Inducing Optimal Attribute Representations for Conditional GANs

Binod Bhattarai¹ and Tae-Kyun Kim^{1,2}

¹ Imperial College London, UK
 ² KAIST, Daejeon, South Korea
 {b.bhattarai,tk.kim}@imperial.ac.uk



Fig. 1: Simultaneous manipulation of target facial attributes along with their auxiliary. From left to right, in each pair, original images change to Old (along with negate of black hair), Female (lipstick/makeup), and Bald (wrinkles). The attribute representations embedding relations are automatically learnt and applied to condition the generator.

Abstract. Conditional GANs (cGANs) are widely used in translating an image from one category to another. Meaningful conditions on GANs provide greater flexibility and control over the nature of the target domain synthetic data. Existing conditional GANs commonly encode target domain label information as hard-coded categorical vectors in the form of 0s and 1s. The major drawbacks of such representations are inability to encode the high-order semantic information of target categories and their relative dependencies. We propose a novel end-to-end learning framework based on Graph Convolutional Networks to learn the attribute representations to condition the generator. The GAN losses, the discriminator and attribute classification loss, are fed back to the graph resulting in the synthetic images that are more natural and clearer with respect to the attributes generation. Moreover, prior-arts are mostly given priorities to condition on the generator side, not on the discriminator side of GANs. We apply the conditions on the discriminator side as well via multi-task learning. We enhanced four state-of-the-art cGANs architectures: Stargan, Stargan-JNT, AttGAN and STGAN. Our extensive qualitative and quantitative evaluations on challenging face attributes manipulation data set, CelebA, LFWA, and RaFD, show that the cGANs enhanced by our methods outperform by a large margin, compared to their counter-parts and other conditioning methods, in terms of both target attributes recognition rates and quality measures such as PSNR and SSIM.

Keywords: Conditional GAN, Graph Convolutional Network, Multitask Learning, Face Attributes

1 Introduction

Someone buying bread is likely to buy butter, blue sky comes with a sunny day. Similarly, some of the attributes of the faces co-occur more frequently than others. Fig. 2 shows co-occurring probabilities of facial attributes. We see some set of attributes such as *wearing lipsticks* and *male* are least co-occurring (0.01) and *male* and *bald* are highly co-related (1.0).

Face attribute manipulation using GAN [9,16,26,42,48,7,46,3,23] is one of the challenging and popular research problem. Since the advent of conditional GAN [31], several variants of conditional GANs (cGANs) have been proposed. For conditioning the GAN, existing methods rely on target domain one-hot vectors [9,16,25,33], synthetic model parameters of target attributes [11], facial action units [39], or key point landmarks [31], to mention a few of them. Recently, [26] proposed to use the difference of one-hot vectors corresponding to the target and source attributes. This trick *alone* boosts Target Attributes Recognition Rate (TARR) on synthetic data compared to [9] by a large margin. Another recent study on GAN's [35] identified conditioning on GAN is co-related with its performance. The major limitation of existing cGANs for arbitrary multiple face attributes manipulation is [9,16,26,37,25] are: hard coded 1 and 0 form, treating every attribute equally different and ignoring the co-existence of the attributes. In reality, as we can see in Figure 2, some attributes are more co-related than others. Moreover, the existing methods are giving less attention to conditioning on the discriminator side except minimising the cross-entropy loss of target attributes.

Another recent work [28] identified the problem of artefacts on synthetic examples due to unnatural transition from source to target. This problem arises due to the ignorance of existing GANs regarding the co-existence of certain sub set of attributes. To overcome this, they propose a hard-coded method to condition both target attribute (aging) and its associated attributes (gender, race) on generator and also on discriminator in order to faithfully retain them after translation. However, this approach is limited to a single attribute and infeasible to hard code such rules in the case like ours where multiple arbitrary attributes are manipulated simultaneously. Recent study on GAN [5] identifies the forgetting problem of discriminator due to the non-stationary nature of the data from the generator. Applying a simple structural identification loss (rotation angle) helps to improve the performance and stability of GAN.

