

# Unsupervised Cross-Domain Conditional Generation via Normalizing Flows and Adversarial Training

Hari Prasanna Das, Ryan Tran, Japjot Singh, Yu-Wen Lin and Costas J. Spanos

Department of Electrical Engineering and Computer Sciences, UC Berkeley

{hpdas,bobotran,calzoom,yuwen.lin,spanos}@berkeley.edu

## ABSTRACT

How to generate *conditional* synthetic data for a domain without utilizing information about its labels/attributes? Our work presents a solution to the above question. We propose a transfer learning-based framework utilizing normalizing flows, coupled with both maximum-likelihood and adversarial training. We model a source domain (labels available) and a target domain (labels unavailable) with individual normalizing flows, and perform domain alignment to a common latent space using adversarial discriminators. Due to the invertible property of flow models, the mapping has exact cycle consistency. We also learn the joint distribution of the data samples and attributes in the source domain by employing an encoder to map attributes to the latent space via adversarial training. During the synthesis phase, given any combination of attributes, our method can generate synthetic samples conditioned on them in the target domain. Empirical studies confirm the effectiveness of our method on benchmarked datasets. We envision our method to be particularly useful for synthetic data generation in label-scarce systems by generating non-trivial augmentations via attribute and component transformations. These synthetic samples will introduce more entropy into the label-scarce domain than their geometric and photometric transformation counterparts, helpful for robust downstream tasks.

## CCS CONCEPTS

• **Computing methodologies** → **Neural networks.**

## KEYWORDS

Unsupervised Style Transfer, Conditional Synthetic Data Generation, Domain Translation

## ACM Reference Format:

Hari Prasanna Das, Ryan Tran, Japjot Singh, Yu-Wen Lin and Costas J. Spanos. 2022. Unsupervised Cross-Domain Conditional Generation via Normalizing Flows and Adversarial Training. In *DLP '22: Workshop on Deep Learning Practice and Theory for High-Dimensional Sparse and Imbalanced Data, August 15, 2022, Washington DC, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/nmnnnnn.nnnnnnn>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*DLP '22, August 15, 2022, Washington DC, USA*

© 2022 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM... \$15.00

<https://doi.org/10.1145/nmnnnnn.nnnnnnn>

## 1 INTRODUCTION

A large majority of the real-world signals obtained are unlabeled, and require significant human effort or machine intelligence for labeling. This has led to a surge in popularity of unsupervised learning algorithms. A prominent branch of such algorithms, generative modeling, has proven to be efficient in transferring knowledge gained from one (or multiple) domain(s) to other domain(s). Variants of such approaches include cross-domain translation [21, 44], domain adaptation for classification [19, 46] etc. By jointly modeling the data samples and their labels/attributes, variations of conditional synthesis methods have been proposed [28, 31, 32], which during inference phase, can generate synthetic conditional samples. We combine both the above avenues of cross-domain translation along with conditional synthesis and propose a framework capable of generating conditional samples for a domain without utilizing its labels/attributes for training.

Prior works on cross-domain translation involve construction of a mapping between two (or more) unpaired domains. The translation consistency is maintained by introducing some form of inductive bias terms such as cycle consistency [44], semantic consistency [35], entropic regulation [5] etc. Most of the proposed models for domain translation are generative adversarial network (GAN) [16] based and involve many-to-one/one-to-many mappings, making the cycle consistency only approximate. A recent work, Alignflow [17] achieves exact cycle consistency by modeling the domains with normalizing flows via a common latent space. Normalizing flows [12, 24] are a class of generative models which map an unknown and complex data distribution to a latent space with a simple (e.g. standard gaussian) prior distribution via invertible mappings. Another benefit with having flow model mappings is that they offer a rich latent space, which is suitable for a number of downstream tasks, such as semi-supervised learning [32], synthetic data augmentation and adversarial training [4], text analysis and model based control etc. Conditional synthesis has been explored by CGAN [31] by augmenting the conditions with the data and processing it via GAN and by ACGAN [32] by introducing an auxiliary classifier for the conditions. This becomes challenging for flow models which are bijective in nature, and hence indirect methods must be adopted to jointly model data and the conditions. Liu et al. [28] propose an encoder-discriminator-classifier-decoder based approach on flow latent space which can generate synthetic samples for a domain by passing its conditions via encoders to the data via a flow network. They show improvements in varying the quality of generated images for handles relating to various features.

