

Cascading Generative Adversarial Networks for Targeted Imagination

Thesis by
Abdullah Jamal Hamdi

In Partial Fulfillment of the Requirements

For the Degree of
Masters of Science
in Electrical Engineering

King Abdullah University of Science and Technology
Thuwal, Kingdom of Saudi Arabia

April, 2018

EXAMINATION COMMITTEE PAGE

The thesis of Abdullah Jamal Hamdi is approved by the examination committee

Committee Chairperson: Dr. Bernard Ghanem

Committee Members: Dr. Wolfgang Heidrich, Dr. Peter Wonka

©April,2018

Abdullah Jamal Hamdi

All Rights Reserved

ABSTRACT

Cascading Generative Adversarial Networks for Targeted Imagination

Abdullah Jamal Hamdi

Abundance of labelled data played a crucial role in the recent developments in computer vision, but that faces problems like scalability and transferability to the wild. One alternative approach is to utilize the data without labels, *i.e.* unsupervised learning, in learning valuable information and put it in use to tackle vision problems. Generative Adversarial Networks (GANs) have gained momentum for their ability to model image distributions in unsupervised manner. They learn to emulate the training set and that enables sampling from that domain and using the knowledge learned for useful applications. Several methods proposed enhancing GANs, including regularizing the loss with some feature matching. We seek to push GANs beyond the data in the training and try to explore unseen territory in the image manifold. We first propose a new regularizer for GAN based on K-Nearest Neighbor (K-NN) selective feature matching to a target set Y in high-level feature space, during the adversarial training of GAN on the base set X , and we call this novel model K-GAN. We show that minimizing the added term follows from cross-entropy minimization between the distributions of GAN and set Y . Then, We introduce a cascaded framework for GANs that try to address the task of imagining a new distribution that combines the base set X and target set Y by cascading sampling GANs with translation GANs, and we dub the cascade of such GANs as the Imaginative Adversarial Network (IAN). Several cascades are trained on a collected dataset Zoo-Faces and generated innovative samples are shown, including from K-GAN cascade. We conduct an objective and subjective evaluation for different IAN setups in the addressed task of generating

innovative samples and we show the effect of regularizing GAN on different scores. We conclude with some useful applications for these IANs, like multi-domain manifold traversing.

ACKNOWLEDGEMENTS

This work was supported by the King Abdullah University of Science and Technology (KAUST) Office of Sponsored Research. I would like to thank Nourah AlJafaari for helping me in preparing some of the figures in this thesis .

TABLE OF CONTENTS

Examination Committee Page	2
Copyright	3
Abstract	4
Acknowledgements	6
List of Abbreviations	11
List of Figures	12
List of Tables	15
1 Introduction	16
1.1 Deep Learning and Labeled Data	16
1.2 Unsupervised Learning in Rescue	17
1.3 Advancing GANs	18
1.4 Objectives and Contributions	18
2 Related Work	20
2.1 Generative Adversarial Network	20
2.2 Feature Matching	21
2.3 Guiding the Generator to New Distributions	22
2.4 Cycle-Consistent Adversarial Networks	22
3 Methodology	25
3.1 K-GAN	26
3.1.1 Deriving Eq (3.2)	27
3.1.2 Reasons behind K-NN Feature Matching?	28
3.2 IAN Cascade	29
3.2.1 Sampling	29
3.2.2 Object geometry	30

3.2.3	Regularization.	30
4	Experiments	31
4.1	Implementation Details	31
4.2	Training	31
4.2.1	Training Data	31
4.2.2	Training Procedure	32
4.2.3	Selecting the Parameter K and μ for Training K-GAN	33
4.3	Sampling	36
5	Results	38
5.1	Objective Evaluation	38
5.1.1	Baselines	38
5.1.2	Evaluation Metrics	38
5.2	Subjective Evaluation	40
5.2.1	Quality Pair-wise Comparison	40
5.2.2	Human Subjective Classification Experiment	41
5.3	Ablation Study	41
5.3.1	Changing the Training Data	41
5.3.2	Abandoning the Search	42
5.3.3	Abandoning the Regularizer	43
5.3.4	Different GAN structures	43
5.4	Discussion	44
6	Applications	45
6.1	Multi-Domain Image Manifold Traversing	45
6.2	Celebrity Zoo	45
7	Concluding Remarks	47
7.1	Summary	47
7.2	Future Research Work	48
	References	49
	Appendices	53
A.1	Deriving the K-NN Loss	54
B.1	Justifying Regressing to the Mean of the K-features	58
C.1	Confusion Matrices for Human classification experiment	60

D.1	Zoo-Faces Dataset used in the Training	63
E.1	Samples of the K-BEGAN, K-BEGAN+CycleGAN Models	69
E.2	Full Celebrity Zoo	80
E.3	image Manifold Traversing	81

LIST OF ABBREVIATIONS

BEGAN	Boundary Equilibrium GAN
CycleGAN	Cycle-Consistent Adversarial Networks
DCGAN	Deep Convolutional GAN
EBGAN	Energy Based GAN
GAN	Generative Adversarial Network
HCA	Human Classification Accuracy
IAN	Imaginative Adversarial Network
K-NN	K-Nearest Neighbor
KDE	Kernel Density Estimation

LIST OF FIGURES

2.1	GAN structure : \mathbf{D}_X tries to determine if a given image is real (from training) or fake (generated by \mathbf{G}). \mathbf{G} on the other hand is trying to change the generated samples to better fool \mathbf{D}_X . \mathbf{z} is used to sample images from \mathbf{G}).	21
2.2	CycleGAN [1] : generator \mathbf{A} is conditioned on the set X as its input and is encouraged to output images in the distribution \mathbf{P}_Y , whereas generator \mathbf{B} is conditioned on the set Y as its input and is encouraged to output images in the distribution \mathbf{P}_X . $\mathbf{D}_X, \mathbf{D}_Y$ are discriminators on their respective domains X and Y	23
3.1	Imaginative Adversarial Network (IAN): addressing the task of generating samples from a domain between base set X and target set Y . The first GAN generate samples from learned distribution and the other GAN is conditioned on the output of the first GAN to translate to Y . The final samples combine characteristics of both sets X and Y	25
3.2	K-GAN as a flavoured sampling GAN : We propose K-NN regularizer for GAN that models a set X to push it further toward target set Y . In red frame left is the standard GAN of Generator G , Discriminator D . We add regularizer (right) with deep features extracted by conditioning network C on target Y and output of G . The target of regression in deep feature space is selected based on K-NN search in the features set of Y . Samples of K-GAN maintain X main appearance but obtain some Y properties (<i>e.g.</i> the pandas eyes and nose above).	26
4.1	The percentage of used features in training K-BEGAN model with $K=1, 4, \text{and } 8$, and using different target classes Y	35
4.2	Samples from K-BEGAN and its corresponding IAN: green-framed columns are K-BEGAN samples, orange-framed columns are full-IAN samples, each row is one X - Y IAN model where X :faces and Y is a class in Zoo-Faces. The far left and far right blue framed columns are the closest training examples to the neighboring K-BEGAN and IAN samples respectively.	37

5.1	Confusion Matrix for K-BEGAN + CycleGAN	42
5.2	Ablation study: (a) <i>Hard Imagination</i> , trying to translate images from set X (<i>e.g.</i> tiger) to very different set Y (<i>e.g.</i> flower) using CycleGAN. (b-d)studying the effect of altering parts of the model on the GAN output and its corresponding cascaded sample with CycleGAN (b)abandoning the K-NN search produce high frequency artifacts. (c)abandoning the K-NN feature matching leads to amplifying blurriness after cascade.(d)BEGAN vs DCGAN as base model . in all part red is failure and green is success.	43
6.1	Multi-domain image manifold traversing: moving from one point to another in different Y classes using the same pair in base set X . The top row is the actual traversing by convexing the codes obtained by the auto-encoder of BEGAN [2] and then decoding the convex codes .The rest of the rows are translations of the first row bu using different CycleGANs.	46
6.2	Celebrity-Zoo: seeing how celebrities look in the animal kingdom by using different IAN models trained in Sec.4.2.	46
C.1	Confusion Matrix for K-BEGAN	60
C.2	Confusion Matrix for DCGAN + CycleGAN	61
C.3	Confusion Matrix for K-DCGAN + CycleGAN	61
C.4	Confusion Matrix for BEGAN + CycleGAN	62
C.5	Confusion Matrix for K-BEGAN + CycleGAN	62
D.1	Samples of class apes the Zoo-Faces dataset	63
D.2	Samples of class chowchow in the Zoo-Faces dataset	64
D.3	Samples of class egyptian-cat in the Zoo-Faces dataset	64
D.4	Samples of class lions in the Zoo-Faces dataset	65
D.5	Samples of class koalas in the Zoo-Faces dataset	65
D.6	Samples of class norwegian elkhound in the Zoo-Faces dataset	66
D.7	Samples of class pandas in the Zoo-Faces dataset	66
D.8	Samples of class samoyed dog in the Zoo-Faces dataset	67
D.9	Samples of class spiders in the Zoo-Faces dataset	67
D.10	Samples of class tigers in the Zoo-Faces dataset	68
E.1	Samples of K-BEGAN model trained on Apes and faces	69
E.2	Samples of K-BEGAN+CycleGAN model trained on Apes and faces	70
E.3	Samples of K-BEGAN model trained on egyptian-cat and faces	70

E.4	Samples of K-BEGAN+CycleGAN model trained on egyptian-cat and faces	71
E.5	Samples of K-BEGAN model trained on chowchow and faces	71
E.6	Samples of K-BEGAN+CycleGAN model trained on chowchow and faces	72
E.7	Samples of K-BEGAN model trained on koalas and faces	72
E.8	Samples of K-BEGAN+CycleGAN model trained on koalas and faces	73
E.9	Samples of K-BEGAN model trained on lions and faces	73
E.10	Samples of K-BEGAN+CycleGAN model trained on lions and faces	74
E.11	Samples of K-BEGAN model trained on norwegian elkhound and faces	74
E.12	Samples of K-BEGAN+CycleGAN model trained on norwegian elkhound and faces	75
E.13	Samples of K-BEGAN model trained on pandas and faces	75
E.14	Samples of K-BEGAN+CycleGAN model trained on pandas and faces	76
E.15	Samples of K-BEGAN model trained on samoyed dog and faces . . .	76
E.16	Samples of K-BEGAN+CycleGAN model trained on samoyed dog and faces	77
E.17	Samples of K-BEGAN model trained on spiders and faces	77
E.18	Samples of K-BEGAN+CycleGAN model trained on spiders and faces	78
E.19	Samples of K-BEGAN model trained on tigers and faces	78
E.20	Samples of K-BEGAN+CycleGAN model trained on tigers and faces	79
E.21	More Celebrity-Zoo: seeing how celebrities look in the animal kingdom by using different IAN models trained in Sec.4.2.	80
E.22	multi-domain image manifold traversing 1: moving from one point to another in different Y classes using the same pair in base set X	81
E.23	multi-domain image manifold traversing 2: moving from one point to another in different Y classes using the same pair in base set X	82
E.24	multi-domain image manifold traversing 3: moving from one point to another in different Y classes using the same pair in base set X	82
E.25	multi-domain image manifold traversing 4: moving from one point to another in different Y classes using the same pair in base set X	83
E.26	multi-domain image manifold traversing 5: moving from one point to another in different Y classes using the same pair in base set X	83

LIST OF TABLES

4.1	The values of parameter K and coefficients of regularizer used in training different K-GAN models.	34
5.1	Objective Evaluation :The percentage scores (more is better) and normalized pixel-error percentage (less is better) for base set X, and target set Y. bold is best and <u>underlined</u> is second best. First three rows are actual data put for reference.The following six rows are different GANs, the last seven rows are cascaded GANs.Scores are averaged over the 10 different Y sets See Sec.5.1 for details	39
5.2	Human Subjective Evaluation :Pair-wise quality comparison and Human Classification Accuracy (HCA) for six different IANs, and GANs . The Pair-wise percentages represent the frequency that a certain model's samples were preferred against their rival's model samples by human subjects .HCA is for a classification task given for the subjects that were asked to classify the samples from different IANs. The higher the HCA the more easily identifiable are the IAN samples. See Sec.5.2 for more details	41

Chapter 1

Introduction

Future AI agents will be interacting with humans and performing more advanced tasks than just classification or detection. Among the qualities of such agents is to push the boundary for human knowledge, using a unique capability that we humans take for granted - namely - *imagination*. Imagination can be roughly defined as the ability to envision what we have not seen based on what we have seen before. In this work we introduce a cascaded framework of generative models we call the Imaginative Adversarial Network (IAN) in which we try to model new imaginative image distributions from base set X and target set Y.

1.1 Deep Learning and Labeled Data

Deep neural networks have shown great success in pushing the envelope with regards to discriminative tasks, such as classification and detection [3, 4]. It is clear that large-scale fully annotated datasets (*e.g.* ImageNet [5]) have played a crucial role in this development. However, they are still influenced by biases, which limit their usefulness in the wild [6]. The cost of laborious labeling of data (*e.g.* activity localization in videos or pixel-level segmentation) leads us to look for alternative ways that can enable further improvements in computer vision tasks. Learning from synthetic data seems a promising direction to tackle this problem in which some notable success is shown due to the fact that full knowledge about the simulated environment is taken

for granted [7, 8, 9]. However, as pointed out by Roberto *et al.* [8], learning on such simulated data suffer from the domain adaptation problem when testing on real data.

1.2 Unsupervised Learning in Rescue

In a different direction, Generative Adversarial Networks (GANs) have emerged as a promising family of unsupervised learning techniques that have recently shown the ability to model simple natural images, including faces and flowers [10]. In the last couple of years, this topic has gained a large amount of attention from the community, and since its inception, many variations of GAN have been put forward to overcome some of the impediments it faces (*e.g.* instability during training) [11, 2, 12, 13, 14, 15, 16, 17]. Interestingly, GANs have proven useful in several core image processing and computer vision tasks, including image inpainting [18], style transfer [19], super-resolution [20, 21], manifold traversing [22], hand pose estimation [23], and face recognition [24]. These success cases indicate that GANs have the potential to learn (in an unsupervised way) important underlying natural image properties that are exploited in turn to synthesize images.

However, one characteristic of the GAN is that it is bounded to the data that it uses for training. In many cases, it might degenerate to simply *memorizing* training data, thus, limiting its ability to discover new “territory” in the manifold of natural images. The ability to discover unseen territory is crucial because it can help to bypass the domain adaptation problem of classification [25]. Although some recent work has been done to maximize innovation in synthesized GAN sample [26], it is limited in scope to the specific domain of art and does not specify the direction of innovation. Our proposed cascaded approach, IAN, goes beyond this limitation by defining a generic base set X and target set Y .

1.3 Advancing GANs

Our proposed regularized K-GAN is inspired by recent work on feature matching in adversarial image generation [27, 19, 28, 29, 7, 16], which aims to match the generated image to another set of images in high-dimensional feature space. However, Johnson *et al.* [19] is optimizing for a single image not a distribution, while Warde-Farley and Bengio [16] matching was bounded to denoising local features of GAN. We utilize more general features of AlexNet [3] trained on ImageNet[5] object classification. Unlike the common practice of modeling base class X and enforcing the generated output to follow certain class label or features of class X [27, 28, 29], we aim to push toward features of target set Y different from X while preserving the GAN modelling of X. Furthermore, most of the methods relied on taking the features of sampled data without validating their quality before matching to them. We propose to follow K-nearest neighbor approach in the high-level feature space in the training of GAN. Recently, CycleGAN [1] showed excellent results in image-to-image translation from source domain X to target domain Y based on cycle consistency regularizer to adversarial loss, but as pointed out in their paper, it fails to model transformation when there are geometrical differences between objects in X and Y. We propose to utilize CycleGAN as the translation GAN in our IAN framework in cascade with different sampling GAN variations (including our novel K-GAN) to achieve the goal of modeling new meaningful distribution between X and Y.

1.4 Objectives and Contributions

The contributions of this thesis folds in the following streams:

- We propose a new regularizer for GAN that uses K-NN selective feature matching to target set Y in high-level feature space during the adversarial training of any GAN, and we call this novel model K-GAN.

- We show that minimizing the added term in K-GAN follows from cross entropy minimization between Y and the GAN distribution.
- We present a cascade framework for GANs (IAN) to push their output samples away from the base domain X to target Y , and we train and test different IANs including K-GAN+CycleGAN.
- We conduct objective and subjective evaluation for different IAN setups in the addressed task and show some useful applications for these IANs , like multi-domain manifold traversing.

