

# Graph2Pix: A Graph-Based Image to Image Translation Framework

Dilara Gokay<sup>1,†</sup> Enis Simsar<sup>1,2,†</sup> Efehan Atici<sup>2</sup>  
 Alper Ahmetoglu<sup>2</sup> Atif Emre Yuksel<sup>2</sup> Pinar Yanardag<sup>2</sup>  
<sup>1</sup>Technical University of Munich <sup>2</sup>Bogazici University

dilara.goekay@tum.de enis.simsar@tum.de efehan.atici@boun.edu.tr  
 alper.ahmetoglu@boun.edu.tr atif.yuksel@boun.edu.tr yanardag.pinar@gmail.com

## Abstract

In this paper, we propose a graph-based image-to-image translation framework for generating images. We use rich data collected from the popular creativity platform Artbreeder<sup>1</sup>, where users interpolate multiple GAN-generated images to create artworks. This unique approach of creating new images leads to a tree-like structure where one can track historical data about the creation of a particular image. Inspired by this structure, we propose a novel graph-to-image translation model called Graph2Pix, which takes a graph and corresponding images as input and generates a single image as output. Our experiments show that Graph2Pix is able to outperform several image-to-image translation frameworks on benchmark metrics, including LPIPS (with a 25% improvement) and human perception studies ( $n = 60$ ), where users preferred the images generated by our method 81.5% of the time. Our source code and dataset are publicly available at <https://github.com/catlab-team/graph2pix>.

## 1. Introduction

Recent breakthroughs in deep learning show that artificial intelligence (AI) can not only be creative, but can also help humans be more creative. Tasks that require human creativity are being taken over by generative approaches as their capabilities increase. Music generation [7], face generation [16], domain translation [43], and image manipulation [38] are just a few of the many creative applications of generative models.

The success of generative methods has also sparked several public tools where researchers and artists express their creativity through interaction with pre-trained models. For example, the GauGAN platform is offered as an interactive tool based on an image-to-image translation framework [30]



Figure 1: A sample lineage data up to 2 levels<sup>3</sup>. The creators of the images are annotated with labels.

where users design their own lifelike landscape images with properties such as rivers, rocks, grasses, and clouds. Websites like *This Person Does Not Exist*<sup>2</sup> use pre-trained GAN models like StyleGAN [17] and provide users with uniquely generated faces.

One of the most popular GAN-based creativity platforms is Artbreeder. The platform’s easy-to-use interface has attracted thousands of users, allowing them to create over 70 million GAN-based images. Artbreeder helps users create new images using BigGAN [4] or StyleGAN-based models where users can adjust parameters or mix different images to create new images. Users can *breed* new images from a single one by editing *genes* such as *age*, *gender*, *ethnicity* or create new ones by *crossbreeding* multiple images together. The unique capability of the crossbreeding function allows the generated images to contain *ancestry* data where the ancestors of the generated image can be traced in a tree-like structure (see Figure 1).

<sup>2</sup><https://www.thispersondoesnotexist.com>

<sup>3</sup>Access the full tree via <https://www.artbreeder.com/lineage?k=1fcdf872ec11c80e955bb5c1>.

<sup>†</sup>Equal contribution. Author ordering determined by a coin flip.

<sup>1</sup><http://artbreeder.com>

The lineage structure of images generated with Artbreeder opens up a wide range of possible applications, but also presents unique challenges. The lineage information provides a tree-like structure in which one can trace the parents, grandparents, and further ancestors of a given image. However, it is not entirely clear how such a structure can be used in GAN-based models. For example, how can we generate a *child* image based on a list of its *ancestors*? One could try to use image-to-image translation methods such as Pix2Pix [14] or CycleGAN [43] and use any *ancestor* to generate the *child* image (e.g., feeding  $Parent_1$  and generating  $Child_1$  in Figure 1). However, this approach results in a significant loss of information since only one ancestor can be used (e.g.,  $Parent_2$  or the grandparents are considered).

In this paper, we propose a novel image-to-image translation method that takes multiple images and their corresponding *lineage* structure and generates the target image as output. To propose a general solution, our framework takes a graph structure, considering the tree structure of lineage data as a sub-case. To the best of our knowledge, this is the first image-to-image translation method that uses a graph-based structure.

The remainder of the paper is organised as follows. Section 2 discusses related work. Section 3 describes our novel graph-based image-to-image translation framework. Section 4 provides details about Artbreeder platform and the collected data, and presents the qualitative and quantitative experiments to demonstrate the effectiveness of our model. Section 5 discusses the limitations and failures of our model. Section 6 concludes the paper.

## 2. Related Work

**Image-to-Image Translation Models** Generative Adversarial Networks (GANs) are two-part networks that model the real world to the generative space using deep learning methods [9]. The generative part of the network tries to create images similar to the dataset while the adversarial part tries to detect if the created image is coming from the training dataset or if it is a generated one. The main goal of GANs is to try to model the image space in a way that makes generated images indistinguishable from the ones in the dataset. Generative networks can generate new realistic images from random noise vectors called latent vectors via internal mapping.

