ART AND MACHINE LEARNING
CMU 2019 SPRING
PROJECT 2

A Strange Terrain

Jenn Choi, Angel Hernandez, Julie Kim, Sophia Yoo

Concept

The evidence for man-made global warming has recently reached a "gold standard," as discussed in the Journal of Nature Climate Change. Their findings show that there is a one in a million chance that the climate change evident in floods, droughts, and rising sea levels is not a result of human activity – particularly, the burning of fossil fuels.[1]

It is now critical that people begin to pay attention to how their behaviors affect and interact with the earth that we live on. With this project, we decided to take images of industrial transportation-based objects in various ecosystems. We then used Pix2Pix modeling to map animals of the same ecosystem to the objects' edges.

Transportation is significant in presenting this transformation of industry to nature, as it is widely used in consuming fossil fuel globally. A typical passenger vehicle emits approximately 4.6 metric tons of carbon dioxide per year.[2] The transportation sector as a whole emits more greenhouse gases than any other sector; about 30% of global warming emissions. This means that out of 6.5 billion metric tons of CO_2 equivalent produced in the US in 2016, approximately 2 billion were the result of transportation.[3]

The animals used in our images are not only representative of the ecosystems that they are placed in, but also beneficial to the environment and contribute to the ecosystem's health in many ways. For example, we decided to use dolphins as the animal of transformation in an underwater environment. Dolphins often prey on old or ailing fish, decreasing the likelihood of spread of disease in populations. By absorbing harmful pollutants in their prey, dolphins not impact fish populations but also provide useful information to marine biologists about ocean toxicity.[4] Much like dolphins, the animals in this project play a key role in their ecosystems.

Techniques

Image Scraping

We built our datasets by scraping all of the animal images we would need to perform Pix2Pix. We used the Google Image Download python script written by *hardikvasa*.[5] By specifying keywords, image formats, and the number of images desired to be scrapped, we were able to pull the necessary datasets.

Edge Extraction

For edge extraction, we used Open CV's implementation Canny Edge Detection. We first tried to use the algorithm for Holistically-Nested Edge Detection. However, we were unable to complete this implementation due to complications with Caffe in the Deep Learning AMI. For the parameters,

---

[1]https://www.reuters.com/article/us-climatechange-temperatures/evidence-for-man-made-global-warming-hits-gold-standard-scientists-idUSKCN1QE1ZU
[2] https://www.epa.gov/greenvehicles/greenhouse-gas-emissions-typical-passenger-vehicle
[3] https://www.ucsusa.org/clean-vehicles/car-emissions-and-global-warming
[4] https://education.seattlepi.com/dolphins-ecological-importance-5511.html
[5] https://github.com/hardikvasa/google-images-download

we used Adrian Rosebrock’s method to automatically optimize lower and upper bounds of pixel intensity gradients.[6] We use the median single channel pixel intensities to calculate our bounds:

$$lower = max(0,\ (1 - \sigma) * median)$$
$$upper = min(255,\ (1 + \sigma) * median)$$

This essentially uses a range of pixel intensity values based off of the median value, or the minimum (0) / maximum (255) if values computed using the median are out of bounds. We also played with various σ values to reduce noisy edges while preserving the main frame of the animals.

Figure 1: Canny Edge Extraction

Because some of the photos had white backgrounds with clear edges and others had various backgrounds with unclear edges, it was a challenge to tune the parameters such that all edges turned out well. Generally, images with white backgrounds performed much better, while natural photos had much more noise.

Pix2Pix

We trained the Pix2Pix model to use our animal images and edges in mapping edge-to-animal. We used approximately 400 images of dolphin, 300 of rabbits, 300 of Namib beetles, and 450 of tang fish – trained separately. Training proved to be difficult, as Jupyter Notebook crashed continually during this stage.

Process

Jenn Choi

Using Google Image Download script, we collected about 1000 images by specific search description along with the animal we were looking for. Then, we filtered through our dataset to filter out any miscellaneous objects that would interfere with our desired training outcomes, such as images containing cartoons, text, or more complex backgrounds. We did keep images of animated animals if they were realistic and had all the features of that animal.