Chapter 2

Related Work

2.1 Generative Adversarial Network

Generative Adversarial Networks (GANs) [10] are generative models consisting of a Discriminator \mathbf{D}_X and a Generator \mathbf{G} that are adversarially trained in a similar manner as a minimax game. The discriminator tries to determine if a given image is real (from training) or fake (generated by \mathbf{G}). The Generator on the other hand is trying to change the generated samples to better fool \mathbf{D}_X . This can be formulated as the following minimax game on the loss function $L_{GAN}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X)$:

$$\min_{\mathbf{G}} \max_{\mathbf{D}_X} L_{GAN}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) = \mathbb{E}_{\mathbf{x} \sim p_x(\mathbf{x})} [\log \mathbf{D}_X(x)] + \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [\log(1 - \mathbf{D}_X(\mathbf{G}(z)))] \quad (2.1)$$

where \mathbf{P}_X is the distribution of images in domain X , $\mathbf{z} \in \mathbb{R}^d$ is the latent uniformly random vector. Fig Zhao *et al.* introduces energy-based loss for GANs [30], and recent GANs uses the Wasserstein distance instead of K-L divergence as objective [11, 2]. Both Energy Based GAN (EBGAN) [30] and Boundary Equilibrium GAN (BEGAN) [2] uses auto-encoder Discriminator network with its loss defined as L2 pixel difference between input and output.

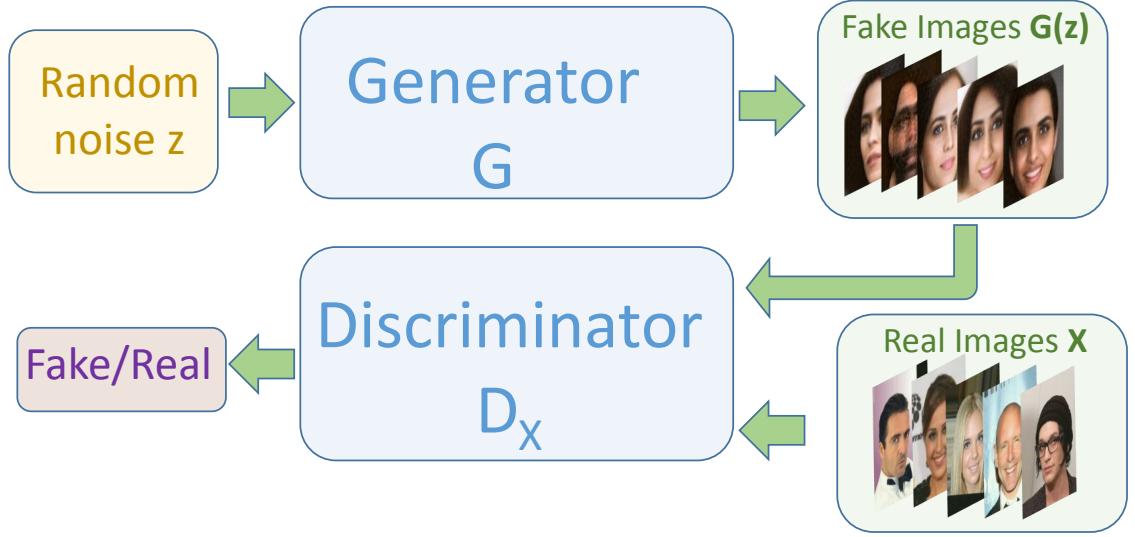


Figure 2.1: GAN structure : \mathbf{D}_X tries to determine if a given image is real (from training) or fake (generated by \mathbf{G}). \mathbf{G} on the other hand is trying to change the generated samples to better fool \mathbf{D}_X . \mathbf{z} is used to sample images from \mathbf{G}).

2.2 Feature Matching

The method of Warde-Farley and Bengio [16] uses Denoising Auto Encoder (DAE) feature matching to enhance the training of the generator and to increase the diversity of the samples. The work of Dosovitskiy A. and Brox T. [27] shows that a conditioning network \mathbf{C} (also called Comparitor) can be incorporated in the GAN loss by adding a discrepancy loss in deep feature space, which measures the difference between the generated image and same class images in this feature space to increase Perceptual similarity. A similar technique is used for style transfer and super resolution by Johnson *et al.* [19]. They try to match covariance matrix of deep features of some image (extracted from VGGNet [31]) and match that to some target domain. Shrivastava *et al.* [7] uses GAN loss to improve the quality of synthetic images and used the deep features of a network trained on classifying eye gaze to maintain the key elements of the synthetic data (*e.g.* the direction of gaze). Plug&Play [29] uses DAE to model the FC7 and Conv5 layers of \mathbf{C} network and shows that it can improve

the quality of the generator network under some conditions.

2.3 Guiding the Generator to New Distributions

As pointed out by Elgammal *et al.* [26], GANs by nature emulate but do not innovate, hence their work introduces the Creative Adversarial Network (CAN) to tackle this issue. They add a loss to the traditional GAN loss to encourage the generator to increase its entropy so as to fool the discriminator, leading the generator to produce more innovative samples. We tackle the same problem but through a different approach. As compared to [26], we define a target set Y with distribution \mathbf{P}_Y (different from the data distribution \mathbf{P}_X) that guides the generator to new territory. In this context, the work of LIU and Tuzel [25] introduces CoGAN to model the joint distribution by coupling two GANs in parallel with shared network weights. This aids domain adaptation when simple distributions are available (*e.g.* MNIST [32]). However, it does not work well for more complex domains as compared to CycleGAN [1]. Unlike image manipulation [33], which manipulates individual image at a time, we manipulate the entire distribution of \mathbf{G} from which we can sample (given the model trained properly).

2.4 Cycle-Consistent Adversarial Networks

Recently, Cycle-Consistent Adversarial Networks (CycleGAN) [1] has shown compelling results in image translation by imposing cycle consistency between two generators (each takes images as input). It produces a sharp translation from one domain X to another domain Y with no paired training. CycleGAN has the following loss on the two distributions \mathbf{P}_X and \mathbf{P}_Y :

$$\begin{aligned} L_{\text{CycGAN}}(\mathbf{A}, \mathbf{B}, \mathbf{D}_X, \mathbf{D}_Y) = \\ L_{\text{GAN}}(\mathbf{A}, \mathbf{D}_Y, \mathbf{P}_Y) + L_{\text{GAN}}(\mathbf{B}, \mathbf{D}_X, \mathbf{P}_X) + \lambda L_{\text{co}}(\mathbf{A}, \mathbf{B}) \end{aligned} \tag{2.2}$$

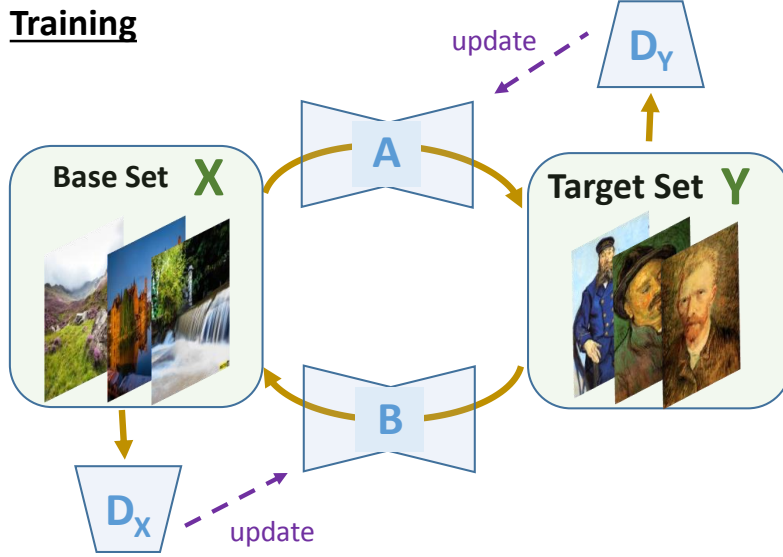
Training**Inference**

Figure 2.2: CycleGAN [1] : generator \mathbf{A} is conditioned on the set X as its input and is encouraged to output images in the distribution \mathbf{P}_Y , whereas generator \mathbf{B} is conditioned on the set Y as its input and is encouraged to output images in the distribution \mathbf{P}_X . $\mathbf{D}_X, \mathbf{D}_Y$ are discriminators on their respective domains X and Y

where \mathbf{A}, \mathbf{B} are the two generators, \mathbf{D}_X and \mathbf{D}_Y are the discriminators for domains X and Y respectively. $L_{co}(\mathbf{A}, \mathbf{B})$ is the cycle consistency loss of the CycleGAN that encourages that $\mathbf{B}(\mathbf{A}(\mathbf{x})) = \mathbf{x}$ and is defined as:

$$L_{co}(\mathbf{A}, \mathbf{B}) = \mathbb{E}_{\mathbf{x} \sim p_{\mathbf{x}}(\mathbf{x})} [\|\mathbf{B}(\mathbf{A}(\mathbf{x})) - \mathbf{x}\|_1] + \mathbb{E}_{\mathbf{y} \sim p_{\mathbf{y}}(\mathbf{y})} [\|\mathbf{A}(\mathbf{B}(\mathbf{y})) - \mathbf{y}\|_1] \quad (2.3)$$

where λ is controls the consistency regularizer. The idea works as follows : generator \mathbf{A} is conditioned on the set X as its input and is encouraged by $L_{GAN}(\mathbf{A}, \mathbf{D}_Y, \mathbf{P}_Y)$ to output images in the distribution \mathbf{P}_Y , whereas generator \mathbf{B} is conditioned on the set Y as its input and is encouraged by $L_{GAN}(\mathbf{B}, \mathbf{D}_X, \mathbf{P}_X)$ to output images in the distribution \mathbf{P}_X . The extra term $L_{co}(\mathbf{A}, \mathbf{B})$ insures that style transfer does not collapse to a single image.

Although CycleGAN shows compelling results when the images translation occurs

in texture or color, it shows less success when the translation requires geometric changes. This is due to the limited features modeling CycleGAN performs. We address this limitation in our proposed method by enforcing change in the object’s deep features, which tend to go beyond mere texture representation. We do this by introducing K-NN feature matching , which encourages the generated sample to be close in deep feature space to its nearest neighbors in the target set. Interestingly, we also see that this type of matching can also be viewed mathematically as a way to reduce the cross-entropy between the target distribution \mathbf{P}_Y and the distribution of the generated images.

Chapter 3

Methodology

We propose Imaginative Adversarial Network (IAN), a two-stage cascade imaginative model that aims to generate samples that resemble base set X and target set Y . First, we sample from a GAN by a latent vector \mathbf{z} , then the second stage is a translation GAN, that takes the samples from the first stage and translate them to a different domain. In the first stage (sampling stage), we use the classical DCGAN[34] and recent BEGAN[2], along with their K-NN regularized versions. We use the recent CycleGAN [1] as a second stage with both stages trained on the same X and Y sets and cascaded during the sampling to obtain an output that both have key properties of X and Y as illustrated in the flowchart in Fig 3.2. For start, we develop the K-NN regularized GAN in the following subsection.

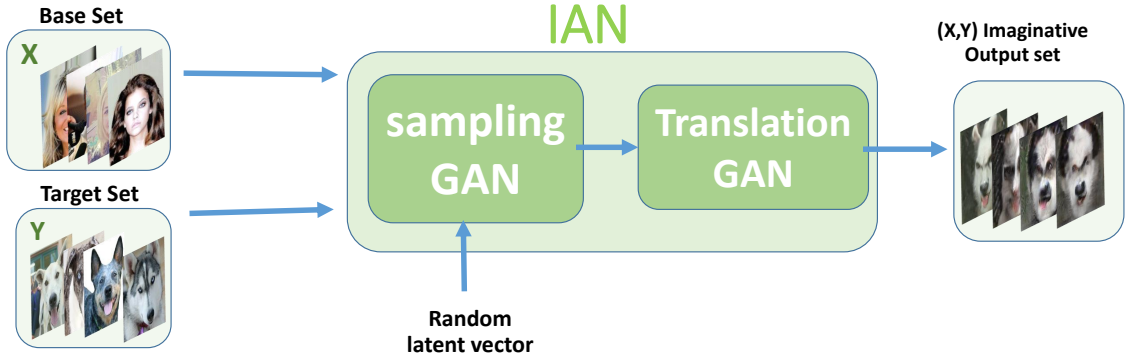


Figure 3.1: Imaginative Adversarial Network (IAN): addressing the task of generating samples from a domain between base set X and target set Y . The first GAN generate samples from learned distribution and the other GAN is conditioned on the output of the first GAN to translate to Y . The final samples combine characteristics of both sets X and Y

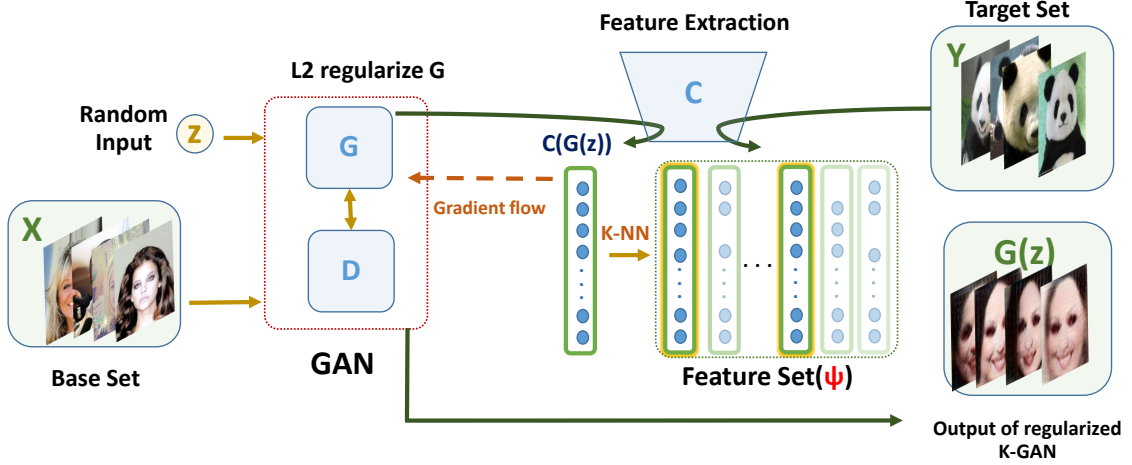


Figure 3.2: K-GAN as a flavoured sampling GAN : We propose K-NN regularizer for GAN that models a set X to push it further toward target set Y . In red frame left is the standard GAN of Generator G , Discriminator D . We add regularizer (right) with deep features extracted by conditioning network C on target Y and output of G . The target of regression in deep feature space is selected based on K-NN search in the features set of Y . Samples of K-GAN maintain X main appearance but obtain some Y properties (*e.g.* the pandas eyes and nose above).

3.1 K-GAN

K-GAN adds to the traditional GAN loss in Eq (2.1) a term that encourages the generator distribution \mathbf{P}_g towards a target distribution \mathbf{P}_Y through cross entropy minimization. The similarity of \mathbf{P}_g to \mathbf{P}_X is also encouraged through the traditional GAN loss. The K-GAN loss is defined as follows:

$$L_{\text{KGAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X, \mathbf{P}_Y) = L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) + \mu \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)] \quad (3.1)$$

where the added term is the cross entropy between the image distribution of our generator \mathbf{P}_g and the image distribution of the target \mathbf{P}_Y , and μ is a regularizing hyperparameter. We replace the cross entropy $H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)$ with the following K-

NN regularizer:

$$L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y) = \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \sum_{n=1}^k \|\mathbf{C}(\mathbf{G}(\mathbf{z})) - \mathbf{r}_n(\mathbf{C}(\mathbf{G}(\mathbf{z})))\|_2^2 \quad (3.2)$$

This is the distance between a generated sample from \mathbf{G} by a latent vector \mathbf{z} and its K-NN samples of the target distribution \mathbf{P}_Y from the total M samples available from \mathbf{P}_Y . This distance is not in pixels, but in high-level feature space defined by the Comparator network $\mathbf{C}(\mathbf{b}) \in \mathbb{R}^s$ which takes an image $\mathbf{b} \in \mathbb{R}^{h \times w}$ as input and outputs feature vector (in our case it was FC7 and FC6 of AlexNet [3] trained on ImageNet [5]). Refer to Fig 3.2 for visualization. $\mathbf{r}_n(\mathbf{c})$ is the K^{th} NN function, a parametric (with parameter n) order selection function that selects an element \mathbf{y}_i from a set of $M - n + 1$ elements based on how close it is to input vector \mathbf{c} . It can be described by the following:

$$\mathbf{r}_n(\mathbf{c}) = \arg \min_{\mathbf{y}_i \in \psi_n} \|\mathbf{c} - \mathbf{y}_i\|_2, \quad (3.3)$$

where ψ_n is the set of deep features of $M - n + 1$ images representing the target distribution \mathbf{P}_Y . The optimization in (3.3) is discrete and the resulting function $\mathbf{r}_k(\mathbf{c})$ is complicated and indifferentiable but fixing it to current estimate of \mathbf{G} solve the problem. See Sec.4.2 for training details.

