

# UNCERTAINTY-GUIDED CONTRASTIVE LEARNING FOR SINGLE SOURCE DOMAIN GENERALISATION

Anastasios Arsenos<sup>1</sup>, Dimitrios Kollias<sup>2</sup>, Evangelos Petrongonas, Christos Skliros<sup>3</sup>, Stefanos Kollias<sup>1</sup>

<sup>1</sup> School of Electrical & Computer Engineering, National Technical University of Athens, Greece

<sup>2</sup> School of Electronic Engineering & Computer Science, Queen Mary University of London, UK

<sup>3</sup> Hellenic Drones S.A., Greece

## ABSTRACT

In the context of single domain generalisation, the objective is for models that have been exclusively trained on data from a single domain to demonstrate strong performance when confronted with various unfamiliar domains. In this paper, we introduce a novel model referred to as Contrastive Uncertainty Domain Generalisation Network (CUDGNet). The key idea is to augment the source capacity in both input and label spaces through the fictitious domain generator and jointly learn the domain invariant representation of each class through contrastive learning. Extensive experiments on two Single Source Domain Generalisation (SSDG) datasets demonstrate the effectiveness of our approach, which surpasses the state-of-the-art single-DG methods by up to 7.08%. Our method also provides efficient uncertainty estimation at inference time from a single forward pass through the generator subnetwork.

**Index Terms**— Domain Generalisation, Augmentation, Adversarial & Contrastive Learning, Uncertainty estimation

## 1. INTRODUCTION

The concept of leveraging diversity for model training has been extensively explored. Previous research [1, 2, 3] demonstrates that employing a wide array of augmentations during training enhances a model’s resilience against shifts in distribution. When the nature of diversity encountered during testing is identifiable, particular augmentations can be applied.

Apart from input diversity, SSDG approaches need to learn domain invariant representations of the data. Many previous works [2, 1] incorporated contrastive learning to acquire domain-invariant representations. These representations ensure that each class forms distinct clusters, thereby enabling the learning of an improved decision boundary that enhances generalisation capabilities. Nevertheless, prior research [4, 5, 6] disregard the potential hazards associated with utilizing augmented data to address out-of-domain generalisation. This omission gives rise to significant apprehensions related to safety and security, particularly in contexts involv-

ing mission-critical applications. For example, in scenarios involving the deployment of self-driving vehicles in unfamiliar surroundings, a comprehensive understanding of predictive uncertainty becomes pivotal for effective risk evaluation.

Recently, [7] proposed a Bayesian meta-learning framework that leverages the uncertainty of domain augmentations to improve the domain generalisation in a curriculum learning scheme and provides fast uncertainty assessment.

Its main limitations included sensitivity to hyperparameters that made the training unstable and high computational demands that made it challenging to scale to complex networks.


Inspired by [7], we aim to leverage the uncertainty of domain augmentations in both input and label spaces. In order to tackle the aforementioned limitations of previous works, we propose a novel framework that contains a task model  $M$  and a domain augmentation generator  $G$ . These components enhance each other through collaborative learning. The domain augmentation generator produces secure and efficient domains, guided by uncertainty assessment. These generated domains are systematically extended to enhance coverage and comprehensiveness. To achieve cross-domain invariant representation across all generated domains, contrastive learning is introduced in the learning process of the task model  $M$ .

The main contributions can be summarized below:

- We propose a novel framework that leverages adversarial data augmentation and style transfer for domain expansion while ensuring semantic information perseverance through contrastive learning.
- Our framework can estimate the uncertainty in a single forward pass while achieving state-of-the-art accuracy.
- We validate our framework’s performance via comparison and ablation study on two SSDG datasets.

## 2. MATERIALS AND METHOD

Here, we present an outline of CUDGNet, as depicted in Figure 1. Our objective involves training a robust model from a solitary domain  $S$ , with the anticipation that the model’s effectiveness will extend across an unfamiliar domain distribution  $\{T_1, T_2, \dots\} \sim p(T)$ . In order to achieve that, we create



**Fig. 1:** The overall framework of the proposed CUDGNet. The Task Model M and the domain augmentation Generator G are jointly trained while the transformation component TC and style mixing (EFDMix) further enrich the augmentation capacity. The contrastive loss will guide semantically similar samples from different domains to be closer in the embedding space.

a sequence of domain augmentations  $\{S_1^+, S_2^+, \dots\} \sim p(S^+)$  to closely approximate  $p(T)$ , allowing the task model to acquire the ability to generalise across domains that it has not encountered before. In addition, we show how we can assess the uncertainty of new domains as a byproduct of the perturbations used for adversarial domain augmentation.

