor, contrasting the Minimization of Inter-Domain Distances with the Maximization of Intra-Domain Distances; the second axis (the discriminability axis) depicts the classifier training process, distinguishing between the Minimization of Intra-Class Distances and Maximization of Inter-Class Distances. This diagram visually organizes domain generalization approaches, illustrating their distinct mechanisms for promoting generalization across unseen domains

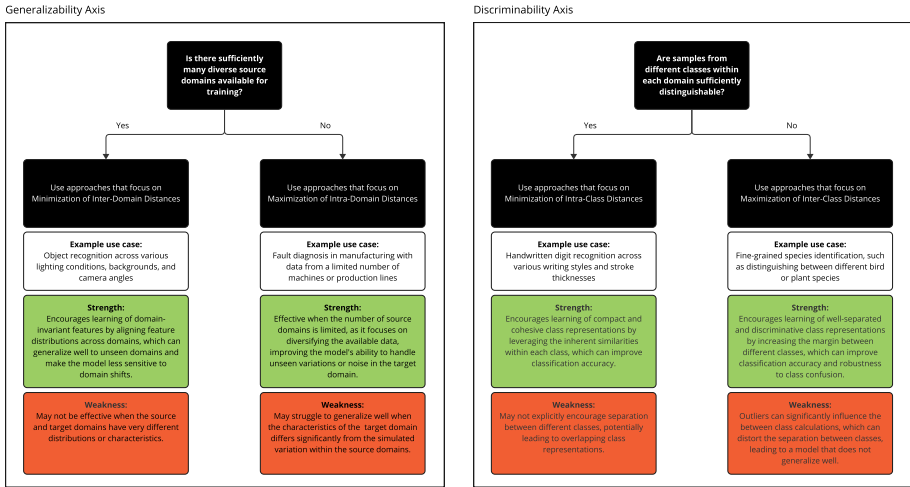


Fig. 2 A decision graph illustrating how to apply the taxonomy of meta-learning approaches for domain generalization. The decision graph categorizes techniques based on the two key aspects: the generalizability axis, which focuses on the feature extractor’s strategy (minimizing inter-domain distances or maximizing intra-domain distances), and the discriminability axis, which focuses on the classifier’s training process (minimizing intra-class distances or maximizing inter-class distances). This decision graph helps the reader navigate through the taxonomy, providing use cases, strengths, and weaknesses for each category to make the concepts more tangible, applicable, and actionable

discuss DG through meta-learning, starting with meta-learning for domain generalization (MLDG) and proceeding accordingly.

4.1 MLDG

Recent meta-learning studies have focused on learning good weight initializations for few-shot learning, particularly Model-Agnostic Meta-Learning (MAML) (Finn et al. 2017). Li et al. (2018) proposed MLDG (meta-learning for domain generalization), which draws inspiration from the MAML approach and trains models capable of performing well on OOD data by transferring knowledge across domains instead of tasks. It is important to note that DG assumes zero training examples from the target domain, in contrast to conventional meta-learning methods like MAML, which address few-shot scenarios using few training examples of new tasks. Through the meta-learning framework, MLDG will be exposed to domain shifts during training, enhancing its ability to handle domain shifts in various situations by learning general representations. Accordingly, they divide data from multi-source domains into meta-train and meta-test sets to simulate domain shifts. In the meta-train phase, parameters θ of the model are updated by gradient descent using the meta-train sets \mathcal{S}_{tr} . As a result, the following steps will be taken:

$$\begin{aligned} \nabla_{\theta} &= \ell'_{\theta}(\mathcal{S}_{tr}; \theta), \\ \theta' &= \theta - \alpha \nabla_{\theta}, \end{aligned} \tag{2}$$

where α is the meta-train (inner loop) learning rate. Following that, in the meta-test phase, we determine the loss using the meta-test sets (virtual-test domains) \mathcal{S}_{te} with the parameters

obtained in the meta-train phase. Consequently, this phase involves calculating gradients over gradients to update the model’s parameters. The parameters update can be formulated as:

$$\theta = \theta - \beta \frac{\partial(\mathcal{L}(\mathcal{S}_{tr};\theta) + \gamma \mathcal{L}(\mathcal{S}_{te};\theta'))}{\partial\theta}, \tag{3}$$

where \mathcal{L} is the loss function, and β is the meta-test (outer loop) learning rate, and γ weights meta-train and meta-test.

MLDG is built on top of MAML and improves the generalization of MAML, which is model-agnostic and easily applicable to reinforcement learning (RL) problems. However, Like MAML, it is computationally expensive since it requires the calculation of second-order derivatives.

4.2 MetaReg

Balaji et al. (2018) introduced MetaReg (domain generalization using meta-regularization) that incorporates a regularization term trained using meta-learning to encourage models to extract domain-invariant features and reduce sensitivity to domain-specific variations to train models that are robust to domain shifts. Their proposed model consists of a feature network ψ along with p task networks θ_i for p source domains. The feature network is designed to extract more general features, and the tasks networks are utilized to enforce domain-specificity in the networks so that they can apply regularizers to make them domain-invariant. Therefore, they implement the meta-learning framework to train the regularizer term and incorporate it into the loss function. The regularizer term will be trained using every pair of source domains (a, b) . This procedure can be expressed by the following set of equations:

$$\begin{aligned} \beta^1 &\leftarrow \theta_a^{(k)} \\ \beta^t &= \beta^{t-1} - \alpha \nabla_{\beta^{t-1}} [L^{(a)}(\psi^{(k)}, \beta^{t-1}) + R_\phi(\beta^{t-1})] \\ &\forall t \in \{2, \dots, l\} \\ \hat{\theta}_a^{(k)} &= \beta^l \end{aligned} \tag{4}$$

$$\phi^{(k+1)} = \phi^{(k)} - \alpha \nabla_\phi L^{(b)}(\psi^{(k)}, \hat{\theta}_a^{(k)})|_{\phi=\phi^{(k)}} \tag{5}$$

where Eq. (4) describes the inner loop of the meta-learning paradigm that performs l steps of gradient descent using meta-train set on a new task network with parameters $\theta_a^{(k)}$ which is the base model’s task network parameters of the a^{th} domain at iteration k to obtain $\hat{\theta}_a^{(k)}$. In addition, Eq. (5) refers to the outer loop of the meta-learning strategy, which incorporates the meta-test set to minimize the unregularized loss with parameters $\hat{\theta}_a^{(k)}$ with respect to the regularizer parameter ϕ . By using the trained regularizer, we fine-tune a new model on the entire source dataset. To further add, weighted L_1 loss is used as the regularization function, i.e., $R_\phi(\theta) = \sum_i \phi_i |\theta_i|$ to build a learnable weight decay mechanism and diminish domain-specific characteristics in task networks.

Although MetaReg increases the training complexity and requires several domains for practical training, it is easily applicable to deep neural networks and helps reduce overfitting.

4.3 Feature-critic networks

Feature-Critic Networks introduced by Li et al. (2019a) is tailored for heterogeneous domain generalization by drawing inspiration from the meta-learning framework. The proposed model meta-learns a loss function that promotes domain robustness. In this way, they train a robust feature extractor that can be used with any classifier for addressing various problems from different domains. More specifically, they introduced a learnable auxiliary loss $\ell_{\omega}^{(Aux)}$ resulting from optimizing the feature-critic network h_{ω} to encourage the base network to extract domain agnostic features, as illustrated in Fig. 3.

Based on the following equations, parameters θ of the base network is computed using the meta-train data \mathcal{D}_{tr} :

$$\begin{aligned} \theta^{(OLD)} &= \theta - \alpha \nabla_{\theta} \ell^{(CE)}(\mathcal{D}_{tr}|\theta), \\ \theta^{(NEW)} &= \theta - \alpha \nabla_{\theta} (\ell^{(CE)}(\mathcal{D}_{tr}|\theta) + \ell_{\omega}^{(Aux)}(\mathcal{D}_{tr}|\theta)), \end{aligned} \tag{6}$$

where $\ell^{(CE)}$ is the cross-entropy loss, and the auxiliary loss $\ell_{\omega}^{(Aux)}$ is defined as follows:

$$\ell_{\omega}^{(Aux)} := \frac{1}{N} \sum_{i=1}^N \text{softplus}(h_{\omega}(f_{\theta}(x_i))), \tag{7}$$

where f_{θ} is the feature extractor. Also, the softplus function is applied to ensure the output is non-negative.

In order to determine parameters ω of the feature-critic network, the following optimization will be used with the meta-test data \mathcal{D}_{val} :

$$\min_{\omega} \tanh(\ell^{(CE)}(\mathcal{D}_{val}|\theta^{(NEW)}) - \ell_{\omega}^{(CE)}(\mathcal{D}_{val}|\theta^{(OLD)})). \tag{8}$$

The tanh function can be considered a softer version of gradient clipping in this context.

Ultimately, we can utilize $g_{\phi} \circ f_{\theta}$, where the trained base network f_{θ} will serve as a fixed feature extractor for target domains, and g_{ϕ} would be the classifier. For heterogeneous DG,

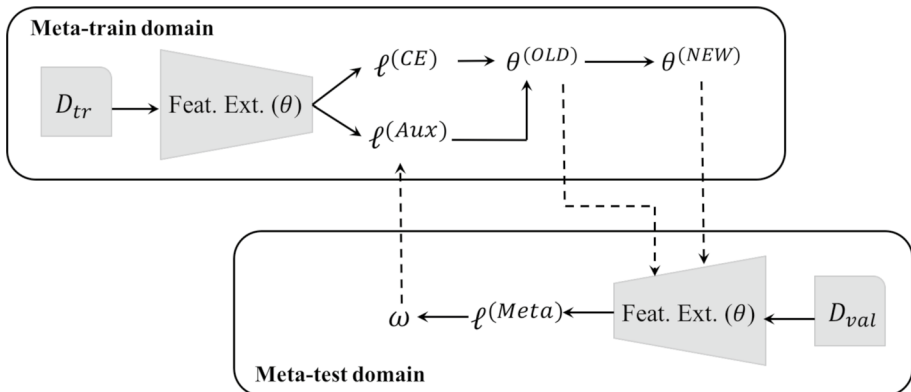


Fig. 3 Visual representation of the Feature-Critic Networks for heterogeneous domain generalization, depicting the base network’s feature extraction guided by an auxiliary loss from the feature-critic networks to promote domain-invariant feature extraction. The domain-invariant features are generated by minimizing the inter-domain distances

we have N distinct classifiers, denoted as $g_{\phi_1}, g_{\phi_2}, \dots, g_{\phi_N}$. On the other hand, in the homogeneous DG, we have a single classifier g_{ϕ} that can be shared across all domains.

