Supplementary for Beyond Boundaries: A Novel Data-Augmentation Discourse for Open Domain Generalization

Shirsha Bose *shirshabosecs@gmail.com*
Technical University of Munich

Ankit Jha *ankitjha16@gmail.com*
Indian Institute of Technology Bombay

Hitesh Kandala *khitesh2000@gmail.com*
Indian Institute of Technology Bombay

Biplab Banerjee *getbiplab@gmail.com*
Indian Institute of Technology Bombay

Reviewed on OpenReview: `https://openreview.net/forum?id=jpZmhiIys1`

In the supplementary materials, we report the following topics:

1. Dataset descriptions and class-splits used for ODG in Section 1.
2. Detailed experimental results for closed-set domain generalization (Table 2-5).
3. Model implementation details and description of the attention modules $\mathcal{A}_d$ and $\mathcal{A}_o$ in Section 3.
4. A t-SNE Van der Maaten & Hinton (2008) plot showing the distribution of the style-space (both original and augmented) is depicted in Fig. 1.
5. Some discussions on model complexity of ODG-NET and DAML Shu et al. (2021).
6. Sensitivity to β in $\mathbf{L_{GAN}}$ in Section 6.
7. System specification for the reported experiments (Section 7).

1 Dataset descriptions

(1) **Office-Home** Venkateswara et al. (2017) - This dataset comprises of 15,500 images distributed among 65 classes spanning four domains: Art, Clipart, Product, and Real.

(2) **PACS** Li et al. (2017) - PACS consists of 9991 images categorized into seven classes, spread across four domains: Artpaint, Cartoon, Sketch, and Photo.

(3) **VLCS** Fang et al. (2013) - VLCS is a combination of images from four image classification datasets: PASCAL VOC 2007 Everingham et al. (2010), Caltech Fei-Fei et al. (2004), LabelMe Russell et al. (2008), and Sun Xiao et al. (2010). The dataset comprises images from five classes: Bird, Car, Chair, Dog, and Person.

(4) **Digits-DG** Zhou et al. (2020b) - This dataset is a combination of several handwritten digit recognition datasets: MNIST LeCun et al. (1998), MNIST-M Ganin & Lempitsky (2015), SVHN Netzer et al. (2011), and SYN Ganin & Lempitsky (2015).

(5) **DomainNet** Peng et al. (2019) - DomainNet is a dataset that consists of images from six different domains, including photos (real), painting, clipart, quickdraw, infograph, and sketch. The dataset has a total of 600K images classified into 345 classes, with each domain having between 48K to 172K images.

(6) **Multi-dataset** Shu et al. (2021) - This dataset is a combination of several public datasets, including Office-31 Saenko et al. (2010), STL-10 Coates et al. (2011), and Visda2017 Peng et al. (2017), with four domains from DomainNet Peng et al. (2019) as the target. Additionally, 20 classes are considered as open classes, which are not present in the joint label set of source domains. The class splits for all the five datasets considered for ODG are mentioned in Table 1, with the classes numbered in alphabetical order.

Table 1: Class splits for PACS, VLS, Office-Home, Digits-DG, and Multi-dataset for the ODG experiments.

Domain	PACS	VLCS	Digits-DG	Office-Home	Multi-dataset
Source 1	3, 0, 1	0, 1	0, 1, 2	$0-2, 3-8, 9-14, 21-31$	$0-30$
Source 2	4, 0, 2	1, 2	2, 3, 4	$0-2, 3-8, 15-20, 32-42$	$1, 31-41$
Source 3	5, 1, 2	2, 3	4, 5, 6	$0, 3-4, 9-10, 15-16$	$31, 33, 34, 41, 42-47$
Target	0–6	0–5	0–9	$21-23, 32-34, 43-45, 54-64$	$31-36, 39-43, 45-46, 48-67$

2 Detailed results for closed-set DG

We show the detailed results and comparisons of the closed-set DG tasks for PACS, Office-Home, VLCS, and Digits-DG, and DomainNet in Tables 2-5. It can be found that ODG-NET produces the new state-of-the-art for all the datasets by outperforming the literature by at least $2-3\%$ for all the domain combinations and the average leave-one-out accuracies.