3.1.1 Deriving Eq (3.2)

To show how we get Eq.(3.2), first we look at \mathbf{P}_Y . Since acquiring a full description of \mathbf{P}_Y in deep feature space is infeasible in most cases, we settle to approximate it using M target samples. We can get a proxy distribution of \mathbf{P}_Y by using Kernel Density Estimation (KDE) on the deep features extracted by deep network $\mathbf{C}(\mathbf{x})$ from those M samples in \mathbf{P}_Y . By picking Gaussian kernel for the KDE, the proxy distribution estimator $\hat{\mathbf{P}}_{Y,\text{proxy}}$ is defined by a Bayesian non-parametric way at any point $\mathbf{C}(\mathbf{b}) \in \mathbb{R}^s$ for $\mathbf{b} \in \mathbb{R}^{h \times w}$ as follows:

$$\hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) = \frac{1}{M\sigma\sqrt{2\pi}} \sum_{i=1}^M \exp\left(-\frac{\|\mathbf{C}(\mathbf{b}) - \mathbf{C}(\mathbf{y}_i)\|_2^2}{\sigma^2}\right) \quad (3.4)$$

\mathbf{y}_i is the i^{th} sample describing \mathbf{P}_Y and σ^2 is the Gaussian kernel variance. For far images from the M samples of $\hat{\mathbf{P}}_{Y,\text{proxy}}$ the probability becomes exponentially negligible. Hence, we can investigate the effect of a subset of K points (out of the M points $\mathbf{C}(\mathbf{y}_i)$) that are closest to $\mathbf{C}(\mathbf{b})$ to compute (3.4). We replace M by K in (3.4), and then use Jansen inequality on that modified KDE to get a bound on the expression. By replacing \mathbf{b} by $\mathbf{G}(\mathbf{z})$ and taking logarithm of expectation of both sides, we obtain upper bound on the cross entropy $\mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})}[H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)]$. The bound found is the one in (3.2), and by minimizing that bound we minimize the cross entropy in Eq.(3.1). For a detailed derivation, please see Appendix A.

The K-GAN objective is therefore defined as follows:

$$\min_{\mathbf{G}} \max_{\mathbf{D}_X} L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) + \mu L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y), \quad (3.5)$$

where the two terms follow from (2.1) and (3.2) respectively.

3.1.2 Reasons behind K-NN Feature Matching?

The goal of employing K-NN feature matching in our loss is to utilize the distribution obtained by deep features of some set Y to guide the generation of the GAN. If GAN generated samples are close to all samples of target set Y in deep features, we become more confident about the "objectiveness" of our samples. However, measuring the distance to all samples features in a set is computationally expensive, especially in an iterative optimization like what is in GAN training. Hence, following the neighbors to guide the generation seems a viable alternative to the expensive global judgment of all samples.

Using K-features rather than a single feature at each training iteration will further

improve the robustness. This is due to the fact that features at layers (*i.e.* FC6 and FC7) of AlexNet are prone to small imperceptible noise. To demonstrate this, we add uniform noise with a maximum pixel corruption of 4% to 10000 randomly selected images from the ImageNet validation set. This resulted in a relative change in intensity values as high as 20%. This shows that the high-level feature space is prone to noise and depending on that solely can lead us astray from the objective we seek. That is why many previous works [29, 16] adapted the DAE in their architectures to mitigate the volatility of deep features. So, using the K-nearest neighbor goes along the same direction with smoothing parameter \mathbf{K} .

3.2 IAN Cascade

As described in Sec.1, an IAN model consists of a sampling GAN and a translation GAN to generate samples that capture the properties of two sets X and Y. We utilize our K-GAN model developed in Sec.3.1, along with other Vanilla GANs, and cascade them with the recently developed cycleGAN [1] after training each separately with some fine-tuning (see Sec.4.2 for training details). The goal is to utilize the ability of CycleGAN to transfer appearance for our sake to generate natural looking samples that capture key properties and appearance of both X and Y. By using K-GAN as input to the CycleGAN, the later gains the following advantages:

3.2.1 Sampling

\mathbf{z} vector sampling is a property of GAN to sample from the modeled distribution. However, CycleGAN is conditioned on input images. Hence, we use the sampling in K-GAN and transform the output with CycleGAN. While this can be done by any GAN, we show in Sec.5.1,5.2 how regularized cascades outperform others in many metrics .

3.2.2 Object geometry

As pointed out by its authors, CycleGAN is limited to the appearance transfer when it comes to cross-class translation (*e.g.* transforming cat to dog) [1]. We try to tackle this by enforcing feature matching to the target set, which results in geometric change that is appended by the appearance transfer of CycleGAN.

3.2.3 Regularization.

While the CycleGAN is trained on natural images, it is used here to translate synthetic images produced by other GANs. This poses the risk of amplifying the imperfections produced by the first GAN. We show in Sec.5.1 that regularizing the GAN (as in K-GAN) limits the possible outputs which helps in mitigating this effect of error amplification when CycleGAN is fine tuned to the regularized samples. The diversity of the final samples is achieved by the cascade effect with CycleGAN.

Chapter 4

Experiments

4.1 Implementation Details

For the K-GAN part, we use the basic network implemented in Deep Convolutional GAN (DCGAN) [34] and the more advanced architecture in BEGAN[2]. The KNN-FM term makes use of deep features extracted from the FC6 and FC7 layers of AlexNet [3] pre-trained on ImageNet[5]. AlexNet was favored over deeper networks because at each iteration of training K-GAN, feature extraction, multiple forward passes, and backpropagations are required for updating \mathbf{G} . We used FC7 because it showed the ability to represent images in many image tasks, including classification [35]. We used FC6 because it helped in transferring the key properties of the object along with FC7. We use the same CycleGAN network as proposed in [1], in which they used the PatchGAN discriminator implemented in [20] and U-Net Generators. The image output samples in all our cascades have 227×227 pixel resolution.

4.2 Training

4.2.1 Training Data

We used CelebA [36] as our base set X and we used 10 different animal classes from ImageNet [5] as our target set Y . The Y dataset was manually filtered to be suitable for the task of imagining humanoid animals by ensuring that the face of the animal appears in the image. Each of the Y target classes contains 100 to 500 images. The Y classes are gorillas, chow chow dog, Egyptian cat, koala ,lion, Norwegian-elkhound

dog, panda, Samoyed-dog, spider, and tiger. They were used because of their distinctive features, the abundance of facial data, and range of difficulty. We call this filtered dataset *Zoo-Faces*, and it will be made publicly available with the code. Refer to for some qualitative examples from this dataset. We tried to use ImageNet [5] without filtering, and that did not help to achieve the objective, and therefore we had to use filtered data collected from the internet. See Sec.5.3 for more details. Before training, we take the Zoo-Faces dataset and extract the FC7 and FC6 features by the Comparator \mathbf{C} (AlexNet) to obtain 20 different feature matrices of FC7 and FC6 for the 10 classes. These feature sets are then passed to the K-GAN to be used in training.

4.2.2 Training Procedure

Since we are running a gradient descent based algorithm (Adam optimizer [37]), by fixing the argument of $\mathbf{r}_k(\cdot)$ in Eq.(3.3) to the current estimate of the generator \mathbf{G}_t at time step t , the term $\mathbf{r}_n(\mathbf{C}(\mathbf{G}_t(\mathbf{z})))$ is just a constant $\mathbf{d}_{n,t}$ at that iteration. For step t and parameter $n = K$, we use the ball-tree algorithm [38] to search for K-NN features in the previously obtained set ψ_n that are close to the current sample from $\mathbf{G}_t(\mathbf{z})$. The term in Eq.(3.2) becomes the familiar ℓ_2 regularizer:

$$\mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \sum_{n=1}^k \|\mathbf{C}(\mathbf{G}(z)) - \mathbf{d}_{n,t}\|_2^2 \quad (4.1)$$

For practical reasons, we minimize the distance between $\mathbf{C}(\mathbf{G}(z))$ and the average of the K neighbors instead of the distances to all of them independently. This is justified and can be seen by expanding the norms and looking at the argument minimizer to realize it is the same in both cases. please refer to Appendix B for detailed mathematical justification. We trained each GAN on CelebA and then used the

pretrained model to fine tune the K-GAN according to KNN-FM on CelebA and each class from Zoo-faces as target Y . At each iteration, we search the KNN features of FC7 and FC6 in Ψ_{FC7}, Ψ_{FC6} , which are the sets of features of class Y and pick the mean of the two K features as two regularization targets for the GAN. Because of the use of the ball-tree search [38] and SciPy library [39], the search is done real-time and delay of the training time of K-GAN is only 70% more compared to GAN with similar structure and for small target class Y of 100-500 instances. In general, the time complexity of the tree search is approximately $\mathcal{O}(s \log M)$, where s is the feature dimensionality, and M is the number of features in Ψ . The CycleGAN was trained independently on a subset of CelebA and each one of the 10 classes of Y and then fine tuned with samples of the cascaded GANs. This procedure was followed to obtain 10 K-DCGAN, 10 K-BEGAN, and 10 different CycleGANs corresponding to the 10 classes in Zoo-Faces. In the following section we look at the procedure to obtain the hyper-parameters and their values.

4.2.3 Selecting the Parameter K and μ for Training K-GAN

To study the effect of changing the K parameter of nearest neighbor search, we look at the cardinality of the set χ (a subset of the set Ψ of a high level feature of the target Y). χ contains different features from class Y that have been selected during the KNN search in K-GAN training. A bigger χ cardinality is an indication of a better hyper-parameter selection, which means we have more effectively utilized the set Ψ and the KNN search did not pick one feature several times during the training. We use this number as a validation metric to assess the quality of the training. We picked $K = 4$ because it compromises both the utilization of Ψ and the stability of the training. The value of μ hyperparameter (coefficient of the KNN regularizer) was picked to be 0.001 for FC7 and 0.0001 for FC6. For χ the subset of features picked by KNN search during the training of K-GAN out of M features, the following Fig

K-GAN model	K	μ_{FC7}	μ_{FC6}
Apes	8	0.012	0.001
egyptian-cat	8	0.010	0.001
Chowchow	8	0.012	0.001
koala	4	0.010	0.001
Lion	8	0.010	0.001
Norwegian elkhound	4	0.012	0.001
panda	4	0.010	0.001
samoyed dog	4	0.010	0.001
spider	8	0.010	0.001
tiger	4	0.010	0.001

Table 4.1: The values of parameter K and coefficients of regularizer used in training different K-GAN models.

D.10 shows percentage of $\text{card}(\chi)$ to M for different choice of target class Y and different values of hyper-parameter K . We used this percentage as quick validation metric to assess the quality of the training. We picked K such that we covers large percentage of the target set while preserving the uniqueness of each image feature. Also , since we are taking the mean of the KNN features , larger K is less prone to noise especially if there is an image in the training that is very unique from others and can be picked by the KNN search several times.

The μ_{FC7}, μ_{FC6} (the coefficients of FC7 and FC6 regularizers of K-GAN) were picked to insure that the key properties were indeed transferred from Y to X , at the same time not over-regularizing the model which can lead to the collapse of the K-GAN model and producing only a single image. Table 4.1 shows the values of K and μ_{FC7}, μ_{FC6} picked to train the final 10 K-BEGAN models that we reported its samples and did all the tests on for the 10 different Zoo-Faces dataset classes.

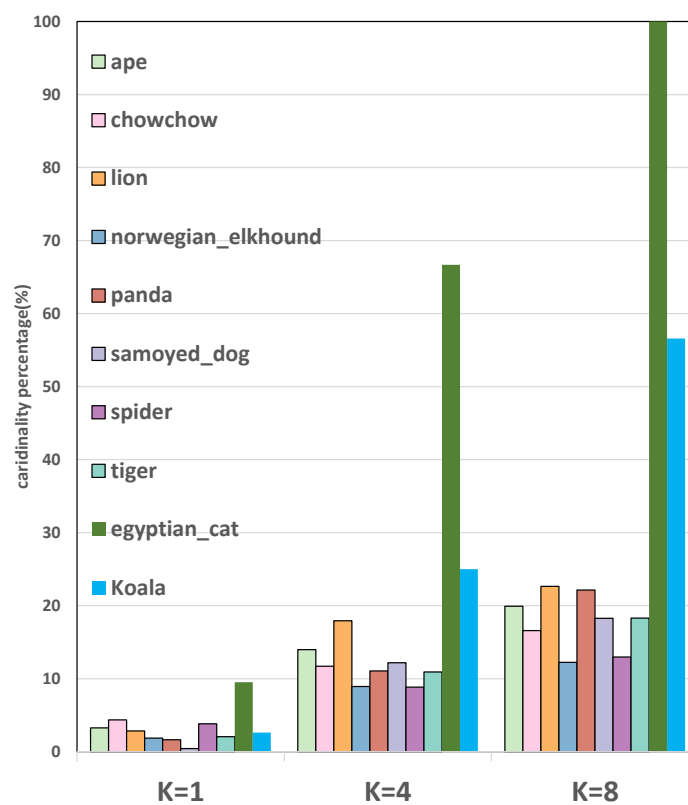


Figure 4.1: The percentage of used features in training K-BEGAN model with $K= 1,4,$ and 8 , and using different target classes Y

4.3 Sampling

After training, the models are cascaded, and sampling takes place in which we sample latent random vector $\mathbf{z} \in [-1, 1]^d, d = 100$ and pass it as input to the GANs and K-GANs similarly (No feature search happens at this stage). The image samples are then passed as input to the CycleGAN trained on the same X and Y sets that K-GAN was trained on. We obtain the output of an IAN model as in Fig 4.2 for K-BEGAN+CycleGAN. We show the closest training example next to some samples, as advised by [40]. More samples from the different models are provided in Appendix E

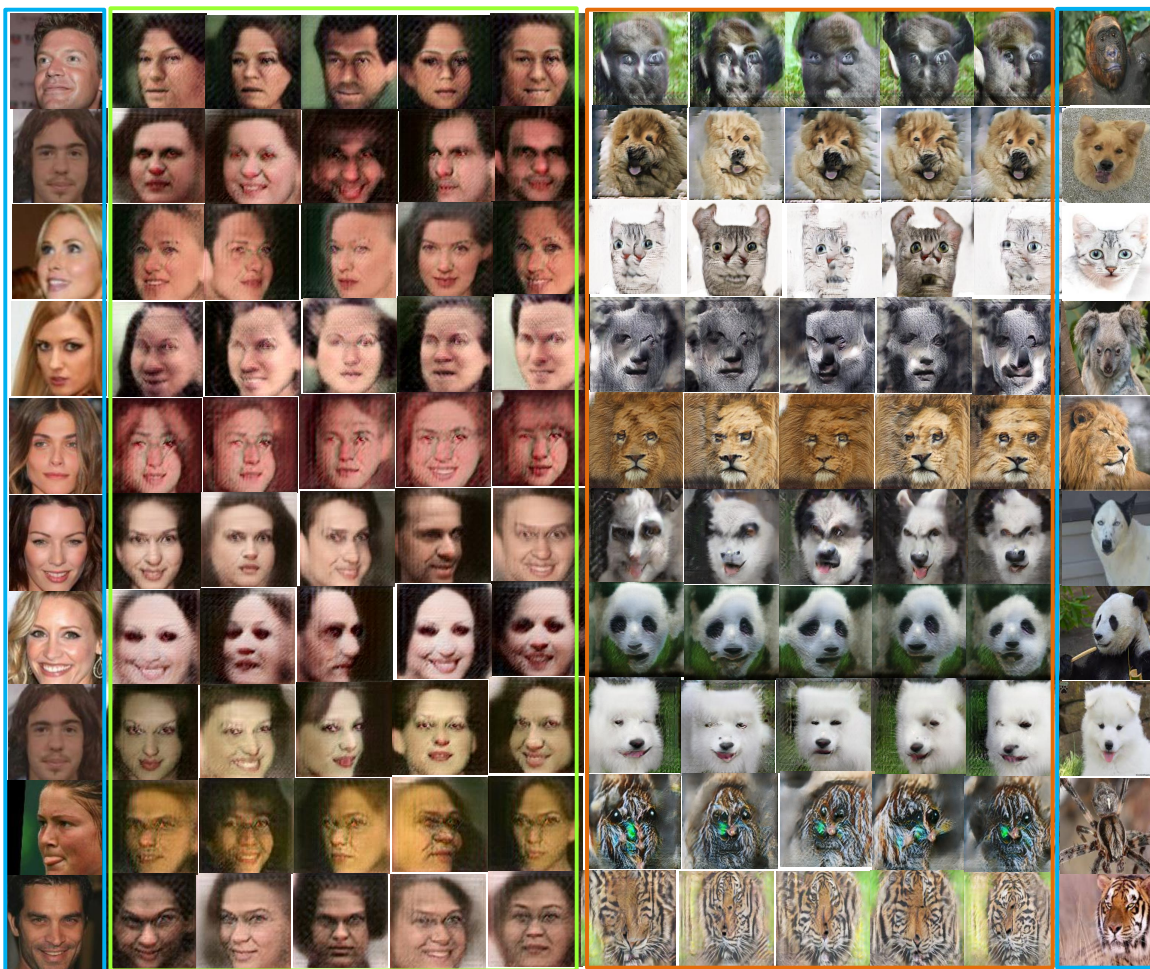


Figure 4.2: Samples from K-BEGAN and its corresponding IAN: green-framed columns are K-BEGAN samples, orange-framed columns are full-IAN samples, each row is one X-Y IAN model where X:faces and Y is a class in Zoo-Faces. The far left and far right blue framed columns are the closest training examples to the neighboring K-BEGAN and IAN samples respectively.