To generate new domains while retaining class-specific details we introduce two auxiliary components, namely the Transformation Component and the domain augmentation generator  $G$  component. The latter is a novel feature perturbation subnetwork that combines style transfer and variational feature perturbations that follow a learnable multivariate Gaussian distribution  $\mathcal{N}(\mu, \sigma^2)$ , for diverse and content-preserved feature augmentations. Additionally, the domain augmentation learning procedure is enriched by incorporating image-structure information generated through our image transformation component which utilizes both affine transformations and fractals to enrich the input space. To effectively organise domain alignment and classification, we incorporate contrastive learning to acquire representations that are invariant to domain shifts as well as to avoid representation collapse caused by extreme domain shifts of feature perturbations. This process facilitates the progressive formation of domain augmentations and well-defined clusters for each class in the representation space.

**2.1. Transformation Component (TC):** it can transform the initial image  $x$  from the original domain  $S$  into a novel image  $\hat{x}$  within the same domain using the following process:

$$\hat{x} = TC(x) = \begin{cases} (x \oplus f_1) \\ (x \oplus f_1) \otimes (x \oplus x_{aff}) \\ (x \oplus f_1) \otimes (x \oplus x_{aff}) \otimes (x \oplus f_2) \end{cases} \quad (1)$$

where: each of the three branches has an equal probability of occurring;  $x_{aff}$  is the resulting image after an affine transformation (like rotation, translation and contrast) is applied to the initial image  $x$ ;  $f_1, f_2$  denote fractal images [8];  $\oplus$  and  $\otimes$  are element-wise additions and multiplications respectively.

The transformation component may not expand the domain space but is a crucial part of our method. It helps to avoid representational collapse in the initial epochs of domain expansion where  $G$  may induce severe noise and thus the whole framework becomes more robust. Fractals exhibit specific structural characteristics, which are notably non-random and improbable to emerge from processes of maximum entropy or gaussian noise (perturbations); thus this transformation is orthogonal to the domain augmentation generator  $G$ . In Eq. 1 the image is passed through  $k$  transformations (with  $k \in [0, 10]$  being a hyperparameter).

**2.2. Domain augmentation generator:** We illustrate the process of producing the unseen domain  $S^+$  from  $S$  through the generator  $G$ , ensuring that the samples generated adhere to the criteria of safety and effectiveness. Safety denotes that the generated samples preserve semantic information, while effectiveness implies that the generated samples encompass a diverse range of unseen domain-specific details.

**Style manipulation.** In conjunction with the Transformation Component, we enrich the input space using style manipulation. We integrate Exact Histogram Matching (EHM) [9] that can represent style information by using high-order feature statistics. While there are several strategies available for EHM, we choose to utilize the Sort-Matching algorithm [10] because of its efficient execution speed. Sort-Matching is implemented by matching two sorted vectors, whose indexes are illustrated in a one-line notation:

$$\begin{aligned} \mathbf{w} : \tau &= (\tau_1 \quad \tau_2 \quad \tau_3 \quad \dots \quad \tau_n) \\ \mathbf{r} : \kappa &= (\kappa_1 \quad \kappa_2 \quad \kappa_3 \quad \dots \quad \kappa_n) \end{aligned} \quad (2)$$

The Sort-Matching output is:  $out_{\tau_i} = r_{\kappa_i}$  (3)

To create a wide range of feature augmentations that combine various styles, we rely on the Exact feature Distribution mixing -outlined in Eq. 3- by incorporating interpolated sorted vectors. This leads to the development of the Exact Feature Distribution Mixing (EFDMix) method [11].

$$EFDMix(w, r)_{\tau_i} = w_{\tau_i} + (1-d)r_{\kappa_i} - (1-d)\langle w_{\tau_i} \rangle \quad (3)$$

We utilize the instance-specific mixing weight denoted as  $d$ , which is obtained by sampling from a Beta distribution( $c, c$ ), where  $c \in (0, \infty)$  serves as a hyper-parameter. where  $\langle \cdot \rangle$  represents the stop-gradient operation [12].

**Learnable mixup with style transfer.** For the purpose of adversarial domain augmentation, we employ feature perturbations. It is assumed that the perturbations, denoted as  $e$ , follow a multivariate Gaussian distribution  $\mathcal{N}(\mu, \sigma^2)$ . This allows for easy access to the uncertainty associated with the perturbations. The parameters of the Gaussian distribution ( $\mu, \sigma$ ) are learned using variational inference.

