CycleGANS for Historical Image Colorization

Lucy Li - Symbolic Systems, **Vinson Luo** - Physics, **Michael Straka** - Mathematical and Computational Sciences

Introduction

We aim to investigate the use of Cycle-Consistent Adversarial Networks (CycleGANs) in translating historical and modern images. These networks have previously been used for a variety of geometric and color transformations, and here, we especially face the challenge of **complete colorization** with the translation from historical to modern using an unpaired training dataset.

Previous work:

- **Isola (2016)** used conditional adversarial networks for colorization, but occasionally produced grayscale or desaturated images.
- **Zhu et al. (2017)** used CycleGANs for image translation between several domains, including artists' styles and photos, apples and oranges, and zebras and horses. They achieved good results on color and texture changes, struggling with geometric changes such as translating between dogs and cats.

Datasets

Our 256 x 256 inputs featured **groups of people**. The **historical grayscale** dataset had 936, 263, and 266 images in our training, test, and validation sets. The **modern colored** dataset had 1,511, 251, and 291 images in our training, test, and validation sets. Aside from color, these datasets also differ in camera quality, racial diversity, and clothing styles.

Problem & Model

The mapping *F*: *X* (color) → *Y* (grayscale) is trivial, so to improve model stability, we only aimed to learn one mapping *G*: *Y* (grayscale) → *X* (color) and hardcoded *F*. A discriminator *D* competes against *G*, where *G* tries to fool *D* into classifying its generated images as real. At the same time, *D* aims to accurately classify *G*'s output as fake. We set our cycle consistency loss weights so that λ2 > λ1 by 30%.