Image-to-image translation models are usually designed as conditional GAN models, where the goal is to transform one image into the style of another image, retaining some features of the first image and applying others from the second image. They can be used in various applications such as transforming maps into aerial photographs, transforming real images into artist drawings, removing or modifying objects, and image quality enhancement. The first approaches to image-to-image translation problem date

back to 2001 [11], where they used more traditional ways like implementing separate, task-specific programs. However, the problem of having domain-specific solutions to a general problem has always been a limiting factor for researchers. The groundbreaking change from task-specific machines to a general-purpose algorithm was made with the discovery of Pix2Pix [14]. This generalized method for automatically generating lifelike images automatically with setting relevant loss functions for models paved the way for other implementations such as GauGAN [30] and CycleGAN [43]. Later, researchers found an impressive solution to image-to-image translation tasks, using a method called U-GAT-IT [19]. Their approach is to develop a new method for unsupervised learning with attention networks that has normalization layers. They were able to get a neural network to creatively transform images in different styles, such as portraits into anime style, or horses into zebra images. UNIT [23] is another unsupervised image-to-image translation network and is built upon the a shared-latent space assumption. This means that UNIT can transform an image from one domain into another and does not require having the corresponding images during training. CC-FPSE [25] is tailored for semantic image synthesis. They propose a conditional convolution to efficiently control the generation. TSIT [15] uses a symmetrical two-stream network to decouple the semantic structure and style information. Recently, pixel2style2pixel [31] has been shown to perform well in tasks such as multi-modal conditional image synthesis, facial frontalization, and super-resolution.

**Graph Convolutional Networks** Graph convolutional networks (GCNs) were introduced by Kipf and Welling [20] and graph structure was used to gain better insight into the hidden representation of data. They showed that GCN usefully incorporates both global (graph level) and local (node level) graph structures for semi-supervised classification. GCN has achieved state-of-the-art results in several application areas, including social analysis [2] and molecular structural analysis [6]. GCNs are also being explored for computer vision [8]. Yao *et al.* [39] proposes a model that combines GCNs with an LSTM architecture to integrate semantic and spatial object relations into an image encoder for image captioning task. Mittal *et al.* [29] proposes a method for generating an image based on a sequence of scene graphs and uses a GCN to handle variable size scene graphs along with generative adversarial image translation networks. Zeng *et al.* [40] uses GCN to represent proposal-proposal relationships and use it for action classification and localization. DGCNet [41] models two graphs, one representing the relationships between pixels in the image and the other dependency along the channels of the feature map of the network.

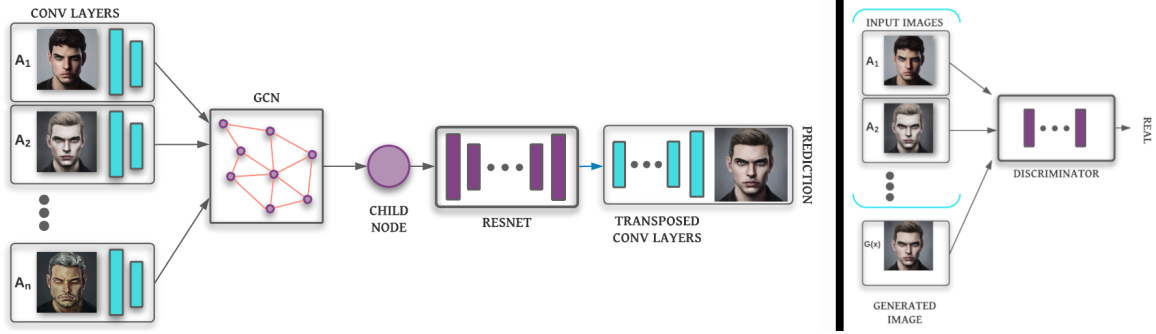


Figure 2: An illustration of Graph2Pix. During the generation process (shown on the left) our 2-layer GCN module takes  $A_1 \dots A_n$ , the ancestors of an image as an input, and generates the prediction. The discriminator (shown on the right) takes the concatenation of the input images  $A_1 \dots A_n$  and the generated image  $G(x)$  and creates a prediction. Note that we placed GCN after the convolutional layer to use it more efficiently.

### 3. Methodology

In this section, we first describe a baseline image-to-image translation framework and then explain details of the proposed Graph2Pix model.

#### 3.1. Pix2Pix and Pix2PixHD

Pix2Pix is an image-to-image translation network based on the conditional GAN structure. Its generator  $G$  uses a network structure based on U-Net [32] while its discriminator  $D$  uses a convolutional classifier called Patch-GAN [22]. It takes a pair of images  $\{(y_i, x_i)\}$  as input (e.g., a label map and the corresponding image) and computes the conditional distribution of the images via  $\min_G \max_D \mathcal{L}_{cGAN}(G, D)$ .

