



Provided by the author(s) and University of Galway in accordance with publisher policies. Please cite the published version when available.

Title	Getting more from your datasets: Data augmentation, annotation and generative techniques
Author(s)	Corcoran, Peter; Bazrafkkan, Shabab; Lemley, Joe
Publication Date	2018-05-18
Publication Information	Corcoran, Peter , Bazrafkan, Shabab , & Lemley, Joe. (2018). Getting more from your datasets: Data augmentation, annotation and generative techniques. Paper presented at the Embedded Vision Summit Santa Clara Convention Center, Silicon Valley, Santa Clara, California, 22-23 May.
Publisher	XPERI Inc.
Link to publisher's version	https://www.dropbox.com/s/bmcjn385y47pgnh/_Corcoran1_Xperi_2018_Embedded_Vision_Summit_Slides_AAG02_MT_JB_PCv6_FINALrr2.pdf?dl=0
Item record	http://hdl.handle.net/10379/7480

Downloaded 2024-05-19T14:46:02Z

Some rights reserved. For more information, please see the item record link above.



embedded
VISION
SUMMIT
2018

*Getting More from Your Datasets: Data
Augmentation, Annotation and Generative
Techniques*

XPERI

Peter Corcoran, Joseph Lemley & Shabab Bazrafkan

18th May 2018



- Portions of this research were funded under the SFI Strategic Partnership Program by Science Foundation Ireland (SFI) and FotoNation Ltd.
 - Project ID: 13/SPP/I2868; *Next Generation Imaging for Smartphone and Embedded Platforms.*
- This work was also supported by an Irish Research Council Employment Based PhD Award.
 - Project ID: EBPPG/2016/280.



Overview of Today's Talk

- Topic #1 – Data Augmentation
 - What is Data Augmentaton?
 - Some Examples & case studies of Augmentation
 - Smart Augmentation – can we 'learn' an optimised augmentation strategy?
- Topic #2 – Generative Adversarial Networks (GANs)
 - Recap – Adversarial Networks; & how a GAN Works ...
- Topic #3 – Generating & Annotating Training Data
 - Boundary Equilibrius GAN (BEGAN) – a GAN to generate Face Images
 - Latent Spaces & learning Annotations
- Summary & Final thoughts - what does it all mean?

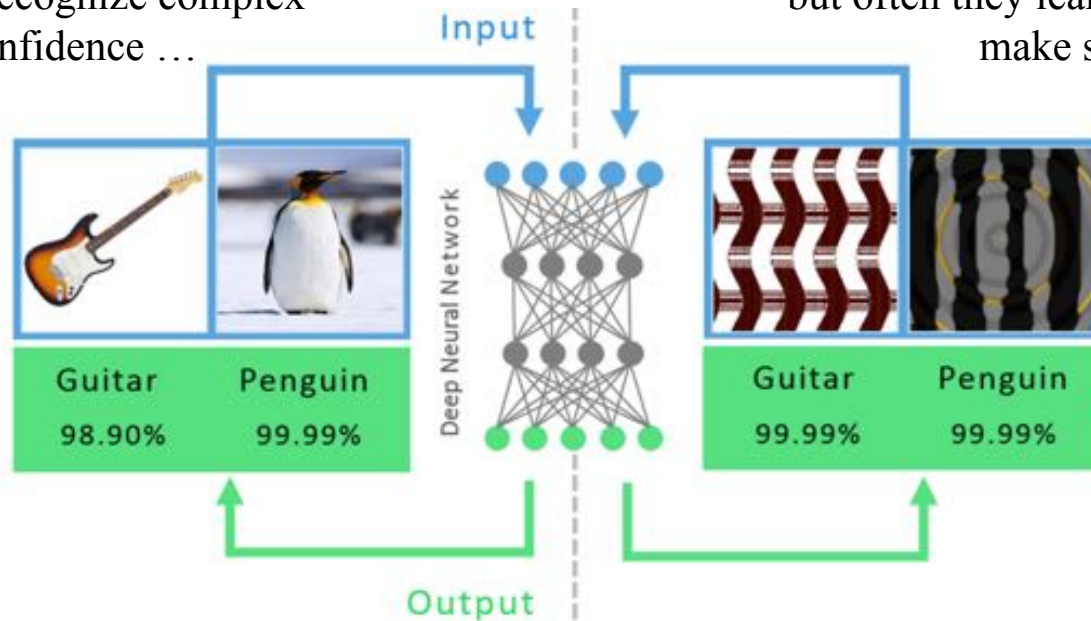
Some Observations on Deep Neural Networks



DNNs are Smart, but not always...

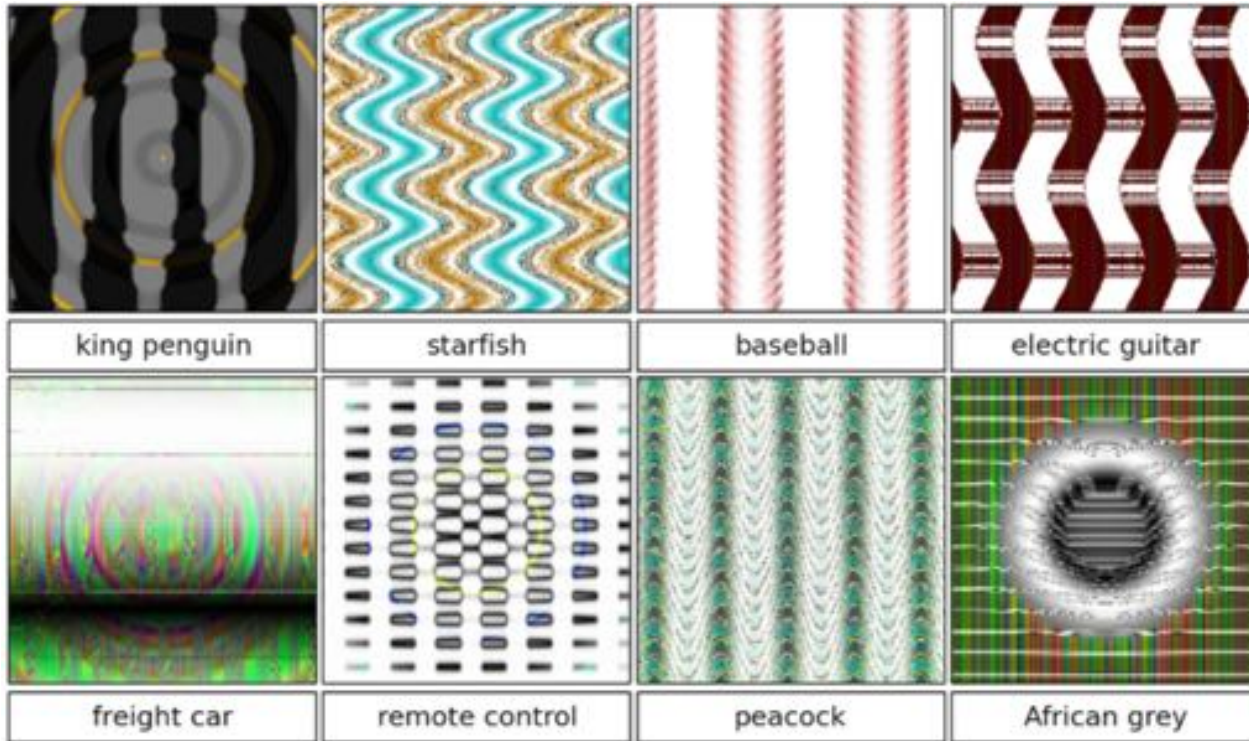
DNNs can learn to recognize complex objects with high confidence ...

but often they learn features that don't make sense to a human ...

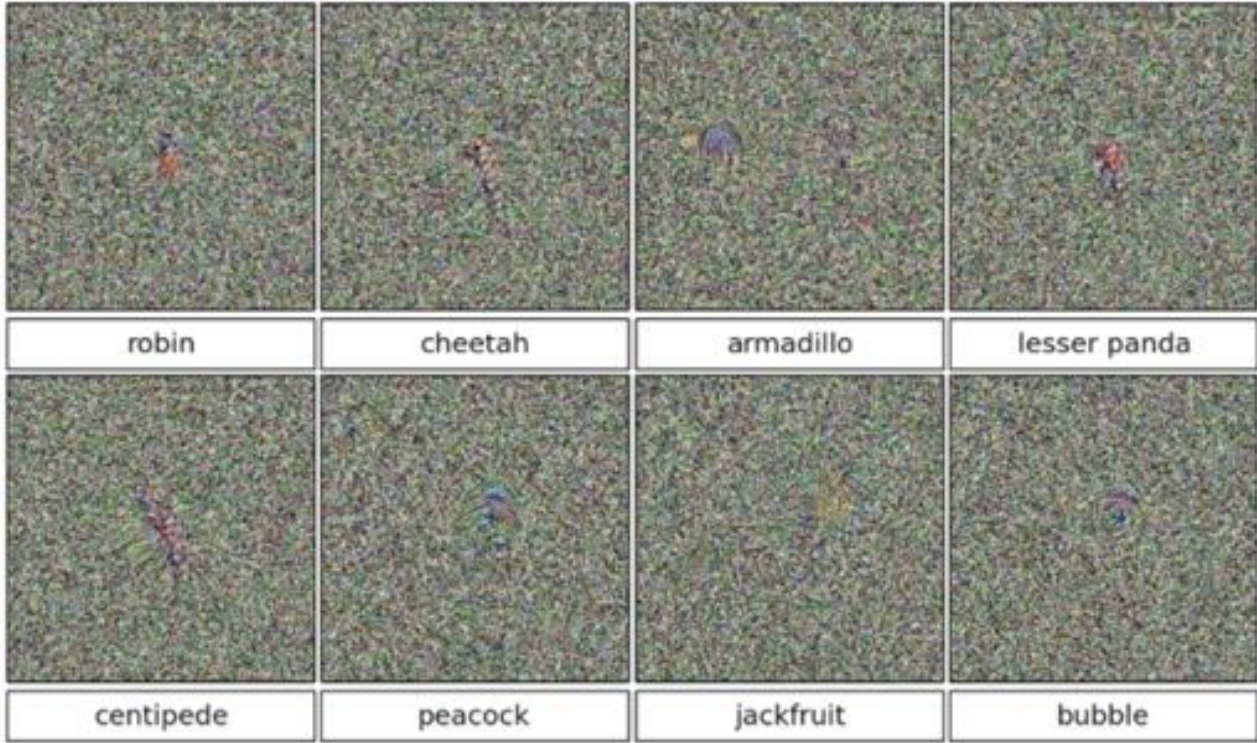


- Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images - *Nguyen, et al - 2014*

Deep Networks can be fooled by certain 'learned' (Adversarial) patterns ...