The updated latent feature  $h^+$  is obtained by adding the perturbations to the original feature  $h$  with interpolated style through EFDMix ( $r$  is obtained by shuffling  $h$  along the batch dimension). This is denoted as:  $h^+ \leftarrow \text{EFDMix}(h, r) + e$ , where  $e$  is sampled from the Gaussian distribution  $\mathcal{N}(\mu, \sigma^2)$ . By employing this approach, we can create a series of feature augmentations in different training iterations.

Our approach also involves blending  $S$  and  $S^+$  via Mixup [13] to achieve intermediate domain interpolations. In particular, we utilize the uncertainty captured in the perturbations ( $\mu, \sigma^2$ ) to forecast adjustable parameters ( $a, b$ ), which guide the direction and intensity of domain interpolations.

$$\begin{aligned} h^{+'} &= \lambda \cdot \text{EFDMix}(h, r) + (1 - \lambda)h^+ \\ y^+ &= \lambda y + (1 - \lambda)\tilde{y} \end{aligned} \quad (4)$$

where:  $\lambda \in \text{Beta}(a, b)$ ;  $\tilde{y}$  is a smoothed version of  $y$  by a chance of lottery  $\tau$ . Beta distribution & lottery are computed by a fully connected layer following the Generator. These parameters integrate the uncertainty of generated domains.

**Adversarial domain augmentation.** In the latent space, we propose an iterative training procedure where two phases are alternated: a maximization phase where new data points are learned by computing the inner maximization problem and a minimization phase, where the model parameters are updated according to stochastic gradients of the loss evaluated on the adversarial examples generated from the maximization phase. The fundamental concept here involves iteratively acquiring "hard" data points from fictitious target distributions while retaining the essential semantic attributes of the initial data points via adversarial data augmentation [14]:

$$\max_G L(M; S^+) - \beta \|z - z^+\|_2^2 \quad (5)$$

where:  $L$  represents the cross-entropy loss and involves the creation of  $S^+$  through the perturbations of  $h^+$ ; the second term is the safety constraint that limits the greatest divergence between  $S$  and  $S^+$  in the embedding space.  $z$  ( $z^+$  when  $G$  activated) denotes the output from the Projection head ( $P$ ) and  $\beta$  is a hyperparameter that controls the maximum divergence. The projection head part of our model transforms the convolutional features into a lower-dimensional feature space  $Z$  suitable for contrastive learning. This is distinct from  $h$  which denotes the outputs emerging from the conv. layers.

**2.3. Learning objective:** Here, our focus is on introducing the process of acquiring cross-domain invariant representations and producing effective domain augmentations  $S^+$ . Initially, we aim to maximize the conditional likelihood of the fictitious domain  $S^+$ :  $\log p(y^+|x, h^+; M)$ . That requires the true posterior  $p(h^+|x; M, G)$  which is intractable. As a consequence, we use amortized variational inference with a tractable form of approximate posterior  $q(h^+|x; M, G)$ . The derived lower bound is expressed as follows:

$$L_{M,G} = \mathbf{E}_{q(h^+|x; M, G)} [\log \frac{p(y^+|x, h^+; M)}{q(h^+|x; M, G)}] \quad (6)$$

We use Monte-Carlo (MC) sampling to maximize the lower bound  $L_{M,G}$  by:

$$\begin{aligned} L_{M,G} &= \min_{M,G} \frac{1}{K} \sum_{k=1}^K [-\log p(y_k^+|x_k, h_k^+; M)] + \\ &KL [q(h^+|x; M, G) || p(h^+|x; M, G)] \end{aligned} \quad (7)$$

where  $h_k^+ \sim q(h^+|x; M, G)$  and  $K$  is the number of MC samples. For Kullback-Leibler divergence (KL) we let  $q(h^+|x; M, G)$  approximate  $p(h^+|x; M, G)$  through adversarial training on  $G$  in Eq. 5 so that the learned adversarial distribution is more flexible to approximate unseen domains.

When the model is faced with a new domain,  $T$ , during testing, we compute its uncertainty with respect to the learned parameters of the task model  $M$ ;  $\sigma(S)$  through  $\frac{\sigma(T) - \sigma(S)}{\sigma(S)}$  [15]. In Section 3, we provide empirical evidence that our estimation aligns with traditional Bayesian methods [16], while substantially reducing the time required, by an order of magnitude.

$$L_{NCE}(z_i, z_i^+) = \min_{M,G} -\log \frac{\exp(z_i \cdot z_i^+)}{\sum_{j=1, j \neq i}^{2N} \exp(z_i \cdot z_j)} \quad (8)$$

$L_{NCE}$  is the InfoNCE loss [17] used for contrastive learning. In the minibatch  $B$ ,  $z_i$  and  $z_i^+$  originate from the task model and perturbed task model (Generator  $G$  was activated) respectively. They share the same semantic content while originating from the source and augmented domain respectively. The utilization of  $L_{NCE}$  will direct the learning process of  $M$  towards acquiring domain-invariant representations by decreasing the distance between  $z_i$  and  $z_i^+$ .