We present CDCGen, a generative framework that is capable of transferring knowledge across multiple domains and using it to generate synthetic samples for domains lacking information about labels/attributes. We model the label/attribute scarce domain as

**Table 1: Comparison of CDCGen with state-of-the-art cross domain translation and conditional synthesis models. Across the board, CDCGen features all the advantages over other models.**

Model	Cross-Domain Translation	Cycle Consistency	Independent Synthesis	Conditional	Availability of Latent Space Embeddings
XGAN [35]	✓	Approximate		✗	✗
CycleGan [44]	✓	Approximate		✗	✗
Taigman et al. [38]	✓	Approximate		✗	✗
Alignflow [17]	✓	Exact		✗	✓
CGAN [31]	✗	-		✓	✗
ACGAN [33]	✗	-		✓	✗
CAGlow [28]	✗	-		✓	✓
<b>CDCGen (ours)</b>	✓	<b>Exact</b>		✓	✓

the target, and a related domain with available information about its labels/attributes as the source. We model the source and target domain via normalizing flows with a common latent space. For conditional synthesis, we introduce a variant of ACGAN by using it on the learned latent space rather than the data space, and train it with only the data and available labels from the source domain. The features can be manipulated easily in the latent space, which is learnt by the conditional synthesis network. During inference phase, CDCGen offers independently specifying conditions, encoding them to a common latent space and moving through the inverse flow to generate conditional synthetic samples in the target domain. Table 1 summarizes the comparison between CDCGen and other related models for different feature availability. CDCGen comes out to be an amalgamation of all features available among the model selections.

We establish the CDCGen framework and conduct empirical evaluations with benchmarked image datasets. CDCGen shows encouraging performance in domain alignment, as well as conditional generation for all source and target combinations.

## 2 PRELIMINARIES

Flow-based generative models and generative adversarial networks constitute the major building blocks for proposed CDCGen. In this section, we will briefly describe the functioning of both types of generative models. We will also iterate through the related works in the space of cross-domain translation and conditional synthesis.

### 2.1 Flow-based Generative Models

Let  $X$  be a high-dimensional random vector with unknown true distribution  $p(x)$ . The following formulation is directly applicable to continuous data, and with some pre-processing steps such as dequantization [18, 37, 41] to discrete data. Let  $Z$  be the latent variable with a known standard distribution  $p(z)$ , such as a standard multivariate gaussian. Using an i.i.d. dataset  $\mathcal{D}$ , the target is to model  $p_\theta(x)$  with parameters  $\theta$ . A flow,  $\mathcal{F}$  is defined to be an invertible transformation that maps observed data  $X$  to the latent variable  $Z$ . A flow is invertible, so the inverse function  $\mathcal{T}$  maps  $Z$  to  $X$ , i.e.

$$Z = \mathcal{F}(X) = \mathcal{T}^{-1}(X) \text{ and } X = \mathcal{T}(Z) = \mathcal{F}^{-1}(Z) \quad (1)$$

The log-likelihood can be expressed as,

$$\log p_\theta(x) = \log p(z) + \log \left| \det \left( \frac{\partial \mathcal{F}(x)}{\partial x^T} \right) \right| \quad (2)$$

where  $\frac{\partial \mathcal{F}(x)}{\partial x^T}$  is the Jacobian of  $\mathcal{F}$  at  $x$ . The training of flow models is accomplished via maximum-likelihood estimation. Confirming with the qualifying properties for a flow as above, different types of flow models have been introduced to efficiently estimate the distribution density and generate synthetic samples [3, 11, 12, 24].

### 2.2 Generative Adversarial Networks (GANs)

GANs [16] are a class of implicit generative models which work based on the principles of a mini-max game. It involves a generator  $\mathcal{G}$  which is tasked to generate synthetic samples from standard noise distribution and a critic  $C$  which learns to discriminate the samples generated by  $\mathcal{G}$  and samples from original data distribution  $p_{data}$ . The training objective for a GAN is given by,

$$\min_{\mathcal{G}} \max_C \mathcal{L}(C, \mathcal{G}) = \mathbb{E}_{x \sim p_{data}} [\log C(x)] + \mathbb{E}_{z \sim \text{standard noise distribution}} [\log(1 - C(\mathcal{G}(z)))]$$

At nash equilibrium, the generator and critic are optimal, and the generator is capable of generating samples resembling original data. GAN based models have been particularly successful in generating high-fidelity images [2, 23], manipulating features of images to generate custom samples [34], audio generation [13, 14], video generation [39] etc. Despite the potential of GANs in generating qualitative samples, they are hard to train due to the mini-max optimization. Unlike flow models, they lack a latent space suitable for a number of downstream applications. Another major problem with GANs is mode collapse, where the generator starts producing the same output (or a small set of outputs) over and over again. A number of remedies have been proposed to tackle this over the years [1, 30, 36].

## 3 RELATED WORK

We discuss the related work from two perspectives relevant to the CDCGen framework, namely cross-domain translation and conditional synthesis.

### 3.1 Cross-Domain Translation

Cross-domain translation involves construction of mappings between two or more domains, by training on unpaired data samples in both the domains. Such a problem is under-constrained and involves aligning the domains in feature space via mappings. A number of research in this space [27, 35, 40, 43, 44] introduce a form of cycle consistency loss which ensures that by translating an image from one domain to another domain via mappings and then applying reverse mappings to translate back yields the same image. XGAN [35] uses additional loss terms to incorporate semantic consistency across domains, to match the subspace for embedding from multiple domains and prior knowledge via pre-trained models. However, since all the above models involve GAN based architectures, they lack a latent space embedding useful for downstream manipulation tasks [24]. Moreover, since the mappings are not guaranteed to be invertible, the cycle consistency is only approximate.