In Fact “noise” will sometimes work ...



Topic #1 – Data Augmentation

- **Topic #1 – Data Augmentation**
 - **What is Data Augmentation?**
 - **Some Examples & case studies of Augmentation**
 - **Smart Augmentation – can we 'learn' an optimised augmentation strategy?**
- Topic #2 – Generative Adversarial Networks (GANs)
 - Recap – Adversarial Networks; & how a GAN Works ...
- Topic #3 – Generating & Annotating Training Data
 - Boundary Equilibrium GAN (BEGAN) – a GAN to generate Face Images
 - Latent Spaces & learning Annotations
- Summary & Final thoughts - what does it all mean?

An Introduction to Data Augmentation



Data Augmentation Basics

- What is Data Augmentation?
- **Data Augmentation** - A regularization technique whereby the dataset is expanded by the creation of artificial variation such as zooming, rotation, shifting, salt/pepper noise, blur, etc.
- The augmentation approach chosen is often more important than the type of network architecture used.



Basic Augmentation - Zooming

Augmenting small datasets is important and challenging. You are not adding much new information into the network, but by augmenting the data you are training the network not to overfit your dataset with regards to the type of augmentation.

Augmentation helps ensure that your Network learns *semantically correct* features!



Basic Augmentation - Rotation

In an image classification task (e.g. dog/cat binary classification), if you rotate the image in various angles you are training the network to be invariant to rotation of the objects in the images.



Augmentation needs to be applied to each data class



Basic Augmentation – Gaussian Blur


R = 5



R = 15



R = 30




So although new “authentic” information isn't added into the network the “synthetic” data augmentation added into the network can both improve the results attained from the network and allow for training with less data.

Basic Augmentation – Noise


Var = 40



Var = 70




Var = 100




The Purpose of Augmentation

- The Main Goal of Augmentation to prevent the deep network (DNN) from ‘learning’ the wrong features (*overfitting*)
 - Here we considered a simple classification problem but augmentation is a valuable tool in other DNN contexts ...



The Augmentation depends on the Problem




- Consider iris authentication on a smartphone
 - Iris resolution is relatively small (<100 pixels)
 - Smartphone is handheld (image blur)
 - Lighting conditions can vary significantly
- A key challenge is to accurately *segment* the iris regions
- High quality iris image datasets with many samples are available
 - but *poor quality images* needed for training ...



High Quality Iris Datasets for Training

Iris diameter = 300+ pixels; **ground truth** (determined from high quality commercial algorithms) is at <https://goo.gl/JVvSyG>

Bath 800



CASIA 1000





Test Databases - Low Quality from Mobile Handsets

Iris diameter < 100 pixels; augmentation code available at:

https://github.com/C3Imaging/Deep-Learning-Techniques/blob/Iris_SegNet/DBaugmentation/DBaug.m

UBIRIS2



MobBio



Augmentation on Training Data – Contrast Reduction

- Contrast reduction


Original image



Low contrast image



Contrast reduction




Iris filled mask

Augmentation on Training Data – Shadowing/Motion Blur

- Shadowing




Shadowing




- Motion Blur



Motion blur



Training Data Augmentation – Mixed Examples



Smart Augmentation

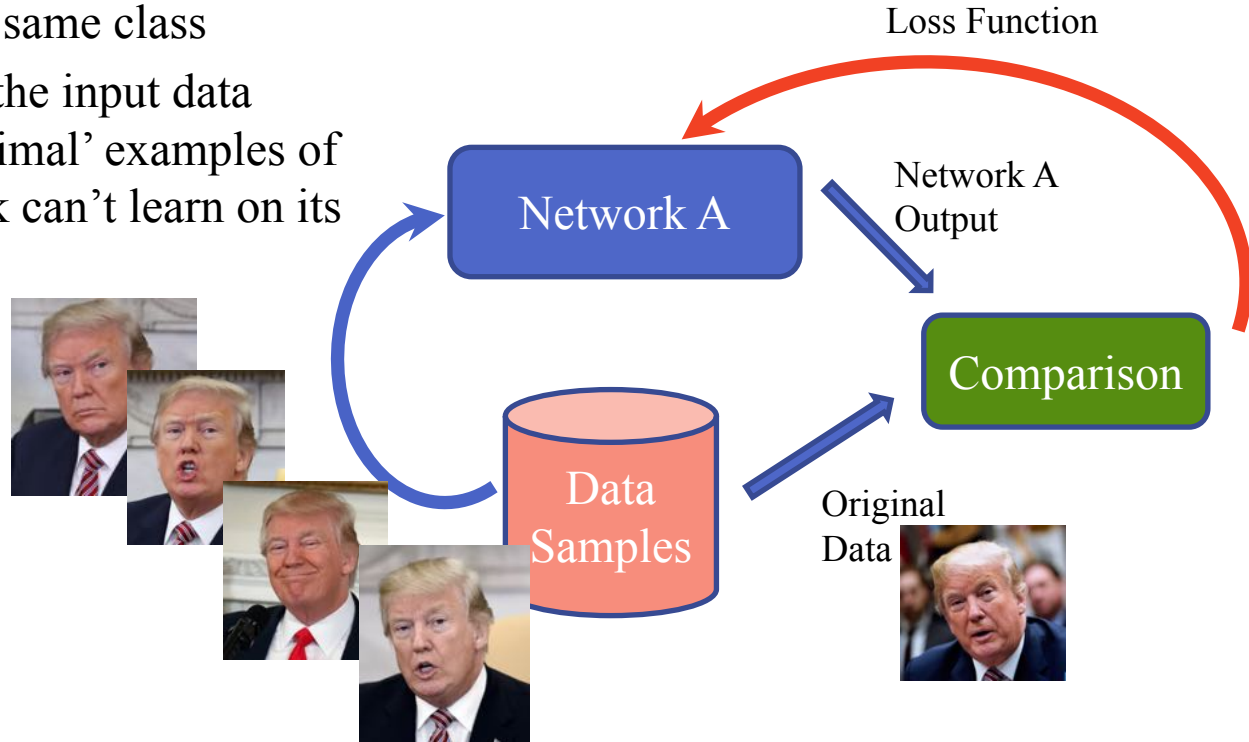


Can We Teach a Deep Network to Learn an Augmentation Strategy?

- Augmenting data is a very time consuming process. Selecting the best augmentation strategy is sometimes a matter of luck, and expertise.
- **Question:** *Can we do better? Can Artificial neural networks learn the augmentation task during training?*
- **Answer:** *Yes!*

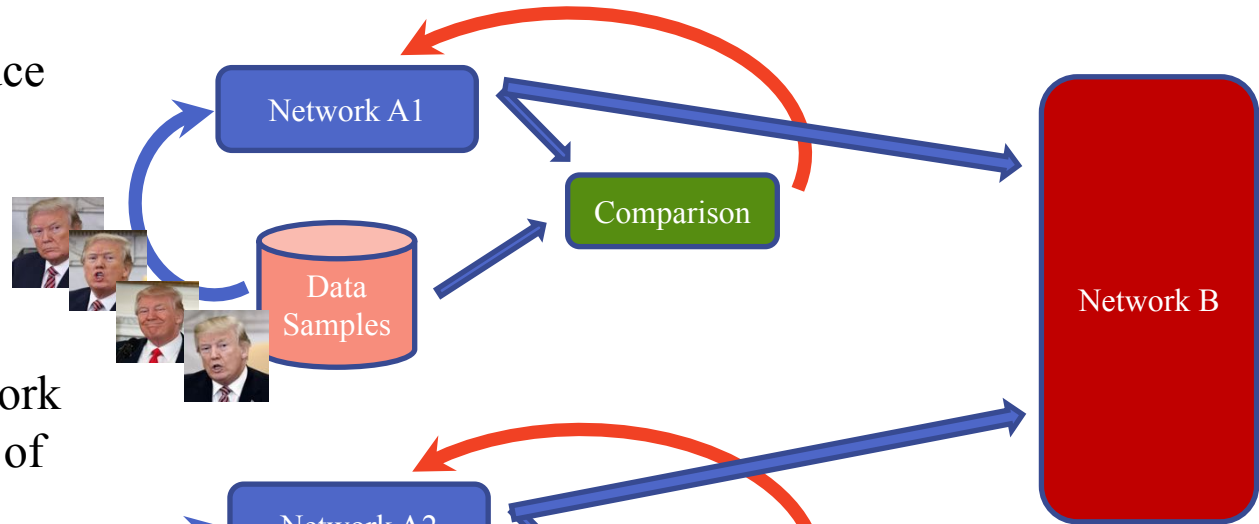
Network A – Learning New Data of a certain Class

- We would like Net A to learn how to make new data of the same class
- The challenge is that the input data provides ‘close to optimal’ examples of class A so the network can’t learn on its own

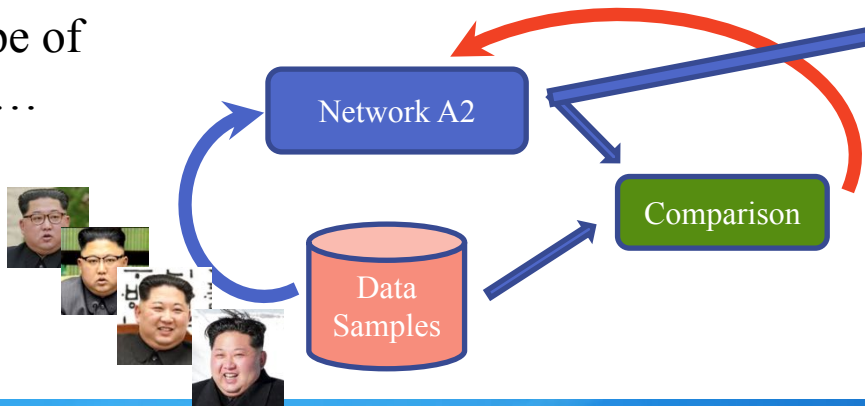


Now add a 2nd Data Class + Classifier Network

- Now suppose we introduce a 2nd Network A for a different Class ...



- Let us feed the two network A outputs into a 2nd type of network – Network B ...