In summary, the loss function used to train the task model  $M$  can be described as follows:

$$L_{total} = L_{M,G} + w_1 \cdot L_{NCE}(z_i, z_i^+) \quad (9)$$

$w_1$ : parameter controlling the importance of contrastive loss.

**2.4. Datasets and Metrics:** CIFA-10-C [18] is a robustness benchmark encompassing 19 diverse types of corruptions applied to the CIFAR-10 test set. These corruptions are classified into four primary categories: weather, blur, noise, and digital. PACS [19] comprises four domains, each accompanied by a distinct number of images: Photo (1670 images),

**Table 1:** The accuracy of single-source domain generalisation (%) on CIFAR-10-C. Models are trained on CIFAR-10 and evaluated on the CIFAR-10-C.

Method	Weather	Blur	Noise	Digital	Avg
ERM [14]	67.28	56.73	30.02	62.30	54.08
M-ADA [5]	75.54	63.76	54.21	65.10	64.65
U-SDG [7]	76.23	65.87	53.05	68.43	65.89
RandConv [6]	76.87	55.36	75.19	77.51	71.23
L2D [2]	75.98	69.16	73.29	72.02	72.61
MetaCNN [4]	77.44	76.80	78.23	81.26	78.45
<b>Ours</b>	<b>89.13</b>	<b>82.94</b>	<b>85.62</b>	<b>84.43</b>	<b>85.53</b>

**Table 2:** The accuracy of single-source domain generalisation (%) on PACS. Models are trained on photo and evaluated on the rest of the target domains.

Method	A	C	S	Avg
ERM [14]	54.43	42.74	42.02	46.39
JiGen [3]	54.98	42.62	40.62	46.07
M-ADA [5]	58.96	44.09	49.96	51.00
L2D [2]	56.26	51.04	58.42	55.24
ALT [22]	<b>68.50</b>	43.50	53.30	55.10
MetaCNN [4]	54.04	<b>53.58</b>	<b>63.88</b>	57.17
<b>Ours</b>	59.30	50.66	62.00	<b>57.32</b>

Art Painting (Art) (2048 images), Cartoon (2344 images), and Sketch (3929 images). The accuracy for each category and the average accuracy across all categories are presented as the main performance metrics.

**2.5. Implementation details:** We have tailored task-specific models and adopt distinct training strategies for the two datasets. We employ WideResNet [20] (16-4) for CIFAR and AlexNet [21] for PACS, for fair comparison [2, 4]. The Generator  $G$  consists of two convolution layers (for  $\mu$  and  $\sigma$ ) and the projection head  $P$  contains one fully connected layer.

### 3. EXPERIMENTAL RESULTS


**Comparison with the state-of-the-art** Tables 1, 2 exhibit the evaluations of single domain generalisation on CIFAR-10-C and PACS, respectively. The results demonstrate that CUDGNet achieves the highest average accuracy compared to other methods. To be specific, as shown in Table 1, there are notable improvements of 11.69%, 6.14%, 7.39%, 3.17% and 7.08% in weather, blur, noise, digital and average categories of CIFAR corruptions respectively. Table 2 demonstrates that CUDGNet outperforms all previous methods in the domain of Art Painting except for [22] which has rather imbalanced results and achieves superior average performance. Sketch and Cartoon in contrast to Art painting have huge domain discrepancies compared to Photo (source domain), but still, our model achieves comparable results in these categories compared to the state-of-the-art.

**Uncertainty estimation.** We compare the domain uncertainty score introduced in Section 2.3 with a more computationally intensive Bayesian approach [16].

**Table 3:** Ablation study of the key components of CUDGNet

Method	Weather	Blur	Noise	Digital	Avg
Baseline	63.43	55.61	31.92	61.01	53.01
+ G	72.19	61.37	51.58	62.71	61.96
+ TC	79.22	75.44	71.46	75.14	75.31
+ Style transfer	87.81	82.10	82.05	80.76	83.18
+ Contrastive learning	89.13	82.94	85.62	84.43	85.53

whereas the latter relies on repeated sampling (30 times) to compute output variance (5.1 ms/batch). In Figure 2, we present the outcomes of uncertainty estimation for CIFAR-10-C. Evidently our estimation consistently aligns with Bayesian uncertainty estimation. This alignment underscores its high efficiency.