Alignflow [17] involves modeling each of the domains via normalizing flow mappings to a common latent space [6, 7, 12]. It has a hybrid training objective constituting both maximum likelihood estimation and adversarial training. Moreover, since flows are invertible mappings, Alignflow achieves exact cycle consistency. However, flow models, by virtue of the training procedure, face a challenge to align domains which are apart in terms of semantics and/or style, apparent from the generated samples quality in comparison with GANs. For CDCGen, we use the best of both worlds: flow model mappings for the domains to a common latent space, along with loss terms useful to align the domains in the embedding space. CDCGen offers a rich latent space, which is further utilized for conditional synthesis in label/attribute scarce domains.

### 3.2 Conditional Synthesis

Conditional generative models have been introduced to generate desired synthetic data by incorporating conditions information in model design. From CGAN [31] which is a modification of conventional GANs and works by feeding the label/attribute information to the generating block, conditional synthesis has seen different algorithmic variations [20, 32, 42]. A notable work, ACGAN [32] employs an auxiliary classifier for the discriminator to classify the class labels. A recent work, CAGlow [28] proposes a variant of ACGAN with an encoder-decoder network, adding ability to model unsupervised conditions. Additionally, above works deal with conditional generation in a single domain. We use a variant of ACGAN over a shared latent space for multiple domains, thereby transferring knowledge from label-rich domains to perform conditional synthesis in label-scarce domains.

## 4 THE CDCGEN FRAMEWORK

In this section, we will present the CDCGen framework capable of generating conditional synthetic samples for a domain in an unsupervised setting. We select a domain with availability of information about the labels/attributes (namely, the source domain) and has shared attributes with the domain for which we don't have information about labels/attributes (namely, the target domain). Under this setting, the framework consists of two major networks: one for domain alignment and one for conditional synthesis. We consider the case of two domains, but under the assumption of

having shared attributes between the source and target domains, the proposed method generalizes to multiple domains seamlessly.

### 4.1 Domain Alignment

The first step in CDCGen is to align the source and target domains. Let the source and target domain be denoted by  $\mathcal{D}_s$  and  $\mathcal{D}_t$  with unknown marginal density  $p_s$  and  $p_t$  respectively. Both the domains are mapped via invertible transformations (normalizing flows)  $\mathcal{F}_s$  and  $\mathcal{F}_t$  to a common latent space  $Z$ , which serves as a shared feature space for alignment. We assume the shared latent space follows a normal gaussian distribution  $p(z)$ , common for training of most of the state-of-the-art flow models. The relationship between the sample space and latent space can be represented as,

$$\mathcal{D}_s \xrightarrow{\mathcal{F}_s} Z \xleftarrow{\mathcal{F}_t} \mathcal{D}_t$$

Note that the invertible nature of the flow model is helpful in two different ways,

- It provides a mechanism to translate between source and target domains, with invertible mappings  $\mathcal{F}_{s \rightarrow t} = \mathcal{F}_t^{-1} \circ \mathcal{F}_s$  and  $\mathcal{F}_{t \rightarrow s} = \mathcal{F}_s^{-1} \circ \mathcal{F}_t$ .
- It helps achieve exact cycle consistency (as introduced in CycleGAN [44] to ensure accurate representation of the mappings) between the domains, since  $\mathcal{F}_{s \rightarrow t} \circ \mathcal{F}_{t \rightarrow s} = \mathcal{F}_t^{-1} \circ \mathcal{F}_s \circ \mathcal{F}_s^{-1} \circ \mathcal{F}_t = I$ , where  $I$  is the identity matrix.

We use a hybrid training objective involving both maximum likelihood estimation and adversarial training. Flow models are trained with an unsupervised maximum likelihood loss, with a normal gaussian prior on the latent space  $Z$ . Since there are two flow models involved for the two domains, the maximum likelihood loss is expressed as,

$$\mathcal{L}_{MLE}(\mathcal{F}_s) + \mathcal{L}_{MLE}(\mathcal{F}_t)$$

For cross-domain mappings, adversarial loss terms are introduced. These terms introduce inductive bias required for cross domain translation [44]. We employ critics  $C_s$  and  $C_t$  for source and target domains respectively, which distinguish between real samples (sampled from the same domain) vs. generated samples (obtained via cross-domain mappings). For example, the adversarial loss for source domain can be expressed as,