Feature-critic networks can effectively deal with heterogeneous DG where label spaces differ. Nevertheless, there is no guarantee that the gradients of the feature-critic networks will not conflict with those from the supervised loss, which could negatively affect the performance of the model.

4.4 Episodic training for DG

A novel framework for DG is proposed by Li et al. (2019b) based on episodic training that resembles the meta-learning training approach. In particular, they used episodic training to train a deep neural network, decomposed into feature extractor and classifier components. Each component is trained by simulating interactions with a poorly tuned partner for the current domain. Moreover, to construct episodes, the framework uses data from domain A to be processed by the classifier trained on domain B , which has not been exposed to data from A , and vice versa. This cross-domain exposure ensures that each component becomes sufficiently robust at handling OOD data, as depicted in Fig. 4.

The intuition of the episodic training using a mismatched classifier is to regularize the feature extractor by asking it to produce features robust enough that even an arbitrary classifier will be able to perform properly. Accordingly, three episodic training strategies are introduced: 1) Train feature extractor to support classifiers from other domains; 2) Train classifier to accept features from other domains; 3) Train feature extractor to support a randomly initialized classifier.

The proposed approach requires only a simple modification of the training procedure and does not rely on special data augmentation or optimization algorithms. It also allows for heterogeneous label spaces across domains since no classifier needs to be shared. However, there is a risk that the feature extractor could learn degenerate representations if the randomness imposed overwhelms the true objectives of the training.

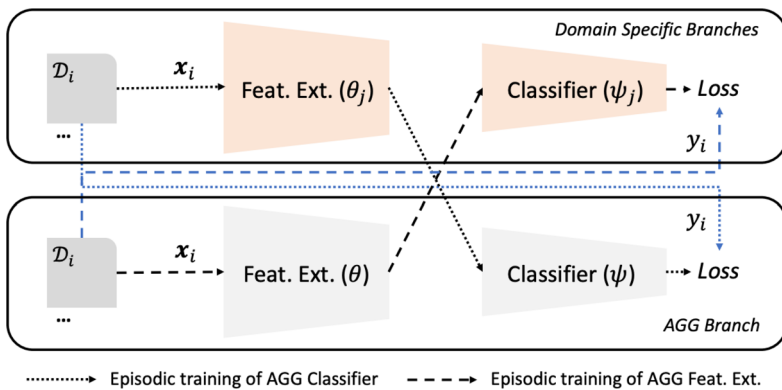


Fig. 4 Overview of Episodic Training for DG framework, illustrating the regularization process where a feature extractor is trained with classifiers from different domains and vice versa, promoting out-of-distribution robustness

4.5 Meta-learning the invariant representation

Learning invariant representations across source domains has shown effectiveness for DG. However, overfitting to source domains can limit generalization to a target domain that differs greatly from the sources. Jia and Zhang (2022) proposed a meta-learning algorithm via bilevel optimization to improve the out-of-domain robustness of the learned invariant representation. Most meta-learning algorithms share a common limitation: they use task objectives directly as the inner-loop and outer-loop objectives. This approach can be suboptimal because it is highly abstracted from the feature representation. To address this issue, the authors focus on a meta-learning approach aimed at reducing the discrepancy between the target domain and source domains, where in each training iteration, any domain can become the meta-target domain while the rest serve as the meta-source domains. Specifically, they developed a bilevel meta-learning procedure based on the first-order MAML framework, which achieves high computational efficiency. The paper employs a \mathcal{Y} -discrepancy measure (Zhang et al. 2012) to quantify domain discrepancy, effectively capturing both covariate and conditional shifts between domains. Notably, \mathcal{Y} -discrepancy has been utilized in previous research for domain invariance learning (Zhang et al. 2021). The primary goal of this algorithm is to minimize the \mathcal{Y} -discrepancy between the source domains and the target domain. A gradient-based meta-learning algorithm is provided to solve the bilevel optimization problem efficiently. The algorithm alternates between sampling meta-tasks, updating the feature embedding and classifiers on the meta-training set, and updating the meta-parameters based on the meta-test performance. The inner-loop objective aims to minimize discrepancy across different source domains, while the outer-loop objective aims to minimize discrepancy between source domains and a potential target domain. In particular, adversarial learning is employed to estimate and minimize the \mathcal{Y} -discrepancy between domains in both the inner and outer loops of the bilevel meta-learning algorithm. In essence, this meta-learning approach reduces the discrepancy between the target domain and the convex hull of the source domains as depicted in Fig. 5. This enables learning a robust invariant representation for improved domain generalization.

The introduced model learns a robust invariant representation that generalizes well to unseen domains by optimizing the feature embedding to minimize the discrepancy between meta-source and meta-target domains during training. However, like other domain generalization methods, the effectiveness of the proposed approach may depend on the diversity of the available source domains. If the source domains lack diversity, the learned invariant representation may not generalize well to significantly different target domains.

4.6 MASF

Model-agnostic learning of semantic features (MASF) has been presented by Dou et al. (2019), in which two complementary losses that explicitly regularize the semantic structure of the feature space are introduced. This approach aims to enhance domain generalization within the meta-learning framework by splitting source domains into meta-train and meta-test sets to simulate domain shift through an episodic training procedure. In more detail, MASF suggests enforcing semantic features through global class alignment and local sample clustering, as illustrated in Fig. 6. These supplementary losses are explicitly derived in an episodic learning procedure.




Fig. 5 Depiction of the meta-learning for invariant representation method. **(a)** Compared to the ERM baseline, **(b)** domain invariance learning decreases the discrepancy among source domains and excels in source-domain classification, yet it may still result in significant errors on the target domain. **(c)** The proposed approach employs bilevel meta-learning to further minimize the discrepancy between the target and source domains, enabling a hypothesis learned from the source domains to generalize effectively to the target domain

The objective of global class alignment is to structure the feature space such that it preserves the learned inter-class relationships on unseen data via explicit regularization. By ensuring consistent class relationships, knowledge transfer across domains is facilitated, thus improving domain generalization. This inter-class alignment is done by exploiting what the model has learned about class ambiguities in the form of per-class soft labels and enforcing their consistency within each two distinct domains by minimizing the symmetric KL divergence between their soft confusion matrix created from the collection of soft labels.

The objective of local regularization is to enhance robustness by promoting feature compactness, grouping features of samples within the same class closely together while being distinct from features of different classes. This is crucial to prevent ambiguous decision boundaries and sensitivity to unseen domain shifts. As a result, an embedding network is utilized to process the extracted features. The output of this network is then used




Fig. 6 Overview of the MASF approach for Domain Generalization, F_{ψ} and T_{θ} represent the feature extractor and task network, respectively, with $F_{\psi'}$ and $T_{\theta'}$ being their updated versions after inner gradient updates on the task loss \mathcal{L}_{task} . M_{ϕ} is the metric embedding network, and D_k represents different source domains. (a) Illustration of gradients flow of episodic training with domain shift simulation; (b) Global alignment for consistent class relationships across domains; (c) Local sample clustering to promote the Maximization of Inter-Class Distances

to measure the distance between the input features of randomly selected pairs of samples from all domains. The training procedure employs either contrastive loss or triplet loss.

The proposed model captures inter-class relationships and ensures intra-class compactness, improving the generalization of the model; however, applicability to large-scale problems with many classes may be a challenge.

4.7 S-MLDG

Li et al. (2020) proposed a framework for sequential learning for domain generalization (S-MLDG) based on the MLDG to enhance its performance by incorporating sequential learning and lifelong learning. They train a base DG model sequentially along a trajectory spanning the source domains, optimizing the cumulative performance across the entire trajectory instead of individual steps. In the case of MLDG, it recursively applies MLDG along the trajectory, simulating lifelong DG learning, as shown in Fig. 7. This approach encompasses more unique DG episodes compared to the base algorithms. While MLDG considers N distinct domain-shot episodes, this method considers $N!$ distinct domain-shift episodes. Furthermore, the approach defines a loss function that aggregates the performance on each step of the trajectory through domains. Additionally, it incorporates shuffling domain orders and sampling new batches for each iteration, ensuring diversity and preventing overfitting.

The fast first-order approximation, denoted as FFO-S-MLDG, constitutes a streamlined variant of the S-MLDG algorithm, drawing inspiration from Reptile to mitigate its computational overhead. FFO-S-MLDG uses a first-order approximation of the original S-MLDG algorithm, enabling it to be trained more quickly and efficiently.

Although the training complexity is higher compared to MLDG due to the optimization over domain trajectories, this approach consistently demonstrates improvements over strong baselines on multiple benchmarks. This is due to its ability to generate more unique DG episodes.

4.8 MetaVIB

Meta variational information bottleneck (MetaVIB) has been introduced by Du et al. (2020) as a probabilistic meta-learning model to handle uncertainty and domain shifts for




Fig. 7 Illustration of the S-MLDG training procedure, showing the sequential training across multiple domain trajectories with random sequences to simulate lifelong learning and increase the diversity of domain-shift episodes

domain generalization by learning domain-invariant representations using variational auto-encoders. In order to address the uncertainty associated with predictions on the unseen target domain, the model represents classifier parameters as distributions inferred from source domains. In a similar manner to other algorithms, source domains are split into meta-train and meta-test sets in each training episode to simulate domain shift, as illustrated in Fig. 8. MetaVIB is derived from novel variational bounds of mutual information by leveraging the meta-learning setting. Also, the meta-learning objective maximizes performance on meta-test while minimizing domain divergence via MetaVIB. As a result, MetaVIB learns to gradually narrow domain gaps to achieve domain invariance, while maximizing prediction accuracy through episodic training. This encourages learning efficient representations that preserve information for prediction but remove domain-specific statistics, thereby enabling the learning of domain-agnostic representations.

