

Date of current version xxxx 00, 0000.

Digital Object Identifier

Barycentric-alignment and Reconstruction Loss Minimization for Domain Generalization

BOYANG LYU¹, THUAN NGUYEN^{1,3}, PRAKASH ISHWAR² (Senior member, IEEE), MATTHIAS SCHEUTZ³ (Senior member, IEEE), and SHUCHIN AERON¹ (Senior member, IEEE)

¹Department of Electrical and Computer Engineering, Tufts University, Medford, MA 02155 USA (e-mail: boyang.lyu@tufts.edu, shuchin@ece.tufts.edu)

²Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA (e-mail: pi@bu.edu)

³Department of Computer Science, Tufts University, Medford, MA 02155 USA (e-mail: thuan.nguyen@tufts.edu, matthias.scheutz@tufts.edu)

Corresponding author: Boyang Lyu (e-mail: boyang.lyu@tufts.edu).

This work was supported in part by the Air Force Office of Scientific Research under award number FA9550-18-1-0465 and NSF CAREER award CCF:1553075, NSF RAISE 1931978, NSF ERC planning 1937057.

ABSTRACT This paper advances the theory and practice of Domain Generalization (DG) in machine learning. We consider the typical DG setting where the hypothesis is composed of a representation mapping followed by a labeling function. Within this setting, the majority of popular DG methods aim to jointly learn the representation and the labeling functions by minimizing a well-known upper bound for the classification risk in the unseen domain. In practice, however, methods based on this theoretical upper bound ignore a term that cannot be directly optimized due to its dual dependence on both the representation mapping and the unknown optimal labeling function in the unseen domain. To bridge this gap between theory and practice, we introduce a new upper bound that is free of terms having such dual dependence, resulting in a fully optimizable risk upper bound for the unseen domain. Our derivation leverages classical and recent *transport inequalities* that link optimal transport metrics with information-theoretic measures. Compared to previous bounds, our bound introduces two new terms: (i) the Wasserstein-2 barycenter term that aligns distributions between domains, and (ii) the reconstruction loss term that assesses the quality of representation in reconstructing the original data. Based on this new upper bound, we propose a novel DG algorithm named Wasserstein Barycenter Auto-Encoder (WBAE) that simultaneously minimizes the classification loss, the barycenter loss, and the reconstruction loss. Numerical results demonstrate that the proposed method outperforms current state-of-the-art DG algorithms on several datasets.

INDEX TERMS Domain generalization, domain alignment, reconstruction loss, Wasserstein barycenter.

I. INTRODUCTION

Modern machine learning applications often encounter the problem that the training (seen) data and the test (unseen) data have different distributions, which can cause a deterioration in model performance. For example, a model trained on data from one hospital may not work well when the test data is from another hospital [1], a drowsiness driving estimator trained on one group of subjects may not perform well for other subjects [2], or a cognitive workload estimator based on fNIRS (functional near-infrared spectroscopy) measurements may not generalize well across sessions and subjects [3]. Methods that aim to mitigate this problem are broadly classified into two categories, namely Domain Adaptation (DA) [4] and Domain Generalization (DG) [5]. Both DA and DG aim to

find a model that can generalize well in scenarios when the training data from the seen domain does not share the same distribution as the test data from the unseen domain. The key difference between DA and DG is that DA allows access to the (unlabeled) unseen domain data during the training process whereas DG does not. This makes DG a more challenging but more practical problem.

To address the problem of DG, motivated by the seminal theoretical works of [4], [6], the hypothesis is typically expressed as the composition of a representation function followed by a labeling function, *e.g.*, see [7]–[10], and the representation and labeling functions are learned by minimizing an upper bound for the classification risk in the unseen domain derived in [4], [6]. The upper bound consists of three terms: (1) the prediction

risk on the mixture of seen domains, (2) the discrepancy or divergence between the data distributions of different domains in the representation space, and (3) a *combined risk* across all domains that implicitly depends on both the representation mapping and the unknown optimal labeling function from the unseen domain. However, most current approaches disregard this dual dependency and treat the third term (*combined risk*) as a constant while developing their algorithms. In fact, the majority of prominent works in DG and DA such as [9], [11], [12] are essentially variations of the following strategy: ignore the *combined risk* term and learn a domain-invariant representation mapping or align the domains in the representation space, together with learning a common labeling function controlling the prediction loss across the seen domains. However, the *combined risk* term is, in fact, a function of the representation mapping and should somehow be accounted for within the optimization process. Additional details of the shortcomings of previous upper bounds are provided in Appendix A.

To address these limitations, we revisit the analysis in [4], [6] and derive a new upper bound that is free of terms with the dual dependence mentioned above. Our new bound consists of four terms: (1) the prediction risk across seen domains in the input space; (2) the discrepancy/divergence between the induced distributions of seen and unseen domains in the representation space, which can be approximated via the Wasserstein-2 barycenter [13] of seen domains; (3) the reconstruction loss term that measures how well the input can be reconstructed from its representation; and (4) a combined risk term that is independent of the representation mapping and labeling function to be learned. Our new bound differs from previous ones in two aspects. Firstly, it introduces two new terms: (a) the Wasserstein-2 barycenter term for domain alignment and (b) the reconstruction loss term for assessing the quality of representation in reconstructing the original data. We note that the Wasserstein-2 barycenter term for controlling the domain discrepancy in our bound is built in the representation space, which is better aligned with the practical implementation than previous Wasserstein-based bounds that are built in the data space. Secondly, the combined risk in our bound is independent of the representation mapping and thus can be ignored during the optimization. Motivated by these theoretical results, we propose an Auto-Encoder-based model that interacts with the Wasserstein barycenter loss to achieve domain alignment.

The contributions of this work can be summarized as follows:

- 1) *Contributions to Theory*: We propose a new upper bound for the risk of the unseen domain using classical and recent transport inequalities that link optimal transport metrics with information-theoretic measures. All terms in our new upper bound are optimizable in practice which overcomes the limitations of previous works and bridges the gap between previous theory and practice.
- 2) *Contributions to Algorithm Development and Practice*: We develop a novel algorithm for domain generaliza-

tion based on our new upper bound. Our algorithm optimizes a new term that controls the domain discrepancy through Wasserstein-2 barycenter. Unlike previous Wasserstein distance-based bounds that form the domain discrepancy term in the data space but optimize it in the representation space, our domain discrepancy term is constructed and optimized in the representation space, making our practical implementation better aligned with the theory.

- 3) *Gains over state-of-the-art methods*: Our algorithm consistently outperforms other theory-guided methods on PACS, VLCS, Office-Home, and TerraIncognita datasets, with a noticeable improvement of 1.7 – 2.8 percentage points on average across all datasets.

II. RELATED WORK

Our work falls within the DG framework wherein domain-invariant features are learned by decomposing the prediction function into a representation mapping followed by a labeling function. A recent example of this framework is [7], where the authors propose a three-part model consisting of a feature extractor, a classifier, and domain discriminators. The feature extractor learns the task-sensitive, but domain-invariant features via minimizing the cross-entropy loss with respect to the task label and maximizing the sum of domain discriminator losses. The domain discriminator loss is based on an estimate of the \mathcal{H} -divergence between all seen domains [6] and has roots in the works [11], [14] on Domain Adaptation. Following a similar idea, the authors of [9] align the representation distributions from different domains by minimizing their Maximum Mean Discrepancy. In [8], the authors adopt a gradient-based episodic training scheme for DG in which the extracted features are driven to simultaneously preserve global class information and local task-related clusters across seen domains by minimizing an alignment loss comprising soft class confusion matrices and a contrastive loss. The authors of [15] propose the Invariant Risk Minimization algorithm to learn features such that the optimal classifiers are matched across domains. In [16], DG is achieved by disentangling style variation across domains from learned features. Among the large body of works on the DG problem, we regard [11], [14], [15], [17], and [18] as recent exemplars of principled algorithms that are guided by theory and compare their performance with our algorithm's.

Our proposed upper bound is based on the Wasserstein barycenters. Related to this context are the works [19], [20], and [10]. In [10], the pairwise Wasserstein-1 distance [13], [21], is used as a measure of domain discrepancy. Using the dual form of the Wasserstein-1 distance, the feature extractor in [10] minimizes a combination of cross-entropy loss, Wasserstein distance loss, and a contrastive loss to achieve DG. The works [19], [20] provide upper bounds for the risk of unseen domain based on the Wasserstein-1 distance. Although they were originally proposed for DA, they can be adapted to the DG set-up. While the bounds from [19], [20] share some similarities with ours, their bounds

are constructed in the input space and therefore do not explicitly motivate the use of representation functions. By contrast, our proposed upper bound measures the discrepancy of domains in the representation space, which naturally justifies the decomposition of the hypothesis in the practical implementation. A detailed analysis and comparison of the bounds in [19], [20] and our proposed bound can be found in Appendix B.

In addition to the domain-invariant feature learning approach, which is the main focus of this paper, there are other noteworthy and emerging directions in domain generalization research. These are data manipulation techniques [22], meta-learning strategies [23], use of pre-trained models [24], and seeking flat minima [25]. For more details, we refer the reader to [26], [27] which are recent survey articles on DG.

III. THEORETICAL ANALYSIS AND PROPOSED METHOD

We consider a domain v as a triple $(\mu^{(v)}, f^{(v)}, g^{(v)})$ consisting of a distribution $\mu^{(v)}$ on the input $\mathbf{x} \in \mathbb{R}^d$, a representation function $f^{(v)} : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$, from the input space to the representation space, and a stochastic labeling function $g^{(v)} : \mathbb{R}^{d'} \rightarrow \mathcal{Y}$ from the representation space to the label space. We denote the unseen domain by $(\mu^{(u)}, f^{(u)}, g^{(u)})$ and S seen domains by $(\mu^{(s)}, f^{(s)}, g^{(s)})$, with $s = 1, \dots, S$.

Let $\mathcal{F} = \{f|f : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}\}$ be the set of *representation functions*, $\mathcal{G} = \{g|g : \mathbb{R}^{d'} \rightarrow \mathcal{Y}\}$ the set of *stochastic labeling functions*, $\mathcal{H} := \mathcal{G} \circ \mathcal{F}$ the set of *hypotheses*, with each hypothesis $h : \mathbb{R}^d \rightarrow \mathcal{Y}$ obtained by composing a $g \in \mathcal{G}$ with an $f \in \mathcal{F}$, i.e., $h = g \circ f$, and $\mathcal{D} = \{\psi|\psi : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d\}$ the set of *reconstruction functions* that map from the representation space back to the input space. In this paper, we limit our theoretical study to binary classification problems, specifically hypothesis functions h such that $h : \mathbb{R}^d \rightarrow \mathcal{Y} = [0, 1]$. Note that a similar set-up is also used in [4] where the hypothesis h occurs non-deterministically and maps a data point to a label between zero and one.

The risk of using a hypothesis h in domain v is then defined by:

$$R^{(v)}(h) := \mathbb{E}_{\mathbf{x} \sim \mu^{(v)}} [\ell(h(\mathbf{x}), h^{(v)}(\mathbf{x}))], \quad (1)$$

where $\mathbb{E}[\cdot]$ denotes the expectation, $h^{(v)} = g^{(v)} \circ f^{(v)}$, and $\ell(\cdot, \cdot)$ is a loss function. We make the following assumptions:

A1: The loss function $\ell(\cdot, \cdot)$ is non-negative, symmetric, bounded by a finite positive number L , satisfies the triangle inequality, and Q -Lipschitz continuous, i.e., for any three scalars a, b, c and positive constant Q ,

$$|\ell(a, b) - \ell(a, c)| \leq Q |b - c|. \quad (2)$$

A2: The optimal hypothesis of the unseen domain $h^{(u)} = g^{(u)} \circ f^{(u)}$ is K -Lipschitz continuous. Specifically, we assume that for any two vectors $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$, and positive constant K ,

$$|h^{(u)}(\mathbf{x}) - h^{(u)}(\mathbf{x}')| \leq K \|\mathbf{x} - \mathbf{x}'\|_2, \quad (3)$$

where $\|\mathbf{x} - \mathbf{x}'\|_2$ denotes the Euclidean distance between \mathbf{x} and \mathbf{x}' .