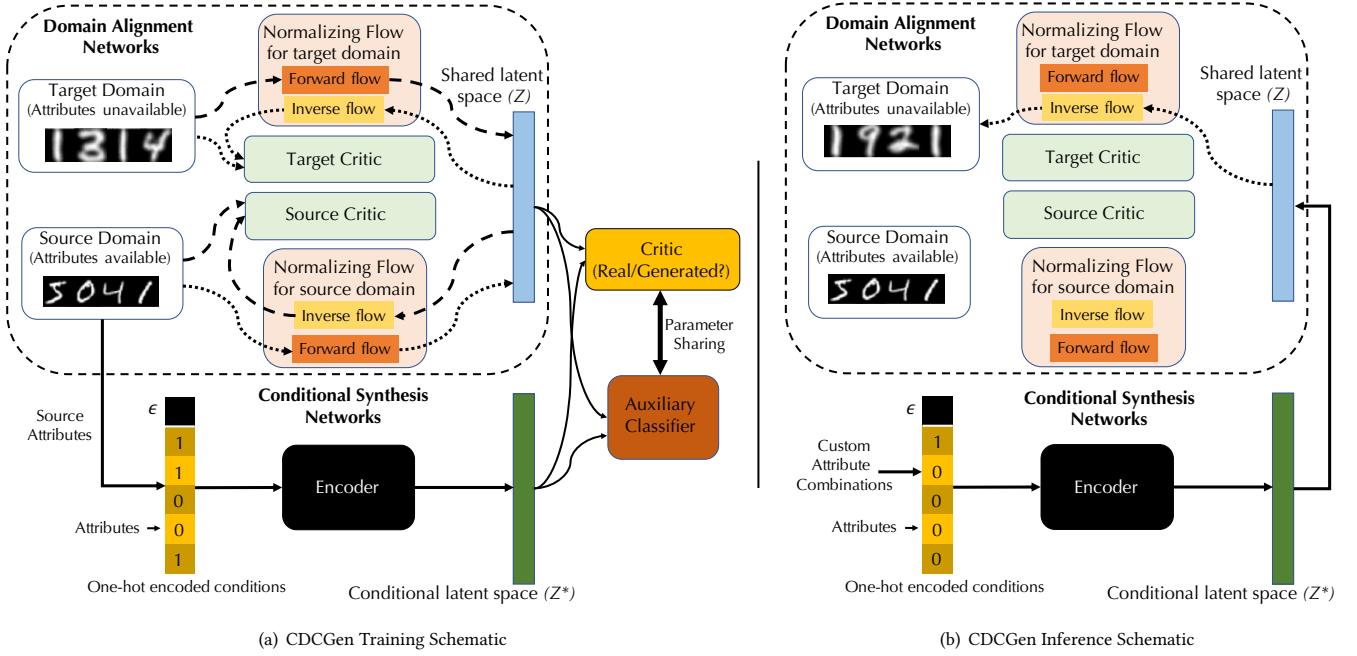
$$\begin{aligned} \mathcal{L}_{ADV}(C_s, \mathcal{F}_{t \rightarrow s}) = & \mathbb{E}_{x_s \sim p_s} [\log C_s(x_s)] \\ & + \mathbb{E}_{x_t \sim p_t} [\log(1 - C_s(\mathcal{F}_{t \rightarrow s}(x_t)))] \end{aligned}$$

We also use a domain-adversarial loss [15] which forces the embeddings learnt by the flow models  $\mathcal{F}_s$  and  $\mathcal{F}_t$  to lie in the same subspace. This is achieved by training a classifier  $C_{DAL}$  which takes the latent space embeddings for each domain and classifies the sample to be coming from  $\mathcal{D}_s$  or  $\mathcal{D}_t$ . It is trained in an adversarial manner, with a classification loss function  $\ell(\cdot, \cdot)$ , such as cross-entropy.  $\mathcal{L}_{DAL}$  can be expressed as,

$$\mathcal{L}_{DAL}(\mathcal{F}_s, C_{DAL}) = \mathbb{E}_{x_s \sim p_s} \ell(\mathcal{D}_s, C_{DAL}(\mathcal{F}_s(x_s)))$$

Finally, for domain alignment, the overall loss term is,

$$\begin{aligned} \mathcal{L}_{Domain\ Alignment}(\mathcal{F}_s, \mathcal{F}_t, C_s, C_t, C_{DAL}; \lambda_s, \lambda_t, \gamma_s, \gamma_t) = \\ \mathcal{L}_{ADV}(C_s, \mathcal{F}_{t \rightarrow s}) + \mathcal{L}_{ADV}(C_t, \mathcal{F}_{s \rightarrow t}) + \gamma_s \mathcal{L}_{DAL}(\mathcal{F}_s, C_{DAL}) \\ + \gamma_t \mathcal{L}_{DAL}(\mathcal{F}_t, C_{DAL}) - \lambda_s \mathcal{L}_{MLE}(\mathcal{F}_s) - \lambda_t \mathcal{L}_{MLE}(\mathcal{F}_t) \end{aligned}$$



**Figure 1: Illustration of training and inference methods in CDCGen. The networks inside the dashed box are for domain alignment (Sec 4.1) and those outside are for conditional synthesis (Sec 4.2).**

where, hyperparameters  $\lambda_s$  and  $\lambda_t$  dictate the relative contribution of maximum likelihood loss, and  $\gamma_s$  and  $\gamma_t$  correspond to contribution of domain adversarial loss, both as compared to the adversarial loss. The objective is minimized w.r.t. the parameters of the flow models  $\mathcal{F}_s$  and  $\mathcal{F}_t$  and maximized w.r.t. parameters of  $C_s$ ,  $C_t$  and  $C_{DAL}$ . This procedure is illustrated in the dashed box in Fig. 1(a).