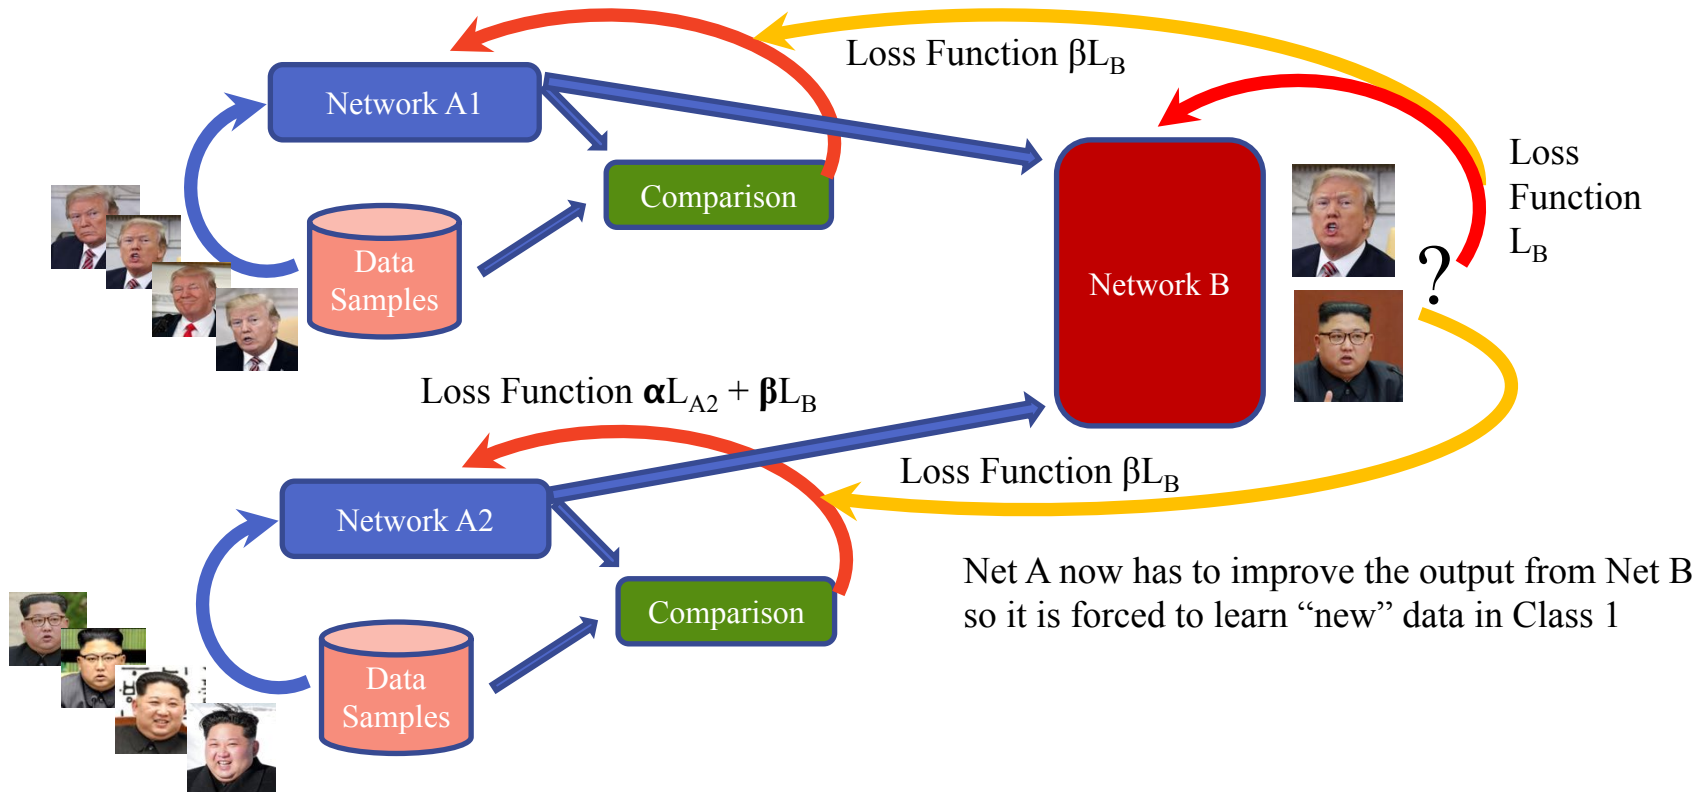


Network B learns to distinguish the 2 classes ...

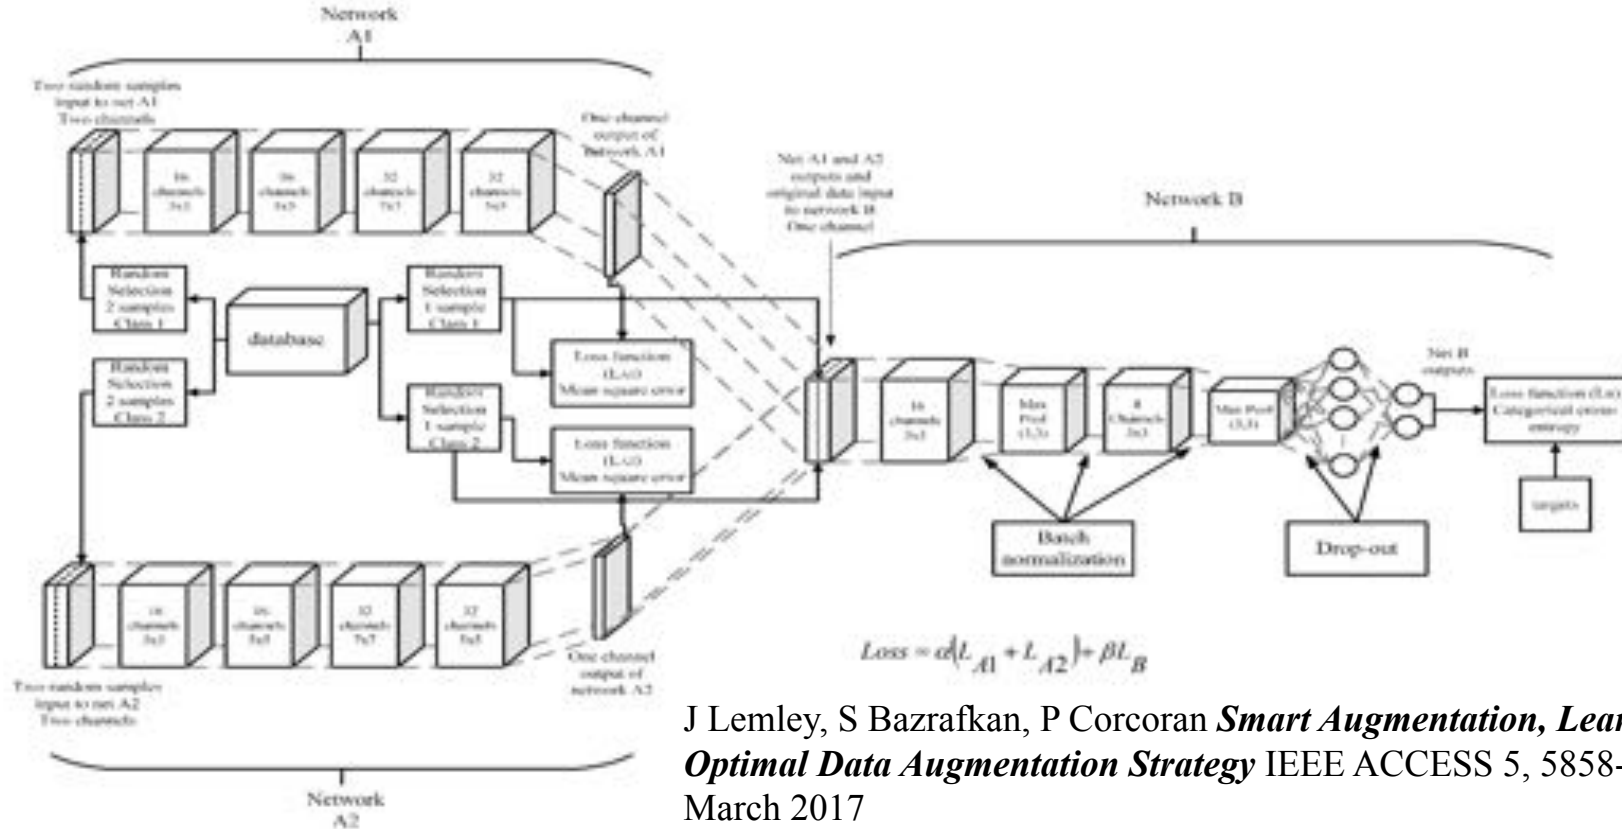
How does Smart Augmentation work with Images? (2)

- Network A learns to merge two or more samples in one class (in a nonlinear way)
 - The merged sample is then used to train a target Network B
 - The loss of the target network B is fed back to inform the augmenter.
 - This generates more images for use by the target network.
- Network A (augmentor) learns to generate images that not only belong to the same class, but that improve the performance of Network B (classifier)
 - The convolutional nature of the merging of samples in Network A generates some surprising results as we'll see later ...

The Full Picture ... Each Network A gets better at generating samples that improve Network B accuracy ...



More details (& complicated figures!) in our paper ...



J Lemley, S Bazrafkan, P Corcoran *Smart Augmentation, Learning an Optimal Data Augmentation Strategy* IEEE ACCESS 5, 5858-5869. March 2017

What kinds of Images Does Network-A Generate?

- The images in the red box are created by a learned combination of the previous two images in that row.



Smart Augmentation was effective on a variety of unconstrained and constrained images.

AR Faces
Highly constrained



FERET
Constrained



Adience
Unconstrained



MIT Places
Highly Unconstrained



Key Results – Significant Improvements with SA

A significant improvement in accuracy resulted from using Smart Augmentation.

(Green cells are Smart-Augmentation (SA) experiments; Orange cells are equivalent methods without SA)

+3.4% up to **+6.7%** on *AR Faces*

+6.1% on *Adiaence*

+5.0% on *FERET dataset*

Citation: Joseph Lemley, Shabab Bazrafkan, and Peter Corcoran. "Smart Augmentation-Learning an Optimal Data Augmentation Strategy." *IEEE Access* (2017).

Dataset	# Net As	Input Channels	Augmented	Test Accuracy
AR Faces	1	2	no	92.50%
AR Faces	1	3	no	95.10%
AR Faces	1	4	no	91.60%
AR Faces	2	2	no	95.70%
AR Faces	2	3	no	94.20%
AR Faces	2	4	no	94.20%
AR Faces	0	NA	no	88.20%
AR Faces	0	NA	yes	89.00%
AR Faces	1	2	yes	95.70%
AR Faces	2	2	yes	95.70%
<hr/>				
Adience	0	NA	no	70.00%
Adience	1	2	no	76.10%
<hr/>				
FERET	0	NA	no	83.50%
FERET	1	2	no	88.50%


Why Use Smart Augmentation?