The first four conditions in Assumption A1 can be easily satisfied by any metric or norm truncated by a finite positive number. Concretely, if $d(a, b)$ is a metric, potentially unbounded like Mean Squared Error (MSE), then $\text{loss}(a, b) := \min(L, d(a, b))$, where L is a positive constant, will satisfy the first four conditions in A1. The Lipschitz condition in A1 and A2 are also widely used in the theory and practice of DG [17], [20], [28].

One may find our assumptions bear some similarities with the assumptions in [19] and [20], but there are some fundamental differences. Specifically, we assume that the loss function is non-negative, symmetric, bounded, Lipschitz, and satisfies the triangle inequality, whereas the loss function in [19] is required to be convex, symmetric, bounded, obey the triangle inequality, and satisfy a specific form. We only assume that the optimal hypothesis function on the unseen domain is Lipschitz, whereas [20] requires all hypotheses to be Lipschitz.

A. BOUND FOR UNSEEN DOMAIN RISK

Our analysis starts by considering a single seen domain. Lemma III.1 below upper bounds the risk $R^{(u)}(h)$ of a hypothesis $h = g \circ f$ in the unseen domain u by four terms: (1) the risk of the seen domain s , (2) the L^1 distance between the distributions of the *data representations* from the seen and unseen domain, (3) the reconstruction loss that quantifies how well the representation can reconstruct its original data input, and (4) an intrinsic risk term that is free of h and is intrinsic to the domains and the loss function. We use the notation $f_{\#}\mu^{(v)}$ to denote the pushforward of distribution $\mu^{(v)}$ under the representation function f , i.e., the distribution of $f(\mathbf{x})$ with $\mathbf{x} \sim \mu^{(v)}$.

Lemma III.1. *Under assumptions A1 and A2, for any hypothesis $h \in \mathcal{H}$ and any reconstruction function $\psi \in \mathcal{D}$, the following bound holds:*

$$\begin{aligned} R^{(u)}(h) &\leq R^{(s)}(h) + L \|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \\ &\quad + QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \\ &\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) + \sigma^{(u,s)} \end{aligned}$$

where $\|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 = \int_{\mathcal{Z}} |f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}| d\mathbf{z}$ denotes the L^1 distance between $(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)})$ in the representation space and:

$$\begin{aligned} \sigma^{(u,s)} &:= \min \{ \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h^{(u)}(\mathbf{x}), h^{(s)}(\mathbf{x}))], \\ &\quad \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\ell(h^{(u)}(\mathbf{x}), h^{(s)}(\mathbf{x}))] \}. \end{aligned}$$

Proof. Please see Appendix C. \square

In typical DG applications, training data from multiple seen domains are available and can be combined in various ways. Therefore, Lemma III.2 below extends Lemma III.1 to a convex combination of distributions of multiple seen domains.

Lemma III.2. *For any convex weights $\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(S)}$ (non-negative and summing to one), any reconstruction*

function $\psi \in \mathcal{D}$, and any hypothesis $h \in \mathcal{H}$, the following bound holds:

$$\begin{aligned} R^{(u)}(h) &\leq \sum_{s=1}^S \lambda^{(s)} R^{(s)}(h) \\ &\quad + L \sum_{s=1}^S \lambda^{(s)} \|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \\ &\quad + QK \left(\sum_{s=1}^S \lambda^{(s)} \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \\ &\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) + \sum_{s=1}^S \lambda^{(s)} \sigma^{(u,s)}. \end{aligned}$$

Proof. Please see Appendix D. \square

The upper bound above relies on the L^1 distances between the pushforwards of seen and unseen distributions. However, accurately estimating L^1 distances from samples is hard [6], [29]. To overcome this practical limitation, we upper bound the L^1 distance by the Wasserstein-2 distance under additional regularity assumptions on the pushforward distributions.

Definition III.3. [30] A probability distribution on \mathbb{R}^d is called (c_1, c_2) -regular, with $c_1, c_2 \geq 0$, if it is absolutely continuous with respect to the Lebesgue measure with a differentiable density $p(\mathbf{x})$ such that

$$\forall \mathbf{x} \in \mathbb{R}^d, \quad \|\nabla \log_2 p(\mathbf{x})\|_2 \leq c_1 \|\mathbf{x}\|_2 + c_2,$$

where ∇ denotes the gradient and $\|\cdot\|_2$ denotes the Euclidean norm.

Lemma III.4. If μ and ν are (c_1, c_2) -regular, then:

$$\begin{aligned} \|\mu - \nu\|_1 &\leq \sqrt{c_1 \left(\sqrt{\mathbb{E}_{\mathbf{u} \sim \mu} [\|\mathbf{u}\|_2^2]} + \sqrt{\mathbb{E}_{\mathbf{v} \sim \nu} [\|\mathbf{v}\|_2^2]} \right) + 2c_2} \\ &\quad \times \sqrt{W_2(\mu, \nu)} \end{aligned}$$

where the Wasserstein- p metric [13], [21] $W_p(\mu, \nu)$ is defined as,

$$W_p(\mu, \nu) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim \pi} [\|\mathbf{u} - \mathbf{v}\|_2^p] \right)^{1/p}$$

where $\Pi(\mu, \nu)$ is the set of joint distributions with marginals μ and ν .

Proof. Please see Appendix E. \square

One may wonder what conditions would guarantee the regularity of the pushforward distributions. Proposition 2 and Proposition 3 in [30] show that any distribution ν for which $\mathbb{E}_{\mathbf{v} \sim \nu} \|\mathbf{v}\|_2$ is finite becomes regular when convolved with any regular distribution, including the Gaussian distribution. Since convolution of distributions corresponds to the addition of independent random vectors having those distributions, it is always possible to make the pushforwards regular by adding a small amount of independent spherical Gaussian noise in the representation space.

Combining Lemma III.2, Lemma III.4, and applying Jensen's inequality, we obtain our main result:

Theorem III.5. If $f_{\#}\mu^{(s)}$, $s = 1, 2, \dots, S$, and $f_{\#}\mu^{(u)}$ are all (c_1, c_2) -regular, then for any convex weights $\lambda^{(1)}, \lambda^{(2)}, \dots, \lambda^{(S)}$, any reconstruction function $\psi \in \mathcal{D}$, and any hypothesis $h \in \mathcal{H}$, the following bound holds:

$$\begin{aligned} R^{(u)}(h) &\leq \sum_{s=1}^S \lambda^{(s)} R^{(s)}(h) \\ &\quad + LC \left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) \right]^{1/4} \\ &\quad + QK \left(\sum_{s=1}^S \lambda^{(s)} \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \\ &\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) + \sum_{s=1}^S \lambda^{(s)} \sigma^{(u,s)} \quad (4) \end{aligned}$$

where:

$$C = \max_s \sqrt{c_1 \left(\sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|f(\mathbf{x})\|_2^2]} + \sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|f(\mathbf{x})\|_2^2]} \right) + 2c_2}.$$

Proof. Please see Appendix F. \square

The upper bound in Theorem III.5 consists of four terms: the first term is the sum of the risk on seen domains, the second term is the Wasserstein distance between the pushforward of seen and unseen domains in the representation space, the third term indicates how well the input can be reconstructed from its corresponding representation, and the fourth term is a combined risk that is independent of both the representation function and the labeling function and only intrinsic to the domain and loss function.

The form of the upper bound derived above shares some similarities with previous bounds in [4], [19], [20]. However, it differs from previous bounds in the following important aspects:

- Firstly, even though Lemma 1 in [19] and Theorem 1 in [20] employ Wasserstein distance to capture domain divergence, the corresponding term is constructed in the *data space*. By contrast, the corresponding term in our bound is constructed in the *representation space*, which not only provides a theoretical justification when decomposing the hypothesis into a representation mapping and a labeling function, but is also consistent with the algorithm implementation in practice. Moreover, the bounds in [19] and [20] are controlled by the Wasserstein-1 distance, while our upper bound is managed by the square root of the Wasserstein-2 distance. There are regimes where one bound is tighter than the other as discussed in Appendix B.
- Secondly, our third term measures how well the input can be reconstructed from its representation. This motivates the use of an encoder-decoder structure in the proposed algorithm in Section IV to minimize the reconstruction loss. This is a novel component absent from [4], [19], [20].
- Finally, the last term in our upper bound is independent of both the representation function f and the labeling function g . This contrasts with the previous results in [4],

where the last term in their upper bound (see Theorem 1 in [4]) depends on the representation function f . We refer the reader to Appendix A for a detailed comparison.

The bound proposed in Theorem III.5 can also be used for the DA problem where one can access the unseen/target domain data and estimate its distribution. However, under the DG setting, the second and third term in (4) are uncontrollable, leading to an intractable upper bound due to the unavailability of the unseen data. This intractability, which cannot be overcome without making additional specific assumptions on the unseen domain, is widely accepted in the literature as a fundamental limitation for all DG methods and analyses.

As a step toward developing a practical algorithm based on our new bound, we decompose both the second term and the third term in (4) into two separate terms where one term completely depends on the unseen distribution and the other fully depends on the seen distributions.

Corollary III.6. *Under the setting and notation of Theorem III.5, for an arbitrary pushforward distribution $f_{\#}\mu$, we have:*

$$\begin{aligned} R^{(u)}(h) \leq & \sum_{s=1}^S \lambda^{(s)} R^{(s)}(h) \\ & + LC \left(\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)}) \right)^{1/4} \\ & + LC \left(W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu) \right)^{1/4} \\ & + QK \left(\sum_{s=1}^S \lambda^{(s)} \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\ & + QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\ & + \sum_{s=1}^S \lambda^{(s)} \sigma^{(u,s)}. \end{aligned} \quad (5)$$

Proof. Please see Appendix G. \square

Motivated by the bound in Corollary III.6, we want to find a suitable representation function f together with a reconstruction function ψ to minimize the second term $\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)})$ and the fourth term $\sum_{s=1}^S \lambda^{(s)} \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2]$ in (5), while ignoring the third term $W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu)$ and the fifth term $\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2]$, as both of them are intractable.

Minimizing the second term $\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)})$ in (5) leads to finding the Wasserstein-2 barycenter of the distributions of seen domains in the representation space. Here, we assume a uniform weight of $\lambda^{(s)} = \frac{1}{S}$ for all s , since there is no additional information for selecting these weights. For this choice, the Wasserstein-2 barycenter of the pushforward distributions of seen domains is defined by:

$$f_{\#}\mu_{\text{barycenter}} := \arg \min_{f_{\#}\mu} \sum_{s=1}^S \frac{1}{S} W_2^2(f_{\#}\mu^{(s)}, f_{\#}\mu). \quad (6)$$

We refer the reader to [31], [32] for the definition and properties (existence, uniqueness) of the Wasserstein barycenter.

On the other hand, minimizing the fourth term $\sum_{s=1}^S \lambda^{(s)} \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2]$ in (5) naturally leads to an auto-encoder mechanism. With a little abuse of notation, we denote the encoder, namely the representation function as f and the decoder, namely the reconstruction function as ψ . The L^2 reconstruction loss should be optimized over all seen domains.

B. PROPOSED METHOD

As the last term in (5) of Corollary III.6 is independent of both the representation function f and the labeling function g , and the third and fifth terms are intractable due to their dependence on unseen domain, we focus on designing f , ψ and g to minimize the first, second, and fourth terms in (5) of Corollary III.6.