## 4.2 Conditional Synthesis

For conditional synthesis, we propose a variant of ACGAN [33]. Instead of using class/attribute conditioning on the sample space as done in ACGAN, we use it in the shared latent space. Under the setting of our problem, we don't have any information about the labels/attributes in the target domain. So, for the conditional synthesis part, only the attributes available from the source domain are used for training.

We denote the available source attributes/conditions as  $c_s \sim p(c_s)$ , represented as one-hot encodings. Our network consists of an encoder to model the conditions, a critic to differentiate between the real and generated latent vectors, and an auxiliary classifier to classify the encoded conditions. We will introduce each of the above components and their associated loss functions separately.

**Encoder:** An encoder network  $E$  encodes the conditions  $(c_s, \epsilon)$  into a latent space  $Z^*$  (separate from the shared latent space  $Z$  for aligned domains), where  $\epsilon$  is sampled from standard gaussian distribution  $(p(\epsilon))$  and is helpful for incorporating stochastic behavior among condition vectors. Let the distribution for the above mentioned latent space be denoted as  $p^*(z)$ . Our objective is to minimize the Jensen-Shannon (JS) divergence between the encoded distribution  $p^*(z)$  and the shared latent distribution  $p(z)$  for aligned

domains  $\mathcal{D}_s$  and  $\mathcal{D}_t$ . So, the encoder loss is represented as,

$$\mathcal{L}_E = \mathbb{E}_{\epsilon \sim p(\epsilon), c_s \sim p(c_s)} [\log C(E(c_s, \epsilon))]$$

where,  $C$  is a critic, more about which we describe now.

**Critic:** A critic  $C$  discriminates between the latent vectors coming from generated conditional distribution  $p^*(z)$  and real shared latent distribution  $p(z)$  for aligned domains. This is an adversarial loss which is trained so as it is unable to distinguish the latent vectors at equilibrium, thus enabling the encoder  $E$  to generate latent vectors close to the real shared latent distribution  $p(z)$ . The loss function for  $C$  is,

$$\mathcal{L}_{CRITIC} = \mathbb{E}_{z \sim p^*(z)} [\log C(z)] + \mathbb{E}_{z \sim p(z)} [1 - \log C(z)]$$


**Classifier:** A classifier takes the latent vectors  $(z \sim p^*(z))$  and  $z \sim p(z)$  as input and classifies the conditions  $(c_s)$ . The classifier loss is a cross entropy loss between the predicted and true conditions. If the class posterior probabilities are  $q(c_s|z)$ , the classifier loss function can be expressed as,

$$\mathcal{L}_{CLASSIFIER} = \mathbb{E}_{z \sim p^*(z), c_s \sim p(c_s)} q(c_s|z) + \mathbb{E}_{z \sim p(z), c_s \sim p(c_s)} q(c_s|z)$$

The overall loss function for the conditional synthesis part is,

$$\mathcal{L}_{Conditional\ Synthesis} = \beta_E \mathcal{L}_E + \beta_{Cr} \mathcal{L}_{CRITIC} + \beta_{CI} \mathcal{L}_{CLASSIFIER}$$

where  $\beta_E, \beta_{Cr}, \beta_{CI}$  are hyperparameters. The critic and the classifier networks share their parameters except for their output blocks. Conditional synthesis procedure is illustrated in Fig. 1(a).



**Figure 2: Results for domain alignment between source and target domains. The top row has original samples from the source domain. The middle row is the corresponding latent space mapping and the bottom row is the sample obtained by translating it to the target domain. The USPS images are slightly blurred due to the upscaling applied as standard pre-processing.**

### 4.3 Inference

CDCGen can generate conditional samples in the target domain, even when the training process does not utilize its class/attribute information. To generate samples with conditions  $\tilde{c}_s$ , a latent vector  $\tilde{z}$  is generated by encoding the one-hot conditions  $\tilde{c}_s$  and  $\tilde{\epsilon} \sim p(\epsilon)$  via the encoder network, i.e.  $\tilde{z} = E(\tilde{c}_s, \tilde{\epsilon})$ . Then the latent vector  $\tilde{z}$  is passed via the inverse flow  $\mathcal{F}_t^{-1}$  to generate the desired sample in the target domain, i.e.  $\mathcal{F}_t^{-1}(\tilde{z})$ . The inference schematic is illustrated in Fig. 1(b).