Chapter 5

Results

5.1 Objective Evaluation

5.1.1 Baselines

Ultimately, we want our generated samples to be indistinguishable from real samples; hence, we include the training images from X and Y as a reference for GAN model evaluation. We include the following GAN models (along with their cascade with CycleGAN) as baselines for fair comparison with our K-GAN and its cascade.

DCGAN , BEGAN: Vanilla DCGAN [34] and BEGAN[2] trained on CelebA.

mxGAN: naive approach of combining X, and one class of Y by mixing the two sets as one training set of the DCGAN.

PerceptualGAN: feature matching GAN without the use of KNN search (random K selection), similar to [27].

rand+Cyc: random noise translated with trained CycleGAN.

5.1.2 Evaluation Metrics

As indicated by Theis *et al.* [40], evaluating the generative models based on log likelihood’s approximations (*e.g.* Parzen window) should be avoided, and the models should be evaluated in the application they are being used in. To evaluate the imagination capability of our IAN model, we use the idea of inception score proposed by Salimans *et al.* [14] on the target class Y. However, since we are developing different models for each class and the goal is not to model the data but to imagine new distri-

model	score(.)			err(.)		
	X	Y	avg	X	Y	avg
real(X)	100	0.03	50.01	0.00	70.12	35.06
real(Y)	10.05	100	55.03	55.87	0.00	27.93
real(X)+Cyc	14.03	66.68	40.35	48.57	56.99	52.78
DCGAN[34]	95.90	0.05	47.97	47.11	67.18	57.14
mxGAN	88.15	0.05	44.10	48.03	67.86	57.95
BEGAN[2]	87.88	0.03	43.95	<u>45.66</u>	60.27	52.97
P.GAN [27]	99.76	0.06	<u>49.91</u>	46.51	67.22	56.86
K-DCGAN	<u>98.55</u>	0.05	49.30	47.27	66.77	57.02
K-BEGAN	97.29	0.36	48.83	44.43	58.10	51.27
rand+Cyc	1.56	0.09	0.83	49.69	54.87	52.28
DCGAN+Cyc	11.76	57.70	34.73	49.46	55.77	52.62
mxGAN+Cyc	10.19	52.15	31.17	49.54	55.98	52.76
P.GAN+Cyc	7.90	61.33	34.62	49.25	55.00	52.12
BEGAN+Cyc	30.05	<u>68.08</u>	49.07	49.04	52.03	50.53
K-DCGAN+Cyc	13.14	61.44	37.29	49.25	55.00	52.12
K-BEGAN+Cyc	22.95	78.09	50.52	49.67	<u>52.80</u>	<u>51.23</u>

Table 5.1: Objective Evaluation :The percentage scores (more is better) and normalized pixel-error percentage (less is better) for base set X, and target set Y. **bold** is best and underlined is second best. First three rows are actual data put for reference. The following six rows are different GANs, the last seven rows are cascaded GANs. Scores are averaged over the 10 different Y sets See Sec.5.1 for details

Contributions between X and Y, we adopt the score for that purpose. We use the inception network[41] softmax probability of the modeled class directly (averaged over all samples), and then average the score over all models learned on the different classes and finally normalize the score by the score of real(Y), and we call this **score(Y)**. For base class assessment (the face), we use face detector accuracy as a metric. We used OpenCV [42] implementation of the famous Viola-Jones detector [43] to measure how our transformed faces retained the properties of the face. The face **score(X)** in table 5.1 is defined as the percentage of images which the detector triggered a face. We use an extra evaluation criterion followed by [27] and [22] by taking the normalized pixel-wise error defined as follows: $\text{err}(\mathbf{x}) = \frac{\|\text{NN}(\mathbf{x})\|_2}{J}$ where $\text{NN}(\mathbf{x})$ is the pixel-wise nearest neighbor to \mathbf{x} in the training, J is the average distance between uniform ran-

dom noise in the training data. An $\text{err}(\mathbf{x}) = 1$ means our samples are no better than random noise, and a $\text{err}(\mathbf{x}) = 0$ indicates we generated a sample from training. For sampling, we follow a similar procedure to [14] by using 1000 samples from different setups. We calculate the scores and error for both X and Y and the average of both on all the samples of each model for all ten classes of Y and the base class X. Table 5.1 summarizes the results.

5.2 Subjective Evaluation

The nature of the task we are addressing, *i.e.* imagination, is subjective and involves human judgment. Therefore, we did extensive human experiments using Amazon Mechanical Turk online tool. We designed two experiments to evaluate the quality and the information conveyed by the samples from the different IAN models respectively.

5.2.1 Quality Pair-wise Comparison

In this experiment, the human subjects were shown human faces and a class from Zoo-Faces and then shown a pair of samples from two different setups. The task was to pick the one which was better looking and represent both the human face and the given class. A total of 15,000 pairs were shown to more than 150 unique human annotators to compare the quality of the generated samples from 6 different setups (4 IANs and 2 GANs) with a total of 15 combinations. The percentages shown in table 5.2 represent the frequency of human subjects picking that model against its rival, averaged over the 10 different classes of Zoo-faces. The higher the percentage of a model the better humans view the samples of that model.

model (id#)	(1)	(2)	(3)	(4)	(5)	(6)	avg	HCA
BEGAN (1)	N/A	25.70	34.50	47.30	34.90	31.80	34.84	N/A
K-BEGAN (2)	74.30	N/A	38.80	38.80	36.90	47.00	47.16	15.35
DCGAN+Cyc (3)	65.50	61.20	N/A	71.60	35.90	32.70	53.38	64.34
K-DCGAN+Cyc (4)	52.70	47.50	28.40	N/A	24.80	20.70	34.82	68.28
BEGAN+Cyc (5)	65.10	63.10	64.10	75.20	N/A	54.40	64.38	57.98
K-BEGAN+Cyc (6)	68.20	53.00	67.30	79.30	45.60	N/A	62.68	63.94

Table 5.2: Human Subjective Evaluation :Pair-wise quality comparison and Human Classification Accuracy (HCA) for six different IANs, and GANs . The Pair-wise percentages represent the frequency that a certain model’s samples were preferred against their rival’s model samples by human subjects .HCA is for a classification task given for the subjects that were asked to classify the samples from different IANs. The higher the HCA the more easily identifiable are the IAN samples. See Sec.5.2 for more details

5.2.2 Human Subjective Classification Experiment

In this experiment , the human subjects were asked to classify the samples generated from each IAN setup .Five setups were trained on 10 different Y target classes from the Zoo-Faces dataset to give a total of 50 IAN models .A hundred samples (like the ones in Fig. 4.2) were drawn from each model to give a total of 5K images as a testing set for the annotators to classify. A total of 25 unique human participants classified the samples to the closest animal from Zoo-Faces classes. The Human Classification Accuracy (HCA) is shown in table 5.2 . The HCA indicates how easily identifiable are these generated samples. Please refer to Fig 5.1 for the Human Confusion Matrix of the K-BEGAN+Cyc setup and Appendix C for other setups.

5.3 Ablation Study

5.3.1 Changing the Training Data

We tried different data other than faces, and by using target class Y that are not animals and very far from the faces. We found that the more significant the difference

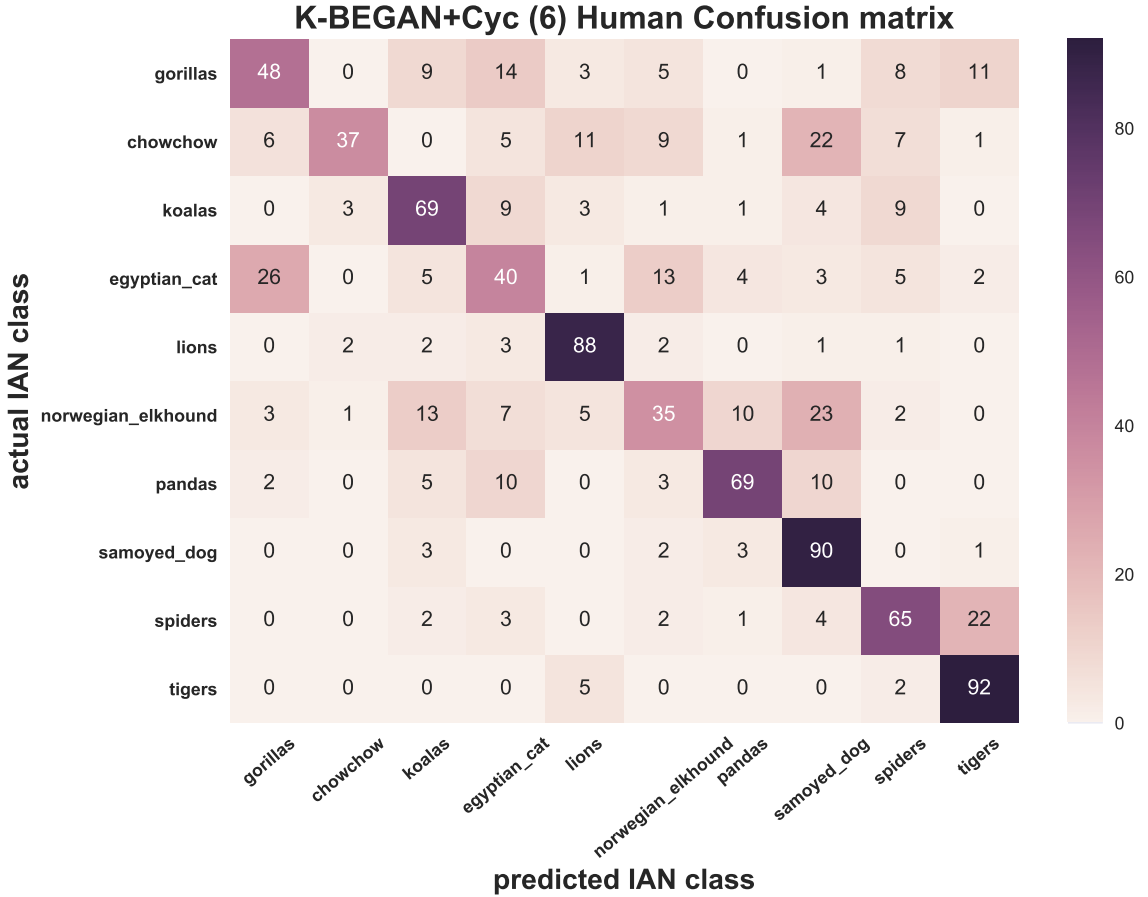


Figure 5.1: Confusion Matrix for K-BEGAN + CycleGAN

is between X and Y the harder it becomes to produce meaningful results from the IAN structure. We show some cases of *hard* imagination (*e.g.* tiger to flower and shoe to tiger) in Fig 5.2(a) with success (maintaining some key X properties after translation to Y) and failure (losing all X properties after translation). We reason that the further the classes in feature space, the possibility for the network to converge to a local minimum increases and hence not reaching to the target distribution.

5.3.2 Abandoning the Search

Picking a random image in the target set and using its features as the target of the regularizer seems to be a valid simpler option. However, randomly picking features

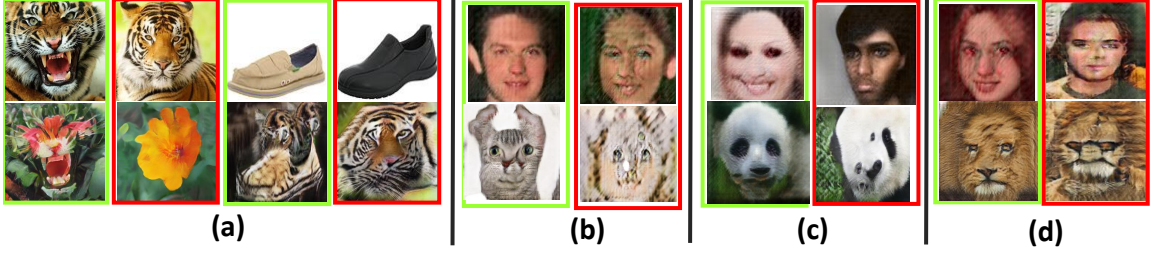


Figure 5.2: **Ablation study:** (a) *Hard Imagination*, trying to translate images from set X (*e.g.* tiger) to very different set Y (*e.g.* flower) using CycleGAN. (b-d) studying the effect of altering parts of the model on the GAN output and its corresponding cascaded sample with CycleGAN (b) abandoning the K-NN search produce high frequency artifacts. (c) abandoning the K-NN feature matching leads to amplifying blurriness after cascade. (d) BEGAN vs DCGAN as base model. in all part red is failure and green is success.

can cause a rapid change in the objective of the GAN and hence introduce high-frequency artifacts as in Fig 5.2(b). These artifacts amplify when the samples pass by CycleGAN, and as results in table 5.1 suggest. Picking KNN features ensures stability in the objective and hence produce better samples.

5.3.3 Abandoning the Regularizer

Using GAN to produce samples of the modeled data X then passing those samples to a trained cycleGAN to transform them and give the final samples seems intuitive. However, as we argued in section 3.2, and as can be seen in Fig 5.2(c), using unregularized GAN and building on top that can lead to systemic failure in which the error can amplify.

5.3.4 Different GAN structures

using the advanced BEGAN[2] instead of DCGAN[34] as base model produced better results from resulting K-GANs and IANs. 5.2(d)

5.4 Discussion

Investigating the results in Fig 4.2 and tables 5.1 and 5.2 , we see that cascading any GAN model by CycleGAN increase the realism of target class Y by 40% in the inception accuracy and be very close from the inception of an actual image .Also, it increases its likeability score in table 5.2 by about 15%. However, it deteriorates the realism of the base class (the human face in our case). adding the K-NN regularizer made it easier for both Inception Network Softmax and humans HCA. The human likeability tends to favour unregularized IAN over regularized IAN , but K-GAN still outperforms GANs.

Chapter 6

Applications

6.1 Multi-Domain Image Manifold Traversing

We follow the standard practice of utilizing the modeled distribution in manifold traversing [22]. But since we have trained different IANs on the same X and different Y , we can do multi-domain traversing for different manifolds with only two images. First, we pick two training samples from X , *e.g.* $\mathbf{x}_1^R, \mathbf{x}_2^R$. Then, we utilize the auto-encoder architecture of the discriminator of the trained BEGAN to encode the two samples and obtain their latent vectors $\mathbf{z}_1^*, \mathbf{z}_2^*$ as follows: $\mathbf{z}^* = \mathbf{D}_{\text{encoder}}(\mathbf{x}^R)$. Then, by taking N convex points between these two latent vectors, we obtain $\mathbf{z}_i, i \in [1, \dots, N]$ and find the corresponding decoded images as $\mathbf{D}_{\text{decoder}}(\mathbf{z}_i)$. By translating these samples by CycleGAN, we can traverse different Y manifolds with only two samples from X as in Fig 6.1. More qualitative results of this application are provided in Appendix E.