Following previous works [4], [6], [7], we optimize the first term by training f together with g using a standard cross-entropy (CE) loss, such that the empirical classification risk on seen domains is minimized. The classification loss function can be written as:

$$\mathcal{L}_c(f, g) = \frac{1}{S} \sum_{s=1}^S \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\text{CE}(h^{(s)}(\mathbf{x}), g(f(\mathbf{x})))] \quad (7)$$

where $\text{CE}(h^{(s)}(\mathbf{x}), g(f(\mathbf{x})))$ denotes the cross-entropy (CE) loss between the output of classifier and the ground-truth label of seen domain s .

As discussed in Corollary III.6, we propose to use the Wasserstein-2 barycenter of representation distributions of seen domains to optimize the second term in (5). Specifically, the barycenter loss is defined by:

$$\mathcal{L}_{\text{bary}}(f) := \sum_{s=1}^S \frac{1}{S} W_2^2(f_{\#}\mu^{(s)}, f_{\#}\mu_{\text{barycenter}}) \quad (8)$$

where $f_{\#}\mu_{\text{barycenter}}$, as defined in (6), denotes the Wasserstein barycenter of pushforward distributions of seen domains.

In contrast to the previous Wasserstein distance-based method [10] where pairwise Wasserstein distance loss is employed, we motivate the use of Wasserstein barycenter loss based on our Corollary III.6 and demonstrate its ability in enforcing domain-invariance in the ablation study of Section V-D. Notably, the barycenter loss (8) only requires computing S Wasserstein distances, whereas using pairwise Wasserstein distance would require $S(S-1)/2$ computations.

Furthermore, to handle the fourth term in (5), we utilize the auto-encoder structure. Specifically, a decoder $\psi: \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$ is adopted, leading to the following reconstruction loss term:

$$\mathcal{L}_r(f, \psi) := \frac{1}{S} \sum_{s=1}^S \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\mathbf{x} - \psi(f(\mathbf{x}))\|_2^2]. \quad (9)$$

From the analysis above, we aim to find a representation function f , a classifier g , and a decoder function ψ to optimize the following objective function:

$$\arg \min_{f, g, \psi} \mathcal{L}_c(f, g) + \alpha \mathcal{L}_{\text{bary}}(f) + \beta \mathcal{L}_r(f, \psi) \quad (10)$$

where weights $\alpha, \beta > 0$ are hyper-parameters. One can observe that the terms in our proposed upper bound are incorporated into our objective function in (10). Specifically, the first term in our objective function aims to determine a good classifier g together with a representation mapping f by minimizing the risk of seen domains, which corresponds to the first term of the upper bound in (5). The second term in (10) acts as a domain alignment tool to minimize the discrepancy between seen domains, aligning with the second term in the proposed bound in (5). Note that although L_{bary} itself requires solving an optimization problem, we leverage fast computation methods, which are also discussed in Section IV, to directly estimate this loss without invoking the Kantorovich-Rubenstein dual characterization of Wasserstein distance [13]. This avoids solving a min-max type problem that is often plagued by unstable numerical dynamics. Finally, the third term in the objective function minimizes the mean squared error between the input and its reconstruction over all seen domains, which directly minimizes the fourth term in (5).

IV. ALGORITHM

Based on the loss function designed above, we propose an algorithm named Wasserstein Barycenter Auto-Encoder (WBAE). The pseudo code of the WBAE algorithm can be found in Algorithm 1 while its block diagram is shown in Fig. 1.

Algorithm 1 Wasserstein Barycenter Auto-Encoder (WBAE)

Input: Data from S seen domains, m samples from each domain, learning rate η , parameters α, β, ϵ . **Output:** Encoder f_{θ_e} , decoder ψ_{θ_d} , classifier g_{θ_c}

```

1: while training is not end do
2:   Randomly choose  $m$  samples from each domain, denoted as  $\mathcal{X}^{(s)} := \{\mathbf{x}_i^{(s)}\}_{i=1}^m \sim \hat{\mu}^{(s)}$  and  $\mathbf{y}^{(s)} := \{y_i^{(s)}\}_{i=1}^m$ 
3:   for  $s = 1 : S$  and  $i = 1 : m$  do
4:      $\mathbf{z}_i^{(s)} \leftarrow f_{\theta_e}(\mathbf{x}_i^{(s)})$  with set  $\mathcal{Z}^{(s)} \sim f_{\#}\hat{\mu}^{(s)}$ 
5:   end for
6:   Calculate the Wasserstein barycenter  $\hat{\mu}_{bary}$  of  $\{f_{\#}\hat{\mu}^{(s)}\}_{s=1}^S$  and its supporting points with  $f_{\theta_e}$  detached from automatic backpropagation
7:    $L_{bary} \leftarrow \frac{1}{S} \sum_{s=1}^S Sinkhorn_{\epsilon}(\hat{\mu}_{bary}, f_{\#}\hat{\mu}^{(s)})$ 
8:    $L_c \leftarrow -\frac{1}{mS} \sum_{s=1}^S \sum_{i=1}^m y_i^s \log p(g_{\theta_c}(f_{\theta_e}(\mathbf{x}_i^{(s)})))$ 
9:    $L_r \leftarrow \frac{1}{mS} \sum_{s=1}^S \sum_{i=1}^m \|\mathbf{x}_i^{(s)} - \psi_{\theta_d}(\mathbf{z}_i^{(s)})\|_2^2$ 
10:   $L \leftarrow L_c + \alpha L_{bary} + \beta L_r$ 
11:   $\theta_c \leftarrow \theta_c - \eta \nabla_{\theta_c} L_c$ 
12:   $\theta_d \leftarrow \theta_d - \eta \nabla_{\theta_d} L_r$ 
13:   $\theta_e \leftarrow \theta_e - \eta \nabla_{\theta_e} L$ 
14: end while

```

As shown in the pseudo code, we use an encoder f and a decoder ψ , which are parameterized by θ_e and θ_d for feature extraction and reconstruction, respectively. Here we denote $\mathcal{X}^{(s)}$ as a set of samples from domain s with empirical distribution $\hat{\mu}^{(s)}$ and with $\mathbf{x}_i^{(s)}$ as one of its element. The

corresponding label set of $\mathcal{X}^{(s)}$ is denoted as $\mathbf{y}^{(s)}$, where $\mathbf{y}^{(s)} := \{y_i^{(s)}\}$ with $y_i^{(s)}$ as the label for sample $\mathbf{x}_i^{(s)}$. The extracted feature $\mathbf{z}_i^{(s)} = f_{\theta_e}(\mathbf{x}_i^{(s)})$ in set $\mathcal{Z}^{(s)}$ follows the empirical distribution of $f_{\#}\hat{\mu}^{(s)}$. The decoder takes the extracted features as input and outputs the reconstructions as $\psi_{\theta_d}(\mathbf{z}_i^{(s)})$ for domain s . The classifier g , parameterized by θ_c is then applied to the extracted features for label prediction.

The proposed algorithm requires calculating Wasserstein-2 barycenter and its supporting points. Here we use an off-the-shelf python package [33] that implements a free-support Wasserstein barycenter algorithm described in [32]. This algorithm is executed in the primal domain and avoids the use of the dual form of Wasserstein distances, which otherwise would turn the problem into an adversarial (min-max) type setting that we want to avoid due to its instability. The barycenter loss is approximated via an average Sinkhorn divergence [34] between the seen domains and the estimated barycenter. Sinkhorn divergence is an unbiased proxy for the Wasserstein distance, which leverages entropic regularization [35] for computational efficiency, thereby allowing for integrating automatic differentiation with GPU computation. We incorporate the implementation from [34] into our algorithm for a fast gradient computation and denote it as $Sinkhorn_{\epsilon}$ in Algorithm 1, where ϵ is the entropic regularization term.

V. EXPERIMENTS AND RESULTS

TABLE 1. Performance of tested methods on PACS dataset in the DomainBed setting, measured by accuracy (%). A, C, P, S are left-out unseen domains.

Algorithm	A	C	P	S	Avg
theory-guided algorithms					
ERM	84.7 ± 0.4	80.8 ± 0.6	97.2 ± 0.3	79.3 ± 1.0	85.5
IRM	84.8 ± 1.3	76.4 ± 1.1	96.7 ± 0.6	76.1 ± 1.0	83.5
DANN	86.4 ± 0.8	77.4 ± 0.8	97.3 ± 0.4	73.5 ± 2.3	83.6
CDANN	84.6 ± 1.8	75.5 ± 0.9	96.8 ± 0.3	73.5 ± 0.6	82.6
MTL	87.5 ± 0.8	77.1 ± 0.5	96.4 ± 0.8	77.3 ± 1.8	84.6
VREx	86.0 ± 1.6	79.1 ± 0.6	96.9 ± 0.5	77.7 ± 1.7	84.9
WBAE (Ours)	86.9 ± 0.3	81.3 ± 0.4	97.2 ± 0.2	80.5 ± 0.4	86.5
best-performing heuristic algorithm					
SagNet	87.4 ± 1.0	80.7 ± 0.6	97.1 ± 0.1	80.0 ± 0.4	86.3

TABLE 2. Performance of tested methods on VLCS dataset in the DomainBed setting, measured by accuracy (%). C, L, S, V are left-out unseen domains.

Algorithm	C	L	S	V	Avg
theory-guided algorithms					
ERM	97.7 ± 0.4	64.3 ± 0.9	73.4 ± 0.5	74.6 ± 1.3	77.5
IRM	98.6 ± 0.1	64.9 ± 0.9	73.4 ± 0.6	77.3 ± 0.9	78.5
DANN	99.0 ± 0.3	65.1 ± 1.4	73.1 ± 0.3	77.2 ± 0.6	78.6
CDANN	97.1 ± 0.3	65.1 ± 1.2	70.7 ± 0.8	77.1 ± 1.5	77.5
MTL	97.8 ± 0.4	64.3 ± 0.3	71.5 ± 0.7	75.3 ± 1.7	77.2
VREx	98.4 ± 0.3	64.4 ± 1.4	74.1 ± 0.4	76.2 ± 1.3	78.3
WBAE (Ours)	98.3 ± 0.2	65.5 ± 1.0	72.8 ± 0.3	78.6 ± 0.4	78.8
best-performing heuristic algorithm					
CORAL	98.3 ± 0.1	66.1 ± 1.2	73.4 ± 0.3	77.5 ± 1.2	78.8

The proposed method was evaluated on four benchmark datasets for DG: PACS [36], VLCS [37], Office-Home [38],

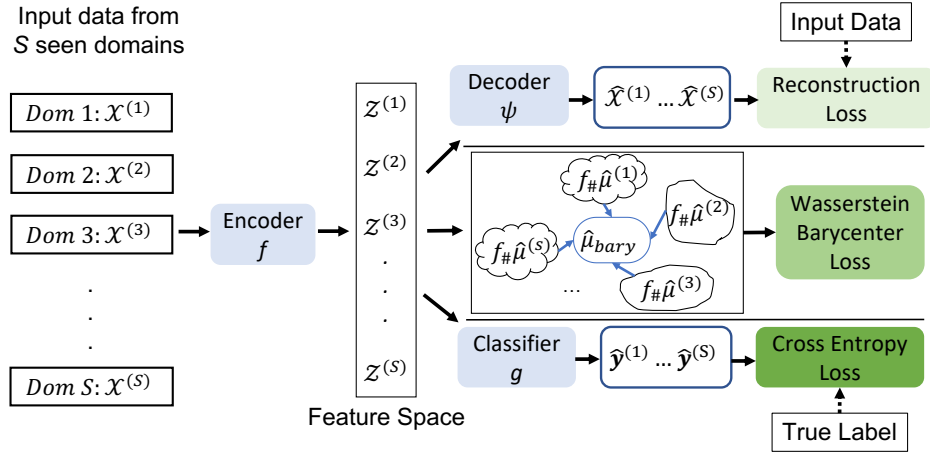


FIGURE 1. An overview of the proposed algorithm. The top, middle, and bottom branches refer to the reconstruction loss term, the Wasserstein barycenter loss term, and the classification risk (from seen domains), respectively.