	strate for her	the second s	and the local diversion of the local diversio	and the second sec	and the second	Contractory of the	Conception in the	the second s	and the second second	and the second se	_	
						Concession in which the local division in th		and the second se				
and the second second							- CORE 1	and the second second	NR. 111			
A 10 Longs Auto and and			and the second second		100.000		1000	and the second se		A COMPANY OF A		
and the second second	the second s			-		and the second second			_	States and s		
	and the second se				100 C	COLUMN TWO IS NOT	1000	-				
								And in case of the	ALC: U.S.			
	Statement of the local division of the local	And the second second	1000		10 million (* 1777)	C		and the second se	10.000	1000000000		
	And in case of the local division of the loc		A REAL PROPERTY.		and the second second		100.00	Contraction of the local division of the loc		COLUMN TWO IS NOT		
and the second	ALC: NOT THE OWNER.	and the second			And in case of the local division of the loc	ALC: NOT THE OWNER OF	1000	A CONTRACTOR OF	ALC: NOT THE			
	Concession of the local division of the loca	THE OWNER WATER OF THE OWNER OF T			A DESCRIPTION OF			and the second second	-100 - 100 ET			
		COLUMN TWO IS NOT			And in case of the local division of the loc		THE OWNER WATER	And an other Designation of the local division of the local divisi	And in case of	ALC: NO DECKS		
	1000	The second s	1.000		A COLORED	and the second second	100 H H	COLUMN TWO IS NOT	CONTRACTOR OF THE OWNER OWNER OF THE OWNER OWNE	A CONTRACTOR OF		
	NAME OF TAXABLE PARTY.		A REAL PROPERTY.			Contraction of the local division of the loc	1 C C C C C C C C C C C C C C C C C C C			1000		
Aug 100, 510	ALC: NO TO BE ADDRESS OF TAXABLE PARTY.	A REAL PROPERTY AND A	A REAL PROPERTY.		and the second second	the second second	100.00	and the second second				
	All second s	THE R. LEWIS CO., LANSING MICH.		THE OWNER WATER		the second second	100000000	All the same				
the second second			And in case of the local division of the loc									
	and the second second	Contract, 199						1000		And the second		
and the second second	the second second					C. Ballings	- Contract (1997)		ALC: NOT THE OWNER OF			
and the second second								and so the second	ALC: U.S.	A		
and second	1000		100 C 100		the second second		1000	A 1 4 4 4 4	Concession of the			
	And the second second	10000000000	and the second second		the second second			And in case of	the second second			
	Contraction of the	And Control of States	and the second second	COLUMN TO A COLUMN				and the second second	All the second T	100010		
Rented Land Links	the second second	And Address of the owner owner owner owner own	and the second second		A COLUMN TWO IS NOT			And in case of				
	State of the local division of the local div	THE REPORT	And the other Distance		And in case of the local division of the loc		_	ALC: NO. OF THE OWNER, NO.	All shares and shares a			
had have being	22 C - C	TRUE TO A			ALC: YOU ARE ADD	10 March 10	1000	And in case of the local division of the loc	COLUMN 1	A REAL PROPERTY.		
free late at \$100	12.2	CONTRACTOR OF THE			And a second second			A REAL PROPERTY.	- Contract of the local division of the loca	CONTRACTOR OF THE OWNER		
	The second s	IN COLUMN 2 IN COLUMN			Address of the local division of the local d	A COLUMN	- 10	COLUMN TWO IS NOT	and the second second	A REAL PROPERTY.		
	THE OWNER WATER	THE OWNER WATCH				the second second		and the second second	COLUMN TWO IS NOT	And in case of the		
And many limits	the second second	100000000000000000000000000000000000000				and the second se		the second second	and the second se	100 100		
NAME OF TAXABLE	Contraction of the local division of the loc		ALC: NOT THE OWNER OF	COLUMN TWO IS NOT			1.00		100 C			
NAME AND ADDRESS	10000-000	100000000000000000000000000000000000000			-		1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1			CONTRACTOR OF A	_	
mapping house	EE00000	-0.000000			Contract and and	the second se		and the second second	-	A	_	
mity (m) 14, 46		-0.000 111			1000 C 1000		100.00				_	
Reading, Lotings (4),742						-					_	
the second second	the second s					-		and the second second	and the second second	A REAL PROPERTY.	_	
Watering, Status & Color		And in case of the local division of the loc				- 	A REAL PROPERTY.	-	100 100		_	
	States and	And Address of the Owner of the O			ALC: NOT THE OWNER.	-	10000	-	- Colored			
Manha Autor (#11)	And in case of the local division of the loc	10001-010-010	and the local division of the local division	COLUMN TWO IS NOT	States of the local division of the local di	The second value of the se						
Tanks I'm Tak	And in case of the local division of the loc	ALC: NO.	COLUMN TWO IS NOT		F 1 1 1	-	And Parks	and the second second	TAXABLE INCOME.	a second second		

Fig. 2: Co-occurrence matrix of facial attributes (zoom in the view).

To address the above mentioned challenges of cGANs, we investigate few continuous representations including semantic label embedding (word2vec) [30], attributes model parameters (attrbs-weights) (see Sec. 4). The advantages of conditioning with such representations instead of 0s and 1s form are mainly two folds: i) carries high-order semantic information, ii) establishes relative relationship between the attributes. These representations are, however, still sub-optimal and less natural as these are computed offline and also do not capture simultaneous existing tendency of different face attributes. Thus, we propose a novel conditioning method for cGAN to induce higher order semantic representations of target attributes, automatically embedding inter-attributes co-occurrence and also sharing the information based on degree of interaction. Towards this goal, we propose to exploit the attributes co-occurrence probability and apply Graph Convolutional Network (GCN) [24] to condition GAN. GCN is capable of generating dense vectors, distilling the higher dimensional data and also capable of convolving the un-ordered data [15]. The conditioning parameters i.e. GCN are optimised via the discriminator and attribute classification losses of GAN in an end-to-end fashion. In order to maintain such semantic structural relationship of the attributes at the discriminator side as well, we adapted on-line multitask learning objectives [4] constrained by the same co-occurrence matrix. The experiments show that the proposed method substantially improve stateof-the-art cGAN methods and other conditioning methods in terms of target attributes classification rates and PSNR/SSIM. The synthesised images by our method exhibit associated multi-attributes and clearer target attributes. Details of the method are explained in Sec. 3, following the literature review in Sec. 2, and experimental results and conclusions are shown in Sec. 4 and Sec. 5.

2 Related Works

Conditional GANs. After the seminal work from Mirza et al. [31] on Conditional GANs (cGANs), several variants such as, conditioning target category labels in addition to the input [35, 6, 9, 16, 26, 50], semantic text or layout representations conditioning [40,49,17,51], image conditioning [19,18,27], facial landmarks [47], have been proposed to solve different vision tasks. These works highlight the importance of semantic representations of target domain as a condition. In this work, we focus on conditioning target category labels especially by continuous and semantic representations. [20] proposes multiple strategies for random continuous representation to encode target label but limits to a single attribute. [10] extended similar approaches for arbitrary attributes. [21] proposes to use decision tree to generate hierarchical codes to control the target attributes. These are some of the works related to ours in terms of inducing continuous representations. Recent works on cGANs for face attribute manipulations [9,37,25,16] encodes in the form of 0s and 1s or their difference [26]. These representations are hardcoded. STGAN [26] also proposes conditional control of the information flow from source to target in the intermediate layers of GANs. This is one of the closest works in terms of the adaptive conditioning target information. Other cGANs, such as StyleGAN [23] propose to condition on the intermediate layers of the generator. Progressive GAN [22] proposed to gradually increase the parameters of the generator and discriminator successively to generate high quality images. Our

method is orthogonal to this line of methods, and can be extended to these works for a higher quality. Recently, attribute aware age progression GAN [28] proposes to condition both associated attributes and target attributes at both generator and discriminator side. This work is closest in terms of conditioning at both the sides and retaining the auxiliary attributes in addition to target attribute. This approach limits to single attribute manipulation i'e aging, whereas, our method supports multiple attributes. Also their method is hard-coded whereas our method is automatic and directly inferred from the co-occurrence matrix.