6.2 Celebrity Zoo

Imagine going to a zoo full of hybrid human-animals. We try doing this by using the IAN models trained in Sec.4.2. We take a collection of celebrities from CelebA and encode them using the BEGAN Discriminator auto-encoder as in Sec.6.1, and passing the decodes to the CycleGAN for different animals from the Zoo-Faces dataset, obtaining animal (and insect) counterpart for each celebrity in Fig 6.2.

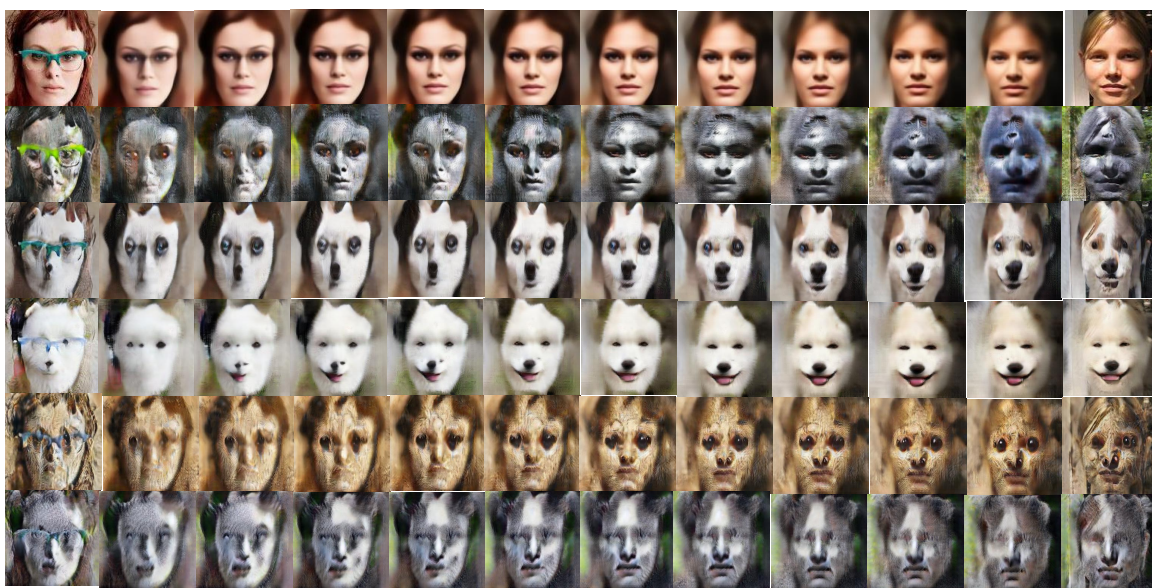


Figure 6.1: Multi-domain image manifold traversing: moving from one point to another in different Y classes using the same pair in base set X . The top row is the actual traversing by convexing the codes obtained by the auto-encoder of BEGAN [2] and then decoding the convex codes. The rest of the rows are translations of the first row but using different CycleGANs.

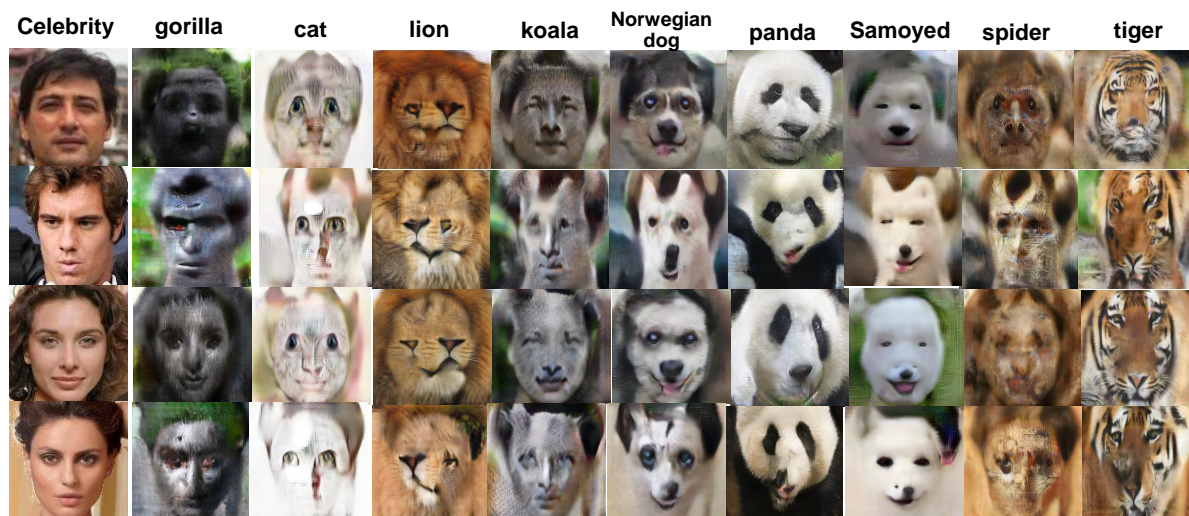


Figure 6.2: Celebrity-Zoo: seeing how celebrities look in the animal kingdom by using different IAN models trained in Sec.4.2.

Chapter 7

Concluding Remarks

7.1 Summary

Deep neural networks, cheap parallel computing , and abundance of labelled data shaped the new wave of AI in the recent years and helped advancing the fields of computer vision -and to lesser extent- natural language processing. The availability of abundant densely labelled low-noise data and transferability of learned models to the wild remain unsolved problems. One approach is to use synthetic data and simulations to learn models that will eventually be used in the real world , but domain adaptation is substantial obstacle to this approach. Another direction is to utilize the data without labels , *i.e.* unsupervised learning , in learning valuable information and put it in use to develop different applications. GAN is one unsupervised approach that attracts the attention of the AI community for its abilities to model image distributions and generate sharp images, along with its uses in many image processing and vision applications.

In this work , a new regularizer is proposed for GANs based on K-NN feature matching in GAN setup with two different sets (base X, target Y) and it was shown how that helps in obtaining key properties of Y while keeping the GAN modeling of X. The added regularizer follows mathematically from the cross entropy distance between the distributions of Y and the GAN distribution.

Also, a framework for cascading GANs (*i.e.* IAN) is presented to address the task of imagining new distributions that combine base set X and target set Y. We

experiment with different IAN setups and show samples of one such setup that uses K-GAN and CycleGAN in which a filtered dataset was collected for the training. A series of objective and subjective evaluations were conducted on different IANs and bench-marked against baselines. From the results, we note that cascading any GAN model by CycleGAN increase the realism of target class Y by 40% in the inception accuracy and by about 15% in the likeability average score in the subjective evaluation. Furthermore, while the regularized IAN comprising of K-GAN and CycleGAN was easily identifiable by Inception networks and humans , humans favoured Vanilla IAN in pair-wise comparison. We conclude our work by showing some applications for the IAN framework, like multi-domain manifold traversing and Celebrity-ZOO.

7.2 Future Research Work

Extending the IAN for more than two two sets , rather a pool of different sets, follow directly from our work. A potential use is to utilize IAN in zero-shot learning in which we learn from the generated samples based on some priors about the new class of unseen samples.

A more interesting direction that can be discovered is to use the innovative adversarial generation as part of training another task , *e.g.* imitation learning or segmentation , in which the performances of both generative and discriminative systems are increasing with time and both systems are optimized to beat each other . If the systems were differentiable , then gradient descent optimization is to be followed . If there is a non-differentiable part of the system ,two approaches can be followed. Montecarlo search for each update of that non-differentiable subpart , or adding another differentiable abstracting system. The goal of the abstracting system is to simplify the output of the generating system and bypass any non-differentiable part and to be optimized to reduce the loss of both discriminating and generating systems.

REFERENCES

- [1] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” *arXiv preprint arXiv:1703.10593*, 2017.
- [2] D. Berthelot, T. Schumm, and L. Metz, “BEGAN: boundary equilibrium generative adversarial networks,” *CoRR*, vol. abs/1703.10717, 2017.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2012, pp. 1097–1105.
- [4] S. Ren, K. He, R. Girshick, and J. Sun, “Faster r-cnn: Towards real-time object detection with region proposal networks,” in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 91–99.
- [5] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. S. Bernstein, A. C. Berg, and F. Li, “Imagenet large scale visual recognition challenge,” *CoRR*, vol. abs/1409.0575, 2014.
- [6] A. Torralba and A. A. Efros, “Unbiased look at dataset bias,” in *CVPR*, 2011.
- [7] A. Shrivastava, T. Pfister, O. Tuzel, J. Susskind, W. Wang, and R. Webb, “Learning from simulated and unsupervised images through adversarial training,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [8] C. R. de Souza, A. Gaidon, Y. Cabon, and A. M. L. Peña, “Procedural generation of videos to train deep action recognition networks,” *CoRR*, vol. abs/1612.00881, 2016.
- [9] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid, “Learning from Synthetic Humans,” in *CVPR*, 2017.
- [10] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes,

- N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672–2680.
- [11] M. Arjovsky, S. Chintala, and L. Bottou, “Wasserstein gan,” 2017, arXiv:1701.07875.
 - [12] E. L. Denton, S. Chintala, a. szlam, and R. Fergus, “Deep generative image models using a laplacian pyramid of adversarial networks,” in *Advances in Neural Information Processing Systems 28*, C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, Eds. Curran Associates, Inc., 2015, pp. 1486–1494.
 - [13] S. Gurumurthy, R. Kiran Sarvadevabhatla, and R. Venkatesh Babu, “Deligan : Generative adversarial networks for diverse and limited data,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
 - [14] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen, “Improved techniques for training gans,” in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 2234–2242.
 - [15] J. T. Springenberg, “Unsupervised and semi-supervised learning with categorical generative adversarial network,” in *ICLR*, 2016.
 - [16] D. Warde-Farley and Y. Bengio, “Improving generative adversarial networks with denoising feature matching,” in *CVPR*, 2017.
 - [17] J. Zhao, M. Mathieu, and Y. LeCun, “Energy-based generative adversarial network,” *arXiv preprint arXiv:1609.03126*, 2016.
 - [18] R. Yeh, C. Chen, T. Lim, M. Hasegawa-Johnson, and M. N. Do, “Semantic image inpainting with perceptual and contextual losses,” *CoRR*, vol. abs/1607.07539, 2016.
 - [19] J. Johnson, A. Alahi, and F. fei Li, “Perceptual losses for real-time style transfer and super-resolution,” in *ECCV*, 2016.
 - [20] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” *arxiv*, 2016.
 - [21] C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi, “Photo-realistic single image super-resolution using a generative adversarial network,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

- [22] J.-Y. Zhu, P. Krähenbühl, E. Shechtman, and A. A. Efros, “Generative visual manipulation on the natural image manifold,” in *Proceedings of European Conference on Computer Vision (ECCV)*, 2016.
- [23] C. Wan, T. Probst, L. Van Gool, and A. Yao, “Crossing nets: Combining gans and vaes with a shared latent space for hand pose estimation,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [24] L. Tran, X. Yin, and X. Liu, “Disentangled representation learning gan for pose-invariant face recognition,” in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.
- [25] M. Liu and O. Tuzel, “Coupled generative adversarial networks,” *CoRR*, vol. abs/1606.07536, 2016.
- [26] A. M. Elgammal, B. Liu, M. Elhoseiny, and M. Mazzone, “CAN: creative adversarial networks, generating ”art” by learning about styles and deviating from style norms,” *CoRR*, vol. abs/1706.07068, 2017.
- [27] A. Dosovitskiy and T. Brox, “Generating images with perceptual similarity metrics based on deep networks,” in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 658–666.
- [28] A. Nguyen, A. Dosovitskiy, J. Yosinski, T. Brox, and J. Clune, “Synthesizing the preferred inputs for neurons in neural networks via deep generator networks,” in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 3387–3395.
- [29] A. Nguyen, J. Yosinski, Y. Bengio, A. Dosovitskiy, and J. Clune, “Plug & play generative networks: Conditional iterative generation of images in latent space,” *CoRR*, vol. abs/1612.00005, 2016.
- [30] J. J. Zhao, M. Mathieu, and Y. LeCun, “Energy-based generative adversarial network,” *CoRR*, vol. abs/1609.03126, 2016.
- [31] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *CoRR*, vol. abs/1409.1556, 2014.
- [32] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” in *Intelligent Signal Processing*. IEEE Press, 2001, pp. 306–351.

- [33] J. Lu, K. Sunkavalli, N. Carr, S. Hadap, and D. A. Forsyth, “A visual representation for editing face images,” *CoRR*, vol. abs/1612.00522, 2016.
- [34] A. Radford, L. Metz, and S. Chintala, “Unsupervised representation learning with deep convolutional generative adversarial networks,” *CoRR*, vol. abs/1511.06434, 2015.
- [35] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, “CNN features off-the-shelf: an astounding baseline for recognition,” *CoRR*, vol. abs/1403.6382, 2014.
- [36] Z. Liu, P. Luo, X. Wang, and X. Tang, “Deep learning face attributes in the wild,” in *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- [37] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014.
- [38] T. Liu, A. W. Moore, and A. Gray, “New algorithms for efficient high-dimensional nonparametric classification,” *Journal of Machine Learning Research*, vol. 7 (2006) 11351158, 2006.
- [39] E. Jones, T. Oliphant, P. Peterson *et al.*, “SciPy: Open source scientific tools for Python,” 2001–, [Online; accessed jtoday]. [Online]. Available: <http://www.scipy.org/>
- [40] L. Theis, A. van den Oord, and M. Bethge, “A note on the evaluation of generative models,” in *ICLR*, 2016.
- [41] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, “Going deeper with convolutions,” in *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [42] Itseez, “Open source computer vision library,” <https://github.com/itseez/opencv>, 2015.
- [43] P. Viola and M. Jones, “Rapid object detection using a boosted cascade of simple features,” in *CVPR*, 2001.
- [44] C. Z., “Convexity, jensens inequality,” 2012, in: *Inequalities*. Springer, Berlin, Heidelberg.

APPENDICES

A Deriving the K-NN Loss in Chapter 3

A.1 Deriving the K-NN Loss

We would like to show how to reach our final objective in (A.1) from the initial expression in (A.2)

$$\min_{\mathbf{G}} \max_{\mathbf{D}_X} L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) + \mu L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y), \quad (\text{A.1})$$

$$\min_{\mathbf{G}} \max_{\mathbf{D}_X} L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) + \mu \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)] \quad (\text{A.2})$$

for Generator \mathbf{G} and discriminator \mathbf{D}_X of modeled base set X with distribution \mathbf{P}_X and latent vector $\mathbf{z} \in \mathbb{R}^d$, the GAN loss $L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X)$ is given by

$$L_{\text{GAN}}(\mathbf{G}, \mathbf{D}_X, \mathbf{P}_X) = \mathbb{E}_{\mathbf{x} \sim p_x(\mathbf{x})} [\log \mathbf{D}_X(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - \mathbf{D}_X(\mathbf{G}(\mathbf{z})))] \quad (\text{A.3})$$

For Target set Y with distribution \mathbf{P}_Y , the K-nearest neighbor loss $L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y)$ is defined as follows:

$$L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y) = \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \sum_{n=1}^k \|\mathbf{C}(\mathbf{G}(\mathbf{z})) - \mathbf{r}_n(\mathbf{C}(\mathbf{G}(\mathbf{z})))\|_2^2 \quad (\text{A.4})$$

This is the distance between a generated sample from \mathbf{G} by a latent vector \mathbf{z} and its K nearest neighbors samples of the target distribution \mathbf{P}_Y . This distance is not pixels, but in high-level feature space defined by the Comparitor network $\mathbf{C}(\mathbf{b}) \in \mathbb{R}^s$ which takes an image $\mathbf{b} \in \mathbb{R}^{h \times w}$ as input and outputs feature vector. $\mathbf{r}_n(\mathbf{c})$ is the

K-nearest neighbor function, a parametric order selection function that selects an element \mathbf{y}_i from a set of $M - n + 1$ elements based on how close it is to input vector \mathbf{c} . It can be described by the following:

$$\mathbf{r}_n(\mathbf{c}) = \arg \min_{\mathbf{y}_i \in \psi_n} \|\mathbf{c} - \mathbf{y}_i\|_2, \quad (\text{A.5})$$

where ψ_n is the set of deep features of $M - n + 1$ images representing the target distribution \mathbf{P}_Y . The total number of samples features we have for \mathbf{P}_Y is M , where the function $\mathbf{r}_n(\mathbf{c})$ selects the nearest n^{th} feature out of the remaining furthest $M-n+1$ features after removing the nearest $n-1$ from the global set of all features \geq .