TABLE 3. Performance of tested methods on Office-Home dataset in the DomainBed setting, measured by accuracy (%). A, C, P, R are left-out unseen domains.

Algorithm	A	C	P	R	Avg
theory-guided algorithms					
ERM	61.3 ± 0.7	52.4 ± 0.3	75.8 ± 0.1	76.6 ± 0.3	66.5
IRM	58.9 ± 2.3	52.2 ± 1.6	72.1 ± 2.9	74.0 ± 2.5	64.3
DANN	59.9 ± 1.3	53.0 ± 0.3	73.6 ± 0.7	76.9 ± 0.5	65.9
CDANN	61.5 ± 1.4	50.4 ± 2.4	74.4 ± 0.9	76.6 ± 0.8	65.8
MTL	61.5 ± 0.7	52.4 ± 0.6	74.9 ± 0.4	76.8 ± 0.4	66.4
VREx	60.7 ± 0.9	53.0 ± 0.9	75.3 ± 0.1	76.6 ± 0.5	66.4
WBAE (Ours)	63.7 ± 0.5	56.4 ± 0.8	76.1 ± 0.3	78.8 ± 0.4	68.8
best-performing heuristic algorithm					
CORAL	65.3 ± 0.4	54.4 ± 0.5	76.5 ± 0.1	78.4 ± 0.5	68.7

TABLE 4. Performance of tested methods on TerraIncognita dataset in the DomainBed setting, measured by accuracy (%). L100, L38, L43, L46 are left-out unseen domains.

Algorithm	L100	L38	L43	L46	Avg
theory-guided algorithms					
ERM	49.8 ± 4.4	42.1 ± 1.4	56.9 ± 1.8	35.7 ± 3.9	46.1
IRM	54.6 ± 1.3	39.8 ± 1.9	56.2 ± 1.8	39.6 ± 0.8	47.6
DANN	51.1 ± 3.5	40.6 ± 0.6	57.4 ± 0.5	37.7 ± 1.8	46.7
CDANN	47.0 ± 1.9	41.3 ± 4.8	54.9 ± 1.7	39.8 ± 2.3	45.8
MTL	49.3 ± 1.2	39.6 ± 6.3	55.6 ± 1.1	37.8 ± 0.8	45.6
VREx	48.2 ± 4.3	41.7 ± 1.3	56.8 ± 0.8	38.7 ± 3.1	46.4
WBAE (Ours)	55.3 ± 0.4	44.3 ± 0.7	56.4 ± 0.5	39.1 ± 0.6	48.8
best-performing heuristic algorithm					
SagNet	53.0 ± 2.9	43.0 ± 2.5	57.9 ± 0.6	40.4 ± 1.3	48.6

and TerraIncognita [39] under two different settings: DomainBed setting [1] and Stochastic Weight Averaging Densely (SWAD) setting [25]. In the DomainBed setting, we implemented our method with the widely used DomainBed package and compared it with various *theory-guided* DG algorithms. Additionally, incorporating the recent advancement in DG-specific optimization, we conducted separate experiments using the SWAD [25] weight sampling strategy with the same experimental setting described in [25]. Furthermore, to

TABLE 5. Performance of theory-guided methods on four datasets in the DomainBed setting, measured by accuracy (%). The average accuracy is reported over different tasks per dataset.

Algorithm	PACS	VLCS	Office-Home	TerraIncognita	Avg
ERM	85.5	77.5	66.5	46.1	68.9
IRM	83.5	78.5	64.3	47.6	68.5
DANN	83.6	78.6	65.9	46.7	68.7
CDANN	82.6	77.5	65.8	45.8	67.9
MTL	84.6	77.2	66.4	45.6	68.5
VREx	84.9	78.3	66.4	46.4	69.0
WBAE (Ours)	86.5	78.8	68.8	48.8	70.7

TABLE 6. Performance of tested methods on four datasets in the SWAD setting, measured by accuracy (%).

Algorithm	PACS	VLCS	Office-Home	TerraIncognita	Avg
ERM + SWAD	88.1 ± 0.1	79.1 ± 0.1	70.6 ± 0.2	50.0 ± 0.3	72.0
CORAL + SWAD	88.3 ± 0.1	78.9 ± 0.1	71.3 ± 0.1	51.0 ± 0.1	72.4
WBAE + SWAD	88.4 ± 0.1	79.5 ± 0.1	71.4 ± 0.2	51.8 ± 0.3	72.8

investigate the impact of different components of the proposed loss function, we conducted an ablation analysis on the PACS, VLCS, and Office-Home datasets and reported the results in Section V-D.

A. DATASETS

The details for the four datasets are described below:

- **PACS dataset [36]:** PACS contains 9,991 images with 7 classes from 4 domains: Art (A), Cartoons (C), Photos (P) and Sketches (S), where each domain represents one type of images.
- **VLCS dataset [37]:** VLCS consists of 10,729 images from 4 different domains: VOC2007 (V), LabelMe (L), Caltech (C), PASCAL (S). A total of 5 classes are shared by all domains.
- **Office-Home dataset [38]:** Office-Home contains 15,500 images from 4 different domains: Artistic (A),

Clipart (C), Product (P), and Real-World (R). Each domain has 65 object categories.

- **TerraIncognita dataset [39]:** TerraIncognita contains four domains {L100, L38, L43, L46} with a total of 24,788 pictures of wild animals belonging to 10 classes.

Example images of the above datasets are shown in Fig. 2, Appendix H.

B. METHODS FOR COMPARISON

In this paper, we compare the empirical performance of our method against the state-of-the-art DG methods reported in [1] under the DomainBed setting. Specifically, the competing methods include:

- Empirical Risk Minimization (**ERM**) [40] which aims to minimize the cumulative training error across all seen domains.
- Domain-Adversarial Neural Networks (**DANN**) [11] which is motivated by the theoretical results from [4]. In particular, to minimize the upper bound of the risk in the unseen domain, **DANN** adopts an adversarial network to enforce that features from different domains are indistinguishable.
- Class-conditional DANN (**C-DANN**) [14] is a variant of **DANN** that aims to match the conditional distributions of feature given the label across domains.
- Invariant Risk Minimization (**IRM**) [15] aims to learn features such that the optimal classifiers applied to these features are matched across domains.
- Risk Extrapolation (**VREx**) [18] is constructed on the assumption from [15] which assumes the existence of an optimal linear classifier across all domains. While **IRM** specifically seeks the invariant classifier, **VREx** aims to identify the form of the distribution shift and propose a variance penalty, leading to the robustness for a wider variety of distributional shifts.
- Marginal Transfer Learning (**MTL**) [5], [17] is proposed based on an upper bound for the generalization error under the setting of an Agnostic Generative Model. Specifically, **MTL** estimates the mean embedding per domain and uses it as a second argument for optimizing the classifier.
- CORrelation ALignment (**CORAL**) [41] is based on the idea of matching the mean and covariance of feature distributions from different domains.
- Style-Agnostic Networks (**SagNet**) [16] minimizes the style induced domain gap by randomizing the style feature for different domains and train the model mainly on the disentangled content feature.

We can categorize the algorithms provided in [1] into two groups: (1) heuristic algorithms, which lack theoretical analysis, and (2) theory-guided algorithms. As the proposed method in this paper falls into the second category, we primarily compare it with the theory-guided methods. Here, **ERM** acts as the baseline theory-guided model and **DANN**, **C-DANN**, **IRM**, **VREx**, **MTL** are five state-of-the-art theory-

guided algorithms. Besides these six methods, for a complete comparison, we also include three heuristic algorithms that achieve the best performances on four evaluated datasets [1]. More specific, **SagNet** [16] is the best-performing algorithm for the PACS and TerraIncognita datasets, and **CORAL** [41] is the best-performing algorithm for both the VLCS and Office-Home datasets. In the SWAD setting, following [25] where **CORAL** was considered as the representative of the previous state-of-the-art methods, we compare our method with both **ERM** and **CORAL**, all of which employed the SWAD strategy. The results for the competing methods above are sourced from [1] and [25].

C. EXPERIMENT SETTINGS

Model Structure: We used the same feature extractor and classifier as used in [1] for all four datasets. Specifically, an ImageNet pre-trained ResNet-50 model with the final (softmax) layer removed is used as the feature extractor. The decoder is a stack of 6 ConvTranspose2d layers for all datasets. The detailed structure of the decoder is described in Table 8, Appendix I. The classifier is a one-linear-layer model with the output dimension the same as the number of classes.

Hyper-parameters: In the DomainBed setting, we performed a random search of 20 trials within the joint distribution of $10^{\text{Uniform}[-3.5, -2]}$ for α and $10^{\text{Uniform}[-3.5, -1.5]}$ for β (see (10)) with other hyper-parameters (*e.g.*, learning rate, batch size, dropout rate, *etc.*) set as the default values recommended in [1]. In the SWAD setting, following [25], we performed a grid search for α in $\{10^{-3.5}, 10^{-3}, 10^{-2.5}, 10^{-2}\}$ and β in $\{10^{-3.5}, 10^{-3}, 10^{-2}, 10^{-1.5}\}$. We chose the value of ϵ for the Sinkhorn loss (line 7, Algorithm 1) as 20, which is the smallest value that can produce stable training processes. A complete description of hyper-parameter tuning in the SWAD setting and a full list of hyper-parameters can be found in Table 9, Appendix I.

Model Selection: We adopted the commonly used training-domain validation strategy in [1], [18] for hyper-parameter tuning and model selection. Specifically, we split the data from each domain into training and validation sets in the proportion 80% and 20%, respectively. During training, we aggregated together the training/validation samples from each seen domain to form the overall training/validation set and selected the model with the highest validation accuracy for testing.

All models were trained on a single NVIDIA Tesla V100 16GB GPU. Experiments on each dataset are repeated three times with different random seeds and the average accuracy together with its standard error are reported.

D. RESULTS AND ABLATION STUDY

As shown in Table 1, 2, 3, and 4, the proposed method (WBAE) performs comparably or better than the state-of-the-art methods. In particular, WBAE achieves the highest accuracy in three out of the four datasets compared to all methods, with a moderate improvement over all theory-guided methods on all four datasets. Additionally, the proposed

method performs equally well as, or slightly better than, the best-performing heuristic DG methods.

In Table 1, it is demonstrated that WBAE outperforms other theory-guided methods by 0.5% and 1.2% points in both Cartoons (C) and Sketches (S) domains, respectively, and by at least 1% point on average on the PACS dataset. Similarly, Table 2 shows that WBAE achieves a performance gain of at least 0.2% points over all theory-guided comparison methods on the VLCS dataset. The effectiveness of the proposed method is further highlighted in Tables 3 and 4, which present results on the larger and more challenging Office-Home and TerraIncognita datasets. Specifically, compared to all theory-guided methods on Office-Home, WBAE boosts the average accuracy by at least 2.3% points on average and at least 2.2%, 3.4%, 0.3%, and 1.9% points on each task. Regarding the TerraIncognita dataset, the proposed algorithm still exhibits superiority by outperforming all theory-guided methods by at least 1.2% points, as shown in Table 4. A summary of evaluation results in the DomainBed setting is reported in Table 5. The proposed method outperforms all theory-guided methods with a noticeable improvement of 1.7-2.8 percentage points on average across all tested datasets.