Graph Convolutional Network (GCN). Frameworks similar to [24] are popular for several tasks including link prediction [12], clustering [38], node classification [45]. Recent works on image classification [8] and face attributes classification [34] propose to use GCN to induce more discriminative representations of attributes by sharing information between the co-occurring attributes. Unlike these works, we propose to apply GCN to induce such higher-order representations of target categories for the generative neural networks and optimise it via end-to-end adversarial learning. To the best of our knowledge, this is the first work to use such embedding as conditions in cGANS. For more details on the work based on Graph Convolutional Networks, we suggest reader to refer to [52]. **Regularizing/Conditioning the Discriminator.** Conditioning on the discriminator side has been shown useful in generating more realistic and diverse images [33,32,5,36]. [35] maximises the distribution of target label in addition to source distribution to improve the quality of synthetic images. [5] introduced rotation loss on the discriminator side to mitigate the forgetting problem of the discriminator. Projecting the target conditional vector to the penultimate representation of the input at discriminator side [33] substantially improved the quality of synthetic images. Another work on Spectral normalisation of weight parameters [32] of every layer of the discriminator stabilises the training of GANs. Recent works on face attribute manipulations [9,16,26] minimise the target label cross entropy loss on discriminator. In this work, we introduce conditioning of the discriminator with multi-task learning framework while minimising the target attributes cross entropy loss.

3 Proposed Method

3.1 Overview on the Pipeline

Fig. 3 shows the schematic diagram of the proposed method, where both the generator G and discriminator D are conditioned. As mentioned in Sec. 1, existing cGANs arts such as Stargan [9], AttGAN [16] or STGAN [26] condition the generator, which can be done either at the encoder or the decoder part of G, to synthesise the image with intended attributes. But the problem with their conditions is that they ignore the intrinsic properties of the attributes and their relationships. They use single digit (0 or 1) for an attribute, and treat every attributes are equally similar to each other. In Fig. 3, the graph on the generator side represents the attributes and their relationships. Each node in the graph represents higher-order semantic representations of attributes, and the edges

between them represent their relationship. We propose to induce the attribute representations which encode attribute properties including their relations, based on how they co-occur in the real world scenario. To induce such representations, we propose to apply GCN [24] with convolutional layers on the generator side. The graph is optimised via end-to-end learning of the entire system of networks. The discriminator and the attribute classifier guide the graph learning such that the learnt conditional representations help the generator synthesise more realistic images and preserve target attributes. Such semantically rich and higher-order conditional representations of the target attributes play an important role in the natural transitioning to the target attribute. This helps to synthesise images with less artefacts, improved quality and better contrast, as also partially observed in StackGAN [49].

We also condition the parameters of attributes on the discriminator side using multi-task learning framework, similar to [4], based on the co-occurrence matrix. Using the learnt representation i.e. the graph to condition the discriminator might also be possible via EM-like alternating optimisation, however due to the complexity and instability, is not considered in this work. Conditioning both target and its associated attributes on generator and on discriminator enabled GAN to retain the target as well as the associated attributes faithfully [28]. Unlike [28] which is hard-coded, limited to a single attribute, our method is automatic, and supports arbitrary multiple attributes. See Section 3.3 for more details. Before diving into in the details of the proposed method, we first introduce attributes co-occurrence matrix which is exploited in both the generator and discriminator of the proposed method.

Co-occurrence matrix: To capture the relationship between the attributes based on how frequently they go together, we constructed a co-occurrence matrix, $C \in \mathbb{R}^{k \times k}$, where k is the total number of attributes. The value at position (i, j) in the matrix gives us an idea about probability of attributes a_j occurring given the attribute a_i . Fig. 2 shows the co-occurrence matrix. We approximate this probability from the training data set as in Eqn. 1.

$$P(a_i|a_j) = \frac{\#\text{images with } a_i \cap a_j}{\#\text{images with attribute } a_j}$$
(1)

3.2 Graph Convolution and Generator

As stated before, we propose to learn the representations via GCN [24], which we simultaneously use to condition the generator. They are in the form of a Graph, $\mathcal{G} = (V, E)$. In our case there are k different facial attributes, thus total nodes in the graph will be k. We represent each node, also called a vertex V of the graph, by their initial representations of the attributes and the edges between the graph encode their relationship. In our case, this is the co-occurrence probability. Since $P(a_i|a_j) \neq P(a_j|a_i)$, the co-occurrence matrix, C is asymmetric in nature. The graph is constructed from co-occurrence information encoded on C and initial continuous representations of the attributes $X \in \mathbb{R}^{k \times d}$. In Fig. 3, we show a



Fig. 3: Schematic diagram showing the pipeline of the proposed method. Each node of the Graph \mathcal{G} represents an attribute and the edge between them encode their co-occurrence defined on \hat{C} . GCN induces the higher-order representations of the attributes (Z) which are further scaled by (t-s) to generate Z_t . We concatenate Z_t with the latent representations of input image and feed to the decoder of (G). At discriminator, we apply Multi-Task Learning (MTL) to share the weights between the tasks constrained on C. During end-to-end learning, we back-propagate the error to the induced representations (Z) to fine-tune their representations. We maintain the color codes among the attributes (best viewed in color).

single un-directed edge between the two nodes for clarity. The thickness of edges is proportional to the probability of co-occurrence.

