ART & ML PROJECT 2

NOVELTY PETS

Caroline Wu, Claire Yuan
github.com/claireyuan/ArtML-project-2

Concept

We chose to generate fish and dogs for this project because it is interesting to see what we could potentially keep as household pets. We experimented with other living things, like tigers and plants, but we ultimately settled on fish and dogs because there were readily available datasets and the results turned out relatively realistic.

Technique

We experimented with different algorithms and datasets. With CycleGAN [1], we used a dataset of fish [3] and a dataset of flowers [4]. With DCGAN [2], we used 20,580 images of dogs from the Stanford Dogs Dataset [5].

Process

We tried a few different algorithms and datasets before deciding on the datasets for our final pieces.

CycleGAN - See Appendix 1

We tried different combinations of datasets using CycleGAN: dogs and flowers, dogs and tigers, fish and flowers. The first two combinations failed to produce meaningful results possibly because dogs do not take well to color change, and tiger stripes were too complex for the GAN. We experimented with changing lambda and image size. We also tried using least squares loss instead of sigmoid loss, which made the models converge faster.

DCGAN - See Appendix 2

We tried to generate a new breed of dog from the full dataset of Stanford Dogs. In our first few models, the discriminator learned too quickly before the generator was able to produce meaningful results. To combat this, we reduced filter size on the discriminator compared to the generator. We also tried to continue training the generator with a reset discriminator. For our final experiment, we added a layer to the DCGAN to be able to generate larger images. The given code was for 64x64, and we modified the model to handle 128x128.

Results

Using patterned fish and bright flowers, we created various colorful fish.

Method	Picture 1	Picture 2	Result
CycleGAN			

CycleGAN	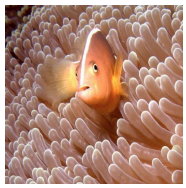		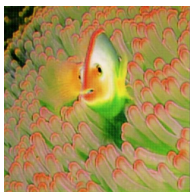
CycleGAN	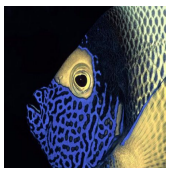		
CycleGAN			
CycleGAN			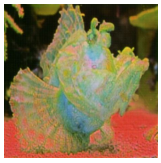

We also used DCGAN to generate new dogs from the full Stanford Dogs Dataset.

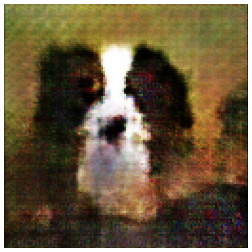


Reflection

Given more time, we would have liked to experiment with Pix2Pix. The DCGAN models on Stanford Dogs dataset was not as successful as we hoped. We hypothesize that this is at least partially because the dataset was not standardized enough. The images were of all different breeds with varying backgrounds, scales, and positions. Some of the images were closeups of dogs' faces, some were incomplete pictures of the dogs' bodies, and some were poorly lit and framed. If we had more time and resources, it would be interesting to see if our method would have worked better with a cleaner dataset, with the images carefully scaled and cropped, and the dogs in a more standard position. Ultimately, we are satisfied with our CycleGAN results and we learned a lot about the importance of dataset curation while training our DCGAN model.

Individual Contributions

Caroline

- Prototyping on CycleGAN
- CycleGAN results

Claire

- Prototyping on DCGAN
- DCGAN results

References

[1] Cycle Generative Adversarial Networks
[2] Pytorch examples. *Deep Convolution Generative Adversarial Networks.*
[3] ImageNet Fish Dataset
[4] 102 Flower Dataset
[5] Aditya Khoosla et al. *Stanford Dogs Dataset.* 2011.

Appendix 1

Method	Dataset 1	Dataset 2	Result
CycleGAN	Dogs	Flowers	
CycleGAN	Dogs	Tigers	
CycleGAN	Fish	Flowers	

Appendix 2

First try, with 64x64 images. No attempt to address discriminator-generator imbalance.

Another try with 64x64 images. Restarted the discriminator while continuing from a trained checkpoint of the generator.

Another try with 64x64 images. Reduced filter size on discriminator compared to generator.

Increased image size to 128x128. Added another convolutional layer to the model. This made training extremely slow and the models converged even slower.