Table 6 presents the results obtained by applying SWAD, a DG-specific optimizer and weight-averaging technique, in combination with our proposed algorithm WBAE. It can be observed that this combination outperforms all comparison methods on all four evaluated datasets, with an average improvement of 0.4% point over the previous best-performing method **CORAL** as reported in [25].

Based on the results above, it is evident that the proposed algorithm has a more significant impact on the PACS, Office-Home, and TerraIncognita datasets compared to the VLCS dataset. One possible explanation for this, as also suggested in [42], is that three out of four domains in the VLCS dataset contain a greater proportion of scenery contents rather than object information. Unlike the scenery background in TerraIncognita dataset, the scenery contents in the VLCS dataset are usually more intricate and sometimes include multiple objects, making it challenging for the feature extractor to obtain useful object information for the downstream classification.

TABLE 7. Ablation study for the proposed algorithm (WBAE) on PACS, VLCS, and Office-Home datasets.

Dataset	no L_{bary}	no L_r	WBAE
PACS	85.3 \pm 0.3	86.0 \pm 0.1	86.5 \pm 0.2
VLCS	77.9 \pm 0.1	78.4 \pm 0.2	78.8 \pm 0.2
Office-Home	65.7 \pm 0.2	67.7 \pm 0.1	68.8 \pm 0.1

To study the impact of different components of the loss function in (10), we conducted an ablation study for WBAE on all datasets except TerraIncognita due to our limited computational resources. In particular, we consider the following variants of our method: (1) no L_{bary} : using the WBAE loss function without the Wasserstein barycenter term L_{bary} ; (2) no L_r : using the WBAE loss function without the reconstruction

term L_r . We re-ran all the experiments three times using the same model architectures, hyper-parameter tuning, and validation method.

Table 7 demonstrates the performance of the model with different loss terms removed from the original WBAE loss function. It can be observed that removing L_r from the WBAE loss function leads to a decrease in the accuracy of 0.5%, 0.4%, and 1.1% points for PACS, VLCS, and Office-Home datasets, respectively. The performance deterioration is more significant when removing L_{bary} from the WBAE loss function, leading to a drop of 1.2%, 0.9%, and 3.1% points for PACS, VLCS, and Office-Home datasets, respectively. Our ablation study demonstrates the importance of the Wasserstein barycenter loss and also highlights the auxiliary role of the reconstruction loss. Specifically, removing the Wasserstein barycenter loss (L_{bary}) will result in diminished performance, and a similar, though less significant, decrease will occur if the reconstruction loss (L_r) is removed.

VI. CONCLUSION AND FUTURE WORK

In this paper, we revisited the theory and methods for DG and provided a new upper bound for the risk in the unseen domain. The proposed upper bound contains four terms: (1) the empirical risk of the seen domains in the input space; (2) the discrepancy between the induced representation distribution of seen and unseen domains, which can be further represented by the Wasserstein-2 barycenter of representation in the seen domains; (3) the reconstruction loss term that measures how well the data can be reconstructed from its representation; and (4) a combined risk term. The proposed upper bound provides valuable insights in three aspects. Firstly, we observed that the combined risk term in previous bounds relied on the representation function, which made optimization challenging. By contrast, our combined risk term in the proposed upper bound is a constant with respect to both the representation and the labeling function, making optimization straightforward, thus bridging the previous gap between theory and practice. Secondly, compared with other upper bounds using Wasserstein distance to measure the domain discrepancy, the proposed bound constructs the discrepancy term in the representation space rather than in the data space. This approach offers a theoretical justification for the decomposition of the hypothesis when bounding the risk and for practical implementation when designing the algorithm. Lastly, motivated by the proposed upper bound, our practical algorithm WBAE demonstrates competitive performance over state-of-the-art DG algorithms, validating the usefulness of the proposed theoretical bound for addressing the DG problem. In addition, our bound encourages minimizing the reconstruction loss term, which theoretically supports the use of (nearly) invertible representation mappings in recent works [43], [44]. In terms of algorithm and numerical implementation, it should be noted that while our theory-guided method is effective in addressing the DG problem, it may become computationally expensive if one wants to use a larger batch size for a more accurate estimation of the Wasserstein-2 barycenter.

To alleviate this constraint, our future works will focus on leveraging the recently proposed large-scale-barycenter and mapping estimators [45], [46] to enable the calculation of barycenters with a larger number of samples.

APPENDIX.

A. LIMITATIONS OF PREVIOUS UPPER BOUNDS

First, let us recall that a *domain* v is defined as a triple $(\mu^{(v)}, f^{(v)}, g^{(v)})$ consisting of a distribution $\mu^{(v)}$ on the input $\mathbf{x} \in \mathbb{R}^d$, a representation function $f^{(v)} : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}$ that maps an input \mathbf{x} from the input space to its representation \mathbf{z} in the representation space, and a stochastic labeling function $g^{(v)} : \mathbb{R}^{d'} \rightarrow \mathcal{Y}$ that maps the representation space $\mathbb{R}^{d'}$ to a label space \mathcal{Y} .

We denote the unseen domain by $(\mu^{(u)}, f^{(u)}, g^{(u)})$ and the seen domain by $(\mu^{(s)}, f^{(s)}, g^{(s)})$. Let $\mathcal{F} = \{f|f : \mathbb{R}^d \rightarrow \mathbb{R}^{d'}\}$ be a set of *representation functions*, $\mathcal{G} = \{g|g : \mathbb{R}^{d'} \rightarrow \mathcal{Y}\}$ a set of *stochastic labeling functions*. Here, the label space \mathcal{Y} is considered as binary. A *hypothesis* $h : \mathbb{R}^d \rightarrow \mathcal{Y}$ is obtained by composing each $g \in \mathcal{G}$ with each $f \in \mathcal{F}$, i.e., $h = g \circ f$. Next, we rewrite Theorem 1 in [4] using our notations.

Theorem 1 in [4]. *Let f be a fixed representation function from the input space to representation space and \mathcal{G} be a hypothesis space of VC-dimension k . If random labeled samples of size m are generated by applying f to i.i.d. samples from the seen domain, then with probability at least $1 - \delta$, for every $g \in \mathcal{G}$:*

$$R^{(u)}(g) \leq R^{(s)}(g) + d_{\mathcal{H}}(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) + \lambda \quad (11)$$

$$\leq \hat{R}^{(s)}(g) + \sqrt{\frac{4}{m} \left(k \log \frac{2em}{k} + \log \frac{4}{\delta} \right)} + d_{\mathcal{H}}(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) + \lambda \quad (12)$$

where e is the base of the natural logarithm, $d_{\mathcal{H}}$ is \mathcal{H} -divergence (please see Definition 1 in [6], Definition 2.1 in [12] or Definition 1 in [29]), $R^{(u)}(g) = \mathbb{E}_{\mathbf{z} \sim f_{\#}\mu^{(u)}} [g(\mathbf{z}) - g^{(u)}(\mathbf{z})]$ denotes the risk in the unseen domain, $R^{(s)}(g) = \mathbb{E}_{\mathbf{z} \sim f_{\#}\mu^{(s)}} [g(\mathbf{z}) - g^{(s)}(\mathbf{z})]$ and $\hat{R}^{(s)}(g)$ denote the risk in the seen domain and its empirical estimation, respectively, and:

$$\lambda = \inf_{g \in \mathcal{G}} (R^{(s)}(g) + R^{(u)}(g)) \quad (13)$$

is the *combined risk*.

Although the bound in Theorem 1 of [4] was originally constructed for the domain adaptation problem, it has significantly influenced past and recent works in domain generalization as discussed earlier in Section I. To highlight the differences between our work and previous theoretical bounds (the bound in Theorem 1 of [4] and Theorem 4.1 of [12]), we provide a detailed comparison below:

- Firstly, [4] defines the risk induced by labeling function g from the representation space to the label space based on the disagreement between g and the optimal labeling function $g^{(u)}$:

$$R^{(u)}(g) = \mathbb{E}_{\mathbf{z} \sim f_{\#}\mu^{(u)}} [g(\mathbf{z}) - g^{(u)}(\mathbf{z})]. \quad (14)$$

On the other hand, we define the risk induced by using a hypothesis h from the input space to the label space by the disagreement between h and the optimal hypothesis $h^{(u)}$ via a general loss function $\ell(\cdot, \cdot)$:

$$R^{(u)}(h) = \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h(\mathbf{x}), h^{(u)}(\mathbf{x}))]. \quad (15)$$

Since the empirical risk measures the probability of misclassification of a hypothesis that maps from the input space to the label space, minimizing $R^{(u)}(g)$ does not guarantee to minimize the empirical risk. Though there are some cases for the causality to hold, for example, if the representation function f is invertible i.e., there is a one-to-one mapping between \mathbf{x} and \mathbf{z} , and the loss function has the form of $\ell(a, b) = |a - b|$, it is possible to verify that $R^{(u)}(g) = R^{(u)}(h)$. In general, the representation mapping might not be invertible. For example, let us consider a representation function f that maps $f(\mathbf{x}_1) = f(\mathbf{x}_2) = \mathbf{z}$, $\mathbf{x}_1 \neq \mathbf{x}_2$, with corresponding labels as $y_1 = 0$ and $y_2 = 1$. In this case, the risk defined in (14) will introduce a larger error than the risk introduced in (15) since $g(\mathbf{z})$ cannot be mapped to both “0” and “1”. That said, the risk defined in (15) is more precise to describe the empirical risk. In addition, the risk defined in (14) is only a special case of (15) when the representation mapping f is invertible and the loss function satisfies $\ell(a, b) = |a - b|$.

- Secondly, using the setting in [4], for a given hypothesis space, the ideal joint hypothesis g^* is defined as the hypothesis which globally minimizes the combined error from seen and unseen domains [4], [6]:

$$g^* = \arg \min_{g \in \mathcal{G}} (R^{(s)}(g) + R^{(u)}(g)).$$

In other words, this hypothesis should work well in both domains. The error induced by using this ideal joint hypothesis is called *combined risk*:

$$\lambda = \inf_{g \in \mathcal{G}} (R^{(s)}(g) + R^{(u)}(g)) = (R^{(s)}(g^*) + R^{(u)}(g^*)).$$

Note that the labeling function g is a mapping from the representation space to the label space, therefore, the ideal labeling function g^* depends implicitly on the representation function f , hence, λ depends on f . Simply ignoring this fact and treating λ as a constant may loosen the upper bound. By contrast, our goal is to construct an upper bound with the *combined risk* term $\sigma^{(u,s)}$ independent of both the representation function and the labeling function, which can be seen from Lemma III.1 and Theorem III.5.

- Finally, it is worth comparing our upper bound with the bound in Theorem 4.1 of [12] which also has the *combined risk* term free of the choice of the hypothesis class. However, note that the result in Theorem 4.1 of [12] does not consider any representation function f , i.e., their labeling function directly maps from the input space to the label space, while our hypothesis is composed of a representation function from the input space to the

representation space followed by a labeling function from the representation space to the label space. Since it is possible to pick a representation function f that maps any input to itself, i.e., $f(\mathbf{x}) = \mathbf{x}$ which leads to $h = g \circ f = g$, the bound in [12] can be viewed as a special case of our proposed upper bound in Lemma III.1.

B. COMPARISON WITH UPPER BOUNDS IN [19] AND [20]

The form of the proposed upper bound derived in Theorem III.5 shares some similarities with Lemma 1 in [19] and Theorem 1 in [20], for example, all of them introduce Wasserstein distance between domain distributions. However, they differ in the following key aspects.