The goal of GCN [24,52] is to learn a function $f(\cdot, \cdot)$ on a graph \mathcal{G} , which takes initial node representations X and an adjacency matrix (\widehat{C}) , which we derive from co-occurrence matrix $C \in \mathbb{R}^{k \times k}$ as inputs. And, it updates the node features as $X^{l+1} \in \mathbb{R}^{k \times d'}$ after passing them through every convolutional layer. Every GCN layer can be formulated as

$$X^{l+1} = f(X^l, \widehat{C}) = \sigma(\widehat{C}X^l\theta^l)$$
⁽²⁾

where $\theta^l \in \mathbb{R}^{d \times d'}$ is a transformation matrix learned during training and $\widehat{C} = D^{-1}C \in \mathbb{R}^{k \times k}$, $D^{ii} = \sum_{ij} C$ is a diagonal matrix and σ denotes a non-linear operation, which is LeakyReLU [29] for our purpose. The induced representations of the attributes from the final convolutional layer of GCN, denoted as $Z \in \mathbb{R}^{k \times d'}$ are the condition at the generator side as a cue to synthesise new synthetic images. Graph convolutions enable combining and averaging information from first-order neighbourhood [24,14] to generate higher-order representations of the attributes. From Eqn. 2, we can see that the node representations at layer l + 1 is induced by adding and averaging the representations of a node itself and its neighbours from layer l. The sharing of information from the neighbouring nodes are controlled by the co-occurrence matrix.

Generator. The higher-order representations of the target attributes induced by graph convolution operations are fed into the generator along with the input image. Recent study [26] has shown that the difference of target and source one-hot vector of attributes helps to generate synthetic images with higher rate of target attributes preservation in comparison to the standard target onehot vectors [9,16]. We propose to feed the generator with the graph induced representations of attributes scaled by the difference of target, t and source, sattributes one-hot vectors as: $Z_t = Z \odot (t - s)$, where, $Z_t \in \mathbb{R}^{k \times d'}$ is a matrix containing the final representations of the target attributes which we feed to the generator as shown in Fig. 3. Given an input image x and the matrix containing continuous representations of target attributes Z_t , we learn the parameters, θ_q of the generator G to generate a new image \hat{x} in an adversarial manner. The generator usually consists of both encoder and decoder or only decoder with few convolutional layers. Also conditions are either concatenated with image at encoder side or concatenated with the image latent representations on the input of decoder side. As we mentioned, our approach is agnostic to the architectures of GANs. Hence, the induced representations from our approach can be fed into the encoder [9] or decoder [16,26] of the generator. In Fig. 3 we present a diagram where target attributes conditioning representations are fed from the decoder part of the generator similar to that of Attgan [16] and STGAN [26]. We flatten Z_t and concatenate it with the latent representations of the input image generated from the encoder and feed it to the decoder. In contrast, in Stargan [9] case, each of the columns in Z_t is duplicated separately and overlaid to match the dimension $128 \times 128 \times 3$ of input RGB image and concatenated with RGB channels.

$$\hat{x} = G(x, Z_t; \theta_q)$$

Loss Functions and End-to-end Learning. The overall loss for the generator is

$$\mathcal{L} = \alpha_1 \mathcal{L}_{G_{adv}} + \alpha_2 \mathcal{L}_{cls} + \alpha_3 \mathcal{L}_{rec}$$

where $\mathcal{L}_{G_{adv}}, \mathcal{L}_{cls}, \mathcal{L}_{rec}$ are the adversarial loss, the classification loss and the reconstruction loss respectively and $\alpha_1, \alpha_2, \alpha_3$ represent the hyper-parameters. We minimise the adversarial loss to make the generated image indistinguishable from the real data. The generator θ_g and discriminator $\theta_d = \{\theta_c, \theta_p, \theta_r\}$ compete to each other in an adversarial manner. Here, θ_c are the parameters of convolutional layers shared by the discriminator and attribute classifier, θ_p is of penultimate layers of the classifier, and θ_r are the parameters of the discriminator. In our case, we use WGAN-GP [1,13]:

$$\max_{[\theta_c;\theta_r]} \mathcal{L}_{D_{adv}} = \mathbb{E}_x[D(x; [\theta_c, \theta_r])] - \mathbb{E}_{\hat{x}}[D(\hat{x}; [\theta_c, \theta_r])] \\ + \lambda \mathbb{E}_{x'=\beta x + (1-\beta)\hat{x}}[(||\nabla_{x'}D(x'; [\theta_c, \theta_r])||_2 - 1)^2] \\ \min_{\theta_a} \mathcal{L}_{G_{adv}} = \mathbb{E}_{\hat{x}}[1 - D(\hat{x}; [\theta_c, \theta_r])]$$

where $\hat{x} = G(x, Z_t; \theta_g)$.

The classification loss here is the standard binary cross-entropy loss in target category label: $\mathcal{L}_{cls} = \sum_k \mathcal{L}^k([\theta_c, \theta_p])$, where θ_c and θ_p form the attribute classifier.

The reconstruction loss is computed setting the target attributes equal to that of the source. This results (t - s) into zero vector and ultimately, Z_t turns to the zero matrix.

$$\mathcal{L}_{rec} = \|x - G(x, O; \theta_g)\|_1$$

The above combined loss, \mathcal{L} trains the generator G, the discriminator D, and the graph CNN in an end-to-end fashion. The optimal attribute representations are learnt to help generate realistic images (by the discriminator loss), and preserve target attributes (by the classification loss). In the process, multiattribute relations are also embedded to the representations. The networks would consider more natural i.e. realistic when the output image has the presence of associated other attributes as well as the target attribute.

3.3 Online Multitask Learning for Discriminator

While minimising the target attribute classification loss on the discriminator side, we propose to share weights between the co-occurring attributes model parameters. We adapted online multitask learning for training multiple linear classifiers [4] to achieve this. The rate of the weights shared between the model parameters of attributes is constrained by the attributes interaction matrix. We derive the interaction matrix from the co-occurrence matrix C.