MetaVIB explicitly addresses the uncertainty associated with predictions for new domains through probabilistic modeling. Also, by limiting the information flow from the inputs to only what is relevant for the classification task through the information bottleneck. This prevents the model from encoding irrelevant domain-specific details. However, it does not explicitly align distributions across domains, implying a certain level of domain shift might remain.

4.9 M-ADA

Qiao et al. (2020) introduced Meta-learning based Adversarial Domain Augmentation (M-ADA) as an effective single domain generalization method. Recognizing that acquiring data from multiple training domains might not be feasible due to constraints such as budget or privacy concerns, they proposed generating fictitious domains through adversarial training to use in the meta-test phase. As shown in Fig. 9, the model consists of a task model and a Wasserstein Autoencoder (WAE).

The task model with parameters θ encompasses a feature extractor F , which maps the input space to embedding space z , and a classifier C that predicts labels from the embedding space z , using cross-entropy as its loss function \mathcal{L}_{task} . Moreover, \mathcal{L}_{const} imposes a semantic consistency constraint on the adversarially augmented samples x^+ . It ensures that the augmented domains S^+ remain close to the source domain S within the embedding space. Accordingly, \mathcal{L}_{const} is defined as:




Fig. 8 Computational graph of the MetaVIB model for domain generalization, depicting the training process where domain-invariant representations are learned. Global parameters (Θ) and classifier parameters (ψ) are optimized over source domains, which are split into meta-train (D^s) and meta-test (D^t) sets to simulate domain shifts within each episodic training phase. During the test phase, Θ generates representations for the target domain, while ψ is used for predictions of data from the source domain. While this is a typical example of the inter-domain minimization and intra-class minimization paradigm, the distinctive feature of MetaVIB lies in its explicit probabilistic modeling approach




Fig. 9 Overview of M-ADA method for single domain generalization using adversarial domain augmentation. The architecture employs a task model comprising a feature extractor F and classifier C , alongside a Wasserstein Autoencoder (WAE), to generate fictitious domains

$$\mathcal{L}_{const} = \frac{1}{2} \|z - z^+\|_2^2 + \infty \cdot 1 \{y \neq y^+\}, \tag{9}$$

where $1 \{ \cdot \}$ is the 0-1 indicator function, so \mathcal{L}_{const} will be ∞ if x and x^+ have different class labels. Essentially, \mathcal{L}_{const} governs the capacity for generalizing beyond the source domain, a measure facilitated by the Wasserstein distance.

In addition, the task model learns from domain augmentations with the assistance of a WAE, denoted as V , with parameters ψ . V consists of an encoder $Q(x|e)$ and a decoder $G(x|e)$, where x and e denote the input and bottleneck embedding respectively. Also, they utilized Maximum Mean Discrepancy (MMD) as the distance metric (\mathcal{D}_e) to minimize the divergence between $Q(x)$ and a priori distribution $P(e)$. Therefore, we can pre-train V using the following optimization on the source domain \mathcal{S} :

$$\min_{\psi} \|G(Q(x)) - x^2\| + \lambda \mathcal{D}_e(Q(x), p(e)), \tag{10}$$

where λ is a hyper-parameter. Following this, the reconstruction error \mathcal{L}_{relax} should be maximized for domain augmentation:

$$\mathcal{L}_{relax} = \|x^+ - V(x^+)\|^2. \tag{11}$$

WAEs use the Wasserstein distance metric to measure the distribution distance between the input and reconstruction. As such, the pre-trained V is better equipped to capture the distribution of the source domain, and maximizing \mathcal{L}_{relax} ensures larger domain transportation. Finally, the overall loss function is:

$$\mathcal{L}_{ADA} = \mathcal{L}_{task}(\theta; x) - \alpha \mathcal{L}_{const}(\theta; z) + \beta \mathcal{L}_{relax}(\psi; x), \tag{12}$$

where α and β are two hyper-parameters used to balance \mathcal{L}_{const} and \mathcal{L}_{relax} .

M-ADA offers a novel adversarial domain augmentation approach using relaxed constraints; however, it requires pre-training a Wasserstein autoencoder, which subsequently adds computational overhead.

4.10 MetaNorm

While batch normalization is essential in training deep neural networks, small batch sizes and distribution shifts can destabilize batch statistics. MetaNorm, introduced by Du et al. (2020), ensures effective batch normalization under such conditions. It uses a meta-learning setting to infer adaptive normalization statistics from limited samples instead of relying on direct calculations of batch statistics, which can be unreliable with small batches or distribution shifts. Unlike the transductive batch normalization (TBN) method, which uses the same statistics for both meta-train and meta-test phases, MetaNorm applies a non-transductive approach. Given that test samples may not always be available, MetaNorm is designed to learn to generate statistics only from the support set, and at the meta-test time, it directly applies the model to infer statistics for new tasks.

MetaNorm employs two separate feed-forward networks known as hypernetworks. One is for inferring μ , denoted as $f_{\mu}^{\ell}(\cdot)$, and the other for inferring σ , denoted as $f_{\sigma}^{\ell}(\cdot)$, in order to find the statistics for each individual channel in each ℓ convolutional layer in the meta-learning model. In essence, the hypernetworks accept sample activations as input and produce the estimated statistics as output. During the meta-training phase, the estimated statistics of the support set are applied for normalization of both support and query samples:

$$a' = \gamma \left(\frac{a - \mu_s}{\sqrt{\sigma_s^2 + \epsilon}} \right) + \beta, \tag{13}$$

where γ and β are jointly learned with parameters of the hypernetworks during the meta-training and directly applied at meta-test time, and ϵ is a small scalar to prevent division by zero.

As a final step, it minimizes the following KL term to learn the normalization statistics:

$$\sum_i^{|\mathcal{D}^s|} D_{KL}[q_{\phi}(m|a_i)||p_{\theta}(m|\mathcal{D}^s)], \tag{14}$$

where $q(m|a_i)$ and $p(m|\mathcal{D}^s)$ are defined as a Gaussian distribution based on their respective estimated statistics, and also a_i is a sample from the meta-source domain \mathcal{D}^s . MetaNorm learns to generate its own appropriate statistics and applies them to the samples in the meta-target domain. Notably, in the meta-target domain, the KL term is not used; instead, each example generates its own normalization statistics.

Although adding hypernetworks to infer normalization statistics slightly increases the overall model complexity, MetaNorm overcomes the issues of batch normalization with small batches and domain shifts, offering a promising approach to domain generalization through meta-learning.

4.11 DADG

Discriminative Adversarial Domain Generalization (DADG) is a novel framework proposed by Chen et al. (2022) that integrates Discriminative Adversarial Learning (DAL) for extracting domain-invariant features and Meta-learning-based Cross-Domain Validation (Meta-CDV) for training a robust classifier. DAL aims to learn a domain-invariant feature representation to minimize domain variance. It achieves this by training a feature extractor to confuse a domain discriminator, which is tasked with predicting the domain of the features. Further, as depicted in Fig. 1, the Gradient Reversal Layer (GRL) is placed between the feature extractor and the domain discriminator in the DAL component. During forward propagation, the GRL functions as an identity transform, conveying the features unchanged to the discriminator. Conversely, during backpropagation, the GRL reverses the gradient by multiplying it by a negative scalar (e.g., $-\lambda$), thereby updating the parameters of the feature extractor θ^m into θ^{m+1} . This reversed gradient, originating from the discriminator, encourages the feature extractor to create domain-invariant representations that deceive the discriminator.

Subsequently, Meta-CDV applies the feature extractor to supervised learning by training a classifier c_φ with parameters φ . This is done in a meta-learning manner, where the classification model will be trained on the domains seen in the previous step to update θ^{m+1} to θ^{m+2} and update φ^m to φ^{m+1} . The model's performance is then validated on cross domains to enhance the classification model. It should be noted that this evaluation is performed using the updated parameters θ^{m+2} and φ^{m+1} . In essence, the optimization of the classification model involves the third derivative with respect to θ and the second derivative with respect to φ . Particularly, Meta-CDV simulates train/test domain shifts by training the model on seen domains and validating it on a held-out domain, thereby making the classifier more robust (Fig. 10).

DADG learns domain-invariant features using adversarial discriminative learning with GRL, which enables the adversarial mechanism in a simple way. However, the effectiveness of the discriminative adversarial learning component can diminish as the number of source domains increases.