Table 2: Results of PACS dataset under closed-set DG. (In $\%$)

Method	A	C	S	P	Avg
CCSA Motiian et al. (2017)	80.50	76.90	**93.60**	66.80	79.40
SFA-A Li et al. (2021)	81.20	77.80	73.70	93.90	81.70
MetaReg Balaji et al. (2018)	83.70	77.20	70.30	95.50	81.70
MixStyle Zhou et al. (2021)	84.10	78.80	75.90	96.10	83.70
DSON Seo et al. (2020)	84.67	77.65	82.23	95.87	85.11
JiGen Carlucci et al. (2019)	79.42	75.25	71.35	96.03	80.51
Epi-FCR Li et al. (2019)	82.10	77.00	73.00	93.90	81.50
SagNet Wu et al. (2019)	83.58	77.66	76.30	95.47	83.25
RSC Huang et al. (2020)	83.43	80.31	80.85	95.99	85.15
DDAIG Zhou et al. (2020a)	84.20	78.10	74.70	95.30	83.10
L2A-OT Zhou et al. (2020b)	83.30	78.20	73.60	96.20	82.80
FACT Xu et al. (2021)	85.37	78.38	79.15	95.15	84.51
STEAM Chen et al. (2021)	85.50	80.60	82.90	97.50	86.60
Style Neo. Kang et al. (2022)	84.41	79.25	83.27	94.93	85.47
DAML Shu et al. (2021)	83.00	78.10	74.10	95.60	82.70
ODG-NET	**91.75**	**85.27**	86.79	**98.83**	**90.66**

3 Implementation details

We adopt a double-convolution layered network (double-conv) for our embedding networks $\mathcal{F}_{im}, \mathcal{F}_v, \mathcal{F}_y$, and $\mathcal{F}_\eta$, where each network consists of the sequence: conv-batchnorm-ReLU repeated twice. The architecture of $\mathcal{F}_G$ is inspired by Ronneberger et al. (2015) and consists of four down-sampling layers, each comprising convolution, max-pooling, batch norm, and ReLU nonlinearity, followed by four up-sampling layers that form a symmetric structure. For each domain-specific network $\mathcal{F}_l^s$, we employ two double-conv blocks followed by global average pooling (GAP) and classification layers to implement the $\mathcal{F}_{ls}^c$s.

The classifiers $\mathcal{F}_d$ and $\mathcal{F}_{disc}$ use the double-conv structure repeated four times, with the only difference being the activation function, where we use a Leaky-ReLU layer with a negative multiplication factor of 0.2 in $\mathcal{F}_{disc}$. Following the convolution blocks, $\mathcal{F}_{disc}$ applies the GAP layer and a linear layer with sigmoid activation, while $\mathcal{F}_d$ employs a combination of GAP and a fully connected layer with softmax activation. The architecture of $\mathcal{F}_o$, which serves as a feature extractor cum classifier, depends on the specific benchmark, as in Zhou et al. (2020b); Kang et al. (2022).

Table 3: Results of Office-Home dataset under closed-set DG. (In %)

Method	A	C	P	R	Avg
CCSA Motiian et al. (2017)	59.90	49.90	74.10	75.70	64.90
D-SAM D'Innocente & Caputo (2018)	58.03	44.37	69.22	71.45	60.77
DSON Seo et al. (2020)	59.37	45.70	71.84	74.68	62.90
Jeon et al. Jeon et al. (2021)	60.24	53.54	74.36	76.66	66.20
MMD-AAE Li et al. (2018b)	56.50	47.30	72.10	74.80	62.70
MixStyle Zhou et al. (2021)	58.70	53.40	74.20	75.90	65.50
JiGen Carlucci et al. (2019)	53.00	47.50	71.50	72.80	61.20
Cross-Grad Shankar et al. (2018)	58.40	49.40	73.90	75.80	64.40
SagNet Wu et al. (2019)	60.20	45.38	70.42	73.38	62.34
RSC Huang et al. (2020)	58.42	47.90	71.63	74.54	63.12
DDAIG Zhou et al. (2020a)	59.20	52.30	74.60	76.00	65.50
L2A-OT Zhou et al. (2020b)	60.60	50.10	74.80	77.00	65.60
FACT Xu et al. (2021)	60.34	54.85	74.48	76.55	66.56
STEAM Chen et al. (2021)	62.10	52.30	75.40	77.50	66.80
Liu et al.Liu et al. (2021)	62.24	54.38	76.12	78.64	67.85
Style Neo. Kang et al. (2022)	59.55	55.01	73.57	75.52	65.89
DAML Shu et al. (2021)	62.47	54.39	76.33	77.65	67.71
ODG-NET	**67.89**	**58.91**	**81.22**	**83.68**	**72.92**