As before, θ_c are the parameters of convolutional layers and θ_p is of the penultimate layers of the classifier. We minimize the objective given in Eqn. 3 for target attribute classification with respect to discriminator. The first term in Eqn. 3 is the standard binary cross entropy loss. The second term is a regularizer which enforces to maintain similar model parameters of frequently co-occurring attributes by sharing the weights. During training, if prediction for any attribute k is wrong, we update not only the parameters of the particular attribute but also the parameters of related attributes. Rate is determined by the co-relation defined on \hat{C} . For more details on regularizer, we suggest to refer the original paper [4].

$$\mathcal{L}_{cls}^{real} = \sum_{k} \mathcal{L}^{k}([\theta_{c}, \theta_{p}]) + \lambda \mathcal{R}(\theta_{p}, \widehat{C})$$
(3)

Note that the multi-task loss is computed on real data. Such updates induce similar model parameters of the attributes which are frequently co-occurring as defined in the co-occurrence matrix.

The multitask attribute classification loss in the above is exploited instead of the conventional single task loss without sharing the parameters. The sharing of parameters between the tasks has advantages over conventional methods: it enforces the discriminator to remember the semantic structural relationship between the attributes defined on the co-occurrence matrix. Such kind of constrains on the discriminator also helps to minimize the risk of forgetting [5] and also retain associated attributes [28]. We can also draw analogy between our method and Label Smoothing. Our difference from one-sided Label Smoothing [41] is randomly softening the labels while our approach is constrained with the meaningful co-occurrence matrix and regularises the parameters of the attributes by sharing the weights. We train the whole system i.e. G, D, and the graph in end-to-end as in the previous section, by replacing the binary classification loss with the multitask loss. Conditioning at the discriminator side also helps improve the generator.

4 Experiments

Data Sets and Evaluation Metrics: To evaluate the proposed method we carried out our major experiments on **CelebA** which has around 200K images annotated with 40 different attributes. In our experiments, we took 13 attributes similar to that of [26] on face attribute editing. Similarly, **LFWA** is another benchmark. This data set contains around 13,233 images and each image is annotated with the same 40 different attributes as CelebA. We took 12K images to train the model and report the performance on the remaining examples. Finally, we use **RaFD** data set annotated with the expressions to do attributes transfer from **CelebA**. This data set consists of 4,824 images annotated with 8 different facial expressions.

For quantitative evaluations, we employed **Target Attributes Recognition Rate (TARR), PSNR, SSIM** which are commonly used quantitative metrics for conditional GANs [9,26,16,43]. For cGANs, it is not sufficient just to have synthetic realistic images, these being recognisable as the target class is also highly important [43]. Thus we choose to compute TARR similar to that of existing works [9,26,16]. TARR measures the generation of conditioned attributes on synthetic data by a model trained on real data. We took a publicly available pre-trained attribute prediction network [26] with a mean accuracy of 94.2% on 13 different attributes on test set of CelebA. Similarly, we employ PSNR (Peak Signal to Noise Ratio) and SSIM (Structural Similarity) to assess the quality of the synthetic examples.

Compared Methods: To validate our idea, we compare the performance of our GCN induced representations (*gcn-reprs*) with wide ranges of both categorical and continuous types of target attributes encoding methods.

• One-hot vector: As mentioned from the beginning, this is the most commonly and widely used conditioning technique for cGANs [9,16,25,26]. Here, the presence and absence of a *target attribute* (t) is encoded by i.e. 1 and 0 respectively.

• Latent Representations (latent-reprs): [10] proposed to represent presence/absence of a target expression by a positive/negative d-dimensional normal random vector for expression manipulation.

• *Word2Vec:* Words embedding [30] to encode target domain information are successfully applied to synthesise image from text [40]. We represented target attributes by the embedding of the attributes labels.

• *Co-occurrence:* We use co-occurrence vectors as representations of target label attributes to obtain an approximate performance comparison to [28]. As, [28] rules were hard coded which is not feasible in our arbitrary attributes manipulation case.

• Attrbs-weights: We use attribute model parameters obtained from [26] to repre-

10 B. Bhattarai and T-K Kim

sent the target attributes.

As [26] demonstrated the effectiveness of conditioning the difference of target and source attributes one-hot vector (t-s) compared to target attributes one-hot vector (t) alone, we mostly perform our experiments on the difference set up. We call conditioning difference of target and source as Difference (**Diff**) mode and conditioning target only as Standard (**Std**) mode. We employed several *types* of attribute encoding on both the *modes* and report their performances on multiple state-of-the-art GAN architectures (*GAN Archs*): Stargan [9], Attgan [16], STGAN [26], Stargan-JNT [9]. In addition to these conditioning on the generator side, we also proposed to apply Multi-task Learning (**MTL**) on the discriminator side.

Implementation Details: We initialise the nodes of the graph with model parameters of attributes (weight vectors) obtained from pre-trained attribute classifiers [26]. The dimensions of input and output nodes of the graph are 1024 and 128 respectively. GCN has 2 convolution layers. For all the data sets in our experiment, we pre-process and crop image to the size of $128 \times 128 \times 3$.



Fig. 4: Perf. comparison on LFWA

Quantitative Evaluations.

	Condition		
Condition Type	Std	Diff	Average
one-hot vec	√		78.6
one-hot vec		\checkmark	80.2
co-occurrence		\checkmark	78.6
word2vec		\checkmark	81.3
attrbs-weights		\checkmark	81.9
gcn-reprs		√	84.0

Table 1: Avg. TARR due to various types of attribute representations

Ablation Studies: We train Stargan [9] on both Std and Diff mode with various types of attribute encodings. We computed TARR on 5 different target attributes viz. hair color (black, blond, brown), gender (male/female), and age (young/old) on CelebA, similar to the original paper. Tab. 1 summarises the performance comparison. Among the four compared condition types (one-hot vec, co-occurrence, word2vec, attrbs-weights), attrbs-weights obtain the best performance. Thus, we chose attrb-weights to initialise nodes of GCN. Please note, nodes can be initialised with any other type of representations. Referring to the same Table, we observe GCN induced representations (gcn-reprs) out-performing all other compared methods.