4.12 Uncertainty-guided model generalization

Qiao and Peng (2021) proposed a unique solution for single domain generalization using Bayesian meta-learning, called uncertainty-guided model generalization. The primary strategy of this method is to augment the source domain's capabilities, not only in the input space, as most current data augmentation methods do, but also in the label space, guided by an uncertainty assessment. As illustrated in Fig. 11, instead of directly augmenting the input space, they introduce an auxiliary network $\psi = \{\phi_p, \phi_m\}$ to create feature perturbations h^+ by adding softplus of Gaussian noise $e \sim \mathcal{N}(\mu, \sigma)$ to generate new domain \mathcal{S}^+ from domain \mathcal{S} . The distribution parameters (μ, σ) represent the uncertainty with respect to the backbone θ . This uncertainty is further utilized to predict learnable parameters (a, b, τ) , which are used to construct learnable label mixup. In learnable label mixup, learnable parameters (a, b) are used to define $\lambda \sim \text{Beta}(a, b)$ to mixup \mathcal{S} and \mathcal{S}^+ to achieve in-between domain interpolations. Notably, learnable parameters (a, b) control the direction and strength of domain interpolations. As a result, the following equations are used:





Fig. 10 Illustration of the DADG framework, which combines Discriminative Adversarial Learning (DAL) featuring a Gradient Reversal Layer (GRL) for creating domain-invariant features with Meta-learning-based Cross-Domain Validation (Meta-CDV) to improve the performance of classifiers across unseen domains

Fig. 11 Schematic representation of the uncertainty-guided model generalization approach, where an auxiliary network ψ introduces feature perturbations to form a new domain S^+ . This process utilizes uncertainty in the form of distribution parameters (μ, σ) to guide the creation of domain-augmented data and label mixup, thereby enabling the model to generalize to unseen domains by interpolating between the original and perturbed features and labels



$$\begin{aligned}
 h^+ &= \lambda h + (1 - \lambda)h^+, \\
 y^+ &= \lambda y + (1 - \lambda)\tilde{y},
 \end{aligned}
 \tag{15}$$

where \tilde{y} represents a label-smoothing version of y . Particularly, label smoothing is carried out with a probability of τ , which means that we allocate a value of $\rho \in (0, 1)$ to the true

class and equally distribute $\frac{1-p}{c-1}$ to the others, where c represents the number of classes. Moreover, by learning a and b , the model can learn to modulate the Beta distribution to produce optimal λ values that result in efficient interpolations between \mathcal{S} and \mathcal{S}^+ . Accordingly, this process results in consistent domain shifts in both input and output via ϕ_p and ϕ_m , enhancing generalization across unseen domains.

The authors proposed a practical Bayesian meta-learning framework that optimizes θ and ψ to maximize the posterior $p(\mathcal{S}^+)$. In other words, they utilized Bayesian inference to maximize the posterior of domain augmentations, approximating the distribution of unseen domains.

While uncertainty-guided augmentation provides a mechanism to approximate unseen target distributions, relying solely on randomized sampling procedures to cover the wide range of target domains may be inadequate for real-world distribution shifts.

4.13 Memory-based multi-source meta-learning

Zhao et al. (2021) proposed a memory-based meta-learning method to address multi-source DG for person re-identification (Re-ID) called Memory-based Multi-source Meta-Learning (M^3L). The purpose of person Re-ID is to match persons with the same identity across multiple camera views. Existing works have shown the effectiveness of meta-learning in classification tasks, but its parametric classifier falls short for Re-ID. This is due to the open-set nature of Re-ID tasks, with each domain having numerous and unique identities. Consequently, the M^3L framework is equipped with a memory-based module that applies the identification loss in a non-parametric way, preventing instability in meta-optimization commonly associated with traditional parametric methods. Specifically, this framework maintains a feature memory \mathcal{M}^i for each source domain \mathcal{D}_S^i consisting of n_i slots, where each slot saves the feature centroid of the corresponding identity. Identification loss is then computed using the similarities between these features and memory centroids as follows:

$$\mathcal{L}_M = -\log \frac{\exp(M[i]^T f(x_i)/\tau)}{\sum_{k=1}^{n_i} \exp(M[k]^T f(x_i)/\tau)}, \quad (16)$$

where $f(\cdot)$ is the feature extractor and τ is the temperature factor. In addition, the following triplet loss is utilized to train the model:

$$\mathcal{L}_{Tri} = [d_p - d_n + \delta]_+, \quad (17)$$

where d_p represents the Euclidean distance between an anchor feature and a hard positive feature, while d_n signifies the Euclidean distance between an anchor feature and a hard negative feature. Also, δ is defined as the margin of the triplet loss and $[\cdot]_+ = \max(\cdot, 0)$. As a result, the sum of \mathcal{L}_{Tri} and \mathcal{L}_M constitutes the meta-train loss (\mathcal{L}_{mtr}) and meta-test loss (\mathcal{L}_{mte}).

Furthermore, M^3L incorporates a meta batch normalization layer (MetaBN) that injects meta-train feature statistics into meta-test features, diversifying them and enabling the model to simulate an expanded range of feature variations. In this way, through iterative generalization processes from meta-train to meta-test domains, the model avoids overfitting due to domain bias and acquires domain-invariant representations that generalize well on unseen domains.

During training, as shown in Fig. 12, source domains are divided into meta-train and meta-test sets at each iteration, employing a meta-learning approach to simulate the train-test process of DG. Therefore, it is necessary to consider the following optimization to build generalizable representations:

$$\min_{\Theta} \mathcal{L}_{mtr}(\Theta) + \mathcal{L}_{mte}(\Theta'), \tag{18}$$

where Θ represents the network parameters, while Θ' denotes the parameters of the model optimized by \mathcal{L}_{mtr} .

Although MetaBN further enhances generalization ability by diversifying meta-test features, the memory module adds computational overhead in maintaining and updating centroid features for each identity.

4.14 MetaBIN

Choi et al. (2021) proposed Meta Batch-Instance Normalization (MetaBIN) as an effective heterogeneous single domain generalization method for person re-identification (Re-ID). The main goal is to simulate unsuccessful generalization scenarios by combining batch-instance normalization layers with meta-learning to address challenging cases caused by both batch and instance normalization layers. A key feature of MetaBIN is its learnable balancing parameters between Batch Normalization (BN) and Instance Normalization (IN), which, depending on their bias, cause the DG model to experience under-style-normalization and over-style-normalization scenarios during meta-learning. Under-style-normalization occurs in the BN model when the model struggles to distinguish identities of samples with unexpected styles from unseen target domains. Conversely, over-style-normalization arises in the IN model, which, while efficient at eliminating instance-specific style information, can inadvertently filter out discriminative information; refer to Fig. 13 for more information.




Fig. 12 Overview of the M^3L framework designed for person Re-ID. The training process includes dividing the source domains into one meta-test domain and multiple meta-train domains in each iteration. The model utilizes memory-based identification and triplet loss for the meta-training phase and then optimizes the meta-test loss on a copy of the model that has been updated with the meta-train loss (maximization of inter-class distance). MetaBN is also applied during the meta-test stage to enhance feature diversity (maximization of intra-domain distance). Finally, the aggregated meta-train and meta-test losses are used to update the original model for improved domain generalization

MetaBIN consists of a classifier g_ϕ with parameter ϕ to predict identities and a feature extractor f_θ with parameters θ where $\theta = (\theta_f, \theta_\rho)$. Here, θ_ρ signifies the balancing parameters between BN and IN, and θ_f represents the remaining parameters.

In the base model updating process, classifier parameters ϕ and feature extractor parameters θ_f are updated by minimizing the cross-entropy loss along with the triplet loss according to the following:

$$\mathcal{L}_{tri} = [d_p - d_n + \delta]_+, \tag{19}$$

where d_p represents the Euclidean distance between an anchor feature and the hardest positive sample, while d_n signifies the Euclidean distance between an anchor feature and the hardest negative sample. Also, δ is defined as the margin of the triplet loss and $[\cdot]_+ = \max(\cdot, 0)$.




Fig. 13 Visualized here are scenarios where generalization is unsuccessful, along with the MetaBIN framework. (a) Under-style-normalization where, the BN model fails to recognize identities in novel domains due to a lack of style variation in the training data. (b) Over-style-normalization, where the IN model excessively removes style information, including identity-discriminative features. (c) The main idea behind the MetaBIN framework is to adjust the balance between BN and IN, simulating and learning from both types of normalization errors within a meta-learning framework, thereby enhancing the model’s ability to generalize to unseen domains without overfitting to source domain styles

In the meta-learning step, source domains are split into meta-train set \mathcal{D}_{mtr} and meta-test set \mathcal{D}_{mte} . First, in the meta-training phase, only θ_ρ is updated to obtain θ'_ρ using \mathcal{D}_{mtr} to simulate over-style-normalization and under-style-normalization cases via the aggregation of the following losses: the intra-domain scatter loss (\mathcal{L}_{scat}), the inter-domain shuffle loss (\mathcal{L}_{shuf}), and the triplet loss (\mathcal{L}_{tri}). The intra-domain scatter loss is used to spread the feature distribution for each domain by minimizing the cosine similarity between each sample and its domain centroid. The inter-domain shuffle loss is introduced to shuffle or mix up the distributions across different source domains by minimizing the distance between anchor samples and inter-domain negatives while maximizing the distance between anchor samples and intra-domain negatives. In summary, the intra-domain scatter loss \mathcal{L}_{scat} and inter-domain shuffle loss \mathcal{L}_{shuf} are used to simulate over-style-normalization scenarios where styles are confused. At the same time, adding the triplet loss enhances intra-class compactness regardless of style differences to simulate under-style-normalization. Next, in the meta-testing phase, the model with parameters θ_f and θ'_ρ is evaluated on \mathcal{D}_{mte} , and θ_ρ will be updated using \mathcal{L}_{tri} . Note that θ_ρ contains channel-wise balancing parameters ρ for each normalization layer, with some channels biasing θ_ρ toward IN to remove unessential style info, while others retain BN properties.

MetaBIN improves generalization to unseen domains without needing extra networks or data augmentation; however, it may not perform well on datasets with complex variations or highly different styles between domains, as it relies on the simulated scenarios to learn a robust representation.

5 Datasets and evaluations

In the DG and domain adaptation (DA) fields, datasets are specifically constructed to represent various 'domains'-each a unique distribution of data that captures a certain variation in the input space. These variations can include changes in visual appearance, sensor modality, environmental conditions, or even different tasks. The purpose of these multi-domain datasets is to simulate real-world scenarios where a model trained on limited domains must perform well on another unseen domain. The richness and diversity of domains within these datasets are pivotal to developing and benchmarking algorithms that can generalize beyond their training data.

Several important datasets, commonly used in the fields of DG and DA, have been examined in studies that leverage meta-learning for DG. Given the fundamental overlap between DA and DG, datasets from DA can also be applied to DG purposes, and vice versa. Each dataset, associated with a specific application, is summarized in Table 3 and will be discussed in this section. This discussion will be followed by an overview of the evaluation strategies commonly employed in DG, which differ from the more straightforward evaluation systems used in DA.