Table 4: Results of VLCS dataset under closed-set DG. (In %)

Method	C	L	V	S	Avg
D-MTAE Ghifary et al. (2015)	89.10	60.10	63.90	61.30	68.60
CIDDG Li et al. (2018c)	88.80	63.10	64.40	62.10	69.60
CCSA Motiian et al. (2017)	92.30	62.10	67.10	59.10	70.20
DBADG Li et al. (2017)	93.60	63.50	70.00	61.30	72.10
MMD-AAE Li et al. (2018b)	94.40	62.60	67.70	64.40	72.30
MLDG Li et al. (2018a)	94.40	61.30	67.70	65.90	72.30
Epi-FCRLi et al. (2019)	94.10	64.30	67.10	65.90	72.90
SFA-A Li et al. (2021)	97.20	62.00	70.40	66.20	74.00
JiGen Carlucci et al. (2019)	96.93	60.90	70.62	64.30	73.19
RSC Huang et al. (2020)	97.61	61.86	73.93	68.32	75.43
MASF Dou et al. (2019)	94.80	64.90	69.10	67.60	74.10
EIS-NET Wang et al. (2020)	97.30	63.50	69.80	68.00	74.65
MetaVIB Du et al. (2020)	97.37	62.66	70.28	67.85	74.54
DGER Zhao et al. (2020)	96.92	58.26	73.24	69.10	74.38
Liu et al. Liu et al. (2021)	97.86	64.33	74.35	69.37	76.48
DAML Shu et al. (2021)	95.51	62.11	67.48	66.72	72.95
ODG-NET	**98.53**	**69.25**	**77.94**	**73.71**	**79.85**

We employ self-attention modules $(\mathcal{A}_d, \mathcal{A}_o)$ to offer spatial-spectral attention using query-key-value-based processing Han et al. (2022). The attention parameters for $\mathcal{A}_d$ and $\mathcal{A}_o$ are disjoint. For a given input feature map $Fe_l(x)$ with shape $C \times W \times H$ (representing the channel, width, and height dimensions), we define the mathematical operations for the spatial and spectral parts of the attention modules.

Spatial attention: We initialize Query (Q), Key (K), and Value (V) with $Fe_l(x)$. We perform 1×1 convolution to transform (V, Q, K) into shapes: $HW \times C$ for V, and $HW \times \frac{C}{4}$ for K and Q, respectively. We further obtain $K \otimes Q^T$ of shape $HW \times HW$. The final self-attention mask with shape $C \times W \times H$ is obtained as,

$$V^T \otimes \text{softmax}(K \otimes Q^T) \tag{1}$$

Spectral attention: Similar to the spatial attention, we initial Query (Q), Key (K), and Value (V) with $Fe_l(x)$. We perform 1×1 convolution to transform (V, Q, K) into shapes: $C \times HW$. We further obtain $Q \otimes K^T$ of shape $C \times C$. The final self-attention mask with shape $C \times W \times H$ is obtained as,

$$V^T \otimes \text{softmax}(Q \otimes K^T) \tag{2}$$

Table 5: Results of Digits-Dg dataset under closed-set DG. (In $\%$)