Discussion: Semantic representations of attributes: word2vec and attrbs-weights outperformed one-hot vec. Co-occurrence also lagged behind the semantic representations as it has no visual information and not optimised for arbitrary attribute manipulations. Please note, [28] designed similar representation for single target attribute. As we know, word2vec are learned from the large corpus and bears syntactic and semantic relationships between the words [30]. These are also being useful for attributes classification [2]. Attrbs-weights are equipped with higher-order visual information including the spatial location of the attributes (See Fig. 4 from [44]). Finally, gcn-reprs benefited from both semantic representations and co-occurrence relationship. Thus, it is essential that conditions hold semantically rich higher-order target attributes characteristics.

MTL at discriminator: We compare our idea to apply MTL at discriminator side. It is evident that, in MTL, interaction between the tasks is important. Thus, we extend the number of target attributes to 13 for this experiments. First two rows from the second block of Table 2 shows the performance comparison between the the baseline w/ and w/o MTL. We observe an overall increase of +1.1% in performance over the baseline. On category level, MTL on the discriminator is outperforming in 9 different attributes out of 13 attributes.

Comparison With Existing Arts: To further validate our idea, we compare our method with multiple state-of-the-art GAN architectures to date in three different quantitative measurements viz. TARR, PSNR and SSIM on three different benchmarks.

		C. 1	Mode Attributes														
GAN Arch.	Condition Type	Std	Diff	B_{ald}	B_{angs}	$Black H_{air}$	$Blonde H_{air}$	Brown Hair	B. Eyebrows	$E_{yeglasses}$	Mth. Sh. Open	M_{ustach_e}	$N_0 \ beard$	$P_{ale\ Skin}$	$M_{\rm ale}$	Y_{Oung}	Average
IcGAN [37]	one-hot vec	1		19.4	74.2	40.6	34.6	19.7	14.7	82.4	78.8	5.5	22.6	41.8	89	37.6	43.2
FaderNet[25]	one-hot vec	\checkmark		1.5	5	27	20.9	15.6	24.2	87.4	44	10	27.2	11.1	48.3	20.3	29.8
Stargan [9]	one-hot vec	√		24.4	92.3	59.4	68.9	55.7	50.1	95.7	96.1	18.8	66.6	84	77.1	83.9	67.2
+ MTL	one-hot vec	√		22.7	95.4	63	62.3	51.9	58	99.2	98.7	24	52.2	90.5	83.7	86.8	68.3
	one-hot vec		\checkmark	41.9	93.6	74.7	75.2	67.4	65.9	99	95.3	26.8	64.3	86.2	89	89.3	74.5
	latent-reprs	1		18.4	93.8	68.5	60.9	62.5	69.4	97.0	97.7	14.0	34.4	91.3	78.5	76.7	66.4
	latent-reprs		\checkmark	32.5	93.2	68.9	79.5	71.5	55.3	97.2	98.4	30.0	58.5	85.1	84.0	75.1	71.4
	attrbs- $weights$		\checkmark	32.7	96.0	74.4	77.5	74.1	66.7	98.8	32.2	78.2	90.9	81.6	98.5	86.7	76.0
	gcn-reprs		 √	28.2	99.4	76.5	77.1	70.9	74.2	99.5	99.4	37.3	89.6	92	93.4	94.9	79.4
+ MTL	gcn-reprs		\checkmark	34.4	98.4	73.3	78.6	70.8	85.5	99.5	99.1	44.2	90	92.3	95.6	91.7	81.0
+ MTL + End2End	gcn-reprs		\checkmark	56.6	98.2	76	81.1	80	73.4	99.4	99.5	50.9	89.1	94.2	92	95	83.4
STGAN [26]	one-hot vec	~		40.7	92.5	69.5	69.7	59.4	65.2	99.2	95	26.8	69.7	90.8	70	52.8	69.3
	one-hot vec		\checkmark	59.9	97.7	93	79	89.9	88.3	99.7	96.7	38.9	93.4	97.0	98.5	86.7	86.1
+ MTL + End2End	gcn-reprs		~	82.0	95.5	92.6	85.4	82.0	86.2	99.9	99.4	55.3	98.4	96.0	98.1	86.9	89.1
AttGAN [16]	one-hot vec	1		22.5	93	46.3	40.4	51	49.2	98.6	97	30.3	81.3	84.4	83.3	67.9	65
	one-hot vec		\checkmark	69.1	97.5	78.8	84.4	76.5	73.4	99.6	95.8	34.2	85.8	96.8	95.8	92.9	83
+ MTL + End2End	gcn-reprs	1	~~	72.8	97.9	93.9	94.4	92.3	86.8	99.8	98.6	48.4	97.1	97.1	98.5	96.4	90.3

Table 2: Comparison of Target Attributes Recognition Rate (TARR) on CelebA with different existing cGANs architectures with different target attribute label conditioning.

Target Attribute Recognition Rate (TARR): Table 2 compares the TARR on CelebA. In the Table, the top block shows the performance of two earlier works: IcGAN [37] and FaderNet [25]. These figures are as reported on [26]. These methods relied on one-hot vec type of the target attributes in Std mode as