## 5 EXPERIMENTS

In this section, we empirically evaluate CDCGen for synthetic generation in label scarce domains.

**Datasets:** We perform experiments on 2 standard image datasets for digits, namely MNIST [26] and USPS. MNIST contains 60,000 training and 10,000 test images with ten classes corresponding to digits from 0 to 9. USPS has 7291 training and 2007 test data with the same classes as MNIST. To address this imbalance, for each domain, we sample 542 images from the original training set for each class to form the new training set. To form the test set, we sample 147 images from the original test sets for each class. We resize all the images to  $32 \times 32$  for training and synthesis.

**Source and Target Domain Combinations:** We consider two cases, first with MNIST as the source and USPS as the target domain, and second, with the roles interchanged, i.e. USPS as the source and MNIST as the target. We report results for domain alignment and as well as subsequent conditional synthesis in the target domain, all while not using any labels from that domain.

**Networks:** We use architecture from RealNVP [12] for each of the domain flows ( $\mathcal{F}_s$  and  $\mathcal{F}_t$ ). Typical configurations for RealNVP can be specified as a tuple comprising  $N_{scales}$  (number of scales),  $N_{channels}$  (number of channels) in the intermediate layers, and  $N_{blocks}$  (number of residual blocks in the scaling and translation networks of the coupling layers). For MNIST  $\leftrightarrow$  USPS case, both  $\mathcal{F}_s$  and  $\mathcal{F}_t$  are set to RealNVP(2, 64, 8). The critics ( $C_s$  and  $C_t$ ) used convolutional discriminators from PatchGAN [21], each with 16 filters in the critic’s first convolutional layer. For conditional synthesis, we concatenate the one-hot vector of labels with components of random noise as input to the encoder. The vector then passes through one fully-connected layer and eight transposed convolutional layers with upsampling scale 2, 2, 2, 2, 2, 1, 1, 1 and channel sizes 256, 1024, 512, 256, 128, 64, 32, 16 respectively. The

supervision block contains four convolutional layers with stride 2 and channel sizes 64, 128, 256, 512. This is followed by two separate fully-connected layers for each network head, one for outputting probabilities of real or fake and the other for classifying the one-hot encoded conditions.

**Optimizer:** For training the domain alignment network, we use the Adam optimizer with  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ , and learning rate  $1 \cdot 10^{-6}$ . For training the conditional synthesis network, we use the Adam optimizer with  $\beta_1 = 0.5$ ,  $\beta_2 = 0.999$ , and learning rate  $2 \cdot 10^{-5}$ .

### 5.1 Domain Alignment

In this section we present the results for domain alignment between source and target combinations. Fig. 2(a) shows the source MNIST samples and corresponding USPS samples by translating it via the forward source and inverse target flows. The middle sample is visualization of corresponding latent space sample. Fig. 2(b) depicts the same with USPS as the source and MNIST as the target. It can be observed that the class identity is preserved with the translation with the style adapted for the target domains. The sharpness of the translated samples are compromised, which is a result of the flow model assigning some probability mass to all the samples it is fed. This is unlike pure GAN based models which selectively assign probability mass to meaningful samples.

Another interesting observation is the appearance of digit class identity in latent space visualizations. This is particularly useful from the perspective of CDCGen, since the conditional synthesis network works based on the latent space mappings from both the domains.

We present the t-SNE embeddings for the shared latent space in our proposed domain alignment network for MNIST and USPS in Fig. 3. It can be observed that the visualization has distinct clusters for each digit class, but the embeddings from both the source and target domain are close and belong to the same cluster for the overall digit class clustering. The visualization allows us to infer that the latent space has learned a subspace corresponding to each digit, and interpolating across this subspace is effectively a conditional feature-preserving domain transfer.

### 5.2 Conditional Synthesis

We trained the conditional synthesis part of CDCGen (Section 4.2) with source labels to generate conditional synthetic samples in the

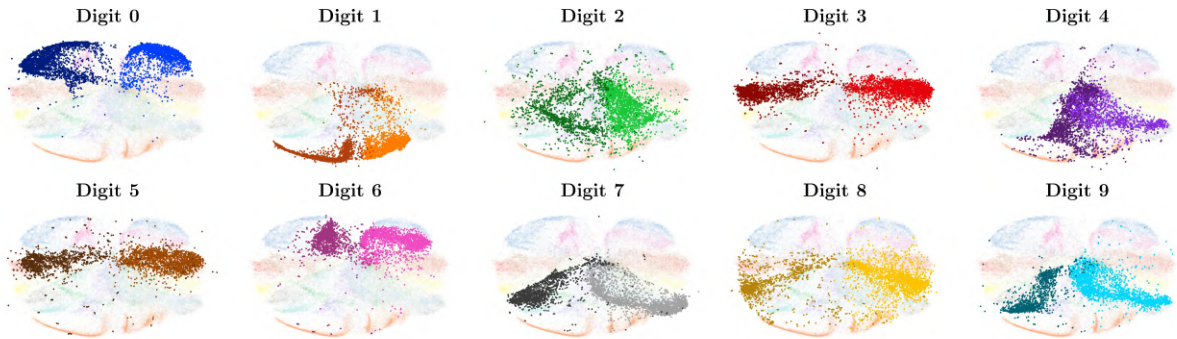



Figure 3: t-SNE representation of shared latent space for MNIST  $\leftrightarrow$  USPS. For each digit, points for USPS are visualized with the darker colors, and points with lighter colors correspond to MNIST.