Method	MNIST	MNIST_M	SVHN	SYN	Avg
CCSAMotiian et al. (2017)	95.20	58.20	65.50	79.10	74.50
MMD-AAELi et al. (2018b)	96.50	58.40	65.00	78.40	74.60
SFA-A Li et al. (2021)	96.50	66.50	70.30	85.00	79.60
JiGen Carlucci et al. (2019)	96.50	61.40	63.70	83.20	76.20
Cross-Grad Shankar et al. (2018)	96.70	61.10	65.30	80.20	75.83
DDAIG Zhou et al. (2020a)	96.60	64.10	68.60	81.00	77.58
L2A-OT Zhou et al. (2020b)	96.70	63.90	68.60	83.20	78.10
Liu et al. Liu et al. (2021)	97.68	66.24	70.97	85.18	80.02
FACT Xu et al. (2021)	97.90	65.60	72.40	90.30	81.55
STEAM Chen et al. (2021)	96.80	67.50	76.00	92.20	83.13
DAML Shu et al. (2021)	96.33	65.38	72.43	85.43	79.89
ODG-NET	**98.53**	**71.45**	**80.34**	**96.71**	**86.75**

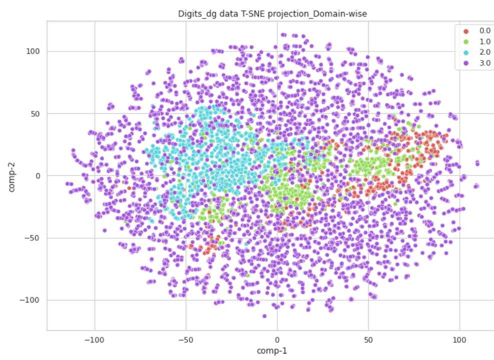

Figure 1: The space of real and generated styles for Digits-DG dataset produced by ODG-NET. The legend '3.0' is for the generated styles, while the other colors represent the source styles per instance.

We use skip connection to add the original $Fe_l(x)$ with both the attended feature maps for better highlighting the important features, which also aids in gradient propagation.

4 Spread of the style space generated by ODG-NET

We show the t-SNE for the original and synthesized styles for the Digits-DG dataset in Fig. 1, which shows that our interpolation technique is able to synthesize a large space of possible styles.

5 Comparison of model complexity between ODG-NET and DAML

It is worth noting that DAML is composed of an ensemble of source domain classification models, each with a different ResNet-18 backbone. In contrast, ODG-NET employs a single feature backbone in $\mathcal{F}_o$, while the other feature extractors for the local classification models $\{\mathcal{F}_l^s\}_{s=1}^{\mathcal{S}}$ are relatively light-weight. Although the parameter sizes of ODG-NET and DAML are comparable, ODG-NET demonstrates better training convergence and faster inference time (0.6-0.75x) in GFLOP than DAML, while achieving significantly improved performance over DAML.

6 Sensitivity analysis of the regularizer weight β for $\mathbf{L}_{\mathbf{GAN}}$

As we observe in Fig. 2, a small beta provide a soft regularizer to the GAN losses by outputting diverse synthesized images for the closed and pseudo-open spaces. As we increase β, the accuracy decreases as it affects the GAN losses. $\beta = 0.5$ is found to provide the best solution.

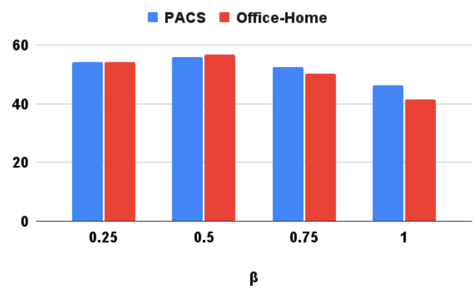

Figure 2: Sensitivity analysis of β in $\mathbf{L_{GAN}}$.

7 System specification

We used a system with 16 CPUs, 128 GB memory and two NVIDIA 2080 Ti graphics cards. The code was written in Pytorch 1.5 with CUDA 10.

References

Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *Advances in neural information processing systems*, 31, 2018.

Fabio M Carlucci, Antonio D'Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2229–2238, 2019.