- Smart Augmentation has been experimentally validated with 30 experiments which show that **Smart Augmentation:**
 - Decreases **overfitting**
 - Increases **accuracy**
 - Increases **generalization capability**
 - Can significantly *reduce the number of parameters* required to perform the same task (i.e. a *smaller network* can work just as well as a large network)


Network Size Reduction via Smart Augmentation

When trained on 2 classes from MIT Places Dataset.


VGG 16 without SA (98.5%)



Small Network without SA (96.5%)

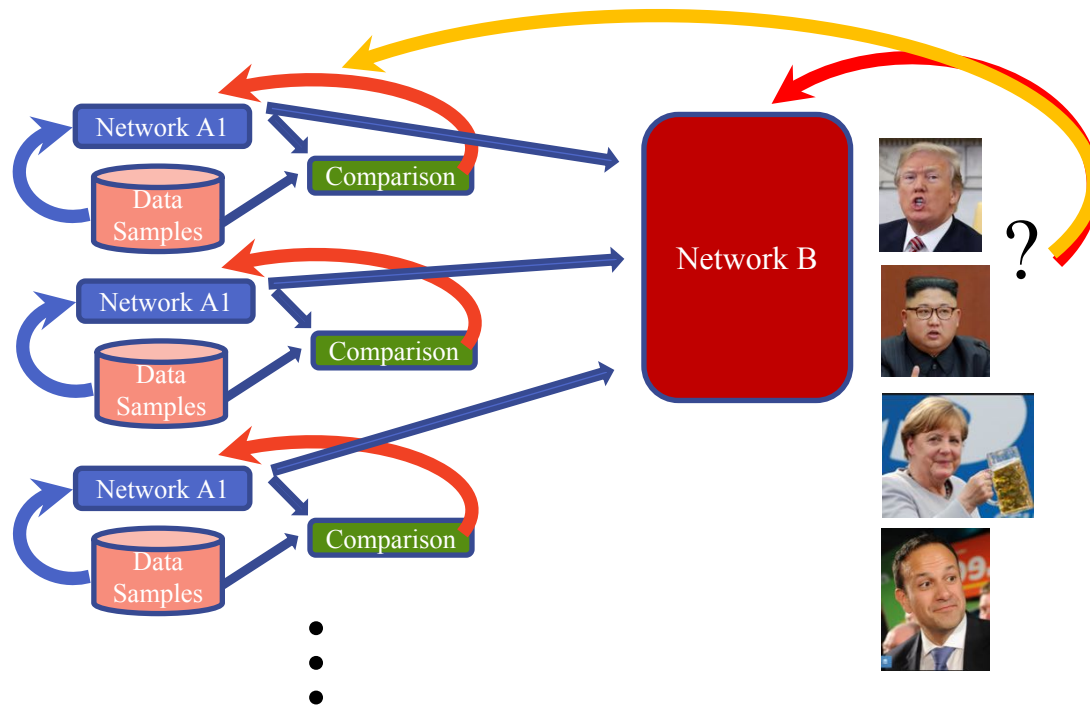


Small Network with SA (99%)



Future Work – Smart Augmentation 2.0

- We can scale Smart Augmentation to work with multi-class problems:



Basics of Generative Adversarial Networks (GANs)

- Topic #1 – Data Augmentation
 - What is Data Augmentaton?
 - Some Examples & case studies of Augmentation
 - Smart Augmentation – can we 'learn' an optimised augmentation strategy?
- **Topic #2 – Generative Adversarial Networks (GANs)**
 - **Recap – Adversarial Networks; & how a GAN Works ...**
- Topic #3 – Generating & Annotating Training Data
 - Boundary Equilibrius GAN (BEGAN) – a GAN to generate Face Images
 - Latent Spaces & learning Annotations
- Summay & Final thoughts - what does it all mean?

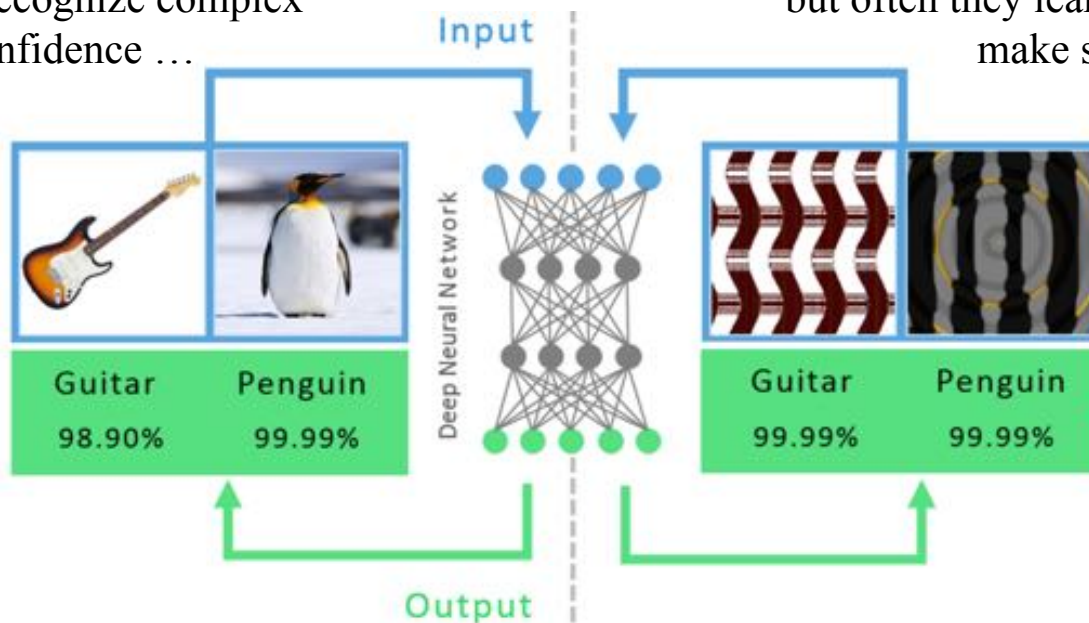
Quick Recap – Adversarial Networks



Quick recap - DNNs are Smart, but not always...

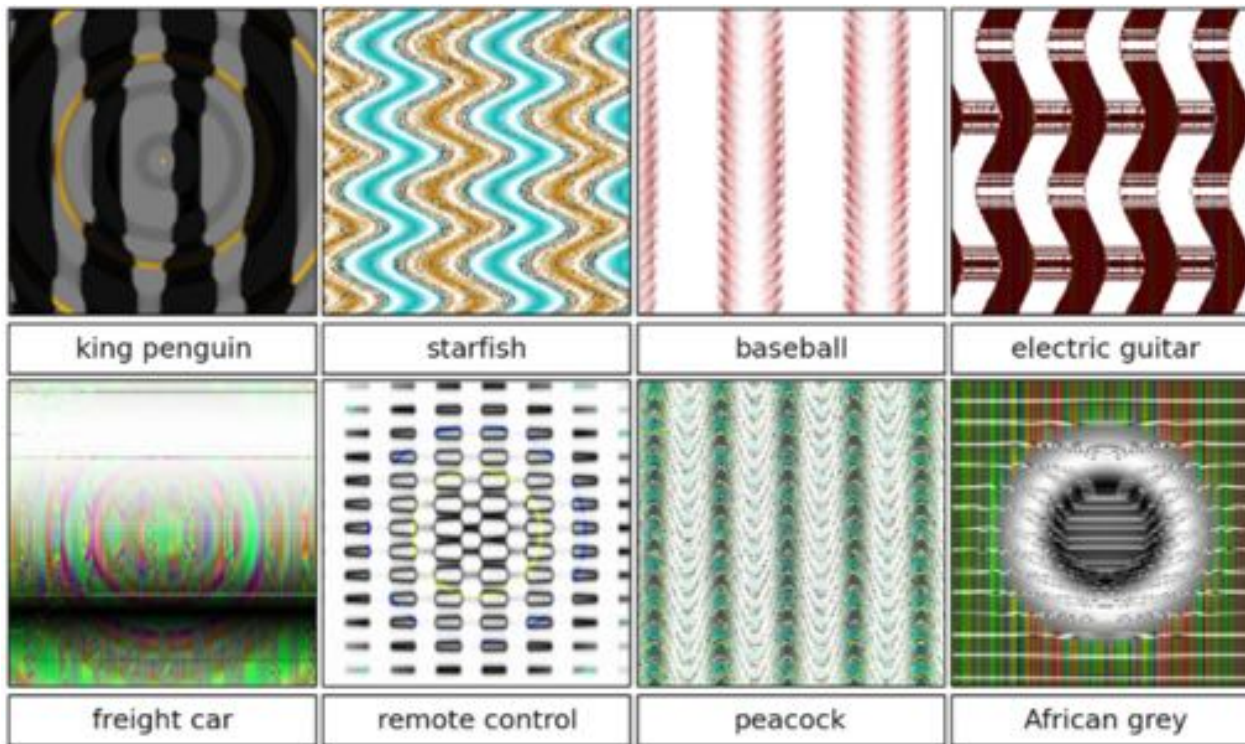
DNNs can learn to recognize complex objects with high confidence ...

but often they learn features that don't make sense to a human ...