Pix2PixHD [37] is built upon the Pix2Pix structure to generate higher resolution images. Compared to Pix2Pix, it divides the generator network  $G$  into two separate parts: the global generator network  $G_1$ , and local enhancer network  $G_2$ . Although they are similar in concept,  $G_1$  is trained on the down-sampled versions of the training images while  $G_2$  acts as a wrapper around  $G_1$  to increase the image resolution of its low-resolution output. This network design supports scalability. To generate higher resolution images, additional generator networks  $G_3 \dots G_n$  can be added. Pix2PixHD also improves the discriminator component with a multi-scale approach where each image scale level has a separate discriminator,  $D_1 \dots D_3$ , respectively. As the images are progressively down-sampled, the discriminator is able to capture the global consistency between real and fake images.

Our proposed model adopts the architecture of Pix2PixHD as a backbone and extends it with a multi-generator that fuses images using a graph convolutional network, which we introduce below.

#### 3.2. Graph2Pix

Pix2PixHD is not suitable for our task because it is limited to a single source and cannot accommodate multiple sources. Therefore, it does not take into account the lineage structure of the input since it cannot use the ancestral information such as *parents* or *grandparents* of the *child* image (see Figure 1). In contrast, a GCN-based approach offers the potential to capture the relationship among the members of a family, which can be represented as a graph.

Graph-based neural networks use a similarity graph that allows the information propagation between similar inputs. In this work, we use a GCN [20] that allows CNNs to work directly with graphs, where each layer in GCNs updates a vector representation of each node in the graph by considering the representations of its neighbors. The forward computation of a graph convolutional layer is as follows:

$$H^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} W^{(l)} H^{(l)}) \quad (1)$$

where  $\tilde{A}$  is the adjacency graph including self-loops,  $\tilde{D}$  is a diagonal matrix that contains the degree of node  $i$  in  $\tilde{D}_{ii}$ , and  $\sigma$  is a non-linear activation function. Here,  $W^{(l)}$  is a trainable weight matrix in layer  $l$ , and  $H^{(l)}$  represents the input activations in layer  $l$ .  $\hat{A} = \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}}$  can be viewed as a normalized adjacency graph. A graph convolutional layer differs from a conventional fully connected layer in the sense that it combines information from different nodes in each layer. This operation turns out to be a first-order approximation of spectral convolutions, which is why the method is called graph convolutional networks [20].

Our proposed global generator consists of a convolutional front-end, a set of residual blocks, and a transposed convolutional back-end (see Figure 2). Graph convolutional layers are inserted between the convolutional block and the residual block. The convolutional layers create low-resolution representations of the images in the graph. These low-resolution representations are inputs to the graph con-

volutional layers. Since we are working in the image domain, we replace the affine transformations in GCN with  $3 \times 3$  convolutions. Then, the graph convolutional layer computes:

$$H^{(l+1)} = \sigma(\tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} \text{Conv}(H^{(l)})) \quad (2)$$

Since we are trying to generate the child image based on its ancestors, we cannot provide the child image as input to the pipeline. However, we want to retain the adjacency information, i.e., the graph, to generate the output with respect to the ancestors. To achieve this, we replace the child image with noise, use the noise as input instead, and predict its original version. This procedure is analogous to denoising autoencoders [36], where we randomly noise the input features. In the end, the network should faithfully reconstruct the child image based on its ancestors. We also noise random nodes in the graph, instead of the child node, to force the model to learn relationships in the graph. For example, if we noise a parent node in the graph, the model should predict that node by combining information from the grandparent node and the child node. Since the global generator exploits the graph structure of the lineage, it is important to have a sufficiently large lineage. Therefore, we only took the children that have at least two parents and at least one grandparent from those two different parents. Thus, the lineages we used had at least four ancestors.

Similar to the global generator, Pix2PixHD’s multi-scale discriminator is not tailored to process multiple source images as it only considers a single input image. Instead, we used the concatenation of all input images as input to the multi-scale discriminator (see Figure 2). In the next section, we show that our graph-based approach performs better than the state-of-the-art image-to-image translation methods both quantitatively and qualitatively.

## 4. Experiments

In this section, we first describe our dataset and present our experimental setup. We then present the results of our quantitative and qualitative experiments.

### 4.1. Dataset

Artbreeder platform offers two ways to generate images: *mutation* and *crossbreed*. The *mutation* option allows users to create a new image by editing various traits, called *genes*, and allows users to directly modify the original images. Users can change high-level features such as *age* and *gender*, as well as low-level features such as the *red*, *green*, and *blue* channels of the image. An alternative method is called *crossbreeding*, which allows users to mix multiple images. In this paper, we are only interested in images created by *crossbreeding*. We collected the images from Art-

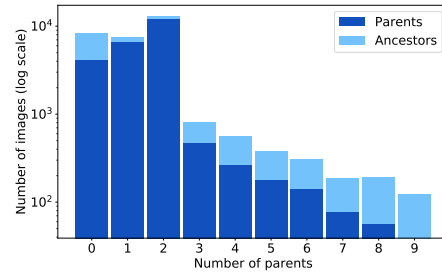


Figure 3: The distribution of parents and ancestors up to 10 is shown. As can be seen from the figures, the majority of the images have parents and ancestors between 1-4.