12 B. Bhattarai and T-K Kim

their conditions. The TARR of these methods are modest. In the Second block, we compare the performance of *attrbs-weights*, *latent-reprs.gcn-repres* with the default conditioning type of Stargan [9] on both the conditioning modes. Simply switching the default type from Std mode to Diff mode improves the TARR from 67.2% to 74.5%. This is the best performance of Stargan reported by [26]. The encoding principle of *latent-reprs* [10] is similar to that of *one hot vec*, as positive and negative random vectors are used instead of 1 and 0. Attribute-specific information rich *attrbs-weights* out-performs these representations. We observe the performance of these representations similar to that of *one-hot vec*. We trained our GCN network to induce gcn-reprs and replaced the default conditioning on the generator of Stargan [9] by gcn-reprs, the performance improved to 79.4%(+4.9%). Another experiment with the same set-up on the generator and MTL on the discriminator improves the performance to 81.0%(+6.5%). Training GCN and GAN on multi-stage fashion make the representations sub-optimal. Thus, we train GCN and GAN and apply MTL simultaneously and attained an average accuracy of 83.4%, the highest performance reported on Stargan [9]. We also applied our method to two other best performing GAN architecture [26,16] for face attribute manipulations. Last two blocks of the Table 2 compares the performance on these architectures. Similar to that with Stargan [9], instead of existing method of conditioning target attributes and applied ours to both the generator and discriminator and train the model from scratch in an end-to-end fashion. After we applied our method on STGAN [26], which is the state-of-the-art method to date, we improve the mean average performance from 86.1% to 89.1%. Similarly, on Attgan [16] we outperformed the best reported performance by +7.3% and attain the new state-of-the-art performance on CelebA. This is +4.2%above the current state-of-the-art [26]. These results show that our method is agnostic to the cGAN architecture. We also evaluated our method on LFWA. We applied our method to Stargan and compared it to the default conditioning type, one-hot vec on Diff mode. Fig. 4 shows the TARR. From this Fig., we can see our approach outperforming the baseline. If we carefully check the performance of individual attributes on both the benchmarks (Tab. 2, Fig. 4), our method is substantially outperforming existing arts in attributes such as *bald*, *moustache*, young, male. These attributes follows law of nature and it is essential to make natural transition to better retain the target label.

PSNR/SSIM: We compute PSNR and SSIM scores between real and synthetic images and compare the performance between with the counter-parts. In Tab. 3, the two columns **Before** and **After** show the scores of GANs before and after applying our method respectively. Our approach consistently improves the performance of the counter-parts.

Qualitative Evaluations: To further validate our idea, we performed extensive qualitative analysis. We compare our method over the existing arts on two different scenarios i'e in and across data set attributes transfer.

In Data Set Attribute Transfer: In this scenario, we train a model on train set and evaluate the performance on the test set of the same data set. Fig. 5 compares the qualitative outcomes of Attgan [16] conditioned with *one-hot vec* on





Fig. 5: Qualitative comparison of Attgan [16] (default conditioned with *one-hot vector*) trained on diff mode vs Attgan trained with *gcn-reprs (ours)* trained on diff-mode on CelebA (best viewed on color).

Diff mode with *qcn-reprs* on the same mode. The left block in the figure shows the result of single target attribute manipulation whereas the right block shows that of multi-attributes manipulation. From the results, we can clearly see that our method is able to generate images with less artefacts and better contrast (see the background). In addition to this, our method is also able to manipulate multiple attributes simultaneously, whenever it is meaningful to do so, to give a natural transition from source to target. For example, for male-to-female transition, our method is able to put on lipsticks, high cheekbones, arched evebrows but the baseline fails to do so. Similarly, wrinkles on face with few remaining grey hair gives natural transition to bald instead just completely removing the hairs from head. As it is highly likely that a person gets bald in his/her old age. Turning grey hair to black hair is making the guy comparatively younger as *black hair* is associated with young attribute. Due to such unique strengths of our method, enabled by GCN on the generator and MTL on the discriminator, we observe substantial improvements over the baselines especially in the recognition of certain attributes: Young, Male, Bald, Moustache where a natural transition is essential as these are naturally occurring attributes associated with different factors.

Cross Data Set Attributes Transfer: Stargan-JNT [9] propose to train a GAN with multiple data sets having disjoint sets of attributes simultaneously to improve the quality of the cross-data set attribute transfer. We applied our conditioning method to train the network on CelebA and RaFD simultaneously.

14 B. Bhattarai and T-K Kim



And, we compare the performance with their default conditioning method which

Fig. 6: Qualitative comparison of Stargan-JNT with default condition vs Stargan-JNT conditioned with *gcn-reprs (ours)* on RaFD (best viewed on color).

is one-hot vec. Fig. 6 shows a few test examples from RaFD and their synthetic images when target attributes are from CelebA. Please note attribute annotations such as *Black Hair, Blonde Hair, Pale Skin* are absent on RaFD train set. From the same Fig., we can clearly see that the synthetic images generated by our method are with less artefacts, better contrast and better preservation of the target attributes.

5 Conclusions

We propose a Graph Convolutional Network enabled novel method to induce target attributes embeddings for Conditional GANs on the Generator part. Similarly, we proposed a MTL based structural regularisation mechanism on the discriminator of the GAN. For both of these, we exploit the co-occurrences of the attributes. Finally, we propose a framework to learn them in an end-to-end fashion. We applied our method on multiple existing target label conditioned GANs and evaluated on multiple benchmarks for face attribute manipulations. From our both extensive quantitative and qualitative evaluations, we observed a substantial improvement over the existing arts, attaining new state-of-the-art performance. As a future work, we plan to design a framework to dynamically adjust the co-occurrence distribution of the attributes to synthesize naturally realistic attributes manipulated images.

Acknowledgements

Authors would like to thank EPSRC Programme Grant 'FACER2VM' (EP/N007743/1) for generous support. We would also like to thank Prateek Manocha, undergraduate student from IIT Guwahati for some of the baseline experiments during his summer internship at Imperial College London.