Yang Chen, Yu Wang, Yingwei Pan, Ting Yao, Xinmei Tian, and Tao Mei. A style and semantic memory mechanism for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9164–9173, 2021.

Adam Coates, Andrew Ng, and Honglak Lee. An analysis of single-layer networks in unsupervised feature learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pp. 215–223. JMLR Workshop and Conference Proceedings, 2011.

Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32, 2019.

Yingjun Du, Jun Xu, Huan Xiong, Qiang Qiu, Xiantong Zhen, Cees GM Snoek, and Ling Shao. Learning to learn with variational information bottleneck for domain generalization. In *European Conference on Computer Vision*, pp. 200–216. Springer, 2020.

Antonio D'Innocente and Barbara Caputo. Domain generalization with domain-specific aggregation modules. In *German Conference on Pattern Recognition*, pp. 187–198. Springer, 2018.

Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2):303–338, 2010.

Chen Fang, Ye Xu, and Daniel N. Rockmore. 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 (ICCV)*, December 2013.

Li Fei-Fei, Rob Fergus, and Pietro Perona. 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, 2004.

Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.

Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pp. 2551–2559, 2015.

Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on vision transformer. *IEEE transactions on pattern analysis and machine intelligence*, 2022.

Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *European Conference on Computer Vision*, pp. 124–140. Springer, 2020.

Seogkyu Jeon, Kibeom Hong, Pilhyeon Lee, Jewook Lee, and Hyeran Byun. Feature stylization and domain-aware contrastive learning for domain generalization. In *Proceedings of the 29th ACM International Conference on Multimedia*, pp. 22–31, 2021.

Juwon Kang, Sohyun Lee, Namyup Kim, and Suha Kwak. Style neophile: Constantly seeking novel styles for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7130–7140, 2022.

Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.

Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 5542–5550, 2017.

Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018a.

Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 1446–1455, 2019.

Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5400–5409, 2018b.

Pan Li, Da Li, Wei Li, Shaogang Gong, Yanwei Fu, and Timothy M Hospedales. A simple feature augmentation for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 8886–8895, 2021.

Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 624–639, 2018c.

Chang Liu, Lichen Wang, Kai Li, and Yun Fu. *Domain Generalization via Feature Variation Decorrelation*, pp. 1683–1691. Association for Computing Machinery, New York, NY, USA, 2021. ISBN 9781450386517. URL https://doi.org/10.1145/3474085.3475311.

Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 5715–5725, 2017.

Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.

Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.

Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Bryan C Russell, Antonio Torralba, Kevin P Murphy, and William T Freeman. Labelme: a database and web-based tool for image annotation. *International journal of computer vision*, 77(1):157–173, 2008.

Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pp. 213–226. Springer, 2010.

Seonguk Seo, Yumin Suh, Dongwan Kim, Geeho Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. In *European Conference on Computer Vision*, pp. 68–83. Springer, 2020.

Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018.

Yang Shu, Zhangjie Cao, Chenyu Wang, Jianmin Wang, and Mingsheng Long. Open domain generalization with domain-augmented meta-learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9624–9633, 2021.

Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9 (11), 2008.

Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5018–5027, 2017.

Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. In *European Conference on Computer Vision*, pp. 159–176. Springer, 2020.

Zhijie Wu, Xiang Wang, Di Lin, Dani Lischinski, Daniel Cohen-Or, and Hui Huang. Sagnet: Structure-aware generative network for 3d-shape modeling. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.

Jianxiong Xiao, James Hays, Krista A Ehinger, Aude Oliva, and Antonio Torralba. 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, 2010.

Qinwei Xu, Ruipeng Zhang, Ya Zhang, Yanfeng Wang, and Qi Tian. A fourier-based framework for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14383–14392, 2021.

Shanshan Zhao, Mingming Gong, Tongliang Liu, Huan Fu, and Dacheng Tao. Domain generalization via entropy regularization. *Advances in Neural Information Processing Systems*, 33:16096–16107, 2020.

Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 13025–13032, 2020a.

Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. In *European conference on computer vision*, pp. 561–578. Springer, 2020b.

Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. *arXiv preprint arXiv:2104.02008*, 2021.