breeder<sup>4</sup> periodically by crawling the *Trending* section of the ‘Portraits’ category daily over a four-week period. The resolution of the images offered by Artbreeder varies from  $512 \times 512$  to  $1024 \times 1024$ . To ensure that we have a consistent dataset, we have resized all images to  $512 \times 512$ . For each image, we collect the ‘lineage’ information that allows us to track the chronological creation of the image. The lineage information includes the parent image, the creation date, and the creator’s name. We collected a total of 101,612 images. Figure 3 illustrates the distribution of ancestors and parents of a child node. Although the number of parents in the figures reaches up to 10 for clarity, the majority of child images have an average of 20.7 ancestors and 1.6 first-degree parents.

**Data exploration.** To gain insight into the creative crowd behind Artbreeder, we visually explore the latent space of images and users. We obtain the latent vectors provided by Artbreeder and project each vector as a point in 2-dimensional space. To achieve this, we first reduced the dimensionality of the latent vectors from 9126 dimensions to 50 by performing principal component analysis (PCA) [27]. After reducing the dimensions of the latent vectors, we apply t-Distributed Stochastic Neighbor Embedding [35] for two-dimensional representation, preserving the proximity in the high-dimensional space.

t-SNE is an algorithm based on Stochastic Neighbor Embedding [13] that aims to visualize high-dimensional data by embedding it in a low-dimensional space. It is tailored to model each data point through a two- or three-dimensional representation such that similar entries are closer together in the final visualization, while different entries are more likely to be farther apart.

To explore the relationship between users in latent space, we used t-SNE to visualize images. We first extracted embeddings learned from a deep neural network [33, 1] and

<sup>4</sup>Note that the images on Artbreeder are public domain (CC0) and we have obtained permission from the owner to crawl the images.

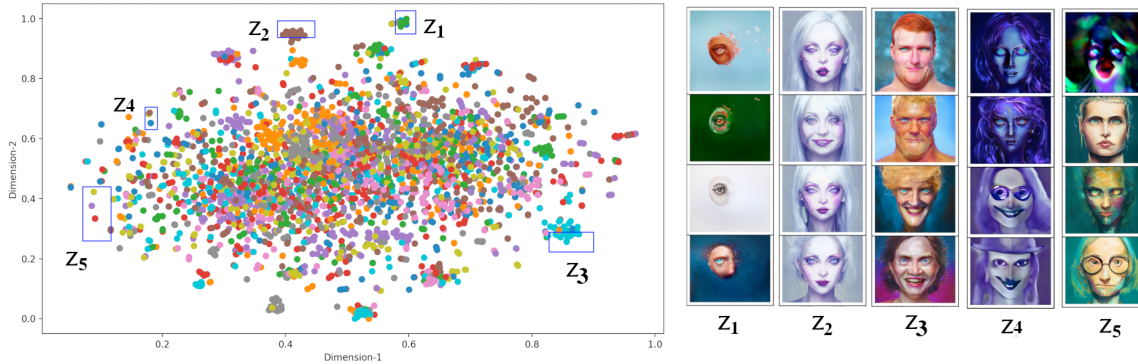


Figure 4: Exploration of the generated images based on user groups. The images are mapped into the 2-dimensional space while maintaining their proximity through t-SNE [35]. The color of the points denotes the user who created the image. The clusters labeled with  $Z_1 \dots Z_5$  on the left show selected image groups whose corresponding images are shown on the right.

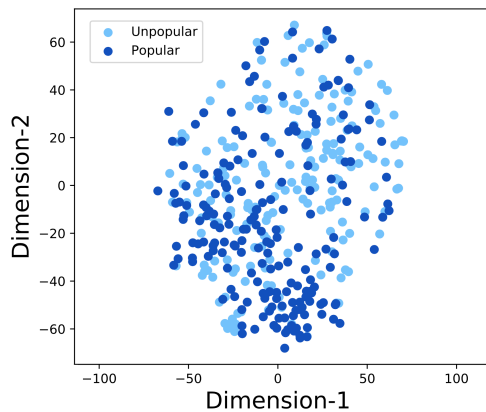


Figure 5: A comparison of popular and unpopular images in 2-dimensions using the t-SNE approach.

used Euclidean distance to compute the distance between the PCA-reduced embedding vectors using the Barnes-Hut algorithm<sup>5</sup>. Figure 4 shows a set of selected image groups with their corresponding visual content, where the colors represent different users. For clarity, only a subset of users who created between 50 and 100 images are shown. While the center of the figure shows that similar images can be created by different users, it can also be observed that certain users have a certain style that distinguishes them from others (such as  $Z_1$ ,  $Z_2$  and  $Z_3$ ). On the other hand, different users with similar styles can also be observed in different clusters, such as  $Z_4$  and  $Z_5$ .