References

- Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein gan. arXiv preprint arXiv:1701.07875 (2017)
- 2. Bhattarai, B., Bodur, R., Kim, T.K.: Auglabel: Exploiting word representations to augment labels for face attribute classification. In: ICASSP (2020)
- 3. Cao, J., Huang, H., Li, Y., Liu, J., He, R., Sun, Z.: Biphasic learning of gans for high-resolution image-to-image translation. CVPR (2019)
- 4. Cavallanti, G., Cesa-Bianchi, N., Gentile, C.: Linear algorithms for online multitask classification. JMLR (2010)
- Chen, T., Zhai, X., Ritter, M., Lucic, M., Houlsby, N.: Self-supervised gans via auxiliary rotation loss. In: CVPR (2019)
- Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In: NIPS (2016)
- Chen, Y.C., Lin, H., Shu, M., Li, R., Tao, X., Shen, X., Ye, Y., Jia, J.: Facelet-bank for fast portrait manipulation. In: CVPR (2018)
- Chen, Z.M., Wei, X.S., Wang, P., Guo, Y.: Multi-label image recognition with graph convolutional networks. In: CVPR (2019)
- Choi, Y., Choi, M., Kim, M., Ha, J.W., Kim, S., Choo, J.: StarGAN: Unified generative adversarial networks for multi-domain image-to-image translation. In: CVPR (2018)
- Ding, H., Sricharan, K., Chellappa, R.: Exprgan: Facial expression editing with controllable expression intensity. In: AAAI (2018)
- Gecer, B., Bhattarai, B., Kittler, J., Kim, T.K.: Semi-supervised adversarial learning to generate photorealistic face images of new identities from 3d morphable model. In: ECCV (2018)
- Grover, A., Leskovec, J.: node2vec: Scalable feature learning for networks. In: SIGKDD (2016)
- 13. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., Courville, A.C.: Improved training of wasserstein gans. In: NIPS (2017)
- Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. In: NIPS (2017)
- Hamilton, W.L., Ying, R., Leskovec, J.: Representation learning on graphs: Methods and applications. arXiv preprint arXiv:1709.05584 (2017)
- He, Z., Zuo, W., Kan, M., Shan, S., Chen, X.: Attgan: Facial attribute editing by only changing what you want. IEEE TIP (2019)
- 17. Hong, S., Yang, D., Choi, J., Lee, H.: Inferring semantic layout for hierarchical text-to-image synthesis. In: CVPR (2018)
- Huang, X., Liu, M.Y., Belongie, S., Kautz, J.: Multimodal unsupervised image-toimage translation. In: ECCV (2018)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: CVPR (2017)
- 20. Kaneko, T., Hiramatsu, K., Kashino, K.: Generative attribute controller with conditional filtered generative adversarial networks. In: CVPR (2017)
- 21. Kaneko, T., Hiramatsu, K., Kashino, K.: Generative adversarial image synthesis with decision tree latent controller. In: CVPR (2018)
- 22. Karras, T., Aila, T., Laine, S., Lehtinen, J.: Progressive growing of gans for improved quality, stability, and variation. ICLR (2018)

- 16 B. Bhattarai and T-K Kim
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: CVPR (2019)
- 24. Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. In: ICLR (2016)
- Lample, G., Zeghidour, N., Usunier, N., Bordes, A., Denoyer, L., et al.: Fader networks: Manipulating images by sliding attributes. In: NIPS (2017)
- Liu, M., Ding, Y., Xia, M., Liu, X., Ding, E., Zuo, W., Wen, S.: Stgan: A unified selective transfer network for arbitrary image attribute editing. In: CVPR (2019)
- Liu, M.Y., Breuel, T., Kautz, J.: Unsupervised image-to-image translation networks. In: NeurIPS (2017)
- Liu, Y., Li, Q., Sun, Z.: Attribute-aware face aging with wavelet-based generative adversarial networks. In: CVPR (2019)
- Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. In: ICML (2013)
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G.S., Dean, J.: Distributed representations of words and phrases and their compositionality. In: NurIPS (2013)
- Mirza, M., Osindero, S.: Conditional generative adversarial nets. arXiv preprint arXiv:1411.1784 (2014)
- Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: ICLR (2018)
- 33. Miyato, T., Koyama, M.: cgans with projection discriminator. In: ICLR (2018)
- Nian, F., Chen, X., Yang, S., Lv, G.: Facial attribute recognition with feature decoupling and graph convolutional networks. IEEE Access (2019)
- Odena, A., Buckman, J., Olsson, C., Brown, T.B., Olah, C., Raffel, C., Goodfellow, I.: Is generator conditioning causally related to gan performance? In: ICML (2018)
- Odena, A., Olah, C., Shlens, J.: Conditional image synthesis with auxiliary classifier gans. In: ICML (2017)
- Perarnau, G., Van De Weijer, J., Raducanu, B., Alvarez, J.M.: Invertible conditional gans for image editing. In: NIPSW (2016)
- Perozzi, B., Al-Rfou, R., Skiena, S.: Deepwalk: Online learning of social representations. In: SIGKDD (2014)
- Pumarola, A., Agudo, A., Martinez, A.M., Sanfeliu, A., Moreno-Noguer, F.: Ganimation: Anatomically-aware facial animation from a single image. In: ECCV (2018)
- 40. Reed, S., Akata, Z., Yan, X., Logeswaran, L., Schiele, B., Lee, H.: Generative adversarial text to image synthesis (2016)
- 41. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. In: NIPS (2016)
- Shen, W., Liu, R.: Learning residual images for face attribute manipulation. In: CVPR (2017)
- 43. Shmelkov, K., Schmid, C., Alahari, K.: How good is my gan? In: ECCV (2018)
- 44. Taherkhani, F., Nasrabadi, N.M., Dawson, J.: A deep face identification network enhanced by facial attributes prediction. In: CVPRW (2018)
- 45. Tang, J., Qu, M., Wang, M., Zhang, M., Yan, J., Mei, Q.: Line: Large-scale information network embedding. In: WWW (2015)
- 46. Xiao, T., Hong, J., Ma, J.: Elegant: Exchanging latent encodings with gan for transferring multiple face attributes. In: ECCV (2018)
- 47. Zakharov, E., Shysheya, A., Burkov, E., Lempitsky, V.: Few-shot adversarial learning of realistic neural talking head models. arXiv preprint arXiv:1905.08233 (2019)
- 48. Zhang, G., Kan, M., Shan, S., Chen, X.: Generative adversarial network with spatial attention for face attribute editing. In: ECCV (2018)

- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: CVPR (2017)
- 50. Zhang, Z., Song, Y., Qi, H.: Age progression/regression by conditional adversarial autoencoder. In: CVPR (2017)
- Zhao, B., Meng, L., Yin, W., Sigal, L.: Image generation from layout. In: CVPR (2019)
- 52. Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Sun, M.: Graph neural networks: A review of methods and applications. arXiv:1812.08434 (2018)