For the Cross-entropy in (A.2), we use the following definition of cross entropy H between two distributions p, q as follows :

$$H(p, q) = \mathbb{E}_p[-\log q] \quad (\text{A.6})$$

To show the derivation, we first look into \mathbf{P}_Y . Since acquiring a full description of \mathbf{P}_Y in deep feature space is infeasible in most cases, we settle to approximate it using M target samples. We can get a proxy distribution of \mathbf{P}_Y by using Kernel Density Estimation (KDE) on the deep features extracted by deep network $\mathbf{C}(\mathbf{x})$ from those M samples in \mathbf{P}_Y . By picking Gaussian kernel for the KDE, the proxy distribution estimator $\hat{\mathbf{P}}_{Y,\text{proxy}}$ is defined by a Bayesian non-parametric way at any point $\mathbf{C}(\mathbf{b}) \in \mathbb{R}^s$ for $\mathbf{b} \in \mathbb{R}^{\mathbf{h} \times \mathbf{w}}$ as follows:

$$\hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) = \frac{1}{M\sigma\sqrt{2\pi}} \sum_{i=1}^M \exp\left(-\frac{\|\mathbf{C}(\mathbf{b}) - \mathbf{C}(\mathbf{y}_i)\|_2^2}{\sigma^2}\right) \quad (\text{A.7})$$

\mathbf{y}_i is the i^{th} sample describing \mathbf{P}_Y and σ^2 is the Gaussian kernel variance. For far images from the M samples of $\hat{\mathbf{P}}_{Y,\text{proxy}}$ the probability becomes exponentially negligible. Hence, we can investigate the effect of a subset of K points (out of the

M points $\mathbf{C}(\mathbf{y}_i)$ that are closest to $\mathbf{C}(\mathbf{b})$ to compute (A.7). We replace M by K in (A.7), and pick $\sigma = 1$ for simplicity to get the following :

$$\hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) \approx \frac{1}{M\sqrt{2\pi}} \sum_{i=1}^K \exp(-\|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{b}))\|_2^2) \quad (\text{A.8})$$

where \mathbf{r}_i is described in (A.5). We use a finite discrete special form of the Jansen inequality described by Theorem 7.3 of ([44]) as follows

$$f\left(\sum_{i=1}^n u_i \lambda_i\right) \leq \sum_{i=1}^n f(u_i) \lambda_i \quad (\text{A.9})$$

where f is convex function for any u_i , λ_i are set of n weights with $\sum_{i=1}^n \lambda_i = 1$.

By Picking $f(\mathbf{x}) = \exp(-\mathbf{x})$ a convex function defined in $[0, \infty)^d$, and picking $u_i(\mathbf{b}) = \|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{b}))\|_2^2$ for $\mathbf{b} \in \mathbb{R}^{h \times w}$, \mathbf{r}_i is just like in (A.5) and \mathbf{C} is like in (A.4), and picking $\lambda_i = \frac{1}{K}$, and $n = K$ inequality (A.9) becomes :

$$\exp\left(-\frac{1}{K} \sum_{i=1}^K \|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{b}))\|_2^2\right) \leq \frac{1}{K} \sum_{i=1}^K \exp(-\|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{b}))\|_2^2) \quad (\text{A.10})$$

Noting that the right-hand side of (A.10) is scaled version of the proxy distribution approximate $\hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b})$ described in (A.8), we reach to the following inequality

$$\exp\left(-\frac{1}{K} \sum_{i=1}^K \|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{y}_i))\|_2^2\right) \leq \frac{M\sqrt{2\pi}}{K} \hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) \quad (\text{A.11})$$

Taking the natural logarithm (a monotonically non-decreasing function) of both sides of (A.11) and then negating both sides results in the following inequality :

$$-\log \left(\frac{M\sqrt{2\pi}}{K} \hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) \right) \leq \frac{1}{K} \sum_{i=1}^K \|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{y}_i))\|_2^2 \quad (\text{A.12})$$

Rearranging (A.12) to get the following:

$$-\log \hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{b}) \leq \log \frac{M\sqrt{2\pi}}{K} + \frac{1}{K} \sum_{i=1}^K \|\mathbf{C}(\mathbf{b}) - \mathbf{r}_i(\mathbf{C}(\mathbf{y}_i))\|_2^2 \quad (\text{A.13})$$

By replacing the \mathbf{b} by the output of the generator $\mathbf{G}(\mathbf{z})$ sampled from latent vector \mathbf{z} and taking the expectation, we get the following expected upper bound

$$\begin{aligned} & \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [-\log \hat{\mathbf{P}}_{Y,\text{proxy}}(\mathbf{G}(\mathbf{z}))] \\ & \leq \log \frac{M\sqrt{2\pi}}{K} + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \left[\frac{1}{K} \sum_{i=1}^K \|\mathbf{C}(\mathbf{G}(\mathbf{z})) - \mathbf{r}_i(\mathbf{C}(\mathbf{y}_i))\|_2^2 \right] \end{aligned} \quad (\text{A.14})$$

We note that the left side is the cross entropy between \mathbf{G}, \mathbf{P}_Y , and the second term in the right hand side is scaled version of the KNN loss $L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y)$ in (A.4), to get :

$$\mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)] \leq \log \frac{M\sqrt{2\pi}}{K} + \frac{M\sqrt{2\pi}}{K} L_{\text{KNN}}(\mathbf{G}, \mathbf{P}_Y) \quad (\text{A.15})$$

The expression in (A.15) gives upper bound for the cross entropy using the KNN loss. By minimizing the KNN loss for parameters of \mathbf{G} as in (A.1) we insure minimizing the cross entropy $\mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [H(\mathbf{G}(\mathbf{z}), \mathbf{P}_Y)]$ as in (A.2) and hence establish mathematical justification for using KNN loss instead of cross entropy.

B Justifying Regressing to the Mean of the K-features Instead of Each Independently During the Training of K-GAN in Chapter 4

B.1 Justifying Regressing to the Mean of the K-features

We would like to show that the equivalence of :

$$\arg \min_{\mathbf{w}} \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \sum_{i=1}^k \|\mathbf{C}(\mathbf{G}(\mathbf{z}, \mathbf{w})) - \mathbf{d}_i\|_2^2 \quad (\text{B.1})$$

$$\arg \min_{\mathbf{w}} \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} \left\| \mathbf{C}(\mathbf{G}(\mathbf{z}, \mathbf{w})) - \frac{1}{k} \sum_{i=1}^k \mathbf{d}_i \right\|_2^2 \quad (\text{B.2})$$

where \mathbf{C} is the feature extraction network , $\mathbf{G}(\mathbf{z}, \mathbf{w})$ is the Generator network with parameters \mathbf{w} and random sampling latent vector \mathbf{z} sampled from uniform distribution, and \mathbf{d}_i 's are the K chosen features in the time step t .

By starting from Eq.(B.1) and using the notation $\mathbf{f}(\mathbf{w})$ instead of $\mathbf{C}(\mathbf{G}(\mathbf{z}, \mathbf{w}))$, and $\bar{\mathbf{d}}$

as the mean $\frac{1}{k} \sum_{i=1}^k \mathbf{d}_i$, Eq.(B.1) can be expanded as the following

$$\begin{aligned}
& \arg \min_{\mathbf{w}} \sum_{i=1}^k \|\mathbf{f}(\mathbf{w}) - \mathbf{d}_i\|_2^2 \\
&= \arg \min_{\mathbf{w}} \frac{1}{k} \sum_{i=1}^k \|\mathbf{f}(\mathbf{w}) - \mathbf{d}_i\|_2^2 \\
&= \arg \min_{\mathbf{w}} \frac{1}{k} \sum_{i=1}^k \langle \mathbf{f}(\mathbf{w}), \mathbf{f}(\mathbf{w}) \rangle - \frac{2}{k} \sum_{i=1}^k \langle \mathbf{f}(\mathbf{w}), \mathbf{d}_i \rangle + \frac{1}{k} \sum_{i=1}^k \langle \mathbf{d}_i, \mathbf{d}_i \rangle \quad (\text{B.3}) \\
&= \arg \min_{\mathbf{w}} \langle \mathbf{f}(\mathbf{w}), \mathbf{f}(\mathbf{w}) \rangle - 2 \langle \mathbf{f}(\mathbf{w}), \bar{\mathbf{d}} \rangle + \frac{1}{k} \sum_{i=1}^k \langle \mathbf{d}_i, \mathbf{d}_i \rangle \\
&= \arg \min_{\mathbf{w}} \langle \mathbf{f}(\mathbf{w}), \mathbf{f}(\mathbf{w}) \rangle - 2 \langle \mathbf{f}(\mathbf{w}), \bar{\mathbf{d}} \rangle
\end{aligned}$$

where $\langle \cdot, \cdot \rangle$ is the inner product and after utilizing the property $\sum_i \langle \mathbf{a}, \mathbf{d}_i \rangle = \langle \mathbf{a}, \sum_i \mathbf{d}_i \rangle$. On the other hand, Eq.(B.2) can be expanded (with the new notation) as the following

$$\begin{aligned}
& \arg \min_{\mathbf{w}} \|\mathbf{f}(\mathbf{w}) - \bar{\mathbf{d}}\|_2^2 \\
&= \arg \min_{\mathbf{w}} \langle \mathbf{f}(\mathbf{w}), \mathbf{f}(\mathbf{w}) \rangle - 2 \langle \mathbf{f}(\mathbf{w}), \bar{\mathbf{d}} \rangle + \langle \bar{\mathbf{d}}, \bar{\mathbf{d}} \rangle \quad (\text{B.4}) \\
&= \arg \min_{\mathbf{w}} \langle \mathbf{f}(\mathbf{w}), \mathbf{f}(\mathbf{w}) \rangle - 2 \langle \mathbf{f}(\mathbf{w}), \bar{\mathbf{d}} \rangle
\end{aligned}$$

We can see that the final expressions in Eq.(B.3) and Eq.(B.4) are equivalent. Hence , the equivalence between Eq.(B.1) and Eq.(B.2) is established , and picking the mean of the K-NN features at each iteration is mathematically justified , while improving the speed of the training of the K-GAN.