- 1) The term containing Wasserstein distance in our upper bound is constructed in the *representation* space, not in the data (ambient) space, which provides a theoretical justification when decomposing the hypothesis into a representation mapping and a labeling function. This is also consistent with the algorithmic implementation in practice.
- 2) The bounds in Lemma 1 of [19] and Theorem 1 of [20] are controlled by the Wasserstein-1 distance while our upper bound is managed by the square-root of the Wasserstein-2 distance. There are regimes where one bound is tighter than the other. It is well-known that $W_1(\mu, \nu) \leq W_2(\mu, \nu)$, if $W_2(\mu, \nu) \leq 1$, then $W_1(\mu, \nu) \leq \sqrt{W_2(\mu, \nu)}$. However, based on Jensen's inequality, it is possible to show that $\sqrt{W_2(\mu, \nu)} \leq [Diam(f(\mathbf{X}))W_1(\mu, \nu)]^{1/4}$ where $Diam(f(\mathbf{X}))$ denotes the largest distance between two points in the representation space $\mathbb{R}^{d'}$ generated by input \mathbf{X} via mapping f . To guarantee $\sqrt{W_2(\mu, \nu)} \leq W_1(\mu, \nu)$, a sufficient condition is $[Diam(f(\mathbf{X}))W_1(\mu, \nu)]^{1/4} \leq W_1(\mu, \nu)$ which is equivalent to $Diam(f(\mathbf{X})) \leq W_1(\mu, \nu)^3$. In fact, for a given $Diam(f(\mathbf{X}))$, the larger the value of $W_1(\mu, \nu)$, the higher the chance that this sufficient condition will hold.

C. PROOF OF LEMMA III.1

Note that in this paper, we assume that any hypothesis function $h(\cdot)$ outputs a value in $[0, 1]$, i.e., $h : \mathbb{R}^d \rightarrow [0, 1]$, and $\ell(\cdot, \cdot)$ is a bounded distance metric. In addition, we assume that $h^{(u)}(\cdot)$ is K -Lipschitz continuous and $\ell(\cdot)$ is Q -Lipschitz continuous. Particularly, we assume that for any two vectors $\mathbf{x}, \mathbf{x}' \in \mathbb{R}^d$ and any three scalars a, b , and c , the following inequalities hold:

$$|h^{(u)}(\mathbf{x}) - h^{(u)}(\mathbf{x}')| \leq K\|\mathbf{x} - \mathbf{x}'\|_2, \quad (16)$$

$$|\ell(a, b) - \ell(a, c)| \leq Q|b - c|, \quad (17)$$

where $\|\mathbf{x} - \mathbf{x}'\|_2$ and $|b - c|$ denote the Euclidean distances between \mathbf{x} and \mathbf{x}' , and b and c , respectively.

Lemma 2.1. If $h^{(u)}(\cdot)$ is K -Lipschitz continuous and $\ell(\cdot)$ is Q -Lipschitz continuous. Then, for any hypothesis $h \in \mathcal{H}$ and

any function (decoder) $\psi : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$, the following bound holds:

$$\begin{aligned} R^{(u)}(h) &\leq R^{(s)}(h) + L\|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \\ &\quad + QK\left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}}[\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2]\right. \\ &\quad \left.+ \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}}[\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2]\right) + \sigma^{(u,s)} \end{aligned}$$

where $\|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 = \int_{\mathbf{z}} |f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}| d\mathbf{z}$ denotes the L^1 distance between $(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)})$ and:

$$\sigma^{(u,s)} := \min \left\{ \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}}[\ell(h^{(u)}(\mathbf{x}), h^{(s)}(\mathbf{x}))], \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}}[\ell(h^{(u)}(\mathbf{x}), h^{(s)}(\mathbf{x}))] \right\}.$$

Proof. First, we want to note that our approach is motivated by the proof of Theorem 1 in [6]. Next, to better demonstrate the relationship between the hypothesis, input distribution, true representation and labeling functions, we use inner product notation $\langle \cdot, \cdot \rangle$ to denote expectations. Specifically,

$$R^{(v)}(h) := \mathbb{E}_{\mathbf{x} \sim \mu^{(v)}}[\ell(h(\mathbf{x}), h^{(v)}(\mathbf{x}))] = \langle \ell(h, h^{(v)}), \mu^{(v)} \rangle. \quad (18)$$

From the definition of risk,

$$\begin{aligned} R^{(u)}(h) &= \langle \ell(h, h^{(u)}), \mu^{(u)} \rangle \\ &= \langle \ell(h, h^{(s)}), \mu^{(s)} \rangle - \langle \ell(h, h^{(s)}), \mu^{(s)} \rangle + \langle \ell(h, h^{(u)}), \mu^{(u)} \rangle \\ &= R^{(s)}(h) + (\langle \ell(h, h^{(u)}), \mu^{(u)} \rangle - \langle \ell(h, h^{(s)}), \mu^{(s)} \rangle) \\ &\quad + (\langle \ell(h, h^{(s)}), \mu^{(u)} \rangle - \langle \ell(h, h^{(s)}), \mu^{(s)} \rangle) \\ &\leq R^{(s)}(h) + \langle \ell(h^{(u)}, h^{(s)}), \mu^{(u)} \rangle \\ &\quad + \langle \ell(h, h^{(s)}), \mu^{(u)} - \mu^{(s)} \rangle \end{aligned} \quad (19)$$

where the inequality of (19) follows from the triangle inequality $\ell(h, h^{(u)}) \leq \ell(h, h^{(s)}) + \ell(h^{(s)}, h^{(u)})$ and $\ell(h^{(s)}, h^{(u)}) = \ell(h^{(u)}, h^{(s)})$.

In an analogous fashion, it is possible to show that:

$$\begin{aligned} R^{(u)}(h) &\leq R^{(s)}(h) + \langle \ell(h^{(u)}, h^{(s)}), \mu^{(s)} \rangle \\ &\quad + \langle \ell(h, h^{(u)}), \mu^{(u)} - \mu^{(s)} \rangle. \end{aligned} \quad (20)$$

Next, we will bound the third term in the right-hand-side of (20). Specifically,

$$\begin{aligned}
& \langle \ell(h, h^{(u)}), \mu^{(u)} - \mu^{(s)} \rangle \\
&= \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h(\mathbf{x}), h^{(u)}(\mathbf{x}))] - \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\ell(h(\mathbf{x}), h^{(u)}(\mathbf{x}))] \\
&\leq \max \left\{ \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) + K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2], \right. \\
&\quad \left. \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) - K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right\} \\
&- \min \left\{ \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) + K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2], \right. \\
&\quad \left. \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) - K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right\} \\
&\leq \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))))] \right. \\
&\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\
&- \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))))] \right. \\
&\quad \left. - \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\
&= \left(\mathbb{E}_{\mathbf{z} \sim f_{\#} \mu^{(u)}} [\ell(g(\mathbf{z}), h^{(u)}(\psi(\mathbf{z})))] \right. \\
&\quad \left. - \mathbb{E}_{\mathbf{z} \sim f_{\#} \mu^{(s)}} [\ell(g(\mathbf{z}), h^{(u)}(\psi(\mathbf{z})))] \right) \\
&+ \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \\
&\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\
&= \langle \ell(g(\mathbf{z}), h^{(u)}(\psi(\mathbf{z}))), f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)} \rangle \\
&+ QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) \\
&\leq L(1, |f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)}|) \\
&+ QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right). \quad (24)
\end{aligned}$$

Note that in this paper, any hypothesis function $h(\cdot)$ is assumed to output a scalar value in $[0, 1]$, *i.e.*, $h : \mathbb{R}^d \rightarrow [0, 1]$, and $\ell(\cdot, \cdot)$ is a distance metric. With all these assumptions, we get (21) due to $\min\{\ell(a, c), \ell(a, d)\} \leq \ell(a, b) \leq \max\{\ell(a, c), \ell(a, d)\}$, $\forall b \in [c, d]$, $a, b, c, d \in \mathbb{R}$ and the fact for Lipschitz function $h^{(u)}(\cdot)$ that:

$$h^{(u)}(\psi(f(\mathbf{x}))) - K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2 \leq h^{(u)}(\mathbf{x}), \quad (25)$$

$$h^{(u)}(\mathbf{x}) \leq h^{(u)}(\psi(f(\mathbf{x}))) + K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2. \quad (26)$$

Next, (22) is due to the Lipschitzness of $\ell(\cdot, \cdot)$:

$$\begin{aligned}
& \max \left\{ \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) + K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2), \right. \\
&\quad \left. \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) - K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2) \right\} \\
&\leq \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x})))) + QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2, \quad (27)
\end{aligned}$$

$$\begin{aligned}
& \min \left\{ \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) + K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2), \right. \\
&\quad \left. \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x}))) - K \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2) \right\} \\
&\geq \ell(h(\mathbf{x}), h^{(u)}(\psi(f(\mathbf{x})))) - QK \|\psi(f(\mathbf{x})) - \mathbf{x}\|_2. \quad (28)
\end{aligned}$$

Finally, we get (23) due to $h = g \circ f$, $f(\mathbf{x}) = \mathbf{z}$, and (24) due to $\ell(\cdot, \cdot)$ is bounded by L .

The proof of Lemma follows by combining (19), (20), (24), and note that:

$$\sigma^{(u,s)} = \min \left\{ \langle \ell(h^{(u)}, h^{(s)}), \mu^{(u)} \rangle, \langle \ell(h^{(u)}, h^{(s)}), \mu^{(s)} \rangle \right\},$$

and

$$\langle 1, |f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)}| \rangle = \|f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)}\|_1. \quad \square$$

D. PROOF OF LEMMA II.2

Apply Lemma III.1 S times for S seen domains, then for any hypothesis $h \in \mathcal{H}$ and function (decoder) $\psi : \mathbb{R}^{d'} \rightarrow \mathbb{R}^d$, the following bound holds:

$$\begin{aligned}
R^{(u)}(h) &\leq R^{(s)}(h) + L \|f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)}\|_1 \\
&+ QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \\
&\quad \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) + \sigma^{(u,s)}, \forall s = 1, \dots, S. \quad (29)
\end{aligned}$$

Next, multiplying (29) with its corresponding convex weight $\lambda^{(s)}$, for $s = 1, 2, \dots, S$, and summing them up, we have:

$$\begin{aligned}
\sum_{s=1}^S \lambda^{(s)} R^{(u)}(h) &\leq \sum_{s=1}^S \lambda^{(s)} \left[R^{(s)}(h) + L \|f_{\#} \mu^{(u)} - f_{\#} \mu^{(s)}\|_1 \right. \\
&\quad \left. + QK \left(\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right. \right. \\
&\quad \left. \left. + \mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|\psi(f(\mathbf{x})) - \mathbf{x}\|_2] \right) + \sigma^{(u,s)} \right]. \quad (30)
\end{aligned}$$

Note that $\sum_{s=1}^S \lambda^{(s)} = 1$, thus, the left-hand side of (30) is $R^{(u)}(h)$, and by re-arranging the terms on the right-hand side of (30), the proof follows.