We expanded our animal dataset to namib desert beetles, rabbits, and yellow/purple tang fish. After gathering a variety of animal dataset, we reduced the images in the dataset to fit a 286 by 286 dimension. We overlaid images over a white background for images that were either smaller than the 286 by 286 dimensions or had non-rectangular edges. This way, we were able to prevent png images from crashing as some images had transparent backgrounds. Using the Open CV’s

[6] https://www.pyimagesearch.com/2015/04/06/zero-parameter-automatic-canny-edge-detection-with-python-and-opencv/

edge extraction, we generated edge images for each respective images and used the Combine_Aand_B.py from the original Pix2Pix implementation.[7]

Sophia Yoo, Angel Hernandez, Jenn Choi

We trained the Pix2Pix algorithm using leaky relu activations and an Adam optimizer for datasets on the dolphin, bunny, beetle, and fish datasets. For all of our training procedures, we used an 80-20 ratio for training and validation. In Figure 2, we can see the dolphin results getting better at each iteration. However, due to the highly variable nature of the dataset we collected, we weren't able to get amazing results. Dolphins have similar colors to water as well, which blurred the clear formulation of dolphin-like figures. Edges2Bunny worked better in some ways than Edges2Dolphins because of the bunnies' drastically different color and texture to grass. In Figure 3, it's apparent that the more clear edges (ones without grass in the background) look a lot better and clearer. We additionally trained our dataset on yellow, purple, and hippo tang. The only difference between these three fish species is their color. We trained Pix2Pix on the combined dataset, but due to the colorful coral reefs and the different colors of the fish, the trained output was extremely hazy and insignificant. (Figure 7).

Generally, for all of the animals that we tried Pix2Pix on, our dataset was a mixture of cartoon visualizations, natural photographs, cutouts, and even stuffed animals. It would have been ideal to have stock images (with white backgrounds) as that would have made all the edges a lot more clear.

Julie Kim

Finally, we drew edges on images of the selected ecosystems. In the example below, the image of an underwater plane wreck[8] was used as a base on which we illustrated three dolphins. (Figure 4) Using the training applied to our dolphin image dataset, we used the illustrated edges to generate images of dolphins:

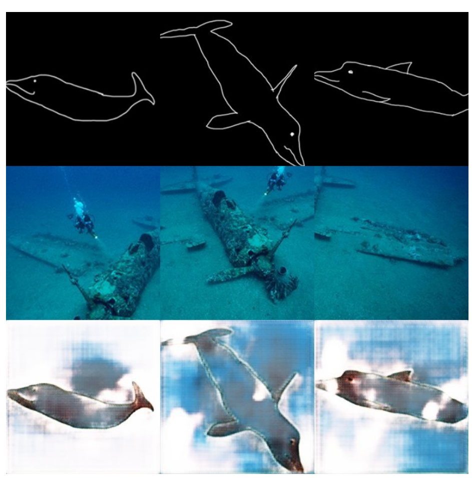

Figure 5: Edge2Dolphin using Sunken Plane Sketches

[7] https://github.com/phillipi/pix2pix/tree/master/scripts

[8] http://diysolarpanelsv.com/a-car-cracked-underwater-clipart.html#gal_post_202_a-car-cracked-underwater-clipart-42.jpg

Another example of the Edge2Animal result is through an urban park[9] in which parked passenger vehicles were replaced with rabbits:

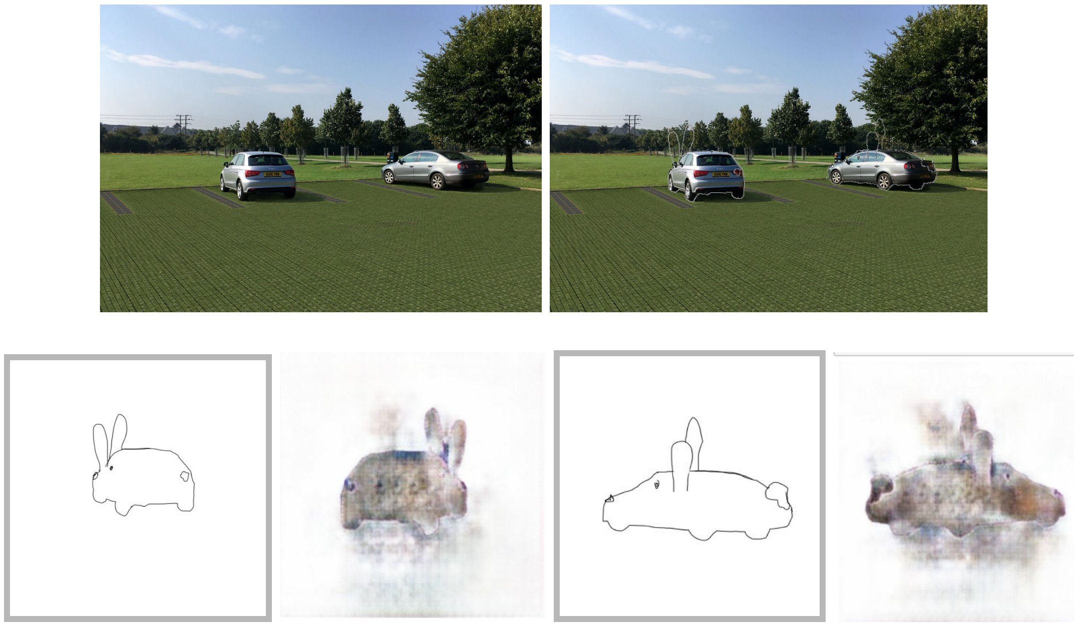

Figure 6: Edge2Bunny using Park Sketches

After the image generation, the final results were trimmed of any excess noise and overlaid on top of the original images in Photoshop. The results section displays the final product achieved through this step.

Results

[9] https://ecogrid.co.uk/popular-uses/grass-parking-surfacing/

 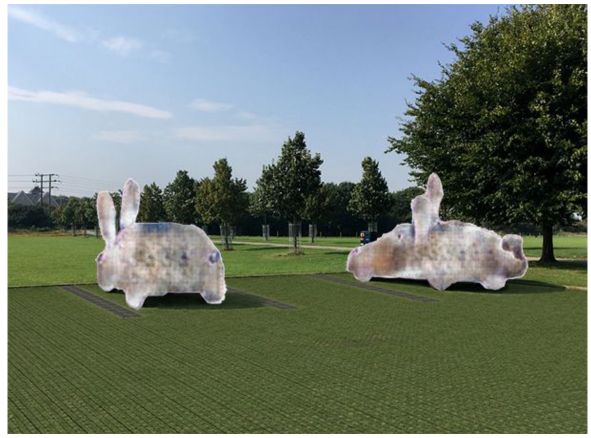

Code: https://github.com/jupiek/artML

Results: https://github.com/jupiek/artML/tree/master/Result

Conclusion

We set out with hope of shedding light on global warming and how it is affecting our environment and our surrounding animals. Along the way, we realized most of our time was spent on gathering and preprocessing data rather than the fine-tuning of the Pix2Pix algorithm for our desired images. We found image scraping is not a trivial task because returned results tend to be noisy and need further preprocessing. Furthermore, an additional algorithm which could have clustered our scrapped images into cartoons, realistic and white backgrounds would have been beneficial, so we could have trained the Pix2Pix model with segregated datasets. Like many other machine learning applications where 80% of your work is spent during data engineering, this project was no different. With that said, our team feels we were still able to generate intriguing images while highlighting how automotive vehicles are negatively affecting our wildlife.