The Artbreeder platform has a *Trending* page that showcases the images currently most used by users. To understand the relationship between popular and unpopular images, we selected the top 200 users whose images are used the most, and 200 unpopular users whose images are used

<sup>5</sup>Learning rate at 200, number of iterations at 1000

the least. Similar to our previous experiment, we used the t-SNE approach and projected the users onto a two-dimensional space after reducing their embedding vectors using PCA. Figure 5 illustrates popular and unpopular images. As can be seen from the figure, unpopular images are scattered over a large area, while popular images seem to be concentrated in a specific region (lower left and bottom), suggesting that they share similar features that appeal to a large number of users.

**Experimental setup.** We train our model with 86,373 training images and evaluate it with a test set of 15,239 images. We used 8 as the batch size to generate images with a resolution of  $256 \times 256$ . While our method requires the ancestor images up to two levels and the corresponding graph structure as input, the competing methods are limited to a single input (e.g., one of the first level parent images). However, since the results can be sensitive to which parent image is used as input, we run all the competing methods twice, once for each parent, and report the average results. For our experiments, we used the PyTorch platform, and each model was trained using its official implementation<sup>7</sup> with the default parameters. All methods are trained for 50 epochs on a Titan RTX GPU, and U-GAT-IT method is run for 100K iterations since it has no epoch option.

**Baselines.** We compared our performance against several state-of-the-art image-to-image translation frameworks, including Pix2PixHD [37], Pix2Pix [14], CycleGAN [43], MS-GAN [26], and U-GAT-IT [18]. Pix2PixHD and Pix2Pix are supervised models that require paired data, as

<sup>6</sup><https://www.artbreeder.com/lineage?k=f8c31131db9a73ac0425>

<sup>7</sup><https://github.com/junyanz/pytorch-CycleGAN-and-pix2pix>, <https://github.com/NVIDIA/pix2pixHD>




Figure 6: A qualitative comparison of our method with the top performing competitors, Pix2PixHD and U-GAT-IT. First-level ancestors ( $Parent_1$  and  $Parent_2$ ) are shown on the left where the ground truth image is denoted with  $Child$ . As can be seen from the results, our method is able to incorporate the ancestor information when generating the images whereas Pix2PixHD and U-GAT-IT are limited to a single ancestor image ( $Parent_1$ ). Due to lack of space, only first-degree parents are shown, however a sample lineage for the bottom-left image can be seen on Artbreeder.<sup>6</sup>

explained in Section 3.1. CycleGAN is one of the most popular image-to-image translation methods and uses a cycle consistency loss to enable training with unpaired data, where mapping an image from the source domain to the target domain and repeating the process in reverse leads the same starting image. However, as we will show later in our experiments, while CycleGAN generates high quality images, the cycle consistency constraint enforces a strong link between domains and does not learn to transform the input parent image into the output child image, but learns to output the parent image itself. A similar trend can be observed for another unsupervised model, U-GAT-IT, which includes an attention module to focus on more important regions that distinguish between source and target domains. We also compare with MS-GAN, which uses a mode seeking regularization term that explicitly maximizes the ratio of the distance between the generated images and the corresponding latent codes. We used the Pix2Pix-Mode Seeking option of MS-GAN.

Note that some of the popular image-to-image translation models are not applicable to our experiments because they specifically require a label map that does not exist in our dataset, such as SPADE [30], or they require the dataset to contain multiple categories during training, such as FUNIT [24].

It should also be noted that the child images on Artbreeder are not simply an average of two parent images: These images are interpolated according to a specific ratio chosen by the user. Therefore, one might consider estimating the synthesis parameters of Artbreeder directly, rather than developing a new method, as a more relevant baseline. Although this would have been an interesting comparison, we would like to emphasise that Artbreeder platform does not keep the synthesis parameters on its servers. Therefore, this information is not available for our experiments. Thus, using a GCN model becomes useful as it implicitly learn how the ancestors are mixed without explicit ratio data.

## 4.2. Quantitative Evaluation

For our quantitative experiments, we used Learned Perceptual Image Patch Similarity (LPIPS) [42], Fréchet Inception Distance (FID) [12], and Kernel Inception Distance (KID) [3]. LPIPS uses features from an ImageNet pre-trained AlexNet and measures the diversity of the generated images using the L1 distance. FID is another method for computing image quality and compares the generated samples with the real images in the dataset using Inception-V3 model [34]. KID is a more recent metric that is similar to FID but is calculated by finding the Maximum Mean Discrepancy [10] between Inception activation layers. One of the advantages of KID compared to FID is that KID compares the skewness of the distributions as well as the mean and variance, and has a faster and unbiased estimator.