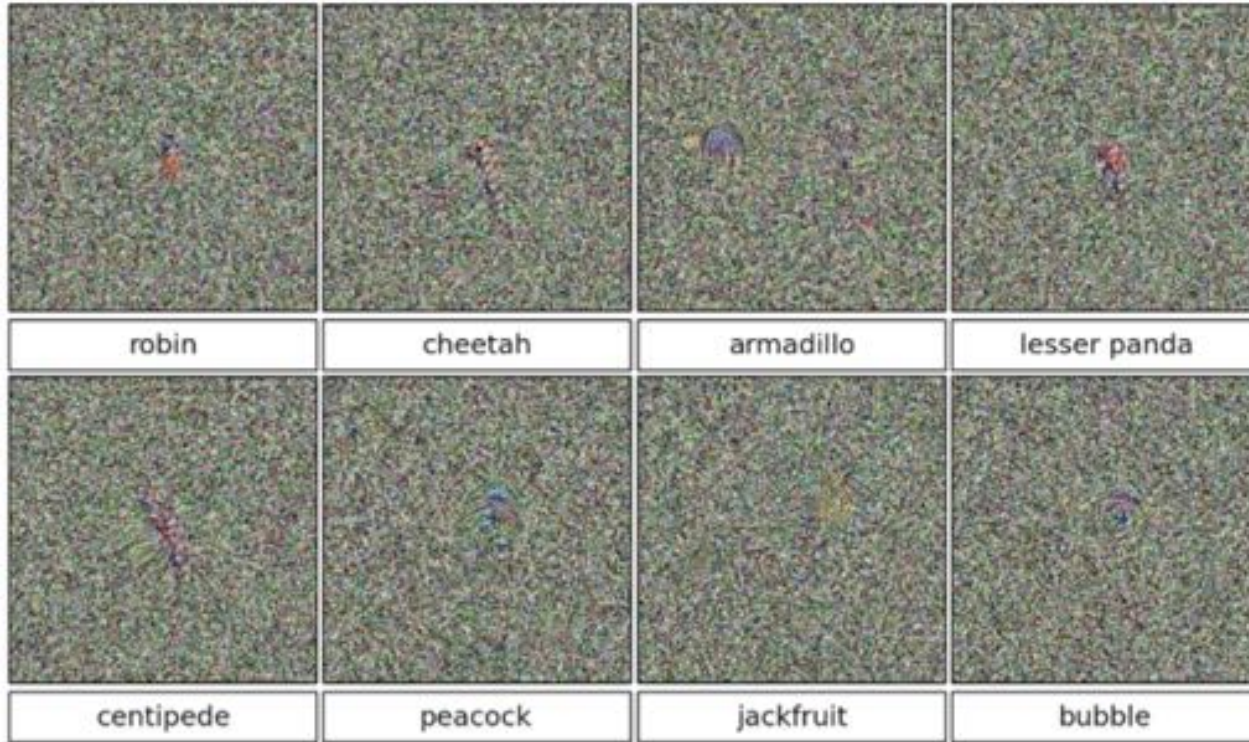


- Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images - *Nguyen, et al - 2014*

Quick recap - DNNs fooled by “Adversarial” patterns ...



In Fact “noise” will sometimes work ...



Training Adversarial Samples (*Goodfellow 2014*) ...




Figure 1: A demonstration of fast adversarial example generation applied to GoogLeNet (Szegedy et al., 2014a) on ImageNet. By adding an imperceptibly small vector whose elements are equal to the sign of the elements of the gradient of the cost function with respect to the input, we can change GoogLeNet's classification of the image. Here our ϵ of .007 corresponds to the magnitude of the smallest bit of an 8 bit image encoding after GoogLeNet's conversion to real numbers.

Goodfellow IJ, Shlens J, Szegedy C. Explaining & harnessing adversarial examples. arXiv:1412.6572. 2014 Dec 20.

Single Pixel Adversarial Examples ...


AllConv



SHIP
CAR(99.7%)



HORSE
DOG(70.7%)



CAR
AIRPLANE(82.4%)


NiN



HORSE
FROG(99.9%)



DOG
CAT(75.5%)



DEER
DOG(86.4%)

VGG



DEER
AIRPLANE(85.3%)



BIRD
FROG(86.5%)



CAT
BIRD(66.2%)

- There are many ways to ‘fool’ DNNs
- These authors showed that a single pixel in the right place can ‘trigger’ a DNN to generate false classifications
- They showed that several well-know object classification frameworks were vulnerable to this simple ‘attack’ ...

Su J, Vargas DV, Kouichi S. One pixel attack for fooling deep neural networks. arXiv:1710.08864. 2017 Oct 24.

Adversarial Data used to improves the Model!

- Generated ‘Adversarial’ samples can be used to fool a discriminative model ...
- **But** adversarial samples can be used in the training process as *False Positive samples* to make the discriminative model more robust!
 - DNN network learns the difference between adversarial & genuine data ...
 - More effort is **then** required to generate *better adversarial samples* ...
- Iterating this process yields an *improved discriminative model* ...
 - and an *improved generator* for adversarial samples!

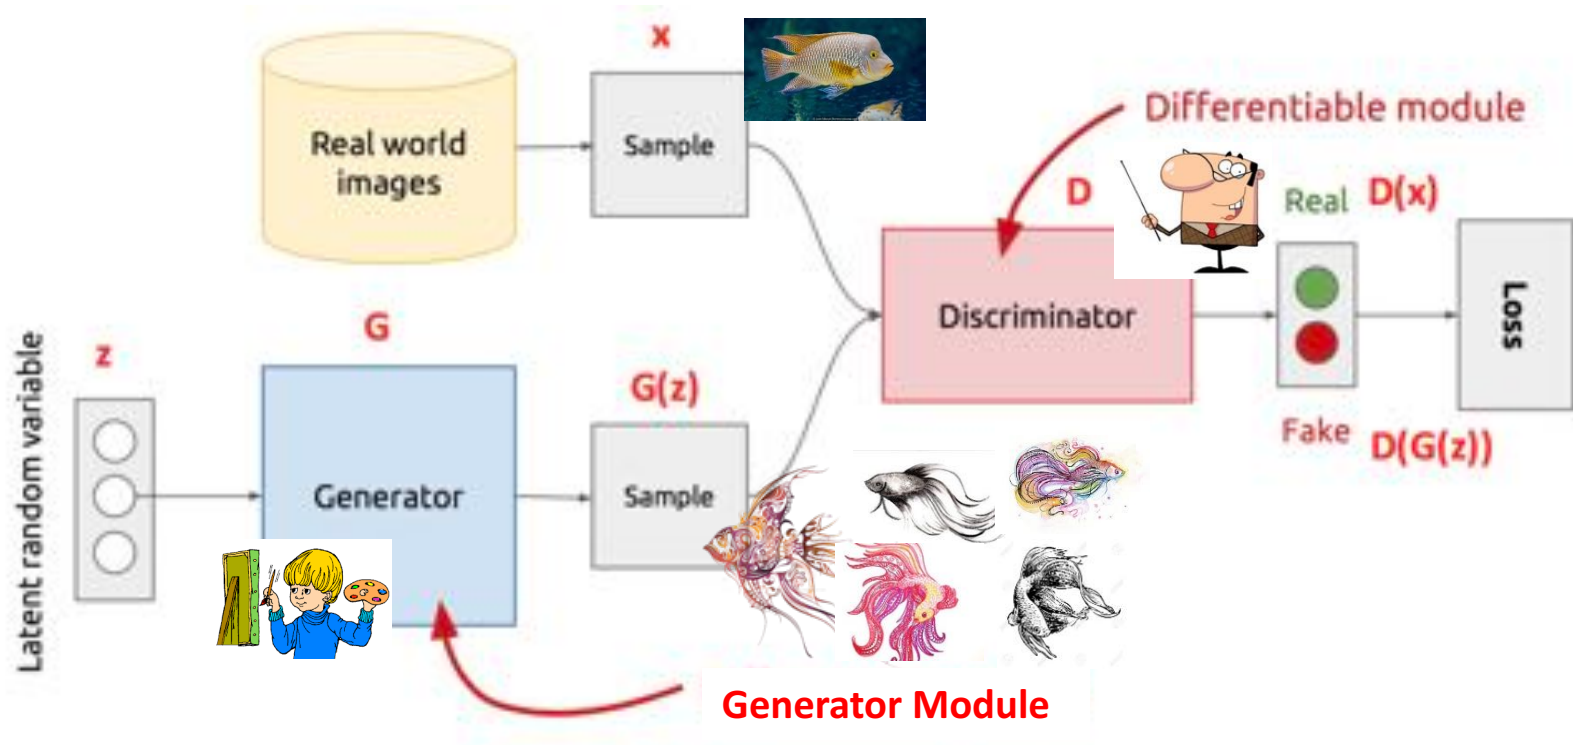
Quick Recap – Generative Adversarial Networks



What is a Generative Adversarial Network? (GAN)

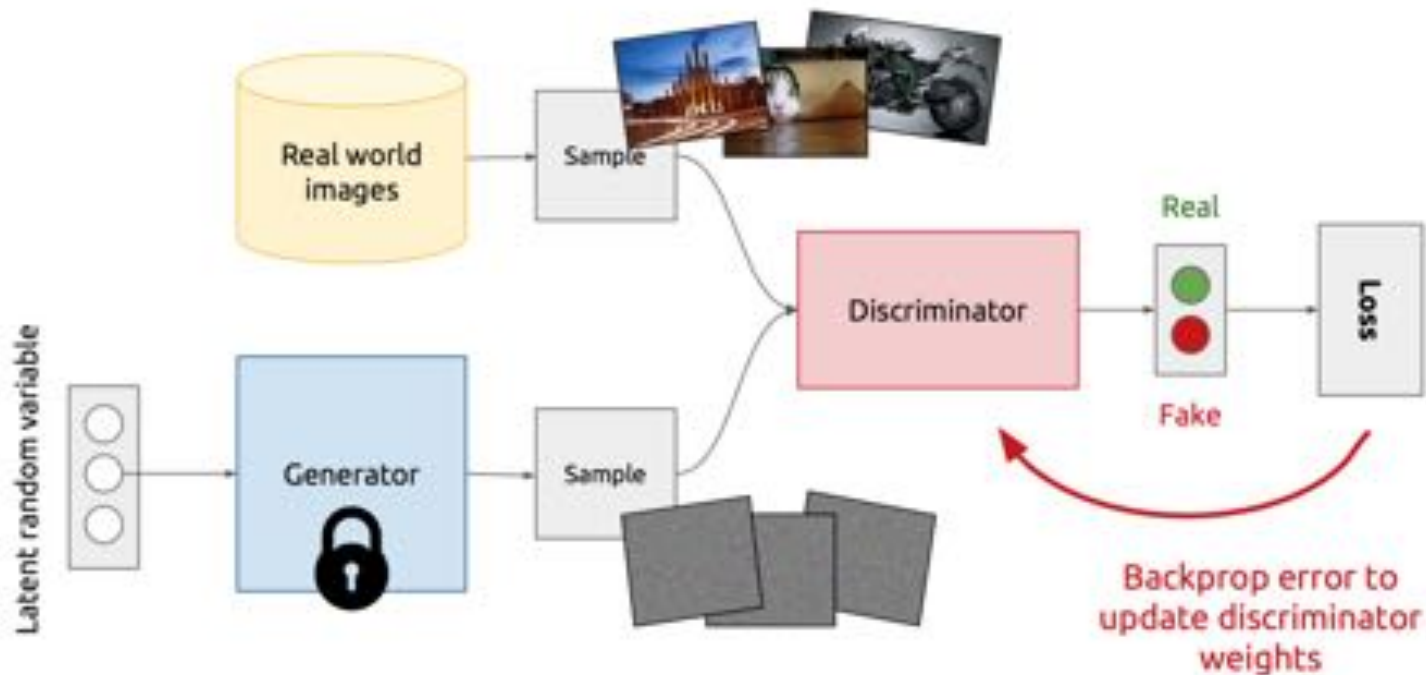
- We saw that a network can learn to make **False Positive** data samples
 - That data can improve the discriminating ability of a classifier
- But suppose the goal is not simply to classify data samples ...
 - but instead to make better **high quality data** of a certain “class” – data that is sufficiently close to the ‘original data’ that it cannot be distinguished from it by a human ...
 - Can we now fomulate a network with two parts ...