(a) Generated samples for USPS as target

(b) Generated samples for MNIST as target

Figure 4: Conditional synthetic samples generated by CDCGen. The rows represent conditioned digit classes (0-9) and the columns include more samples for each class.

target domain. Fig. 4(a) shows the samples generated with USPS as the target domain and Fig. 4(b) shows the samples generated with MNIST as the target. Each row corresponds to the digit classes which are assigned as conditions. It can be observed that CDCGen is able to generate synthetic samples belonging to the digit class as conditioned. There are also variations among the samples across different columns which shows the stochastic nature of generation by CDCGen. The compromise in sharpness of the samples generated can be observed in the generated samples too, and is owed from the domain alignment mappings by flow models.

## 6 CONCLUSIONS

CDCGen, a generative framework capable of generating conditional synthetic samples for domains without the requirement of obtaining its labels/attributes was presented. We also conducted empirical studies with standard image datasets to observe feature transfer and independent conditional generation. In the future, making the conditional generation models across multiple domains can be studied with varying levels of label availability (few-shot learning) for

target domain. CDCGen can also be adapted for other modalities of data including audio and tabular data. It can also be used alongside real-world applications [8–10, 22, 25, 29, 45] where having access to diverse conditional data is important, but is hard to obtain, hence the need for synthetic data.

## REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein gan. *arXiv preprint arXiv:1701.07875* (2017).
- [2] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2019. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=B1xsqj09Fm>
- [3] Ricky TQ Chen, Jens Behrmann, David Duvenaud, and Jörn-Henrik Jacobsen. 2019. Residual Flows for Invertible Generative Modeling. *arXiv preprint arXiv:1906.02735* (2019).
- [4] Moustapha Cisse, Piotr Bojanowski, Edouard Grave, Yann Dauphin, and Nicolas Usunier. 2017. Parseval networks: Improving robustness to adversarial examples. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 854–863.
- [5] Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. 2017. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems*. 3730–3739.