Among these metrics, we argue that LPIPS is better suited to evaluate the performance of our system because it computes the score for corresponding image pairs. In other words, it considers the relationship between the ground truth and the generated image, while FID and KID scores are computed using mean and variance statistics on the entire set of images instead of pairwise comparison. The comparison is shown in Table 1. Our method significantly improves the LPIPS metric (lower score is better) compared to its competitors. Our method achieves an LPIPS score of 0.25, while the second best score is achieved by Pix2PixHD with 0.32. This result shows that our method is able to improve the LPIPS metric by almost 25%. The second metric we use to compare different methods is the FID score, where lower is better. The lowest FID score is obtained by our model, albeit with a slight improvement. This result indicates that both our method and its closest competitors, CycleGAN and U-GAT-IT, are able to generate realistic results. However, as we will show later in our qualitative experiments (Section 4.3), both methods fail to generate images that resemble ground truth. Instead, they only learn to

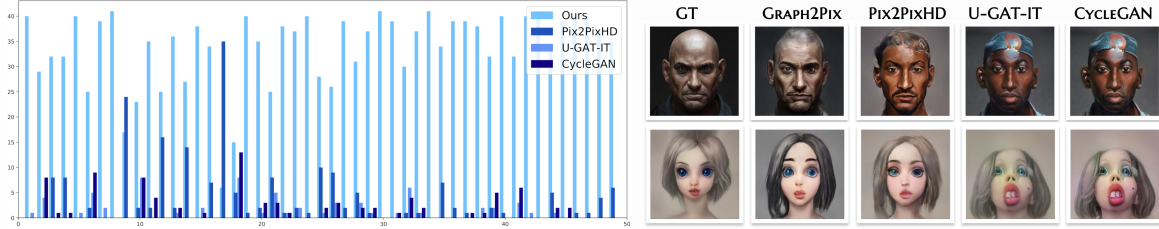


Figure 7: The results for each question in our human evaluation survey are shown on the left (the x-axis represents the image IDs, the y-axis represents the number of votes received for each method). The top-performing image, where participants unanimously found our method was most successful, is shown in the upper right. Our method is the least successful on the image shown on the bottom right, where participants preferred the Pix2PixHD result.

Method	LPIPS↓	FID↓	KID↓	Human↑
Pix2PixHD	0.32	23.89	0.71	0.097
CycleGAN	0.37	19.35	1.36	0.043
U-GAT-IT	0.33	19.35	-1.00	0.044
MS-GAN	0.46	85.62	5.18	N/A
Pix2Pix	0.52	112.01	9.73	N/A
<b>Graph2Pix</b>	<b>0.25</b>	<b>19.29</b>	<b>-1.12</b>	<b>0.815</b>

Table 1: Comparison of our method with baselines in quantitative and qualitative experiments. For LPIPS, FID and KID lower is better. For human evaluations (denoted by **Human**), higher is better.

copy the input image (in this case, one of the parent images). On the other hand, we note that our method significantly improves the FID score of the baseline Pix2PixHD method by 19.25%. Our method also outperforms its competitors in terms of KID score, which compares the skewness of the distributions and the mean and variance statistics.

**Ablation Study.** In addition to comparing our method to existing methods, we investigated whether using a GCN helps the learning process as follows. We created a random adjacency matrix and trained our model without using any lineage information. When we compare the generated images of the random adjacency based model with our model, we observe that FID and LPIPS scores degraded by 53.17% and 37.84%, respectively. This ablation study verifies that our network indeed learns the underlying topology information.

### 4.3. Qualitative Evaluation

In this section, we first present randomly selected images and compare them with the top two competitors, Pix2PixHD and U-GAT-IT. For clarity, we omit the CycleGAN images from the comparison, as it generates images that are very similar to those of U-GAT-IT (see Figure 7 for a comparison between the results of CycleGAN and U-

GAT-IT). Figure 6 lists the two parent images, the child image (i.e., the *ground truth*), and the images generated by Graph2Pix, Pix2PixHD, and U-GAT-IT, respectively. To allow consistent comparison, we only feed the first parent ( $Parent_1$ ) to the competitors to assess how well they were able to use this information. The first observation is that U-GAT-IT generates almost identical images to the parent and does not learn how to transform from the parent image domain to the child image domain. We think this is due to the unsupervised nature of the model. On the other hand, we can see that Pix2PixHD is able to generate images that are different from the parent, but it fails to generate images that are close to the ground truth. Since our method is able to consider both the parent and its ancestor, it achieves a better result that is closer to the ground truth.

We further conducted a qualitative analysis with 60 human participants and asked them to compare our method with the baseline approaches. Our experimental setup follows the standard approach employed by popular image-to-image translation methods [5, 37], where participants are given a pair of source and reference images and instructed to choose between candidate images. For brevity, we selected the top 3 competitors according to our quantitative analysis in Section 4.2, namely Pix2PixHD, U-GAT-IT, and CycleGAN. We randomly select 50 images for our experiment. For each image, we show the ground truth as the reference image, followed by four images (generated by our method and the competitors) whose order is randomly shuffled, and we ask which image is most similar to the reference image. Participants took an average of 9 minutes to complete the survey. As shown in Table 1, our method obtains the majority of votes (81.5%), followed by Pix2PixHD (9.7%).