5.1 Datasets

In the DG landscape, datasets serve as critical benchmarks for evaluating algorithm performance across varied scenarios. We begin with Rotated MNIST (Ghifary et al. 2015), a synthetic dataset based on the original MNIST. It comprises six domains, each containing images of various digits modified by rotations ranging from 0 to 90 degrees at 15-degree intervals. This dataset is instrumental in assessing DG algorithms for handwritten digit

Table 3 Overview of Key Datasets Used in Domain Generalization Based on Their Application, Highlighting the Number of Domains, Number of Classes, Number of Samples, and Their Description

	# Domains	# Classes	# Samples	Description
Handwritten Digit Recognition				
Rotated MNIST (Chifary et al. 2015)	6	10	70,000	Rotations (0, 15, 30, 45, 60, and 75)
Digits-Five (Zhou et al. 2020)	5	10	215,695	MNIST (LeCun et al. 1998), MNIST-M (Ganin and Lempitsky 2015), SVHN (Netzer et al. 2011), SYN (Ganin and Lempitsky 2015), and USPS (Hull 1994)
Object Recognition				
VLCS (Fang et al. 2013)	4	5	10,729	VOC2007 (Everingham et al. 2010), LabelMe (Russell et al. 2008), Caltech-101 (Fei-Fei et al. 2004), and SUN09 (Xiao et al. 2010)
Office-Home (Fang et al. 2013)	4	65	15,588	Art, Clipart, Product, and Real-World
PACS (Li et al. 2017)	4	7	9,991	Photo, Art painting, Cartoon, and Sketch
CIFAR-10-C (Hendrycks and Dietterich 2019)	-	10	60,000	Artificial Corruptions of CIFAR-10 (Krizhevsky et al. 2009)
Visual Decathlon (Rebuffi et al. 2017)	10	-	1,659,142	10 Different Datasets (Maji et al. 2013; Krizhevsky et al. 2009; Muder and Gavrilu 2006; Cimpoi et al. 2014; Stalkamp et al. 2012; Russakovsky et al. 2015; Lake et al. 2015; Netzer et al. 2011; Soomro et al. 2012; Bilen et al. 2016; Nilsback and Zisserman 2008)
Action Recognition				
IXMAS (Weinland et al. 2006)	5	11	1,650	5 Different Camera Views, and 11 Different Subjects (Li et al. 2018)
Semantic Segmentation				
SYNTIA (Ros et al. 2016)	15	13	2,700	5 Different Locations, and 5 Different Weather Conditions (Volpi et al. 2018)
Person Re-Identification				
Market-Duke (Ristani et al. 2016)	2	-	69,079	Market-1501 (Zheng et al. 2015) and DukeMTMC-ReID (Zheng et al. 2017)
Sentiment Classification				
Amazon Reviews (Blitzer et al. 2006)	4	2	> 340,000	Books, DVDs, Electronics, and Kitchen Appliances

recognition. Additionally, to serve a similar purpose, another dataset, known as Digits-Five (Zhou et al. 2020), has been introduced, which includes five sub-datasets: MNIST (LeCun et al. 1998), MNIST-M (Ganin and Lempitsky 2015), SVHN (Netzer et al. 2011), SYN (Ganin and Lempitsky 2015), and USPS (Hull 1994), where each of them can be considered a different domain.

For object recognition, a task particularly sensitive to domain shifts, notable benchmarks include VLCS (Fang et al. 2013), Office-Home (Fang et al. 2013), PACS (Li et al. 2017), CIFAR-10-C (Hendrycks and Dietterich 2019), and VD (Rebuffi et al. 2017). VLCS contains images from four domains (VOC2007 Everingham et al. 2010, LabelMe Russell et al. 2008, Caltech-101 Fei-Fei et al. 2004, and SUN09 Xiao et al. 2010) (see Fig. 14) and five classes. Office-Home includes images of objects from Art, Clipart, Product, and Real-World domains, spanning 65 categories. PACS covers four contrasting domains (Photo, Art painting, Cartoon, and Sketch as depicted in Fig. 14) with seven object classes. VD contains ten diverse domains, including handwritten characters, pedestrians, traffic signs, etc., with varying image categories and sizes suitable for heterogeneous DG. Also, CIFAR-10-C (Hendrycks and Dietterich 2019) is a robustness benchmark comprising CIFAR-10 (Krizhevsky et al. 2009) test images corrupted through 19 distortion types across five severity levels.

IXMAS (Weinland et al. 2006) dataset addresses the domain generalization challenge in action recognition, which relies on learning generalizable representations within video understanding. It includes recordings of 11 actions performed by various actors, captured from multiple angles and with different cameras, to introduce domain shifts.

In autonomous vehicle development, semantic segmentation is essential, yet deep neural networks have not fully bridged the performance gap in unexplored environments. SYNTHIA (Ros et al. 2016) dataset, with its synthetic images of different locations and weather conditions, supports research to overcome this limitation.

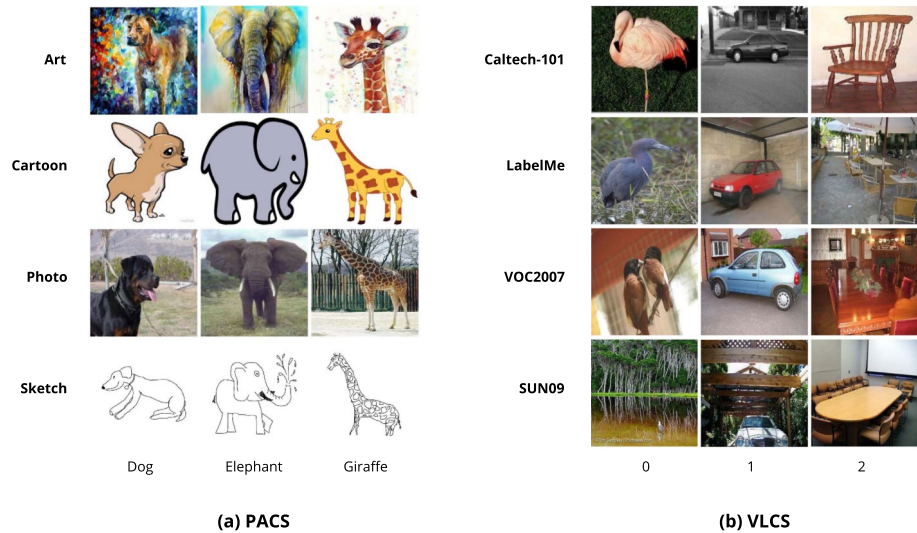


Fig. 14 Example images from two prominent domain generalization benchmarks, illustrating different types of domain shifts. In (a), the PACS dataset highlights domain shifts primarily due to changes in image style. In (b), the VLCS dataset reveals domain shifts caused by variations in environment, scene, and viewpoint, reflecting dataset-specific biases

Person Re-Identification (Re-ID) is another area where DG is paramount, especially for security and surveillance. While traditional Re-ID has been confined to the same camera views, cross-dataset Re-ID seeks to generalize from source to target views across varied resolutions, viewpoints, and lighting conditions. Key datasets like Market-1501 (Zheng et al. 2015) and DukeMTMC-ReID (Zheng et al. 2017) are essential in the advancement of this field.

Lastly, the Amazon Reviews (Blitzer et al. 2006) dataset, containing reviews from domains such as books, DVDs, electronics, and kitchen appliances, is crucial for sentiment classification in natural language processing.

5.2 Evaluations

For evaluating domain generalization algorithms, three principal strategies are commonly employed: (i) *Leave-one-domain-out validation*, which is the most prevalent strategy that is applicable when multiple source domains are available during training. It involves leaving out one source domain as the validation set while using the others for training. However, the final results heavily rely on the selection of the validation domain, leading to results that may lack stability. (ii) *Training-domain validation*, where a subset of the training data is held out for validation to select the best model. However, since there is still a divergence between real-world unseen data and the training subset, it may not achieve optimal performance. (iii) *Test-domain validation*: which involves using a random subset of the target domain data for model selection. This method can lead to the best performance since the validation and test data share the same distribution. However, it presupposes access to the target domain data, which may not be feasible in real-world applications. Using test domain data for model selection can yield misleading results, as an unrepresentative or excessively easy subset may result in optimistic performance, or a challenging subset may lead to pessimistic results (Gulrajani and Lopez-Paz 2020). Careful interpretation of results obtained using this method is therefore necessary.

In general, the average accuracy is used as the primary metric to report the performance of a DG model across held-out domains. For image classification on Cifar-10-C (Hendrycks and Dietterich 2019), the mean Corruption Error (mCE) metric determines model robustness to corruptions, while Relative mCE (RmCE) compares corruption robustness relative to a baseline by accounting for clean data performance. This enables fairer comparisons between models. For semantic segmentation models evaluated on the SYNTHIA (Ros et al. 2016) dataset, the standard mean Intersection Over Union (mIoU) is calculated for each unseen domain to quantify segmentation accuracy. Additionally, the VD-score metric assesses models trained on the diverse VD (Rebuffi et al. 2017) dataset by evaluating performance across its ten distinct image classification tasks.

6 Applications

In real-world scenarios, finding sufficient labeled data to effectively train a model often proves challenging. This difficulty is compounded when the model encounters OOD data during testing. As such, domain generalization through meta-learning has emerged as a highly desirable solution. It enables a model to be trained on one or multiple domains and to be applied to unseen domains without the need for retraining. This approach not only alleviates the risk of overfitting but also diminishes the costs and time associated

with domain-specific retraining, thus boosting the model's adaptability and operational efficiency.

The applicability of domain generalization models is evident across various disciplines, such as medical imaging, intelligent fault diagnosis, computer vision, and natural language processing. For example, Liu et al. (2020) utilized meta-learning for domain generalization to segment prostate MRI images, demonstrating its potential in medical analysis. Ren et al. (2023) proposed Meta-GENE, a model-agnostic meta-learning framework designed for fault diagnosis in industrial prognostics and health management scenarios, which can operate in both homogeneous and heterogeneous DG settings. Wang et al. (2020) applied the technique for cross-lingual semantic parsing, and Balaji et al. (2018) utilized it for sentiment analysis within the Amazon Review dataset. In the area of person re-identification (Re-ID), both Choi et al. (2021) and Zhao et al. (2021) leveraged meta-learning for DG to recognize individuals across varying postures and perspectives. Moreover, meta-learning for DG has been applied in cross-view action recognition by Li et al. (2019b) and robot control through reinforcement learning by Li et al. (2018). Qiao et al. (2020) showcased its utility in semantic segmentation for street scenes, and Dou et al. (2019) explored its application in multi-site brain tissue segmentation.