Architecture of GAN #1 - *Learning to Generate Samples*



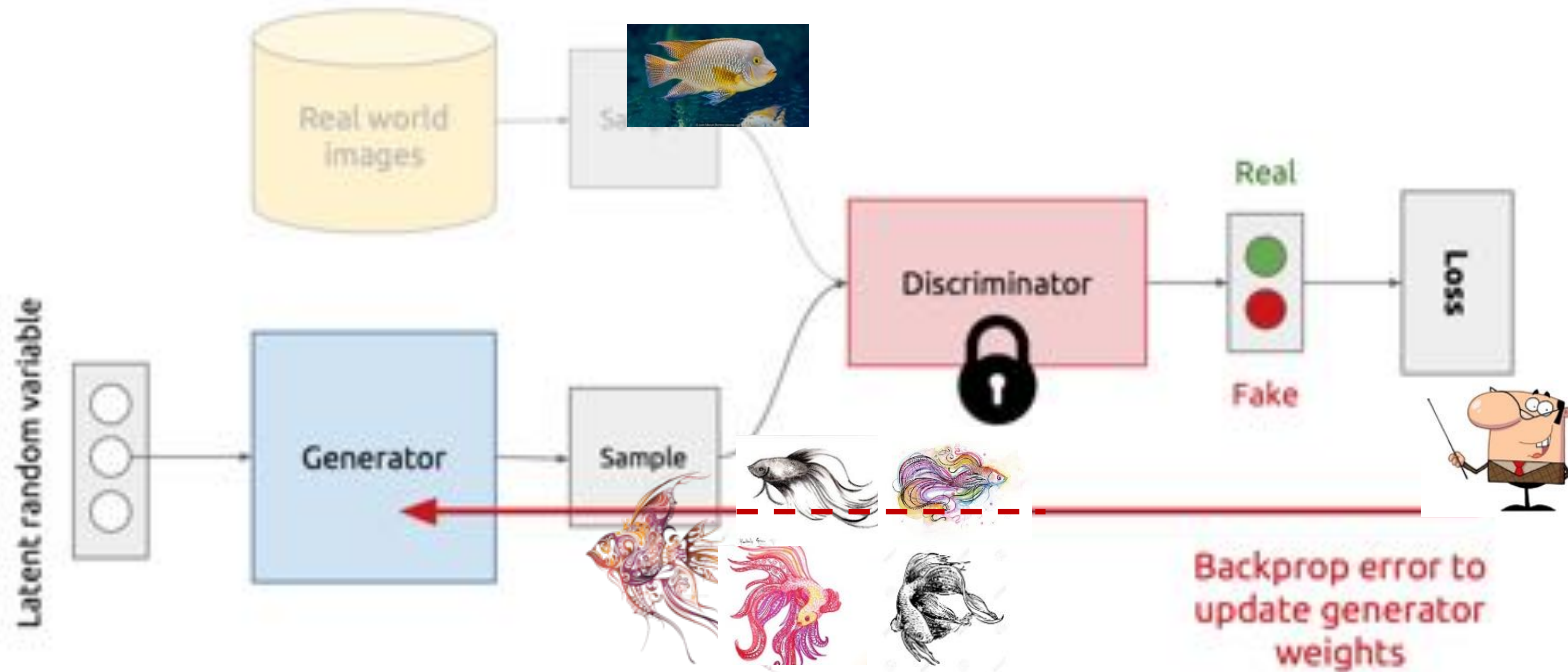
- Generator: generates fake samples, tries to fool the Discriminator
- Discriminator: tries to distinguish between real and fake samples
 - Now suppose we train them against each other
- Repeat & optimise this cycle and we get improved Generator and Discriminator modules

Separate Training needed for the Discriminator ...



<https://www.slideshare.net/kavigiro/deep-learning-for-computer-vision-generative-models-and-adversarial-training-upc-2016>

... and the Generator modules ...



<https://www.slideshare.net/xavigiro/deep-learning-for-computer-vision-generative-models-and-adversarial-training-upc-2016>

GAN is Formulated as a min/max game ...

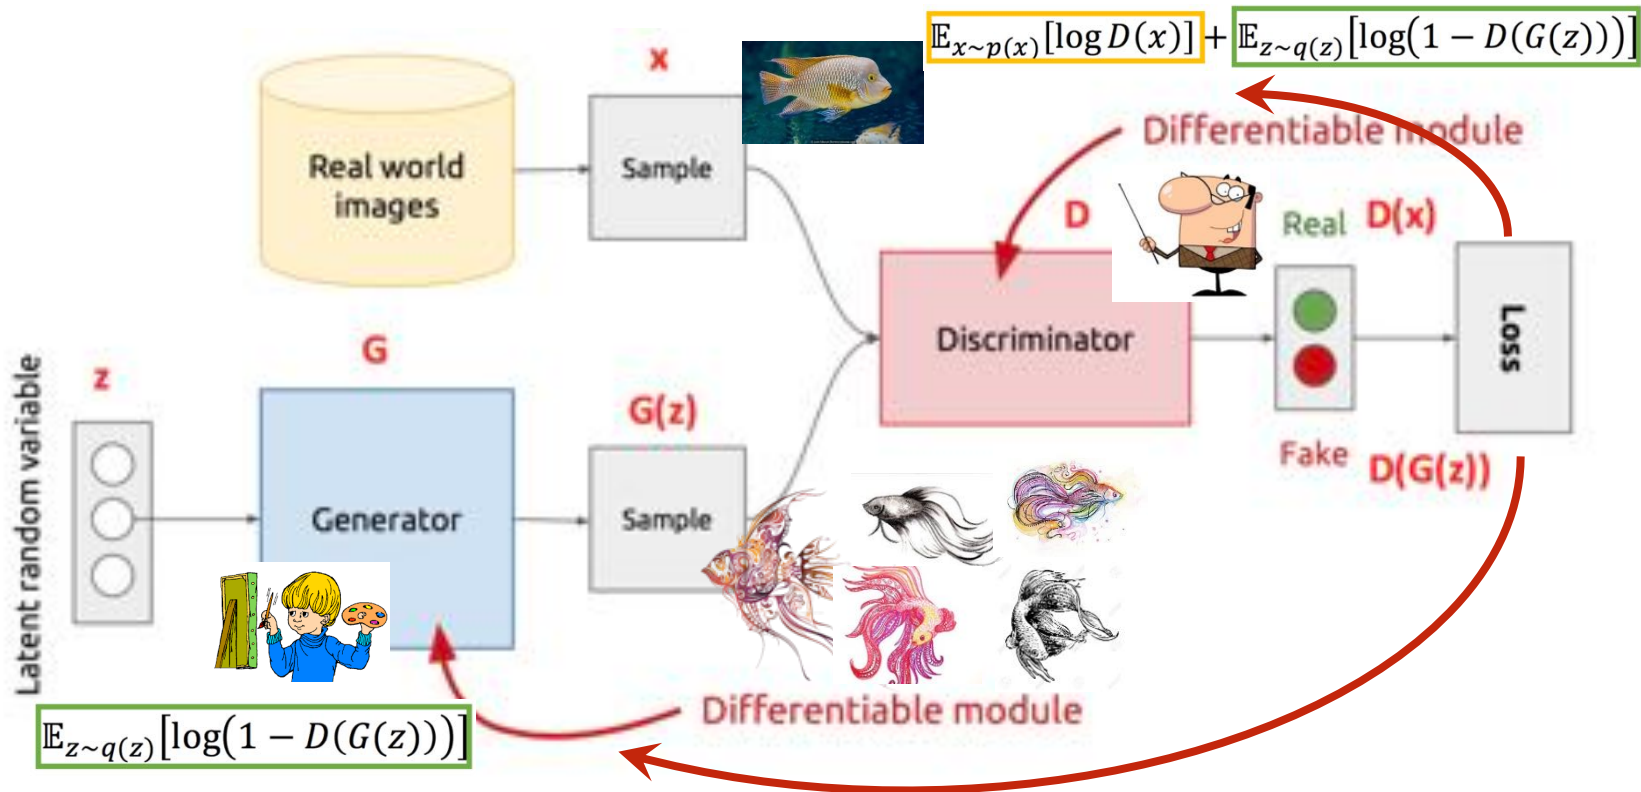
GAN's formulation

$$\min_G \max_D V(D, G)$$

- It is formulated as a **minimax game**, where:
 - The Discriminator is trying to maximize its reward $V(D, G)$
 - The Generator is trying to minimize Discriminator's reward (or maximize its loss)

$$V(D, G) = \mathbb{E}_{x \sim p(x)} [\log D(x)] + \mathbb{E}_{z \sim q(z)} [\log(1 - D(G(z)))]$$

Architecture of GAN



Making Data (& Annotations) with GANs

- Topic #1 – Data Augmentation
 - What is Data Augmentaton?
 - Some Examples & case studies of Augmentation
 - Smart Augmentation – can we 'learn' an optimised augmentation strategy?
- Topic #2 – Generative Adversarial Networks (GANs)
 - Recap – Adversarial Networks; & how a GAN Works ...
- **Topic #3 – Generating & Annotating Training Data**
 - **Boundary Equilibrius GAN (BEGAN) – a GAN to generate Face Images**
 - **Latent Spaces & learning Annotations**
- Summay & Final thoughts - what does it all mean?