We also examined the votes on each image in Figure 7. As can be seen from the figure, our method consistently receives the majority of votes on all images with a few exceptions. It is noticeable that our method is unanimously chosen as the best approach for some images, such as the top right image in Figure 7, where we can see that the skin color and facial structure are much more similar to the ground truth than the competitors. This example shows that the an-



Figure 8: The images in the first three columns are inputs, while the image in the last column is generated by our framework. All input images are treated as a first-degree parent, and the expected cloth image in this case is the child in the lineage.

cestor information is very important for some images and including this structure adds value to the image generation process. In contrast, our method performed poorly on some images, such as the image on the bottom right of Figure 7.

#### 4.4. Further Use Cases

Our method is not limited to generating images from Artbreeder. Due to the flexibility of the graph structure, our method is also applicable to any scenario where multiple images need to be translated into one, such as multi-source human image generation [21]. To show the validity of our method in different use cases, we applied our method to this case as well. We used the images of a cloth on a model from three different views (front, back and close) as input and generated the image of the cloth itself<sup>8</sup> (Figure 8). To see if having multiple parents are indeed helpful, we applied our framework to an image from a single view (front) and compared it to the multiple view scenario based on FID score. We obtained 45.26 in the single view case, while obtaining 43.43 in the multiple view case. Moreover, the comparison between these scores and FID score in Table 1 shows that the network learned Artbreeder dataset better, implying that a deeper lineage structure improves the performance of our graph-based image-to-image translation framework.

#### 5. Limitations and Failure Cases

In Figure 9, we show some of the failure cases of our method. It is noticeable that although our method is able to take into account the ancestral information and generate images that resemble the ground truth, it fails when the colors of the parent images vary significantly. Another reason for this effect could be that some users *crossbreed* images with different ratios. Therefore, the child image may carry more features from one of the parents (like the top image in Figure 9) and our model is not able to capture such dramatic ratios.

An important limitation of our method is that the input images should have a graph- or tree-based structure. However, we argue that due to the flexibility of the graph struc-

<sup>8</sup>Source: <https://modesens.com/product/prada-asymmetric-floralprint-stretchsilk-midi-dress-navy-9969054/>




Figure 9: Some cases where our method fails in learning coloring when the parent images have extreme differences.

ture, our method is also applicable to any scenario where multiple images need to be translated into one, such as human image generation [21] from multiple sources. In addition, we believe that our work can be explored in 3D reconstruction settings or in multi-agent communication settings where multiple drones share their image captures from different locations and construct a single coherent image (see [28] for a similar setting).

#### 6. Conclusion & Future Work

In conclusion, we propose Graph2Pix, an image-to-image translation method that considers multiple images and their corresponding graph structure as input and generates a single image as output. To the best of our knowledge, our method is the first to propose a GCN-based image-to-image translation method. Our quantitative experiments show that our method outperforms the competitors and improves LPIPS, FID, and KID scores by 25%, 0.3%, and 12%, respectively. Furthermore, our qualitative experiments show that human participants prefer the images generated by our method 81.5% of the time compared to its competitors. For future work, we consider using metadata such as popularity as a way to indicate the importance of the ancestors and we plan to extend our GCN module to capture hierarchical graphs. In addition, we plan to extend our method to other multi-source problems, such as generating human images from multiple sources.

**Acknowledgments.** This publication has been produced benefiting from the 2232 International Fellowship for Outstanding Researchers Program of TUBITAK (Project No: 118c321). We also acknowledge the support of NVIDIA Corporation through the donation of the TITAN X GPU and GCP research credits from Google. We also would like to thank Joel Simon for their support in collecting the dataset.