E. PROOF OF LEMMA III.4

From Pinsker's inequality [47], the L^1 distance can be bounded by Kullback–Leibler (KL) divergence as follows:

$$\|\mu - \nu\|_1^2 \leq 2d_{KL}(\mu, \nu) \quad (31)$$

where $\|\mu - \nu\|_1$ and $d_{KL}(\mu, \nu)$ denote L^1 distance and Kullback–Leibler divergence between two distributions μ and ν , respectively. Since $\|\mu - \nu\|_1 = \|\nu - \mu\|_1$, applying Pinsker's inequality to (μ, ν) and (ν, μ) ,

$$2\|\mu - \nu\|_1^2 = \|\mu - \nu\|_1^2 + \|\nu - \mu\|_1^2 \leq 2d_{KL}(\mu, \nu) + 2d_{KL}(\nu, \mu) \quad (32)$$

which is equivalent to,

$$\|\mu - \nu\|_1 \leq \sqrt{d_{KL}(\mu, \nu) + d_{KL}(\nu, \mu)}. \quad (33)$$

Next, if μ and ν are (c_1, c_2) -regular distributions, their Kullback–Leibler divergences can be bounded by their Wasserstein-2 distance as follows (please see equation (10), Proposition 1 in [30]),

$$\begin{aligned}
& d_{KL}(\mu, \nu) + d_{KL}(\nu, \mu) \\
&\leq 2W_2(\mu, \nu) \left(\frac{c_1}{2} \sqrt{\mathbb{E}_{\mathbf{u} \sim \mu} [\|\mathbf{u}\|_2^2]} + \frac{c_1}{2} \sqrt{\mathbb{E}_{\mathbf{v} \sim \nu} [\|\mathbf{v}\|_2^2]} + c_2 \right). \quad (34)
\end{aligned}$$

Combining (33) and (34), we have:

$$\begin{aligned} & \|\mu - \nu\|_1 \\ & \leq [W_2(\mu, \nu)]^{1/2} \sqrt{c_1 \left(\sqrt{\mathbb{E}_{\mathbf{u} \sim \mu} [\|\mathbf{u}\|_2^2]} + \sqrt{\mathbb{E}_{\mathbf{v} \sim \nu} [\|\mathbf{v}\|_2^2]} \right) + 2c_2}. \end{aligned} \quad (35)$$

F. PROOF OF THEOREM III.5

Under the assumption that $f_{\#}\mu^{(s)}$ and $f_{\#}\mu^{(u)}$ are (c_1, c_2) -regular, $\forall s = 1, 2, \dots, S$, we can derive the following inequality from Lemma III.4,

$$\begin{aligned} & \|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \\ & \leq \sqrt{c_1 \left(\sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|f(\mathbf{x})\|_2^2]} + \sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|f(\mathbf{x})\|_2^2]} \right) + 2c_2} \\ & \times [W_2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)})]^{1/2}. \end{aligned} \quad (36)$$

Let:

$$C := \max_s \sqrt{c_1 \left(\sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(s)}} [\|f(\mathbf{x})\|_2^2]} + \sqrt{\mathbb{E}_{\mathbf{x} \sim \mu^{(u)}} [\|f(\mathbf{x})\|_2^2]} \right) + 2c_2}. \quad (37)$$

Multiplying (36) by $\lambda^{(s)}$ and summing over all s , we get:

$$\sum_{s=1}^S \lambda^{(s)} \|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \leq C \sum_{s=1}^S \lambda^{(s)} [W_2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)})]^{1/2}. \quad (38)$$

By Jensen's inequality,

$$\sum_{s=1}^S \lambda^{(s)} [W_2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)})]^{1/2} \leq \left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) \right]^{1/4}. \quad (39)$$

From (38) and (39),

$$\sum_{s=1}^S \lambda^{(s)} \|f_{\#}\mu^{(u)} - f_{\#}\mu^{(s)}\|_1 \leq C \left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) \right]^{1/4}. \quad (40)$$

Finally, combining the upper bound in Lemma III.2 and (40), the proof follows.

G. PROOF OF COROLLARY III.6

We begin with the second term in the upper bound of Theorem III.5. Indeed, for any arbitrary pushforward distribution $f_{\#}\mu$, we have:

$$\left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu^{(s)}) \right]^{1/4} \quad (41)$$

$$\leq \left[\sum_{s=1}^S \lambda^{(s)} \left(W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu) + W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)}) \right) \right]^{1/4} \quad (42)$$

$$= \left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu) + \sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)}) \right]^{1/4} \quad (43)$$

$$= \left[W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu) + \sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)}) \right]^{1/4} \quad (44)$$

$$\leq \left[\sum_{s=1}^S \lambda^{(s)} W_2^2(f_{\#}\mu, f_{\#}\mu^{(s)}) \right]^{1/4} + \left[W_2^2(f_{\#}\mu^{(u)}, f_{\#}\mu) \right]^{1/4} \quad (45)$$

with (42) due to the triangle inequality, (44) due to $\sum_{s=1}^S \lambda^{(s)} = 1$, (45) due to the fact that for any $a, b \geq 0$ and $0 < p \leq 1$, $(a + b)^p \leq a^p + b^p$.

Combining (4) in Theorem III.5 and (45), the proof of Corollary III.6 follows.

H. EXAMPLE IMAGES OF FOUR TESTED DATASETS

Example images of each dataset are shown in Fig. 2.

I. ARCHITECTURE AND HYPER-PARAMETERS

- The model structure of the decoder used in all four datasets can be found in Table 8.
- A list of hyper-parameters used in our proposed method is shown in Table 9.

TABLE 8. Model structure of the decoder.

Layer
ConvTranspose2d (in=2048, out=512, kernel_size=4, stride= 1, padding=0)
BatchNorm2d + ReLU
ConvTranspose2d (in=512, out=256, kernel_size=4, stride=2, padding=1)
BatchNorm2d + ReLU
ConvTranspose2d (in=256, out=128, kernel_size=4, stride=2, padding=1)
BatchNorm2d + ReLU
ConvTranspose2d (in=128, out=64, kernel_size=4, stride=2, padding=1)
BatchNorm2d + ReLU
ConvTranspose2d (in=64, out=32, kernel_size=4, stride=2, padding=1)
BatchNorm2d + ReLU
ConvTranspose2d (in=32, out=3, kernel_size=4, stride=2, padding=1)
Tanh + Interpolate (size=(224, 224))

TABLE 9. Hyper-parameters of the proposed method.

Parameters	DomainBed Setting	SWAD Setting
Optimizer	Adam [48]	Adam [48]
Learning rate	5×10^{-5}	$\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}\}$
Batch size	32	32
ResNet dropout	0	$\{0.0, 0.1, 0.5\}$
Weight decay	0	$\{10^{-4}, 10^{-6}\}$
Training steps	2000	5000
ϵ	20	20
α	$10^{\text{Uniform}(-3.5, -2)}$	$\{10^{-3.5}, 10^{-3}, 10^{-2.5}, 10^{-2}\}$
β	$10^{\text{Uniform}(-3.5, -1.5)}$	$\{10^{-3.5}, 10^{-3}, 10^{-2}, 10^{-1.5}\}$

Following [25], in the SWAD setting, we first fixed all algorithm-agnostic hyper-parameters (HPs) and only tuned the algorithm-specific HPs. Specifically, we first fixed the learning rate as 5×10^{-5} , Resnet dropout rate and weight decay both as 0, and grid searched α, β in $\{10^{-3.5}, 10^{-3}, 10^{-2.5}, 10^{-2}\}$ and $\{10^{-3.5}, 10^{-3}, 10^{-2}, 10^{-1.5}\}$ with the batch size as 32. Then we searched learning rate, Resnet dropout rate, and weight decay in $\{10^{-5}, 3 \times 10^{-5}, 5 \times 10^{-5}\}$, $\{0.0, 0.1, 0.5\}$ and $\{10^{-4}, 10^{-6}\}$, with the selected α, β , as performed in [25]. We used the same values for SWAD-specific hyper-parameters as those used in [25], without any further tuning.

J. CODE AVAILABILITY

The code used to generate the results and tables is available in the GitHub repository: https://github.com/boyanglyu/DG_via_WB.

REFERENCES

- [1] I. Gulrajani and D. Lopez-Paz, "In search of lost domain generalization," in *Int. Conf. Learn. Representations*, 2021.
- [2] Y. Cui, Y. Xu, and D. Wu, "EEG-based driver drowsiness estimation using feature weighted episodic training," *IEEE Trans. Neural Syst. Rehabil. Eng.*, vol. 27, no. 11, pp. 2263–2273, 2019.

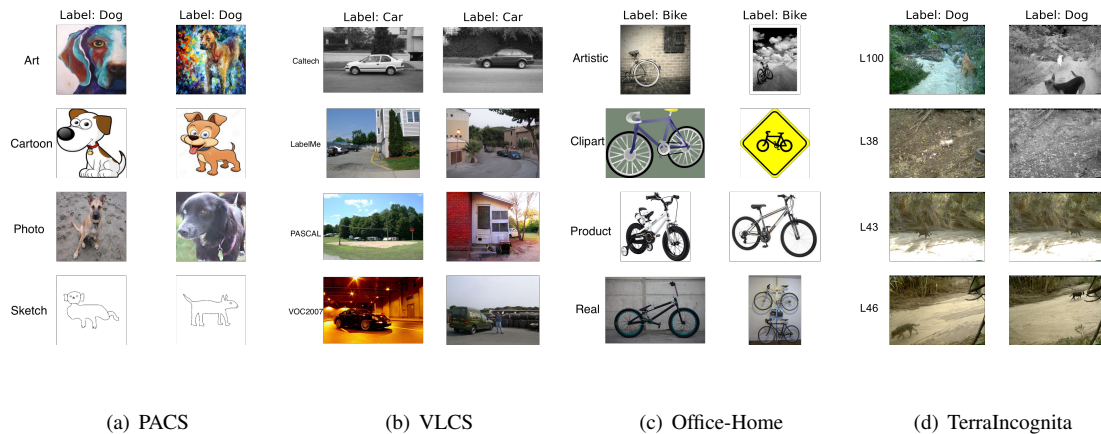


FIGURE 2. Example images of four tested datasets.