C Confusion Matrices for Human Classification Experiment

C.1 Confusion Matrices for Human classification experiment

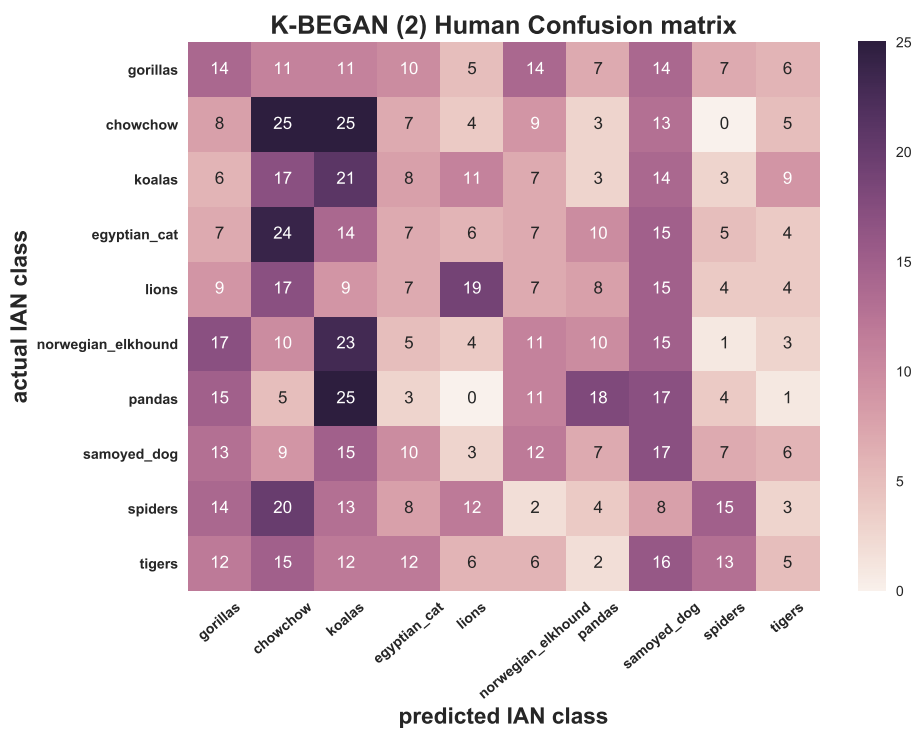


Figure C.1: Confusion Matrix for K-BEGAN

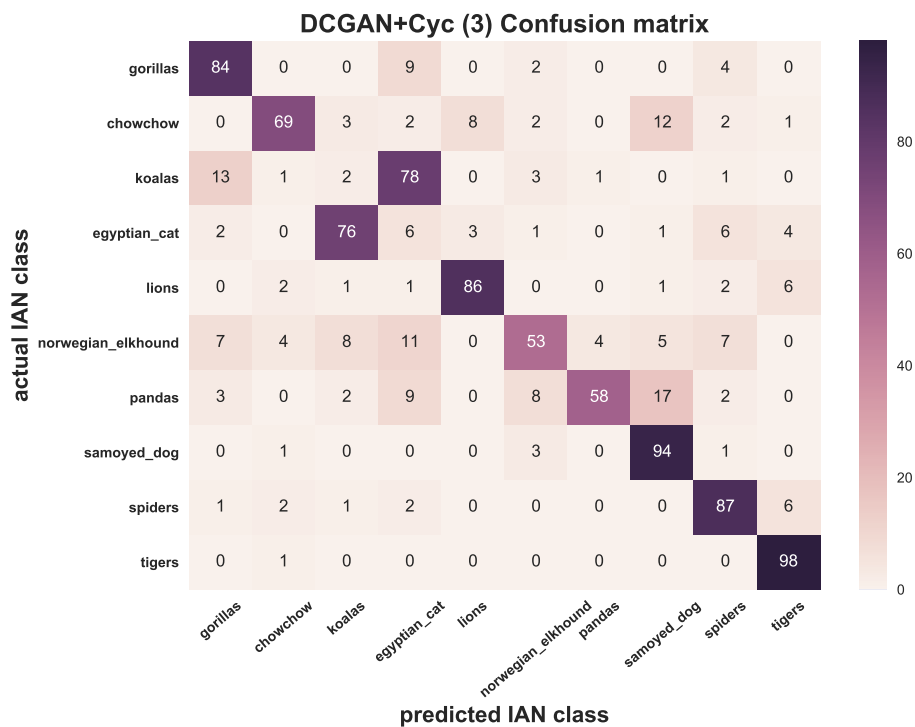


Figure C.2: Confusion Matrix for DCGAN + CycleGAN

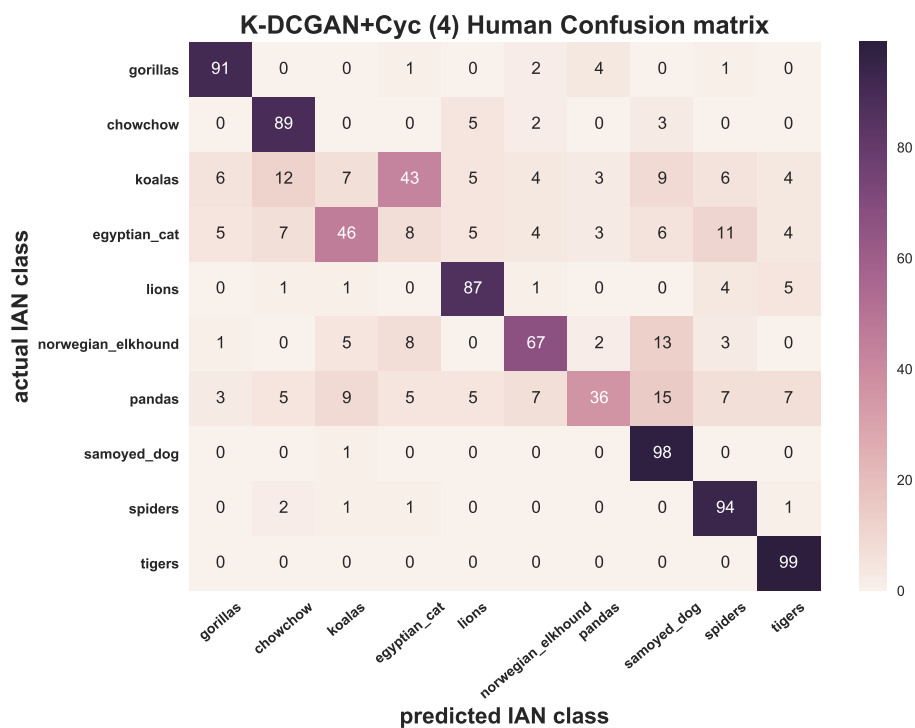


Figure C.3: Confusion Matrix for K-DCGAN + CycleGAN

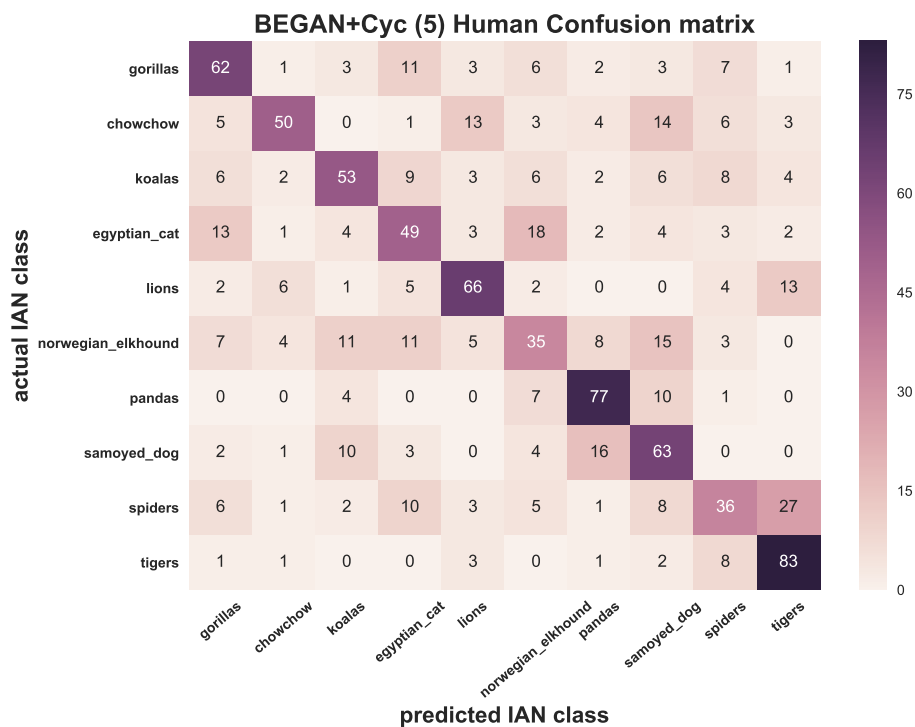


Figure C.4: Confusion Matrix for BEGAN + CycleGAN

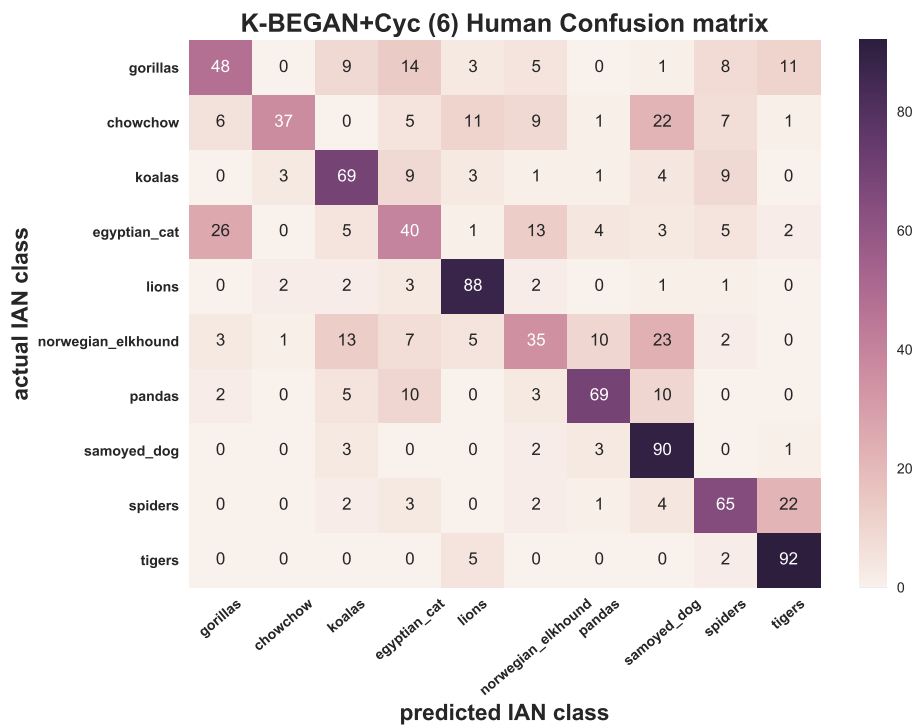


Figure C.5: Confusion Matrix for K-BEGAN + CycleGAN

D Samples of the Zoo-Faces Dataset used in the Training

D.1 Zoo-Faces Dataset used in the Training



Figure D.1: Samples of class apes the Zoo-Faces dataset



Figure D.2: Samples of class chowchow in the Zoo-Faces dataset

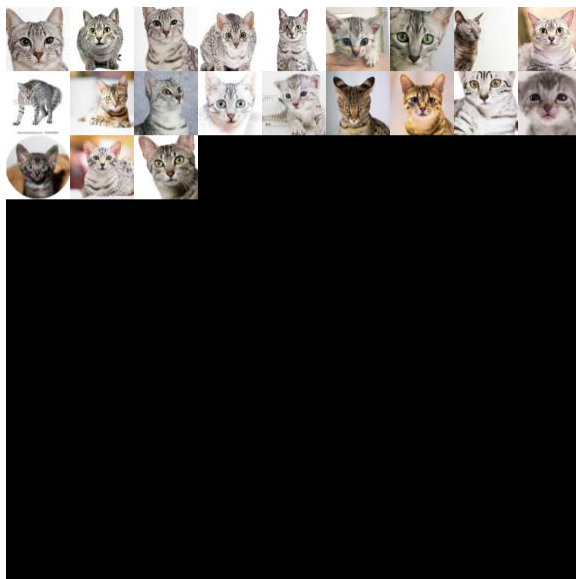


Figure D.3: Samples of class egyptian-cat in the Zoo-Faces dataset

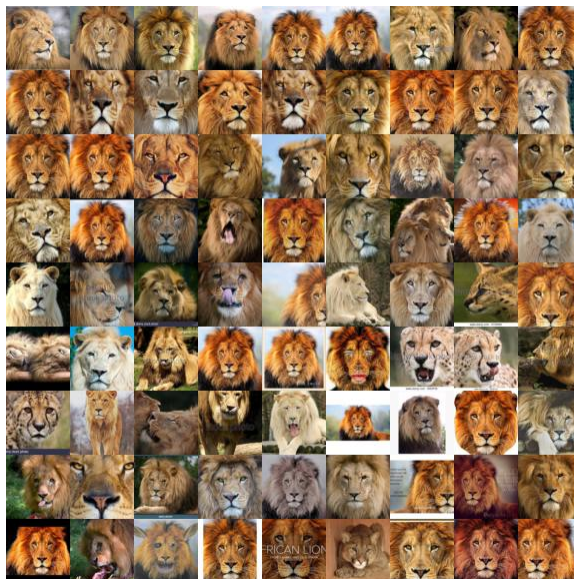


Figure D.4: Samples of class lions in the Zoo-Faces dataset



Figure D.5: Samples of class koalas in the Zoo-Faces dataset



Figure D.6: Samples of class norwegian elkhound in the Zoo-Faces dataset



Figure D.7: Samples of class pandas in the Zoo-Faces dataset



Figure D.8: Samples of class samoyed dog in the Zoo-Faces dataset

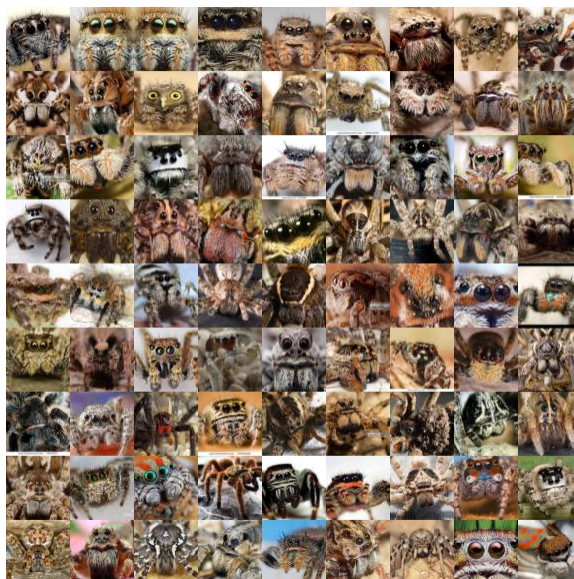


Figure D.9: Samples of class spiders in the Zoo-Faces dataset



Figure D.10: Samples of class tigers in the Zoo-Faces dataset

E More Qualitative Results

E.1 Samples of the K-BEGAN, K-BEGAN+CycleGAN Models



Figure E.1: Samples of K-BEGAN model trained on Apes and faces

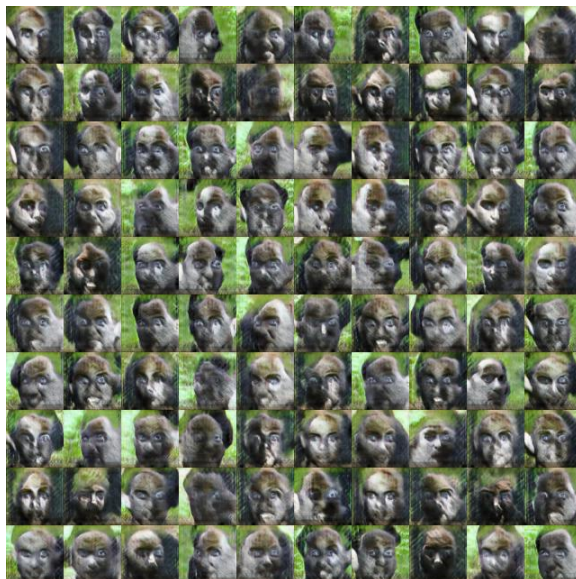


Figure E.2: Samples of K-BEGAN+CycleGAN model trained on Apes and faces



Figure E.3: Samples of K-BEGAN model trained on egyptian-cat and faces



Figure E.4: Samples of K-BEGAN+CycleGAN model trained on egyptian-cat and faces

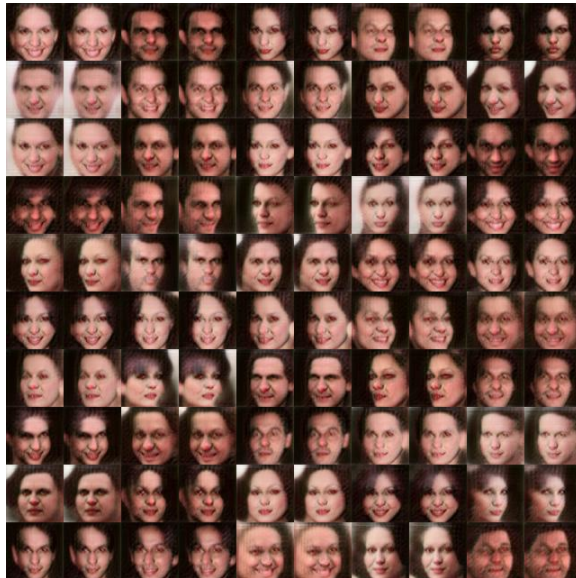


Figure E.5: Samples of K-BEGAN model trained on chowchow and faces

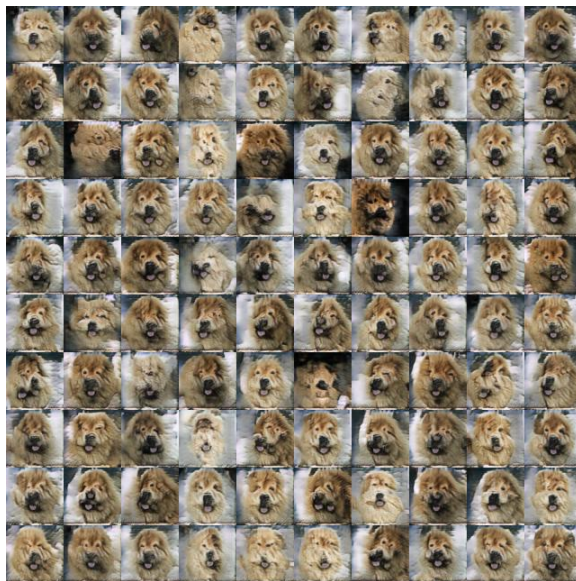


Figure E.6: Samples of K-BEGAN+CycleGAN model trained on chowchow and faces



Figure E.7: Samples of K-BEGAN model trained on koalas and faces

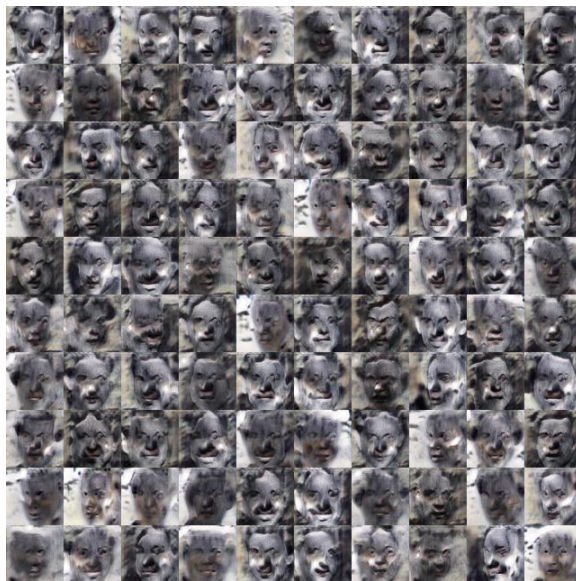


Figure E.8: Samples of K-BEGAN+CycleGAN model trained on koalas and faces



Figure E.9: Samples of K-BEGAN model trained on lions and faces

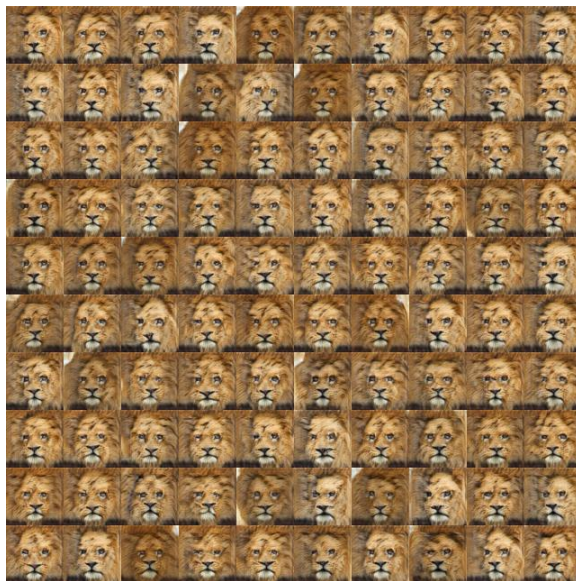


Figure E.10: Samples of K-BEGAN+CycleGAN model trained on lions and faces

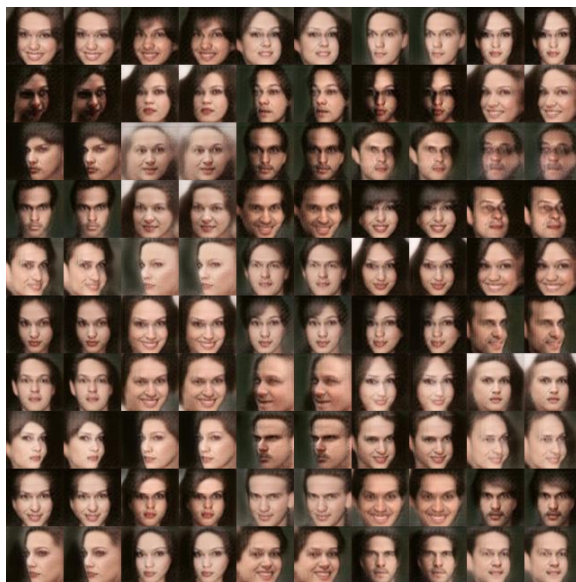


Figure E.11: Samples of K-BEGAN model trained on norwegian elkhound and faces

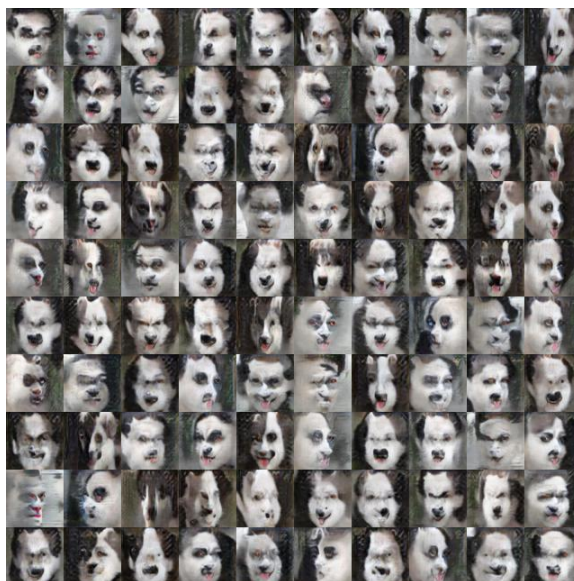


Figure E.12: Samples of K-BEGAN+CycleGAN model trained on norwegian elkhound and faces