Further extending the applications, DG has been employed in face anti-spoofing (Shao et al. 2019), image compression (Zhang et al. 2021), and various medical diagnostics, including Parkinson's disease detection (Muandet et al. 2013), activity recognition (Erfani et al. 2016), chest X-ray analysis (Mahajan et al. 2021), and EEG-based seizure detection (Ayodele et al. 2020). Other studies have used domain generalization for speech utterance recognition (Shankar et al. 2018; Piratla et al. 2020), fault diagnosis (Li et al. 2020; Zheng et al. 2020; Liao et al. 2020), and brain-computer interaction (Han and Jeong 2021). Additionally, DG has found its place in brain-computer interfacing (Han and Jeong 2021), and promising advancements have been made in time series forecasting for financial markets (Du et al. 2021).

The extensive applications of meta-learning for DG position it as a focal point of interest for future research and commercial endeavors. The algorithms developed under this paradigm equip models with rapid generalization and adaptation capabilities, even in the absence of target domain data. Such a capacity is a stepping stone towards creating more powerful AI systems capable of tackling diverse real-world problems. By integrating domain generalization with meta-learning, we can unlock zero-shot learning capabilities for novel, unseen domains, enabling generalization across various tasks and domains without requiring additional data.

7 Discussion and future directions

This survey has examined diverse strategies for DG using meta-learning. The presented taxonomy provides a structured view of the field, categorizing approaches based on the treatment of the input domain and classifier training strategies, as discussed in Sect. 3. Initially, DG models employing meta-learning predominantly focused on the minimization of inter-domain distances and the minimization of intra-class distances, aiming to find a common representation across domains and cohesive grouping of class instances, as this approach is the most intuitive way to develop DG methods using meta-learning. However, recent algorithms have been exploring the effectiveness of the Maximization of Inter-Class Distances or the Maximization of Intra-domain Distances to fully exploit the advantages of

meta-learning. For instance, MASF (Dou et al. 2019) leverages the Maximization of Inter-Class Distances. It is the first model in meta-learning for DG that follows this paradigm, resulting in improved class discriminability and enhanced generalization to unseen domains. Furthermore, M-ADA first introduced the concept of maximizing intra-domain distances through meta-learning, laying the groundwork for future models that employ this strategy. This approach aims to diversify the input feature space, thereby enhancing the robustness of the feature extractor, particularly when there are limited training domains available. Recently, methods such as Memory-Based Multi-Source Meta-Learning (M^3L) (Zhao et al. 2021) and MetaBIN (Choi et al. 2021) have been developed, leveraging the Maximization of both Intra-Domain and Inter-Class Distances. This strategic combination is designed to diversify the feature space within each domain while simultaneously ensuring that features remain discriminative for classification tasks. These methods are beneficial for use cases where different classes are not easily distinguishable, such as Re-Identification (Re-ID), by further diversifying input features and enhancing class discriminability.

The choice of strategy in DG models tends to vary depending on the nature and number of available training domains. When multiple diverse domains are available for training, employing a technique that minimizes inter-domain distances is a natural choice, as it exploits the diversity of the available data to extract domain-invariant features. On the other hand, if the training data is limited to a few domains, models that maximize inter-domain distances are preferable, as they focus on augmenting the input or feature space to create a generalizable model capable of performing well across various domains during inference. In particular, recent advances in generative AI allows more realistic synthetic data generation and diversification, which further enhance algorithms that maximize intra-domain distances.

Additionally, in scenarios where data points within a dataset are highly similar to each other, it is beneficial to utilize models that emphasize the Maximization of Inter-Class Distances. This approach encourages a clear separation between classes, enabling the model to learn effective embedding representations that facilitate discrimination between classes based on distance.

To alleviate the weaknesses of existing meta-learning approaches for DG which has been explored in Fig. 2, a promising future direction is to create a novel distance metric that captures the intra- and inter-domain characteristics. This metric can be utilized as a guidance for selecting appropriate methods to enable rapid adaptation to new domains.

It is also worth mentioning that the exploration of biologically plausible models in meta-learning has gained some attention and aims to bridge the gap between artificial learning systems and the mechanisms underlying learning in biological organisms. While artificial neural networks have achieved impressive results on meta-learning benchmarks, they differ substantially from the learning processes in biological systems, such as the human brain. As of recent, there has been growing interest in developing meta-learning models that are more biologically grounded and mimic the human learning system. One such model incorporates a memory system designed to prevent catastrophic forgetting, functioning similarly to an episodic memory system (Khoei et al. 2024). This model enhances the generalization capabilities of spiking neural networks by simulating an efficient episodic memory capable of storing vast amounts of information and linking similar underlying patterns. A promising direction for future development is to create innovative policies for memory adaptation within this model to handle more complex tasks, such as DG. By implementing these policies, the model can better ignore noise and outlier data and map similar classes from different domains to the same memory representation. This would empower the model to effectively tackle data from diverse domains, enhancing its generalization capabilities.

DG is a critical challenge in federated learning, where multiple clients collaboratively train a model without sharing their raw data, enhancing privacy and security (Zhang et al. 2021). However, the domain of the target test data on the server can differ greatly from the training data of each client. This discrepancy leads to decreased performance of the federated model on the target domain. RFDG (Guan et al. 2023) has integrated domain generalization into federated learning via a reinforcement learning-based feature decorrelation policy that enables clients to reweight samples from a global perspective to learn domain-invariant knowledge. The replay mechanism also addresses challenges posed by mini-batch training. This field represents a potential area where we should incorporate meta-learning for DG to enhance adaptability to different clients with limited labeled data. By leveraging meta-learning techniques, the federated model can quickly adapt to new clients by learning from a small number of labeled examples, even when there are significant distributional shifts between the source and target domains. Accordingly, meta-learning can be used to learn a shared feature extractor that is robust to domain shifts, allowing the model to extract transferable knowledge from the source domains and effectively apply it to the target domain, thereby improving its ability to generalize to diverse and unseen domains encountered in real-world federated learning scenarios.

Although the presented approaches have made notable progress in domain generalization using meta-learning, there are still interesting avenues for future research. One such direction is Generalizable Label Distribution Learning (GLDL), which aims to learn a Label Distribution Learning (LDL) model that can generalize well to unseen target domains. LDL is a machine learning paradigm that assigns a label distribution to each instance, indicating the relative importance or description degree of each label. By reflecting the varying degrees of association between labels and instances, LDL provides a more detailed representation of label information. Consequently, leveraging the full distribution of labels can enhance model performance on complex tasks. A recent work by Zhao et al. (2023) has expanded the scope of domain generalization research by addressing the specific challenges of LDL. In this work, the DICE framework was proposed, which learns to extract domain-invariant feature-label correlations and label-label correlations. It achieves this by aligning the prior distributions and label correlation matrices across different source domains. This work opens up new possibilities for applying meta-learning techniques to the GLDL problem, potentially enabling more robust and generalizable LDL models that can handle distribution shifts between source and target domains. Exploring the integration of meta-learning strategies with GLDL could lead to further advancements in this emerging area of research.

It is also important to note another key challenge that still needs to be addressed in this area of research: enabling algorithms to go beyond pattern recognition to understand and exploit data causality (Rothenhäusler et al. 2021). This understanding is essential for models to make consistent predictions across domains, particularly when existing correlations are unreliable (Sheth et al. 2022; Lv et al. 2022; Chen et al. 2023). Therefore, it is essential to provide meta-learning models with proper inductive biases that can guide the learning process toward causal inference. Inductive biases help in shaping the hypothesis space, enabling the model to prioritize learning generalizable patterns that reflect underlying causal relationships as opposed to spurious correlations. Despite challenges, meta-learning has valuable potential for DG if steered to grasp causal structures rather than merely mimic data. Developing meta-learning frameworks that inherently capture causal relationships will be pivotal for achieving true domain-agnostic capabilities.

8 Conclusions

Meta-learning has experienced rapid growth in interest in recent years and has been applied to many significant research topics in machine learning, including DG. In this paper, we focus on meta-learning for DG, a promising field that has attracted many researchers. We provide a comprehensive review of existing methods and present a detailed taxonomy based on two primary axes crucial for designing effective models, as well as an overview of relevant datasets, benchmarks, and applications. The taxonomy provided in this survey offers a roadmap for understanding the diverse strategies that underpin current research in the field. It is clear that both the feature extraction and classifier training processes are critical to the development of DG models. Furthermore, we share insights from our analysis of these methods and identify potential research challenges that could guide future research directions. We hope that this survey will assist newcomers and practitioners in navigating this growing field and will also highlight opportunities for future research.

Author contributions Arsham composed the primary manuscript text and prepared the figures. Yinan significantly contributed by revising the entire manuscript, composing the discussion section, and refining the figures. Robert focused on revising the manuscript text and figure captions. All authors collaboratively developed the taxonomy for this survey, reviewed the manuscript with constructive comments, and contributed to the final refinement.