- [6] Hari Prasanna Das, Pieter Abbeel, and Costas J Spanos. 2019. Dimensionality reduction flows. *arXiv preprint arXiv:1908.01686* (2019), 1–10.
- [7] Hari Prasanna Das, Pieter Abbeel, and Costas J Spanos. 2019. Likelihood Contribution based Multi-scale Architecture for Generative Flows. *arXiv preprint arXiv:1908.01686* (2019).
- [8] Hari Prasanna Das, Ioannis C Konstantakopoulos, Aummul Baneen Manasawala, Tanya Veeravalli, Huihan Liu, and Costas J Spanos. 2019. A novel graphical lasso based approach towards segmentation analysis in energy game-theoretic frameworks. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*. IEEE, 1702–1709.
- [9] Hari Prasanna Das, Stefano Schiavon, and Costas J Spanos. 2021. Unsupervised personal thermal comfort prediction via adversarial domain adaptation. In *Proceedings of the 8th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. 230–231.
- [10] Hari Prasanna Das, Ryan Tran, Japjot Singh, Xiangyu Yue, Geoffrey Tison, Alberto Sangiovanni-Vincentelli, and Costas J Spanos. 2022. Conditional synthetic data generation for robust machine learning applications with limited pandemic data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 36. 11792–11800.
- [11] Laurent Dinh, David Krueger, and Yoshua Bengio. 2014. Nice: Non-linear independent components estimation. *arXiv preprint arXiv:1410.8516* (2014).
- [12] Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using Real NVP. *CoRR abs/1605.08803* (2016). [arXiv:1605.08803](https://arxiv.org/abs/1605.08803)
- [13] Chris Donahue, Julian McAuley, and Miller Puckette. 2018. Adversarial audio synthesis. *arXiv preprint arXiv:1802.04208* (2018).
- [14] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. 2019. Gansynth: Adversarial neural audio synthesis. *arXiv preprint arXiv:1902.08710* (2019).
- [15] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. 2016. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* 17, 1 (2016), 2096–2030.
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [17] Aditya Grover, Christopher Chute, Rui Shu, Zhangjie Cao, and Stefano Ermon. 2019. AlignFlow: Cycle Consistent Learning from Multiple Domains via Normalizing Flows. *arXiv preprint arXiv:1905.12892* (2019).
- [18] Jonathan Ho, Xi Chen, Aravind Srinivas, Yan Duan, and Pieter Abbeel. 2019. Flow++: Improving Flow-Based Generative Models with Variational Dequantization and Architecture Design. *arXiv preprint arXiv:1902.00275* (2019).
- [19] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei A Efros, and Trevor Darrell. 2017. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213* (2017).
- [20] Seunghoon Hong, Dingdong Yang, Jongwook Choi, and Honglak Lee. 2018. Inferring semantic layout for hierarchical text-to-image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7986–7994.
- [21] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1125–1134.
- [22] Ming Jin, Ruoxi Jia, Hari Prasanna Das, Wei Feng, and Costas Spanos. 2018. BISCUIIT: Building Intelligent System CUsTomer Investment Tool. In *10th International Conference on Applied Energy (ICAE)*.
- [23] Tero Karras, Samuli Laine, and Timo Aila. 2018. A style-based generator architecture for generative adversarial networks. *arXiv preprint arXiv:1812.04948* (2018).
- [24] Durk P Kingma and Prafulla Dhariwal. 2018. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*. 10215–10224.
- [25] Ioannis C Konstantakopoulos, Hari Prasanna Das, Andrew R Barkan, Shiyang He, Tanya Veeravalli, Huihan Liu, Aummul Baneen Manasawala, Yu-Wen Lin, and Costas J Spanos. 2019. Design, benchmarking and explainability analysis of a game-theoretic framework towards energy efficiency in smart infrastructure. *arXiv preprint arXiv:1910.07899* (2019).
- [26] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998), 2278–2324.
- [27] Ming-Yu Liu and Oncel Tuzel. 2016. Coupled Generative Adversarial Networks. *arXiv:1606.07536* [cs.CV]
- [28] Rui Liu, Yu Liu, Xinyu Gong, Xiaogang Wang, and Hongsheng Li. 2019. Conditional adversarial generative flow for controllable image synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 7992–8001.
- [29] Shichao Liu, Stefano Schiavon, Hari Prasanna Das, Ming Jin, and Costas J Spanos. 2019. Personal thermal comfort models with wearable sensors. *Building and Environment* 162 (2019), 106281.
- [30] Luke Metz, Ben Poole, David Pfau, and Jascha Sohl-Dickstein. 2016. Unrolled generative adversarial networks. *arXiv preprint arXiv:1611.02163* (2016).
- [31] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [32] Augustus Odena. 2016. Semi-supervised learning with generative adversarial networks. *arXiv preprint arXiv:1606.01583* (2016).
- [33] Augustus Odena, Christopher Olah, and Jonathon Shlens. 2017. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2642–2651.
- [34] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434* (2015).
- [35] Amélie Royer, Konstantinos Bousmalis, Stephan Gouws, Fred Bertsch, Inbar Mosseri, Forrester Cole, and Kevin Murphy. 2020. Xgan: Unsupervised image-to-image translation for many-to-many mappings. In *Domain Adaptation for Visual Understanding*. Springer, 33–49.
- [36] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*. 2234–2242.
- [37] Tim Salimans, Andrej Karpathy, Xi Chen, and Diederik P. Kingma. 2017. PixelCNN++: Improving the PixelCNN with Discretized Logistic Mixture Likelihood and Other Modifications. *CoRR abs/1701.05517* (2017). [arXiv:1701.05517](https://arxiv.org/abs/1701.05517)
- [38] Yaniv Taigman, Adam Polyak, and Lior Wolf. 2016. Unsupervised cross-domain image generation. *arXiv preprint arXiv:1611.02200* (2016).
- [39] Sergey Tulyakov, Ming-Yu Liu, Xiaocong Yang, and Jan Kautz. 2018. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1526–1535.
- [40] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. 2017. Adversarial Discriminative Domain Adaptation. [arXiv:1702.05464](https://arxiv.org/abs/1702.05464) [cs.CV]
- [41] Benigno Uria, Iain Murray, and Hugo Larochelle. 2013. RNADE: The real-valued neural autoregressive density-estimator. In *Advances in Neural Information Processing Systems*. 2175–2183.
- [42] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8798–8807.
- [43] Zili Yi, Hao Zhang, Ping Tan, and Minglun Gong. 2017. Dualgan: Unsupervised dual learning for image-to-image translation. In *Proceedings of the IEEE international conference on computer vision*. 2849–2857.
- [44] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. 2017. Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks. [arXiv:1703.10593](https://arxiv.org/abs/1703.10593) [cs.CV]
- [45] Han Zou, Jianfei Yang, Hari Prasanna Das, Huihan Liu, Yuxun Zhou, and Costas J Spanos. 2019. Wifi and vision multimodal learning for accurate and robust device-free human activity recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. 0–0.
- [46] Han Zou, Yuxun Zhou, Jianfei Yang, Huihan Liu, Hari Prasanna Das, and Costas J Spanos. 2019. Consensus Adversarial Domain Adaptation. In *AAAI Conference on Artificial Intelligence 2019*.