**Fig. 2:** Estimation of Uncertainty CIFAR-10-C. Our domain uncertainty prediction aligns with Bayesian uncertainty, while our approach is significantly faster.

**Ablation Study** As depicted in Table 3 when incorporating the Transformation Component the model exhibits significant improvements over the baseline with adversarially augmented Generator. This underscores that slightly augmenting the diversity of the source domain and changing the image structure through fractals before the adversarial minmax optimization of the Generator is an effective strategy. Furthermore, the addition of style transfer results in a further performance boost of 7.87%. This suggests that this component empowers our model to generalise more effectively to novel domains. At last, with the integration of contrastive loss, we achieve a new state-of-the-art performance, boasting an average performance score of 85.53%.

### 4. CONCLUSION

In this paper, we introduce a novel framework that enhances domain generalisation capabilities of the model and explainability through uncertainty estimation. Our comprehensive experiments and analysis have demonstrated the effectiveness of our method.

## 5. REFERENCES

- [1] Lei Li, Ke Gao, Juan Cao, Ziyao Huang, Yepeng Weng, Xiaoyue Mi, Zhengze Yu, Xiaoya Li, and Boyang Xia, “Progressive domain expansion network for single domain generalization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 224–233.
- [2] Zijian Wang, Yadan Luo, Ruihong Qiu, Zi Huang, and Mahsa Baktashmotlagh, “Learning to diversify for single domain generalization,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 834–843.
- [3] 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*, 2019, pp. 2229–2238.
- [4] Chaoqun Wan, Xu Shen, Yonggang Zhang, Zhiheng Yin, Xinmei Tian, Feng Gao, Jianqiang Huang, and Xian-Sheng Hua, “Meta convolutional neural networks for single domain generalization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4682–4691.
- [5] Fengchun Qiao, Long Zhao, and Xi Peng, “Learning to learn single domain generalization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 12556–12565.
- [6] Zhenlin Xu, Deyi Liu, Junlin Yang, Colin Raffel, and Marc Niethammer, “Robust and generalizable visual representation learning via random convolutions,” *ICLR*, 2021.
- [7] Xi Peng, Fengchun Qiao, and Long Zhao, “Out-of-domain generalization from a single source: An uncertainty quantification approach,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [8] Kataoka Hirokatsu, Matsumoto Asato, Yamagata Eisuke, Yamada Ryosuke, Inoue Nakamasa, Akio Nakamura, and Satoh Yutaka, “Pre-training without natural images,” *International Journal of Computer Vision*, vol. 130, no. 4, pp. 990–1007, 2022.
- [9] Dinu Coltuc, Philippe Bolon, and J-M Chassery, “Exact histogram specification,” *IEEE Transactions on Image processing*, vol. 15, no. 5, pp. 1143–1152, 2006.
- [10] Jannick P Rolland, V Vo, B Bloss, and Craig K Abbey, “Fast algorithms for histogram matching: Application to texture synthesis,” *Journal of Electronic Imaging*, vol. 9, no. 1, pp. 39–45, 2000.
- [11] Yabin Zhang, Minghan Li, Ruihuang Li, Kui Jia, and Lei Zhang, “Exact feature distribution matching for arbitrary style transfer and domain generalization,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 8035–8045.
- [12] Xinlei Chen and Kaiming He, “Exploring simple siamese representation learning,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 15750–15758.
- [13] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017.
- [14] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese, “Generalizing to unseen domains via adversarial data augmentation,” *Advances in neural information processing systems*, vol. 31, 2018.
- [15] Fengchun Qiao and Xi Peng, “Uncertainty-guided model generalization to unseen domains,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 6790–6800.
- [16] Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra, “Weight uncertainty in neural network,” in *International conference on machine learning*. PMLR, 2015, pp. 1613–1622.
- [17] Aaron van den Oord, Yazhe Li, and Oriol Vinyals, “Representation learning with contrastive predictive coding,” *arXiv preprint arXiv:1807.03748*, 2018.
- [18] Dan Hendrycks and Thomas Dietterich, “Benchmarking neural network robustness to common corruptions and perturbations,” *ICLR*, 2019.
- [19] 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*, 2017, pp. 5542–5550.
- [20] Sergey Zagoruyko and Nikos Komodakis, “Wide residual networks,” *arXiv preprint arXiv:1605.07146*, 2016.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, “Imagenet classification with deep convolutional neural networks,” *Advances in neural information processing systems*, vol. 25, 2012.
- [22] Tejas Gokhale, Rushil Anirudh, Jayaraman J Thiagarajan, Bhavya Kailkhura, Chitta Baral, and Yezhou Yang, “Improving diversity with adversarially learned transformations for domain generalization,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 434–443.