Funding Open access funding provided by Chalmers University of Technology.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ayodele KP, Ikezogwo WO, Komolafe MA, Ogunbona P (2020) Supervised domain generalization for integration of disparate scalp EEG datasets for automatic epileptic seizure detection. *Comput Biol Med* 120:103757
- Balaji Y, Sankaranarayanan S, Chellappa R (2018) Metareg: towards domain generalization using meta-regularization. *Advances in Neural Information Processing Systems* 31
- Ben-David S, Blitzer J, Crammer K, Pereira F (2006) Analysis of representations for domain adaptation. *Advances in Neural Information Processing Systems* 19

- Bilen H, Fernando B, Gavves E, Vedaldi A, Gould S (2016) Dynamic image networks for action recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3034–3042
- Blitzer J, McDonald R, Pereira F (2006) Domain adaptation with structural correspondence learning. In: Proceedings of the 2006 conference on empirical methods in natural language processing, pp. 120–128
- Bousmalis K, Trigeorgis G, Silberman N, Krishnan D, Erhan D (2016) Domain separation networks. *Advances in Neural Information Processing Systems* **29**
- Castro FM, Marín-Jiménez MJ, Guil N, Schmid C, Alahari K (2018) End-to-end incremental learning. In: Proceedings of the European conference on computer vision (ECCV), pp. 233–248
- Chattopadhyay P, Balaji Y, Hoffman J (2020) Learning to balance specificity and invariance for in and out of domain generalization. In: Computer Vision–ECCV 2020: 16th European conference, Glasgow, UK, Proceedings, Part IX 16, pp. 301–318. Springer, 23–28 August 2020
- Chen K, Zhuang D, Chang JM (2022) Discriminative adversarial domain generalization with meta-learning based cross-domain validation. *Neurocomputing* 467:418–426
- Chen J, Gao Z, Wu X, Luo J (2023) Meta-causal learning for single domain generalization. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 7683–7692
- Choi S, Kim T, Jeong M, Park H, Kim C (2021) Meta batch-instance normalization for generalizable person re-identification. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 3425–3435
- Cimpoi M, Maji S, Kokkinos I, Mohamed S, Vedaldi A (2014) Describing textures in the wild. In: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3606–3613
- Csurka G (2017) A comprehensive survey on domain adaptation for visual applications. *Domain adaptation in computer vision applications*, 1–35
- De Lange M, Aljundi R, Masana M, Parisot S, Jia X, Leonardis A, Slabaugh G, Tuytelaars T (2021) A continual learning survey: defying forgetting in classification tasks. *IEEE Trans Pattern Anal Mach Intell* 44(7):3366–3385
- De Lange M, Aljundi R, Masana M, Parisot S, Jia X, Leonardis A, Slabaugh G, Tuytelaars T (2019) Continual learning: a comparative study on how to defy forgetting in classification tasks. 2(6):2 arXiv preprint [arXiv:1909.08383](https://arxiv.org/abs/1909.08383)
- Dou Q, Castro D, Kamnitsas K, Glocker B (2019) Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*. <https://doi.org/10.48550/arXiv.1910.13580>
- Du Y, Wang J, Feng W, Pan S, Qin T, Xu R, Wang C (2021) Adarnn: Adaptive learning and forecasting of time series. In: Proceedings of the 30th ACM international conference on information & knowledge management, pp. 402–411
- Du Y, Xu J, Xiong H, Qiu Q, Zhen X, Snoek CG, Shao L (2020) Learning to learn with variational information bottleneck for domain generalization. In: Computer Vision–ECCV 2020: 16th European conference, Glasgow, Proceedings, Part X 16, pp. 200–216. Springer, 23–28 August 2020
- Du Y, Zhen X, Shao L, Snoek CG (2020) Metanorm: learning to normalize few-shot batches across domains. In: International conference on learning representations
- Erfani S, Baktashmotlagh M, Moshtaghi M, Nguyen X, Leckie C, Bailey J, Kotagiri R (2016) Robust domain generalisation by enforcing distribution invariance. In: Proceedings of the twenty-fifth international joint conference on artificial intelligence (IJCAI-16). AAAI Press, Washington, pp. 1455–1461
- Everingham M, Van Gool L, Williams CK, Winn J, Zisserman A (2010) The pascal visual object classes (voc) challenge. *Int J Comput Vision* 88:303–338
- Fang C, Xu Y, Rockmore DN (2013) Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: Proceedings of the IEEE international conference on computer vision, pp. 1657–1664
- Farahani A, Voghoei S, Rasheed K, Arabnia HR (2021) A brief review of domain adaptation. *Advances in data science and information engineering: proceedings from ICDATA 2020 and IKE 2020*, 877–894
- Fei-Fei L, Fergus R, Perona P (2004) Learning generative visual models from few training examples: an incremental bayesian approach tested on 101 object categories. In: 2004 conference on computer vision and pattern recognition workshop, pp. 178–178. IEEE
- Finn C, Abbeel P, Levine S (2017) Model-agnostic meta-learning for fast adaptation of deep networks. In: International conference on machine learning, pp. 1126–1135. PMLR
- Ganin Y, Ustinova E, Ajakan H, Germain P, Laroche H, Laviolette F, March M, Lempitsky V (2016) Domain-adversarial training of neural networks. *J Mach Learn Res* 17(59):1–35
- Ganin Y, Lempitsky V (2015) Unsupervised domain adaptation by backpropagation. In: International conference on machine learning, pp. 1180–1189. PMLR

- Ghifary M, Kleijn WB, Zhang M, Balduzzi D (2015) Domain generalization for object recognition with multi-task autoencoders. In: Proceedings of the IEEE international conference on computer vision, pp. 2551–2559
- Guan Z, Li Y, Pan Z, Liu Y, Xue Z (2023) RFDG: reinforcement federated domain generalization. *IEEE Trans Knowl Data Eng.* <https://doi.org/10.1109/TKDE.2023.3301036>
- Gulrajani I, Lopez-Paz D (2020) In search of lost domain generalization. arXiv preprint [arXiv:2007.01434](https://arxiv.org/abs/2007.01434)
- Han D-K, Jeong J-H (2021) Domain generalization for session-independent brain-computer interface. In: 2021 9th international winter conference on brain-computer interface (BCI). IEEE, New York, pp. 1–5
- He H, Chen S, Li K, Xu X (2011) Incremental learning from stream data. *IEEE Trans Neural Netw* 22(12):1901–1914
- He J, Mao R, Shao Z, Zhu F (2020) Incremental learning in online scenario. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pp. 13926–13935
- Hendrycks D, Dietterich T (2019) Benchmarking neural network robustness to common corruptions and perturbations. arXiv preprint [arXiv:1903.12261](https://arxiv.org/abs/1903.12261)
- Hoi SC, Sahoo D, Lu J, Zhao P (2021) Online learning: a comprehensive survey. *Neurocomputing* 459:249–289
- Hospedales T, Antoniou A, Micaelli P, Storkey A (2021) Meta-learning in neural networks: a survey. *IEEE Trans Pattern Anal Mach Intell* 44(9):5149–5169
- Huisman M, Van Rijn JN, Plaata A (2021) A survey of deep meta-learning. *Artif Intell Rev* 54(6):4483–4541
- Hull JJ (1994) A database for handwritten text recognition research. *IEEE Trans Pattern Anal Mach Intell* 16(5):550–554
- Jia C, Zhang Y (2022) Meta-learning the invariant representation for domain generalization. *Mach Learn*, 1–21
- Khoee AG, Javaheri A, Kheradpisheh SR, Ganjtabesh M (2024) Meta-learning in spiking neural networks with reward-modulated stdp. *Neurocomputing* 600:128173
- Khosla A, Zhou T, Malisiewicz T, Efros AA, Torralba A (2012) Undoing the damage of dataset bias. In: *Computer Vision–ECCV 2012: 12th European conference on computer vision, Florence, Italy, Proceedings, Part I 12*, pp. 158–171. Springer, 7–13 October 2012
- Krizhevsky A, Hinton G, et al. (2009) Learning multiple layers of features from tiny images. In: *Technical Report TR-2009, University of Toronto, Toronto*
- Lake BM, Salakhutdinov R, Tenenbaum JB (2015) Human-level concept learning through probabilistic program induction. *Science* 350(6266):1332–1338
- LeCun Y, Bottou L, Bengio Y, Haffner P (1998) Gradient-based learning applied to document recognition. *Proceed IEEE* 86(11):2278–2324
- Li X, Zhang W, Ma H, Luo Z, Li X (2020) Domain generalization in rotating machinery fault diagnostics using deep neural networks. *Neurocomputing* 403:409–420
- Liao Y, Huang R, Li J, Chen Z, Li W (2020) Deep semisupervised domain generalization network for rotary machinery fault diagnosis under variable speed. *IEEE Trans Instrum Measure* 69(10):8064–8075
- Li H, Pan SJ, Wang S, Kot AC (2018) Domain generalization with adversarial feature learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5400–5409
- Liu Q, Dou Q, Heng P-A (2020) Shape-aware meta-learning for generalizing prostate mri segmentation to unseen domains. In: *Medical image computing and computer assisted intervention–MICCAI 2020: 23rd International conference, Lima, Proceedings, Part II 23*, pp. 475–485. Springer, 4–8 October 2020
- Liu J, Shen Z, He Y, Zhang X, Xu R, Yu H, Cui P (2021) Towards out-of-distribution generalization: a survey. arXiv preprint [arXiv:2108.13624](https://arxiv.org/abs/2108.13624)
- Li D, Yang Y, Song Y-Z, Hospedales T (2018) Learning to generalize: meta-learning for domain generalization. In: *Proceedings of the AAAI conference on artificial intelligence*, vol. 32
- Li D, Yang Y, Song Y-Z, Hospedales T (2020) Sequential learning for domain generalization. In: *European conference on computer vision*, pp. 603–619. Springer
- Li D, Yang Y, Song YZ, Hospedales TM (2017) Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE international conference on computer vision*, pp. 5542–5550
- Li Y, Yang Y, Zhou W, Hospedales T (2019) Feature-critic networks for heterogeneous domain generalization. In: *International conference on machine learning*, pp. 3915–3924. PMLR
- Li D, Zhang J, Yang Y, Liu C, Song Y-Z, Hospedales TM (2019) Episodic training for domain generalization. In: *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1446–1455
- Long M, Cao Y, Wang J, Jordan M (2015) Learning transferable features with deep adaptation networks. In: *International conference on machine learning*, pp. 97–105. PMLR