GANs to Generate Faces

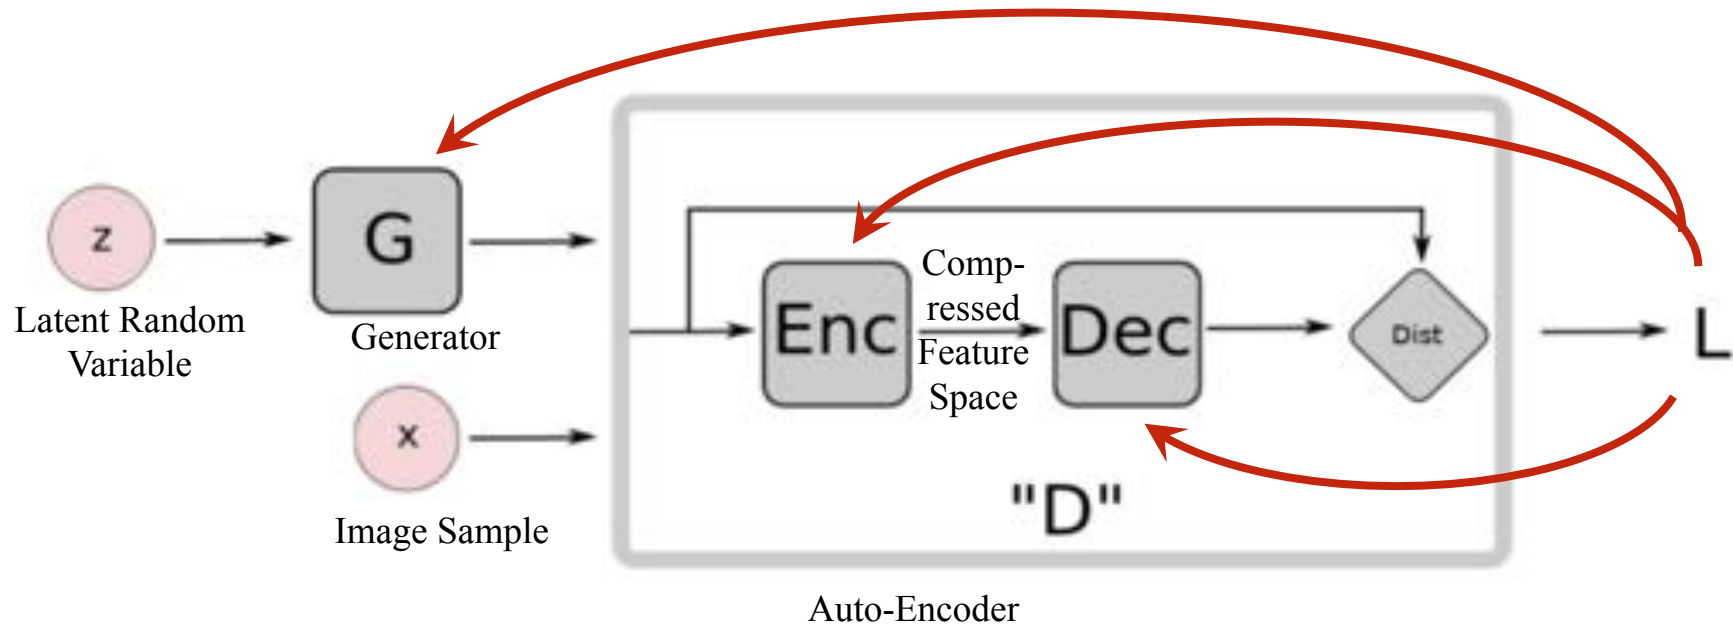


Boundary Equilibrium GAN - BEGAN

- Boundary Equilibrium GAN successfully generates anatomically coherent faces at a resolution of 128×128 pixels
- As of early May 2017, this was the state of the art



BEGAN replaces Discriminator with an Auto-Encoder



- The training goes like this:
 - D (*autoencoder*) learns to reconstruct real images with improving accuracy.
 - i.e. weights of D are updated to minimize the reconstruction loss on *real* images;
 - D **simultaneously** works to *increase* the reconstruction loss of *generated* images; [Min/Max goal]
 - And G works *adversarially* to minimizing the reconstruction loss of generated images
- As long as both D and G continue to improve, then the quality of *generated* images will grow closer to that of real images in the training dataset ...

The Loss Functions have complex relationships ...

<https://blog.heuritech.com/2017/04/11/began-state-of-the-art-generation-of-faces-with-generative-adversarial-networks/>

Here is the complete BEGAN objective:

$$\mathcal{L}_D = \mathcal{L}(x) - k_t \cdot \mathcal{L}(G(z))$$

[Max/Min goal - +ve reconstruction for real images; -ve for generated]

$$\mathcal{L}_G = \mathcal{L}(G(z))$$

$$k_{t+1} = k_t + \lambda * (\gamma \cdot \mathcal{L}(x) - \mathcal{L}(G(z)))$$

[balancing of each iteration based on diversity & gain of learning rate]


- \mathcal{L}_D and \mathcal{L}_G are the respective losses for D and G (what they try to minimize).
 - \mathcal{L}_D is only used to optimize θ_D and \mathcal{L}_G is only used to optimize θ_G .
 - $\mathcal{L}(x)$ and $\mathcal{L}(G(z))$ are the losses of reconstruction of real and generated images.
 - γ is the diversity ratio (in $[0, 1]$) defined before as:
 - k_t is the adaptive term that will allow us to balance the losses
 - λ is the proportional gain for k_t (aka the learning rate for k_t).
- $$\gamma = \mathbb{E}[\mathcal{L}(G(z))] / \mathbb{E}[\mathcal{L}(x)].$$

From BEGAN to Data Annotation?



The Challenge we posed ...

- Can we generate random faces that are accurately “annotated” ... ?
 - Why would you want to do that?
 - It becomes possible to generate **random faces** for **training data** ...
 - Can ‘**prove**’ they are not the ‘original’ faces – solves **Privacy** issues!
 - In theory can generate **more faces** than in the original dataset ... (???)



Dataset and Initial Landmark Detection

- The **CelebA dataset** - **202,599** original images with 40 unique attributes is used for training our GAN framework (<http://mmlab.ie.cuhk.edu.hk/projects/CelebA.html>)
- **OpenCV** frontal face cascade classifier detects facial regions
 - cropped and resized to 128×128 pixels
- Initial landmark detection is performed using the method developed by Astana *et. al* due to its ability to be effective on unconstrained faces
 - Landmark detector estimates a set of **49 landmark points** defined on the contours of eyebrows, eyes, mouth and the nose


*Asthana A, Zafeiriou S, Cheng S, Pantic M. **Incremental face alignment in the wild.** In proceedings of the IEEE conference on computer vision and pattern recognition 2014 (pp. 1859-1866).*

Landmark Examples on CelebA Dataset




Mapping of CelebA data samples into Z-Space

- Generator, G , is trained that can produce random faces using BEGAN

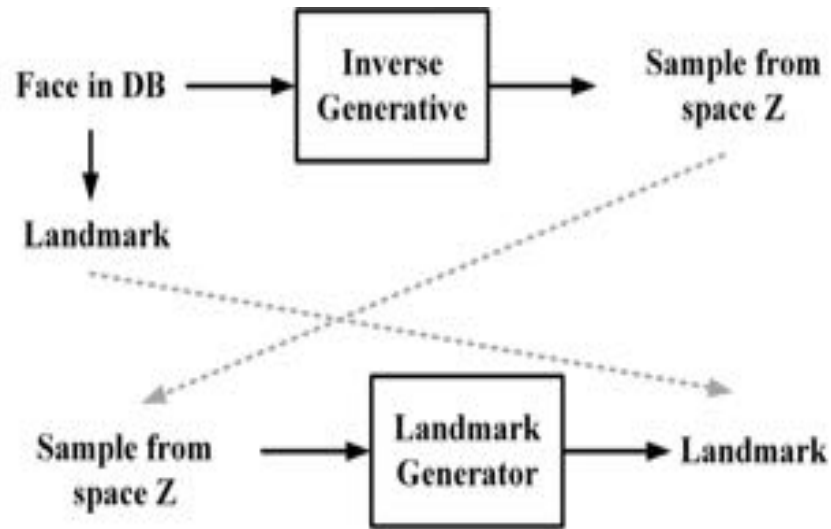


- An Inverse generator maps a face sample back into the latent Z -space



Training the Landmark Generator ...

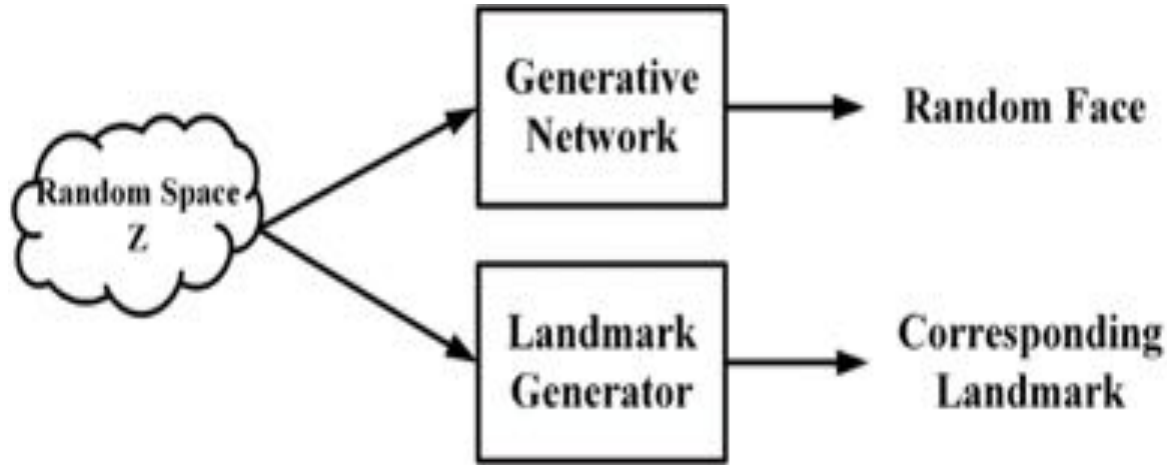
- CelebA data samples are mapped to Z-space to Train the Landmark Generator
- Landmark Generator is trained to learn annotations from latent Z-space



1. A Facial image sample is mapped into the latent space by the *Inverse Generator*
2. The corresponding facial *landmark data* becomes a training sample associated with corresponding location in Z-space; (repeat 200k times)
3. A *generator* is trained to map Z-space locations to corresponding landmark point set
4. If landmark data is accurate *ground truth* the generator can handle wide variations in illumination & pose


Finally, Faces generated together with Annotations ...

- Generate a random Face image from the Latent Z-space ...
- Matching annotations are generated directly from the originating Z-space vector values




Example #1 of Generated Dataset

Multiple differing facial images



Example #2 of Generated Dataset

Same base face with differing facial pose



Some observations

- The learned mapping into latent Z-Space will vary according to the training dataset and the DNN structures employed in the GAN
 - The learned Z-Space characteristics depend on the training data, ***but can encapsulate pose & lighting variations*** – key challenges for face generation
 - Examples of pose variation step incrementally between 2 end-points in the Z-Space but the finer details still elude us ... our understanding is still a ‘work-in-progress’ ...
- The latent Z-space that is trained comprises 64 x 32bit floating point numbers has, potentially 2^{38} discriminating capability ...
 - 274B potential samples or c.30 times the global population ...