Appendix

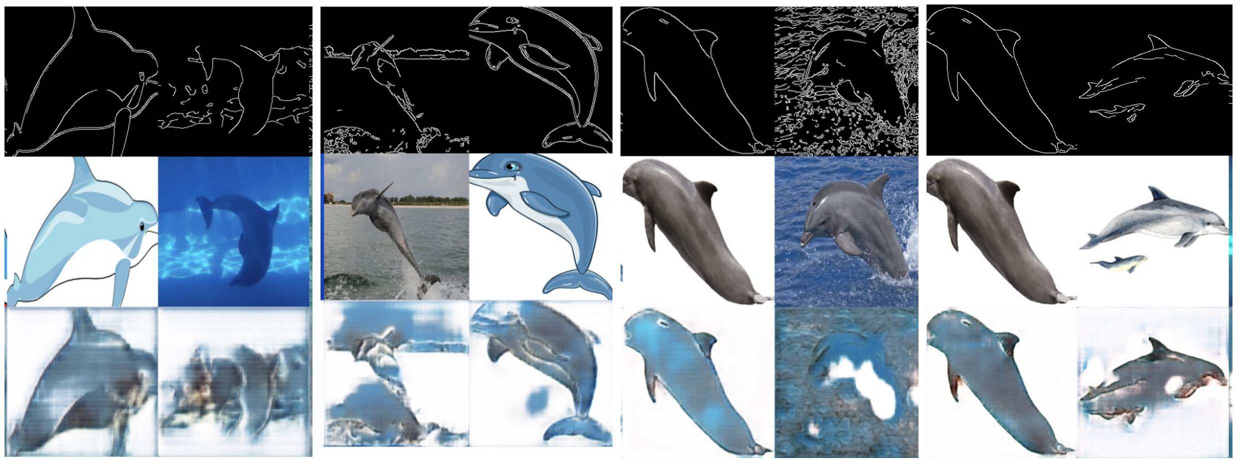

Iteration 15 Iteration 25 Iteration 36 Iteration 49

Figure 2: Edge2Dolphin Progress

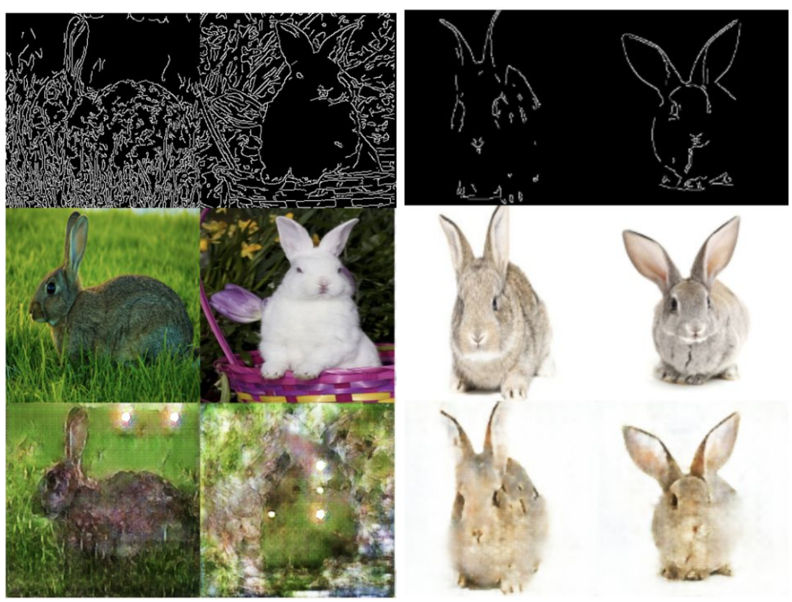

Figure 3: Edge2Bunny Final Iteration

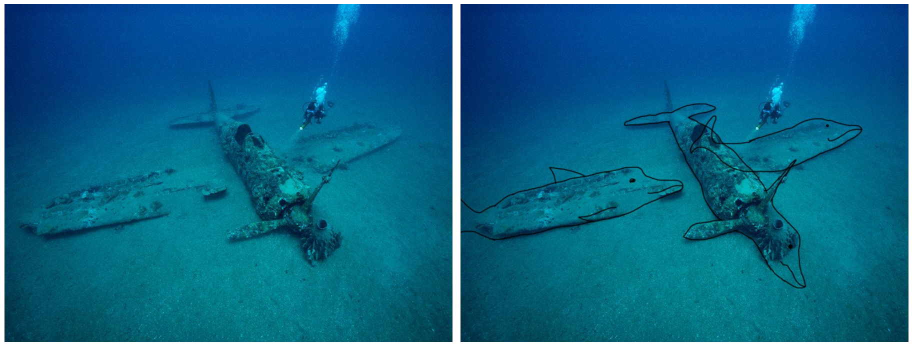

Figure 4: Sunken Plane Dolphin Sketches

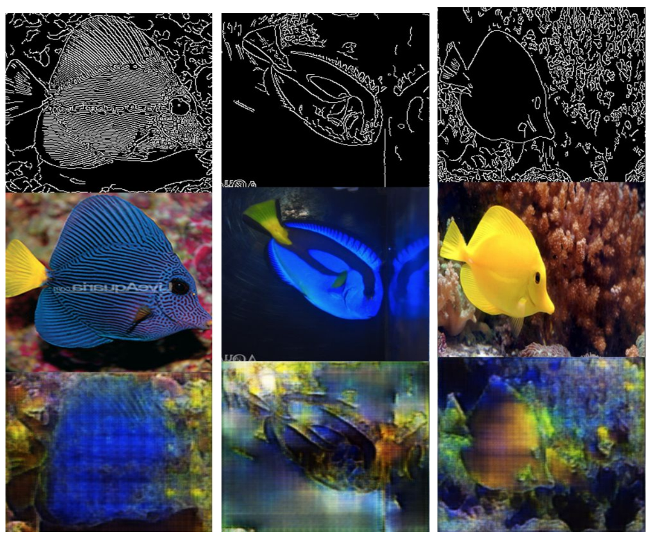

Figure 7: Edge2Fish