- Long M, Zhu H, Wang J, Jordan MI (2016) Unsupervised domain adaptation with residual transfer networks. *Advances in Neural Information Processing Systems* **29**, pp. 136–144
- Lv F, Liang J, Li S, Zang B, Liu CH, Wang Z, Liu D (2022) Causality inspired representation learning for domain generalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8046–8056
- Mahajan D, Tople S, Sharma A (2021) Domain generalization using causal matching. In: *International conference on machine learning*, pp. 7313–7324. PMLR
- Maji S, Rahtu E, Kannala J, Blaschko M, Vedaldi A (2013) Fine-grained visual classification of aircraft. arXiv preprint [arXiv:1306.5151](https://arxiv.org/abs/1306.5151)
- Muandet K, Balduzzi D, Schölkopf B (2013) Domain generalization via invariant feature representation. In: *International conference on machine learning*, pp. 10–18. PMLR
- Munder S, Gavrilu DM (2006) An experimental study on pedestrian classification. *IEEE Trans Pattern Anal Mach Intell* 28(11):1863–1868
- Netzer Y, Wang T, Coates A, Bissacco A, Wu B, Ng AY (2011) Reading digits in natural images with unsupervised feature learning. In: *NIPS workshop on deep learning and unsupervised feature learning*. Granada
- Nilsback ME, Zisserman A (2008) Automated flower classification over a large number of classes. In: *2008 sixth Indian conference on computer vision, graphics & image processing*, pp. 722–729. IEEE
- Pan SJ, Yang Q (2009) A survey on transfer learning. *IEEE Trans Knowl Data Eng* 22(10):1345–1359
- Patel VM, Gopalan R, Li R, Chellappa R (2015) Visual domain adaptation: a survey of recent advances. *IEEE Signal Process Mag* 32(3):53–69
- Piratla V, Netrapalli P, Sarawagi S (2020) Efficient domain generalization via common-specific low-rank decomposition. In: *International conference on machine learning*, pp. 7728–7738. PMLR
- Qiao F, Peng X (2021) Uncertainty-guided model generalization to unseen domains. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6790–6800
- Qiao F, Zhao L, Peng X (2020) Learning to learn single domain generalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 12556–12565
- Ravi S, Larochelle H (2016) Optimization as a model for few-shot learning. In: *International conference on learning representations*
- Rebuffi S-A, Bilen H, Vedaldi A (2017) Learning multiple visual domains with residual adapters. *Advances in Neural Information Processing Systems* **30**, pp. 506–516
- Ren L, Mo T, Cheng X (2023) Meta-learning based domain generalization framework for fault diagnosis with gradient aligning and semantic matching. *IEEE Trans Ind Inf* 20:754–764
- Ristani E, Solera F, Zou R, Cucchiara R, Tomasi C (2016) Performance measures and a data set for multi-target, multi-camera tracking. In: *European conference on computer vision*, pp. 17–35. Springer
- Ros G, Sellart L, Materzynska J, Vazquez D, Lopez AM (2016) The synthia dataset: a large collection of synthetic images for semantic segmentation of urban scenes. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3234–3243
- Rothenhäusler D, Meinshausen N, Bühlmann P, Peters J (2021) Anchor regression: heterogeneous data meet causality. *J Royal Stat Soc Ser B: Stat Methodol* 83(2):215–246
- Russakovsky O, Deng J, Su H, Krause J, Satheesh S, Ma S, Huang Z, Karpathy A, Khosla A, Bernstein M et al (2015) Imagenet large scale visual recognition challenge. *Int J Comput Vision* 115:211–252
- Russell BC, Torralba A, Murphy KP, Freeman WT (2008) Labelme: a database and web-based tool for image annotation. *Int J Comput Vision* 77:157–173
- Sener O, Koltun V (2018) Multi-task learning as multi-objective optimization. *Advances in Neural Information Processing Systems* **31**, pp. 527–538
- Shankar S, Piratla V, Chakrabarti S, Chaudhuri S, Jyothi P, Sarawagi S (2018) Generalizing across domains via cross-gradient training. arXiv preprint [arXiv:1804.10745](https://arxiv.org/abs/1804.10745)
- Shao R, Lan X, Li J, Yuen PC (2019) Multi-adversarial discriminative deep domain generalization for face presentation attack detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10023–10031
- Sheth P, Moraffah R, Candan KS, Raglin A, Liu H (2022) Domain generalization—a causal perspective. arXiv preprint [arXiv:2209.15177](https://arxiv.org/abs/2209.15177)
- Shu Y, Cao Z, Wang C, Wang J, Long M (2021) Open domain generalization with domain-augmented meta-learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9624–9633
- Snell J, Swersky K, Zemel R (2017) Prototypical networks for few-shot learning. *Advances in Neural Information Processing Systems* **30**, pp. 4080–4090

- Soomro K, Zamir AR, Shah M (2012) Ucf101: a dataset of 101 human actions classes from videos in the wild. arXiv preprint [arXiv:1212.0402](https://arxiv.org/abs/1212.0402)
- Stallkamp J, Schlipsing M, Salmen J, Igel C (2012) Man vs. computer: benchmarking machine learning algorithms for traffic sign recognition. *Neural Netw* 32:323–332
- Sugiyama M, Storkey AJ (2006) Mixture regression for covariate shift. *Advances in Neural Information Processing Systems* **19**, pp. 1337–1344
- Tripurani N, Jordan M, Jin C (2020) On the theory of transfer learning: the importance of task diversity. *Advances in Neural Information Processing Systems* 33:7852–7862
- Tzeng E, Hoffman J, Zhang N, Saenko K, Darrell T (2014) Deep domain confusion: maximizing for domain invariance. arXiv preprint [arXiv:1412.3474](https://arxiv.org/abs/1412.3474)
- Vettoruzzo A, Bouguelia M-R, Vanschoren J, Rognvaldsson T, Santosh K (2024) Advances and challenges in meta-learning: a technical review. *IEEE transactions on pattern analysis and machine intelligence*
- Vilalta R, Drissi Y (2002) A perspective view and survey of meta-learning. *Artif Intell Rev* 18:77–95
- Volpi R, Namkoong H, Sener O, Duchi JC, Murino V, Savarese S (2018) Generalizing to unseen domains via adversarial data augmentation. *Advances in Neural Information Processing Systems*. <https://doi.org/10.48550/arXiv.1805.12018>
- Wang M, Deng W (2018) Deep visual domain adaptation: a survey. *Neurocomputing* 312:135–153
- Wang J, Lan C, Liu C, Ouyang Y, Qin T, Lu W, Chen Y, Zeng W, Yu P (2022) Generalizing to unseen domains: a survey on domain generalization. *IEEE transactions on knowledge and data engineering*
- Wang B, Lapata M, Titov I (2020) Meta-learning for domain generalization in semantic parsing. arXiv preprint [arXiv:2010.11988](https://arxiv.org/abs/2010.11988)
- Wang Y, Li H, Kot AC (2020) Heterogeneous domain generalization via domain mixup. In: *ICASSP 2020-2020 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, pp. 3622–3626. IEEE
- Wang Q, Lv Y, Xie Z, Huang J, et al. (2024) A simple yet effective strategy to robustify the meta learning paradigm. *Advances in Neural Information Processing Systems* **36**
- Weinland D, Ronfard R, Boyer E (2006) Free viewpoint action recognition using motion history volumes. *Comput Vision Image Underst* 104(2–3):249–257
- Wu Y, Chen Y, Wang L, Ye Y, Liu Z, Guo Y, Fu Y (2019) Large scale incremental learning. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 374–382
- Xiao J, Hays J, Ehinger KA, Oliva A, Torralba A (2010) Sun database: large-scale scene recognition from abbey to zoo. In: *2010 IEEE computer society conference on computer vision and pattern recognition*, pp. 3485–3492. IEEE
- Yang Y, Hospedales TM (2014) A unified perspective on multi-domain and multi-task learning. arXiv preprint [arXiv:1412.7489](https://arxiv.org/abs/1412.7489)
- Zhang C, Zhang L, Ye J (2012) Generalization bounds for domain adaptation. *Advances in Neural Information Processing Systems*. <https://doi.org/10.48550/arXiv.1304.1574>
- Zhang C, Xie Y, Bai H, Yu B, Li W, Gao Y (2021) A survey on federated learning. *Knowl Based Syst* 216:106775
- Zhang G, Zhao H, Yu Y, Poupard P (2021) Quantifying and improving transferability in domain generalization. *Advances in Neural Information Processing Systems* 34:10957–10970
- Zhang M, Zhang A, McDonagh S (2021) On the out-of-distribution generalization of probabilistic image modelling. *Advances in Neural Information Processing Systems* 34:3811–3823
- Zhao X, Qi L, An Y, Geng X (2023) Generalizable label distribution learning. In: *Proceedings of the 31st ACM international conference on multimedia*, pp. 8932–8941
- Zhao Y, Zhong Z, Yang F, Luo Z, Lin Y, Li S, Sebe N (2021) Learning to generalize unseen domains via memory-based multi-source meta-learning for person re-identification. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6277–6286
- Zheng H, Yang Y, Yin J, Li Y, Wang R, Xu M (2020) Deep domain generalization combining a priori diagnosis knowledge toward cross-domain fault diagnosis of rolling bearing. *IEEE Trans Instrum Measure* 70:1–11
- Zheng L, Shen L, Tian L, Wang S, Wang J, Tian Q (2015) Scalable person re-identification: a benchmark. In: *Proceedings of the IEEE international conference on computer vision*, pp. 1116–1124
- Zheng Z, Zheng L, Yang Y (2017) Unlabeled samples generated by gan improve the person re-identification baseline in vitro. In: *Proceedings of the IEEE international conference on computer vision*, pp. 3754–3762
- Zhou K, Liu Z, Qiao Y, Xiang T, Loy CC (2022) Domain generalization: a survey. *IEEE transactions on pattern analysis and machine intelligence*

Zhou K, Yang Y, Hospedales T, Xiang T (2020) Learning to generate novel domains for domain generalization. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, Proceedings, Part XVI 16, pp. 561–578. Springer, 23–28 August 2020

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.