What does it all mean ... ?


- Topic #1 – Data Augmentation
 - What is Data Augmentaton?
 - Some Examples & case studies of Augmentation
 - Smart Augmentation – can we 'learn' an optimised augmentation strategy?
- Topic #2 – Generative Adversarial Networks (GANs)
 - Recap – Adversarial Networks; & how a GAN Works ...
- Topic #3 – Generating & Annotating Training Data
 - Boundary Equilibrius GAN (BEGAN) – a GAN to generate Face Images
 - Latent Spaces & learning Annotations
- **Summary & Final thoughts - what does it all mean?**

Conclusions – Recap & Final Thoughts



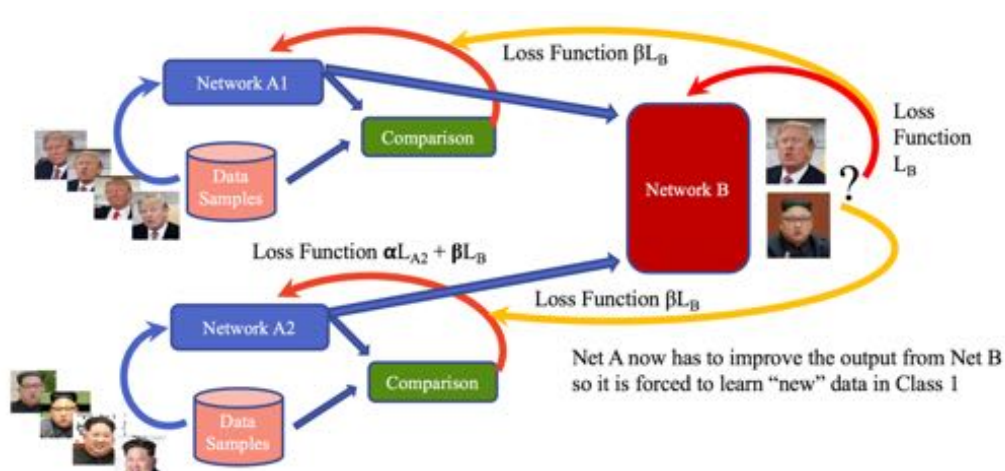
Conclusions #1a – Data Augmentation

- Augmentation is a powerful tool when training DNN
 - Correct choice of augmentation is important & determines robustness and accuracy of final DNN
 - Augmentation is specific to a particular DNN problem



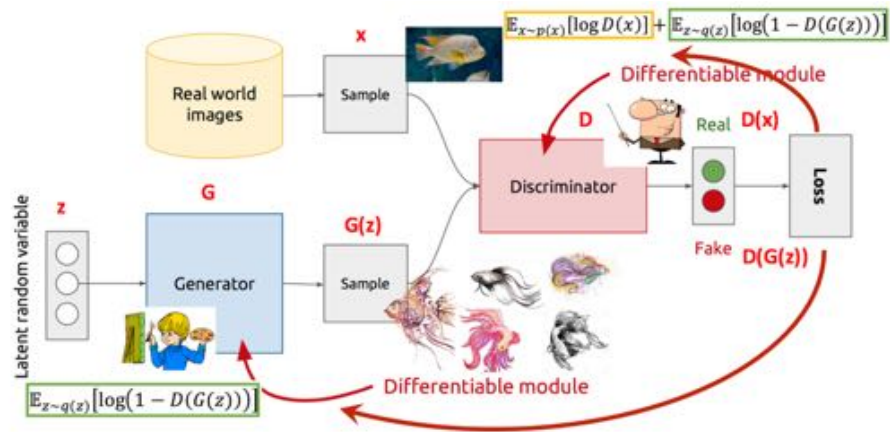
Conclusions #1b – Data Augmentation

- Smart-Augmentation (SA) is a new tool that uses a 2nd DNN to generate new data from an existing training dataset
 - SA can improve DNN accuracy and key performance metrics
 - Data augmentations are “learned” by the SA network and vary in ways that could not be achieved through manual augmentation





Conclusions #2a – GANs & Latent Spaces

- Combining Generative & Adversarial components provides another powerful tool when training DNN
 - Enables training of very accurate data generators; data quality is now equivalent to the original data in many cases
 - Complex loss functions can pose challenges for initial convergence



Conclusions #2b – GANs & Latent Spaces

- An Inverse Generators allows metadata associated with the original training data to be mapped into the latent Z-space
 - Generated data samples can have metadata created from Z-space
 - One key example is to create new facial data with corresponding landmarks for the generated samples



Concluding Thoughts




- We presented two powerful tools for *augmenting & generating* training data
 - *Smart Augmentation* is a relatively new approach ...
 - There are many variations on **GANs** – interesting new papers with new approaches/refinements every week ...
 - Next step is perhaps to combine these?
- We are at the ‘tip’ of an iceberg as researchers continue to refine these techniques and find new ways to scale, refine and develop these approaches ...

- Generated data has some interesting advantages:
 - Avoids *privacy issues* & side-steps *new regulations* (e.g. GDPR)
 - Significant cost reductions in *data annotation & recording metadata* for very large datasets
 - Lots of potential to subtly control/refine how data is generated, augmented & annotated ... we are still learning new ‘tricks’ here ...
 - ... an interesting new field for research & experimentation ...
- Are we entering a new era of “*Generated Data*”? [post *Big-Data* era?]

Questions?



Example of Video Slide Layout



<https://www.youtube.com/watch?v=PWdT3Q5T5U8>

- **General Deep Learning:**
 - Lemley J, Bazrafkan S, Corcoran P.; *Deep Learning for Consumer Devices and Services: Pushing the limits for machine learning, artificial intelligence, and computer vision*. IEEE Consumer Electronics Magazine. 2017 Apr;6(2):48-56.
 - Bazrafkan S, Nedelcu T, Filipczuk P, Corcoran P.; *Deep learning for facial expression recognition: A step closer to a smartphone that knows your moods*. In Consumer Electronics (ICCE), 2017 IEEE International Conference on 2017 Jan 8 (pp. 217-220). IEEE.
 - Bazrafkan S, Corcoran P. *Enhancing iris authentication on handheld devices using deep learning derived segmentation techniques*. In Consumer Electronics (ICCE), 2018 IEEE International Conference on 2018 Jan 12 (pp. 1-2). IEEE.

- **Data Augmentation:**

- J Lemley, S Bazrafkan, P Corcoran *Smart Augmentation Learning an Optimal Data Augmentation Strategy* IEEE ACCESS 5, 5858-5869. March 2017
- J Lemley, S Bazrafkan, P Corcoran *Transfer Learning of Temporal Information for Driver Action Classification* Proceedings of the 28th Modern Artificial Intelligence and Cognitive Science Conference. April 2017
- J Lemley, S Bazrafkan, P Corcoran *Learning Data Augmentation for Consumer Devices and Services* Consumer Electronics (ICCE), 2018 IEEE International Conference on. January 2018
- Bazrafkan S, Thavalengal S, Corcoran P. *An End to End Deep Neural Network for Iris Segmentation in Unconstrained Scenarios*. arXiv preprint arXiv:1712.02877. 2017 Dec 7.

- **Generative Adversarial Networks:**

- Bazrafkan S, Javidnia H, Corcoran P. *Face Synthesis with Landmark Points from Generative Adversarial Networks and Inverse Latent Space Mapping*. arXiv preprint arXiv:1802.00390. 2018 Feb 1.
- Bazrafkan S, Javidnia H, Corcoran P. *Versatile Auxiliary Classifier+ Generative Adversarial Network (VAC+ GAN); Training Conditional Generators*. arXiv preprint arXiv:1805.00316. 2018 May 1.

- **Other Techniques (SPDNN, etc):**

- Bazrafkan S, Javidnia H, Lemley J, Corcoran P. *Depth from Monocular Images using a Semi-Parallel Deep Neural Network (SPDNN) Hybrid Architecture*. arXiv preprint arXiv:1703.03867. 2017 Mar 10.
- Bazrafkan S, Corcoran PM. *Pushing the AI Envelope: Merging Deep Networks to Accelerate Edge Artificial Intelligence in Consumer Electronics Devices and Systems*. IEEE Consumer Electronics Magazine. 2018 Mar;7(2):55-61.

- **Other Key Papers (Adversarial Networks, GANs, Boundary Equilibrium-GAN, etc):**
 - Nguyen A, Yosinski J, Clune J. *Deep neural networks are easily fooled: High confidence predictions for unrecognizable images*. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition 2015 (pp. 427-436).
 - Goodfellow I, Pouget-Abadie J, Mirza M, Xu B, Warde-Farley D, Ozair S, Courville A, Bengio Y. *Generative adversarial nets*. In Advances in neural information processing systems 2014 (pp. 2672-2680).
 - **Generative Adversarial Networks (GANs)** From Ian Goodfellow et al. A short tutorial by :- Binglin, Shashank & Bhargav
 - http://slazebni.cs.illinois.edu/spring17/lec11_gan.pdf
 - Berthelot D, Schumm T, Metz L. Began: Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717. 2017 Mar 31.

- **Other Key Papers (Adversarial Networks, GANs, Boundary Equilibrium-GAN, etc):**
 - Berthelot D, Schumm T, Metz L. Began: Boundary equilibrium generative adversarial networks. arXiv preprint arXiv:1703.10717. 2017 Mar 31.
 - Rosca M, Lakshminarayanan B, Warde-Farley D, Mohamed S. Variational approaches for auto-encoding generative adversarial networks. arXiv preprint arXiv:1706.04987. 2017 Jun 15.
 - Huang B, Chen W, Wu X, Lin CL, Suganthan PN. High-Quality Face Image Generated with Conditional Boundary Equilibrium Generative Adversarial Networks. Pattern Recognition Letters. 2018 Apr 19.
 - Asthana A, Zafeiriou S, Cheng S, Pantic M. Incremental face alignment in the wild. In Proceedings of the IEEE conference on computer vision and pattern recognition 2014 (pp. 1859-1866).
 - Chrysos GG, Antonakos E, Snape P, Asthana A, Zafeiriou S. A comprehensive performance evaluation of deformable face tracking “in-the-wild”. International Journal of Computer Vision. 2018 Apr 1;126(2-4):198-232.
 - Goodfellow IJ, Shlens J, Szegedy C. Explaining and harnessing adversarial examples. arXiv:1412.6572. 2014 Dec 20.
 - Su J, Vargas DV, Kouichi S. One pixel attack for fooling deep neural networks. arXiv preprint arXiv:1710.08864. 2017 Oct 24.