- [3] B. Lyu, T. Pham, G. Blaney, Z. Haga, A. Sassaroli, S. Fantini, and S. Aeron, "Domain adaptation for robust workload level alignment between sessions and subjects using fNIRS," *J. Biomed. Opt.*, vol. 26, no. 2, pp. 1–21, 2021.
- [4] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira, "Analysis of representations for domain adaptation," in *Adv. Neural Inf. Process. Syst.*, B. Schölkopf, J. Platt, and T. Hoffman, Eds., vol. 19, 2007, pp. 137–144.
- [5] G. Blanchard, G. Lee, and C. Scott, "Generalizing from several related classification tasks to a new unlabeled sample," in *Proc. 24th Int. Conf. Neural Inf. Process. Syst.*, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, Eds., vol. 24, 2011, pp. 2178–2186.
- [6] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. W. Vaughan, "A theory of learning from different domains," *Mach. Learn.*, vol. 79, no. 1, pp. 151–175, 2010.
- [7] I. Albuquerque, J. Monteiro, M. Darvishi, T. H. Falk, and I. Mitliagkas, "Generalizing to unseen domains via distribution matching," *arXiv preprint arXiv:1911.00804*, 2019.
- [8] Q. Dou, D. Coelho de Castro, K. Kamnitsas, and B. Glocker, "Domain generalization via model-agnostic learning of semantic features," in *Proc. 32nd Int. Conf. Neural Inf. Process. Syst.*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, Eds., vol. 32, 2019.
- [9] H. Li, S. J. Pan, S. Wang, and A. C. Kot, "Domain generalization with adversarial feature learning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2018, pp. 5400–5409.
- [10] F. Zhou, Z. Jiang, C. Shui, B. Wang, and B. Chaib-draa, "Domain generalization via optimal transport with metric similarity learning," *Neurocomputing*, vol. 456, pp. 469–480, 2021.
- [11] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 2096–2030, 2016.
- [12] H. Zhao, R. T. D. Combes, K. Zhang, and G. Gordon, "On learning invariant representations for domain adaptation," in *Proc. 36th Int. Conf. Mach. Learn.*, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun. 2019, pp. 7523–7532.
- [13] F. Santambrogio, "Optimal transport for applied mathematicians," *Birkhäuser, NY*, vol. 55, no. 58–63, p. 94, 2015.
- [14] Y. Li, X. Tian, M. Gong, Y. Liu, T. Liu, K. Zhang, and D. Tao, "Deep domain generalization via conditional invariant adversarial networks," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 624–639.
- [15] M. Arjovsky, L. Bottou, I. Gulrajani, and D. Lopez-Paz, "Invariant risk minimization," *arXiv preprint arXiv:1907.02893*, 2019.
- [16] H. Nam, H. Lee, J. Park, W. Yoon, and D. Yoo, "Reducing domain gap by reducing style bias," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2021, pp. 8690–8699.
- [17] G. Blanchard, A. A. Deshmukh, Ü. Dogan, G. Lee, and C. Scott, "Domain generalization by marginal transfer learning," *J. Mach. Learn. Res.*, vol. 22, no. 1, pp. 46–100, 2021.
- [18] D. Krueger, E. Caballero, J.-H. Jacobsen, A. Zhang, J. Binas, D. Zhang, R. L. Priol, and A. Courville, "Out-of-distribution generalization via risk extrapolation (rex)," in *Proc. 38th Int. Conf. Mach. Learn.*, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 5815–5826.
- [19] I. Redko, A. Habrard, and M. Sebban, "Theoretical analysis of domain adaptation with optimal transport," in *ECML PKDD*, M. Ceci, J. Hollmén, L. Todorovski, C. Vens, and S. Džeroski, Eds. Cham: Springer International Publishing, 2017, pp. 737–753.
- [20] J. Shen, Y. Qu, W. Zhang, and Y. Yu, "Wasserstein distance guided representation learning for domain adaptation," in *Proc. 32nd AAAI Conf. Artif. Intell.*, vol. 32, no. 1, 2018.
- [21] G. Peyré and M. Cuturi, "Computational optimal transport: With applications to data science," *Foundations and Trends® in Machine Learning*, vol. 11, no. 5–6, pp. 355–607, 2019.
- [22] F. C. Borlino, A. D'Innocente, and T. Tommasi, "Rethinking domain generalization baselines," in *25th Int. Conf. Pattern Recognit.* IEEE, 2021, pp. 9227–9233.
- [23] Y. Du, J. Xu, H. Xiong, Q. Qiu, X. Zhen, C. G. M. Snoek, and L. Shao, "Learning to learn with variational information bottleneck for domain generalization," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 200–216.
- [24] Z. Li, K. Ren, X. Jiang, B. Li, H. Zhang, and D. Li, "Domain generalization using pretrained models without fine-tuning," *arXiv preprint arXiv:2203.04600*, 2022.
- [25] J. Cha, S. Chun, K. Lee, H.-C. Cho, S. Park, Y. Lee, and S. Park, "Swad: Domain generalization by seeking flat minima," in *Adv. Neural Inf. Process. Syst.*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34, 2021, pp. 22 405–22 418.
- [26] J. Wang, C. Lan, C. Liu, Y. Ouyang, T. Qin, W. Lu, Y. Chen, W. Zeng, and P. Yu, "Generalizing to unseen domains: A survey on domain generalization," *IEEE Trans. Knowl. Data Eng.*, pp. 1–1, early access, May 26, 2022.
- [27] K. Zhou, Z. Liu, Y. Qiao, T. Xiang, and C. C. Loy, "Domain generalization: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 45, no. 4, pp. 4396–4415, 2023.
- [28] Y. Wu, E. Winston, D. Kaushik, and Z. Lipton, "Domain adaptation with asymmetrically-relaxed distribution alignment," in *Proc. 36th Int. Conf. Mach. Learn.*, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. PMLR, 09–15 Jun 2019, pp. 6872–6881.
- [29] D. Kifer, S. Ben-David, and J. Gehrke, "Detecting change in data streams," in *VLDB*, vol. 4. Toronto, Canada, 2004, pp. 180–191.
- [30] Y. Polyanskiy and Y. Wu, "Wasserstein continuity of entropy and outer bounds for interference channels," *IEEE Trans. Inf. Theory*, vol. 62, no. 7, pp. 3992–4002, 2016.
- [31] M. Agueh and G. Carlier, "Barycenters in the Wasserstein space," *SIAM J. Math. Anal.*, vol. 43, no. 2, pp. 904–924, 2011.
- [32] M. Cuturi and A. Doucet, "Fast computation of wasserstein barycenters," in *Proc. 31st Int. Conf. Mach. Learn.*, E. P. Xing and T. Jebara, Eds., vol. 32, no. 2. Beijing, China: PMLR, 22–24 Jun 2014, pp. 685–693.
- [33] R. Flamary, N. Courty, A. Gramfort, M. Z. Alaya, A. Boisbunon, S. Chamon, L. Chapel, A. Corenflos, K. Fatras, N. Fournier, L. Gautheron,

- N. T. Gayraud, H. Janati, A. Rakotomamonjy, I. Redko, A. Rolet, A. Schutz, V. Seguy, D. J. Sutherland, R. Tavenard, A. Tong, and T. Vayer, "Pot: Python optimal transport," *J. Mach. Learn. Res.*, vol. 22, no. 78, pp. 1–8, 2021. [Online]. Available: <http://jmlr.org/papers/v22/20-451.html>
- [34] J. Feydy, T. Séjourné, F.-X. Vialard, S.-i. Amari, A. Trounev, and G. Peyré, "Interpolating between optimal transport and MMD using sinkhorn divergences," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, K. Chaudhuri and M. Sugiyama, Eds., vol. 89. PMLR, 16–18 Apr 2019, pp. 2681–2690.
- [35] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport," in *Proc. 26th Int. Conf. Neural Inf. Process. Syst.* Red Hook, NY, USA: Curran Associates Inc., 2013, p. 2292–2300.
- [36] D. Li, Y. Yang, Y.-Z. Song, and T. M. Hospedales, "Deeper, broader and artier domain generalization," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5542–5550.
- [37] C. Fang, Y. Xu, and D. N. Rockmore, "Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 1657–1664.
- [38] H. Venkateswara, J. Eusebio, S. Chakraborty, and S. Panchanathan, "Deep hashing network for unsupervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 5018–5027.
- [39] S. Beery, G. Van Horn, and P. Perona, "Recognition in terra incognita," in *Proc. Eur. Conf. Comput. Vis.*, 2018, pp. 456–473.
- [40] V. N. Vapnik, "An overview of statistical learning theory," *IEEE Trans. Neural Netw.*, vol. 10, no. 5, pp. 988–999, 1999.
- [41] B. Sun and K. Saenko, "Deep coral: Correlation alignment for deep domain adaptation," in *Computer Vision – ECCV 2016 Workshops*, G. Hua and H. Jégou, Eds. Cham: Springer International Publishing, 2016, pp. 443–450.
- [42] S. Zhao, M. Gong, T. Liu, H. Fu, and D. Tao, "Domain generalization via entropy regularization," in *Proc. 34th Int. Conf. Neural Inf. Process. Syst.*, 2020, pp. 16 096–16 107.
- [43] F. D. Johansson, D. Sontag, and R. Ranganath, "Support and invertibility in domain-invariant representations," in *Proc. 22nd Int. Conf. Artif. Intell. Statist.*, K. Chaudhuri and M. Sugiyama, Eds., vol. 89. PMLR, 16–18 Apr 2019, pp. 527–536.
- [44] T. Nguyen, B. Lyu, P. Ishwar, M. Scheutz, and S. Aeron, "Trade-off between reconstruction loss and feature alignment for domain generalization," in *2022 21st IEEE Int. Conf. Mach. Learn. Appl.*, 2022, pp. 794–801.
- [45] J. Fan, A. Taghvaei, and Y. Chen, "Scalable computations of Wasserstein barycenter via input convex neural networks," in *Proc. 38th Int. Conf. Mach. Learn.*, M. Meila and T. Zhang, Eds., vol. 139. PMLR, 18–24 Jul 2021, pp. 1571–1581.
- [46] A. Korotin, V. Egiastian, L. Li, and E. Burnaev, "Wasserstein iterative networks for barycenter estimation," in *Adv. Neural Inf. Process. Syst.*, A. H. Oh, A. Agarwal, D. Belgrave, and K. Cho, Eds., 2022.
- [47] I. Csizsár and J. Körner, *Information theory: coding theorems for discrete memoryless systems*. Cambridge University Press, 2011.
- [48] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *3rd Int. Conf. Learn. Representations*, Y. Bengio and Y. LeCun, Eds., 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>



workload identification.

BOYANG LYU is a Ph.D. student at Tufts University, Medford, MA, USA, under Prof. Shuchin Aeron. She received her B.S. degree in electrical engineering from Harbin Institute of Technology, Weihai, China, in 2016 and her M.S. degree in electrical engineering from Tufts University, Medford, MA, USA, in 2019. Her research involves domain generalization, domain adaptation, and their applications to Brain-Computer Interface. She has applied some of these techniques to mental



THUAN NGUYEN received the B.S. degree in electrical engineering (honors program) from the Post and Telecommunication Institute of Technology, Vietnam and the Ph.D. degree from Oregon State University, USA. He is currently a postdoctoral researcher at Tufts University. His research interests include information theory, signal processing, and machine learning.



PRAKASH ISHWAR received the B.Tech. degree in electrical engineering from the Indian Institute of Technology Bombay, India, in 1996 and the M.S. and Ph.D. degrees in electrical and computer engineering from the University of Illinois Urbana-Champaign, USA in 1998 and 2002, respectively. After working two years as a Post-Doctoral Researcher in the Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, he joined the faculty of Boston University, where he is currently a Professor with the Department of Electrical and Computer Engineering and the Division of Systems Engineering and an Affiliate Professor with the Department of Computer Science, College of Arts and Sciences. His current research centers on data science to advance statistical and computational tools for learning and inference problems using both model-based and data-driven methods.

He was a recipient of a 2005 NSF CAREER Award, a co-recipient of the AVSS10 Best Paper Award, a co-winner of the ICPR10 Aerial View Activity Classification Challenge, and a co-author of the Student Best Paper Award in the NeurIPS'14 Workshop on Analysis of Rank Data. He has published over 140 peer-reviewed articles and received numerous grants. He has served as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING for two terms and served as an Elected Member for the IEEE Signal Processing Theory and Methods Technical Committee and the IEEE Image, Video, and Multidimensional Signal Processing Technical Committee



MATTHIAS SCHEUTZ received a Ph.D. degree in philosophy from the University of Vienna, Austria, and a joint Ph.D. in cognitive science and computer science from Indiana University, USA. He currently a full professor of computer and cognitive science in the Department of Computer Science at Tufts University, Senior Gordon Faculty Fellow in the School of Engineering, and Director of the Human-Robot Interaction Laboratory and the Human-Robot Masters and PhD programs. He has over 400 peer-reviewed publications in artificial intelligence, artificial life, agent-based computing, natural language understanding, cognitive modeling, robotics, human-robot interaction and foundations of cognitive science. His current research focuses on complex ethical cognitive robots with natural language interaction, problem-solving, and instruction-based learning capabilities.



SHUCHIN AERON received the B.Tech. degree in electrical engineering from the Indian Institute of Technology Kanpur, Kanpur, India, in 2002, and the M.S. and Ph.D. degrees in electrical engineering from Boston University, Boston, MA, USA, in 2004 and 2009, respectively. He is currently an Associate Professor with the Department of ECE, Tufts University, Medford, MA, USA. From 2009 to 2011, he was a Postdoctoral Fellow with Schlumberger Doll Research and has a number of

patents with applications to acoustic logging while drilling. His area of expertise and research interests include information theory, mathematical statistics, tensor data analytics, compressed sensing, statistical machine learning, and optimal transport.

In 2016, he was the recipient of the NSF CAREER Award. He is currently an Associate Editor for the IEEE TRANSACTIONS ON GEOSCIENCES AND REMOTE SENSING and is on the technical committee of Machine Learning for Signal Processing, IEEE Signal Processing Society. He is also the Co-Director of the Data Science programs at the School of Engineering, Tufts University

...