## References

- [1] Visualizing representations: Deep learning and human beings - colah's blog. [4](#)
- [2] Lars Backstrom and Jure Leskovec. Supervised random walks: predicting and recommending links in social networks. In *WSDM*, 2011. [2](#)
- [3] Mikolaj Bińkowski, Dougal J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying mmd gans. 2018. [6](#)
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. *CoRR*, 2018. [1](#)
- [5] Yunjey Choi, Youngjung Uh, Jaejun Yoo, and Jung-Woo Ha. Stargan v2: Diverse image synthesis for multiple domains. In *CVPR*, 2020. [7](#)
- [6] Nicola De Cao and Thomas Kipf. Molgan: An implicit generative model for small molecular graphs. 2018. [2](#)
- [7] Jesse Engel, Kumar Krishna Agrawal, Shuo Chen, Ishaan Gulrajani, Chris Donahue, and Adam Roberts. GANSynth: Adversarial neural audio synthesis. In *International Conference on Learning Representations*, 2019. [1](#)
- [8] Junyu Gao, Tianzhu Zhang, and Changsheng Xu. Graph convolutional tracking. In *CVPR*, 2019. [2](#)
- [9] Ian et. al. Goodfellow. Generative adversarial nets. In *NeurIPS*. 2014. [2](#)
- [10] Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *JMLR*, 2012. [6](#)
- [11] Aaron Hertzmann, Charles Jacobs, Nuria Oliver, Brian Curless, and David Salesin. Image analogies. *SIGGRAPH*, 2001. [2](#)
- [12] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, Günter Klambauer, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a nash equilibrium. [6](#)
- [13] Geoffrey Hinton and Sam Roweis. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15:833–840, 2003. [4](#)
- [14] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A. Efros. Image-to-image translation with conditional adversarial networks. 2016. [2](#), [5](#)
- [15] Liming Jiang, Changxu Zhang, Mingyang Huang, Chunxiao Liu, Jianping Shi, and Chen Change Loy. Tsit: A simple and versatile framework for image-to-image translation. In *European Conference on Computer Vision*, pages 206–222. Springer, 2020. [2](#)
- [16] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of gans for improved quality, stability, and variation. 2017. [1](#)
- [17] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *CoRR*, abs/1812.04948, 2018. [1](#)
- [18] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-GAT-IT: unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. [5](#)
- [19] Junho Kim, Minjae Kim, Hyeonwoo Kang, and Kwanghee Lee. U-gat-it: Unsupervised generative attentional networks with adaptive layer-instance normalization for image-to-image translation. *arXiv preprint arXiv:1907.10830*, 2019. [2](#)
- [20] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. 2016. [2](#), [3](#)
- [21] Stéphane Lathuilière, Enver Sangineto, Aliaksandr Siarohin, and Nicu Sebe. Attention-based fusion for multi-source human image generation. 2020. [8](#)
- [22] Chuan Li and Michael Wand. Precomputed real-time texture synthesis with markovian generative adversarial networks. *CoRR*, abs/1604.04382, 2016. [3](#)
- [23] Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Unsupervised image-to-image translation networks. In *Advances in neural information processing systems*, pages 700–708, 2017. [2](#)
- [24] Ming-Yu Liu, Xun Huang, Arun Mallya, Tero Karras, Timo Aila, Jaakko Lehtinen, and Jan Kautz. Few-shot unsupervised image-to-image translation. In *ICCV*, 2019. [6](#)
- [25] Xihui Liu, Guojun Yin, Jing Shao, Xiaogang Wang, and Hongsheng Li. Learning to predict layout-to-image conditional convolutions for semantic image synthesis. *arXiv preprint arXiv:1910.06809*, 2019. [2](#)
- [26] Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. Mode seeking generative adversarial networks for diverse image synthesis. In *CVPR*, 2019. [5](#)
- [27] Andrzej Maćkiewicz and Waldemar Ratajczak. Principal components analysis (pca). *Computers & Geosciences*, 19(3):303–342, 1993. [4](#)
- [28] Rupert Mitchell, Jan Blumenkamp, and Amanda Prorok. Gaussian process based message filtering for robust multi-agent cooperation in the presence of adversarial communication. *arXiv preprint arXiv:2012.00508*, 2020. [8](#)
- [29] Gaurav Mittal, Shubham Agrawal, Anuva Agarwal, Sushant Mehta, and Tanya Marwah. Interactive image generation using scene graphs. 2019. [2](#)
- [30] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. *CoRR*, 2019. [1](#), [2](#), [6](#)
- [31] Elad Richardson, Yuval Alaluf, Or Patashnik, Yotam Nitzan, Yaniv Azar, Stav Shapiro, and Daniel Cohen-Or. Encoding in style: a stylegan encoder for image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2287–2296, 2021. [2](#)
- [32] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. 2015. [3](#)
- [33] Dominik Sacha, Michael Sedlmair, Leishi Zhang, John Lee, Jaakko Peltonen, Daniel Weiskopf, Stephen North, and Daniel Keim. What you see is what you can change: Human-centered machine learning by interactive visualization. *Neurocomputing*, 268, 04 2017. [4](#)
- [34] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567, 2015. [6](#)
- [35] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9:2579–2605, 11 2008. [4](#), [5](#)

- [36] Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. Extracting and composing robust features with denoising autoencoders. In *ICML*, pages 1096–1103, 2008. [4](#)
- [37] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. *CoRR*, abs/1711.11585, 2017. [3](#), [5](#), [7](#)
- [38] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans, 2017. [1](#)
- [39] Ting Yao, Yingwei Pan, Yehao Li, and Tao Mei. Exploring visual relationship for image captioning. In *ECCV*, 2018. [2](#)
- [40] Runhao Zeng, Wenbing Huang, Mingkui Tan, Yu Rong, Peilin Zhao, Junzhou Huang, and Chuang Gan. Graph convolutional networks for temporal action localization. In *ICCV*, 2019. [2](#)
- [41] Li Zhang, Xiangtai Li, Anurag Arnab, Kuiyuan Yang, Yunhai Tong, and Philip HS Torr. Dual graph convolutional network for semantic segmentation. 2019. [2](#)
- [42] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018. [6](#)
- [43] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *CoRR*, 2017. [1](#), [2](#), [5](#)