Figure E.13: Samples of K-BEGAN model trained on pandas and faces



Figure E.14: Samples of K-BEGAN+CycleGAN model trained on pandas and faces

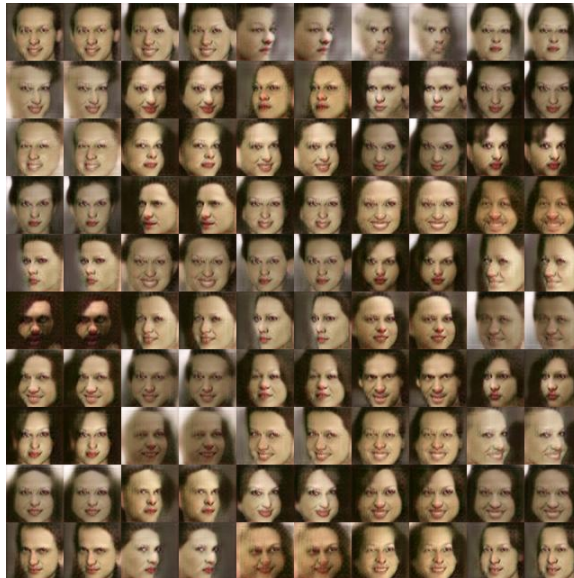


Figure E.15: Samples of K-BEGAN model trained on samoyed dog and faces

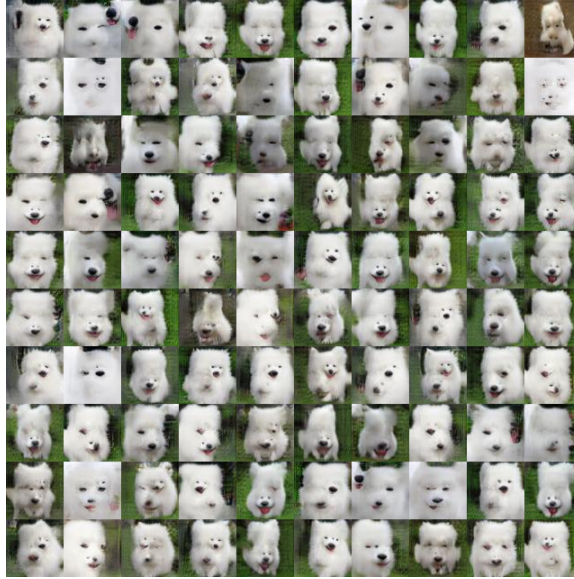


Figure E.16: Samples of K-BEGAN+CycleGAN model trained on samoyed dog and faces

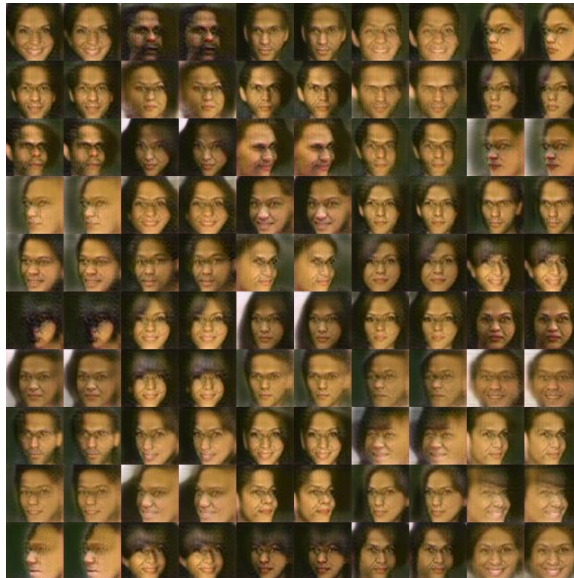


Figure E.17: Samples of K-BEGAN model trained on spiders and faces

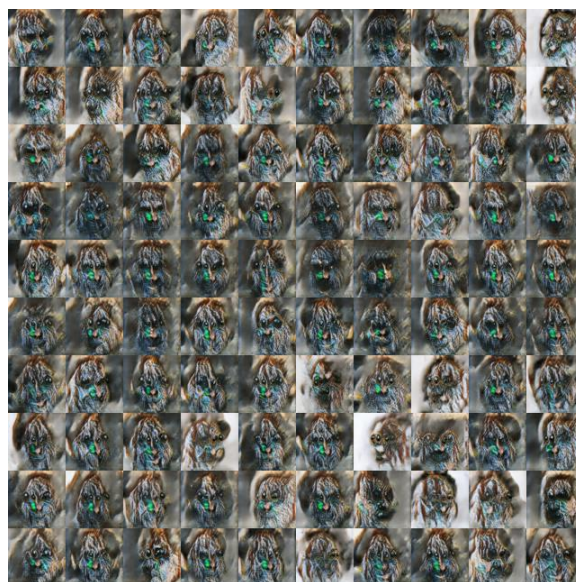


Figure E.18: Samples of K-BEGAN+CycleGAN model trained on spiders and faces

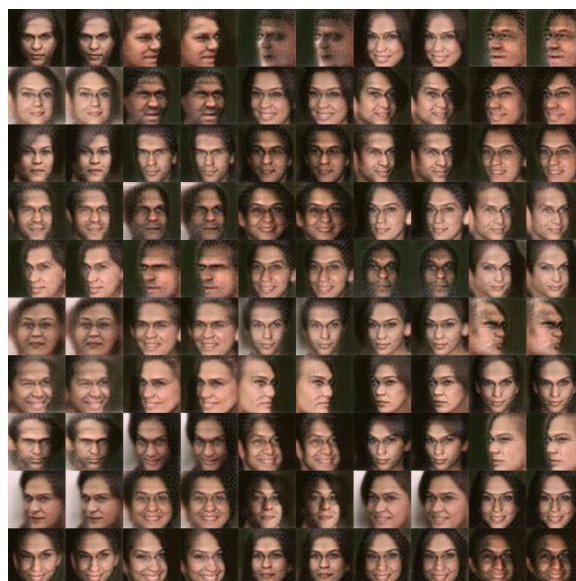


Figure E.19: Samples of K-BEGAN model trained on tigers and faces



Figure E.20: Samples of K-BEGAN+CycleGAN model trained on tigers and faces

E.2 Full Celebrity Zoo

We show more examples of celebrity zoo in the following page.

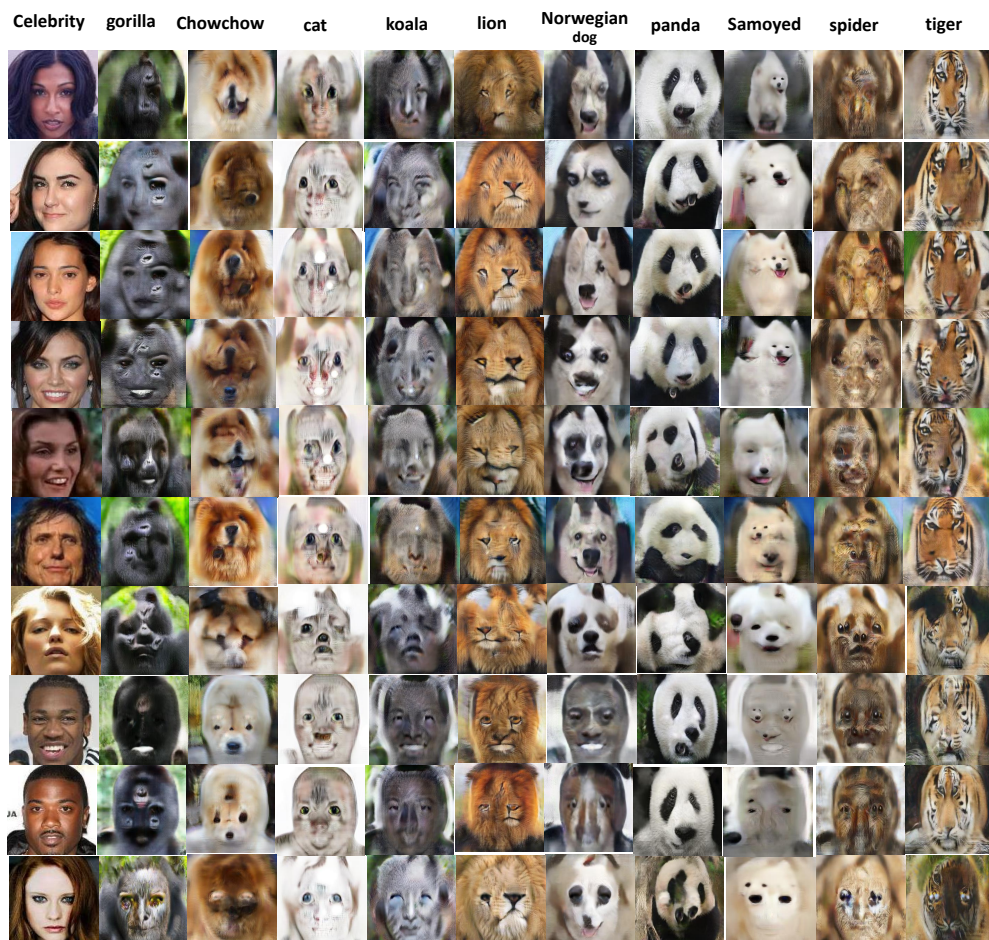


Figure E.21: More Celebrity-Zoo: seeing how celebrities look in the animal kingdom by using different IAN models trained in Sec.4.2.

E.3 image Manifold Traversing

We do image manifold traversing for more classes.

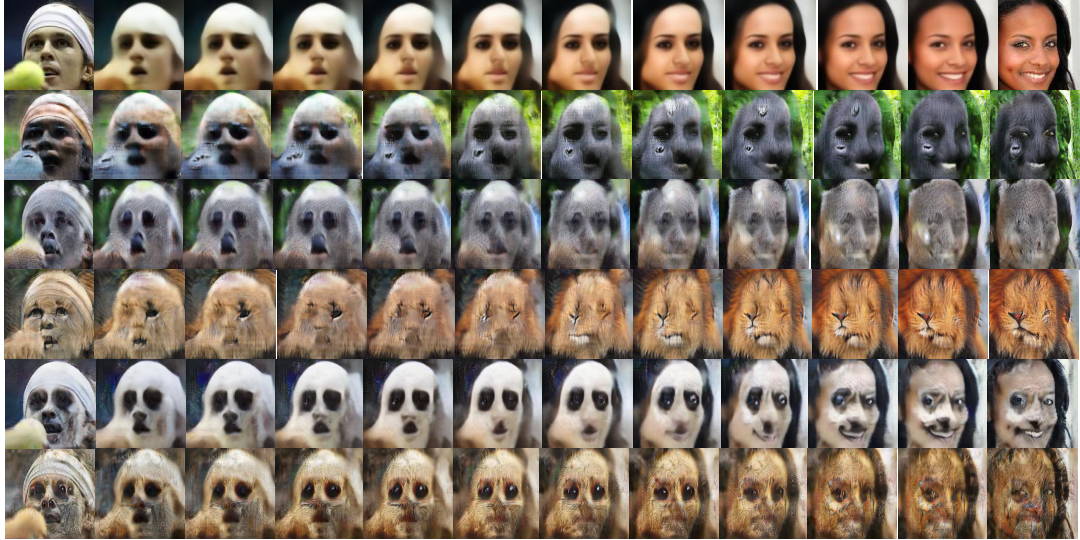


Figure E.22: multi-domain image manifold traversing 1: moving from one point to another in different Y classes using the same pair in base set X

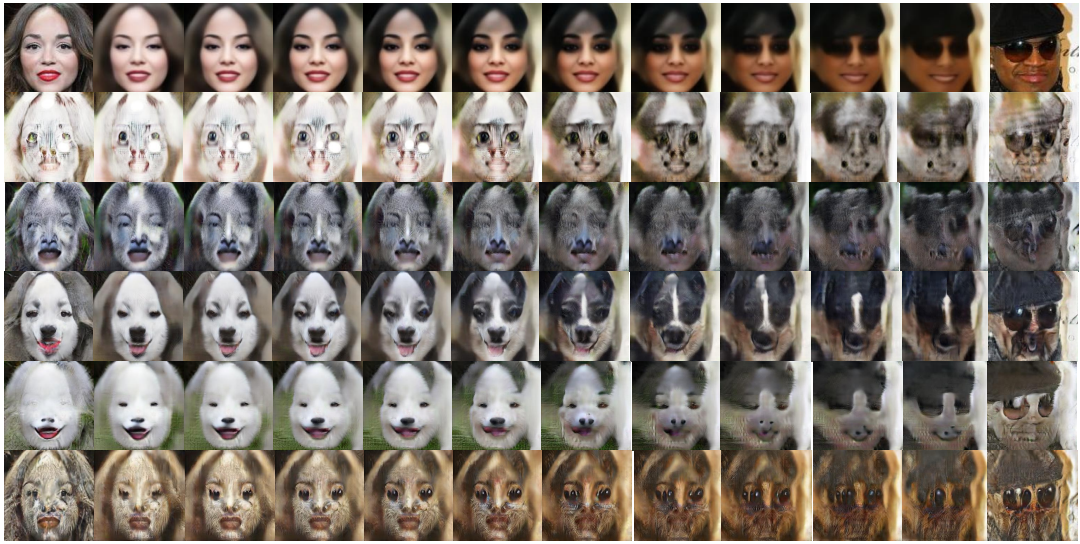


Figure E.23: multi-domain image manifold traversing 2: moving from one point to another in different Y classes using the same pair in base set X

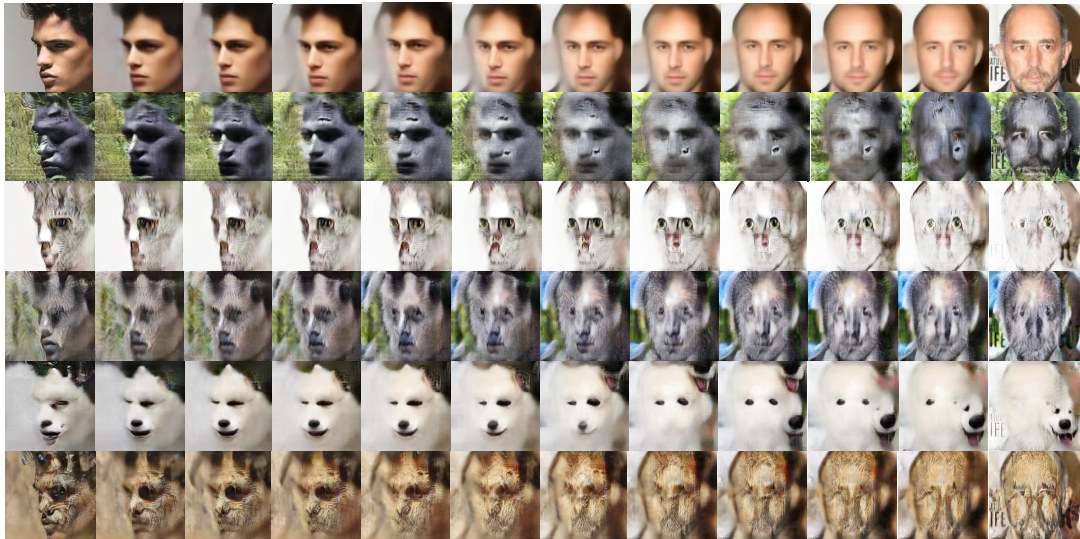


Figure E.24: multi-domain image manifold traversing 3: moving from one point to another in different Y classes using the same pair in base set X



Figure E.25: multi-domain image manifold traversing 4: moving from one point to another in different Y classes using the same pair in base set X



Figure E.26: multi-domain image manifold traversing 5: moving from one point to another in different Y classes using the same pair in base set X

F Papers Submitted and Under Preparation

- Abdullah Hamdi, Bernard Ghanem “IAN: Cascading Generative Adversarial Networks for Targeted Imagination”, *Submitted to European Conference of Computer Vision (ECCV), March. 2018*