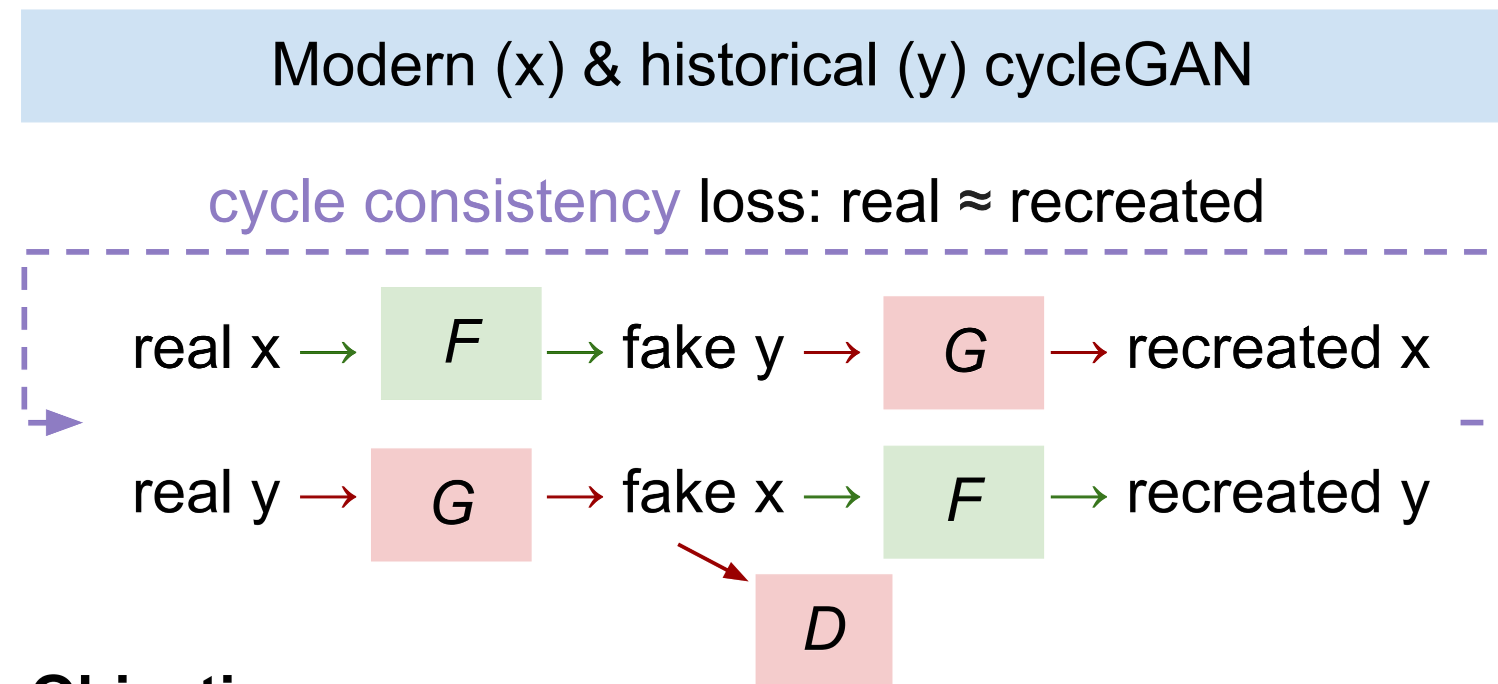

Objective:

$$\mathcal{L}_G(G,D,x,y) = \lambda_1 \mathcal{L}_{\text{cyc}}(F(G(x)),x) + \lambda_2 \mathcal{L}_{\text{cyc}}(G(F(y)),y) + \mathcal{L}_{\text{gan}}(G,y,D)$$

Evaluation and Findings

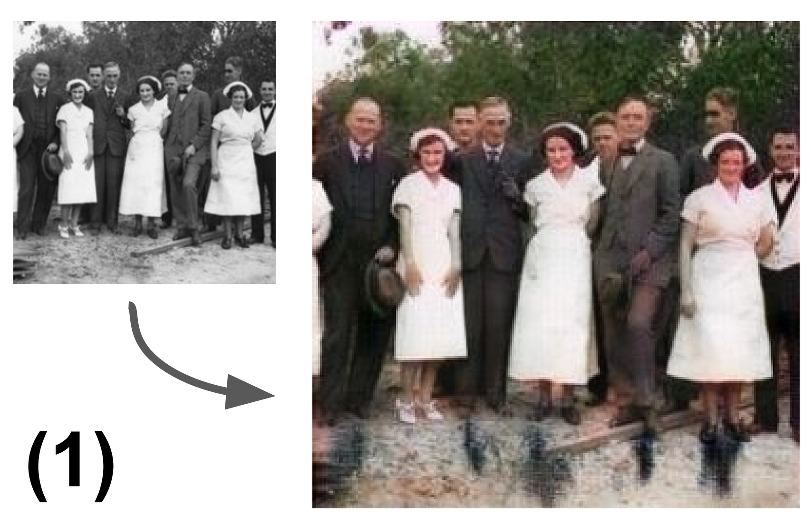

(1)

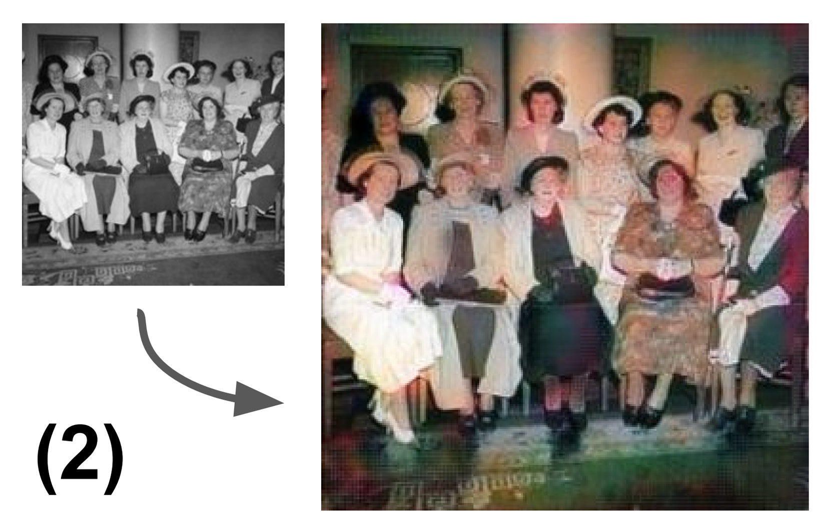

(2)

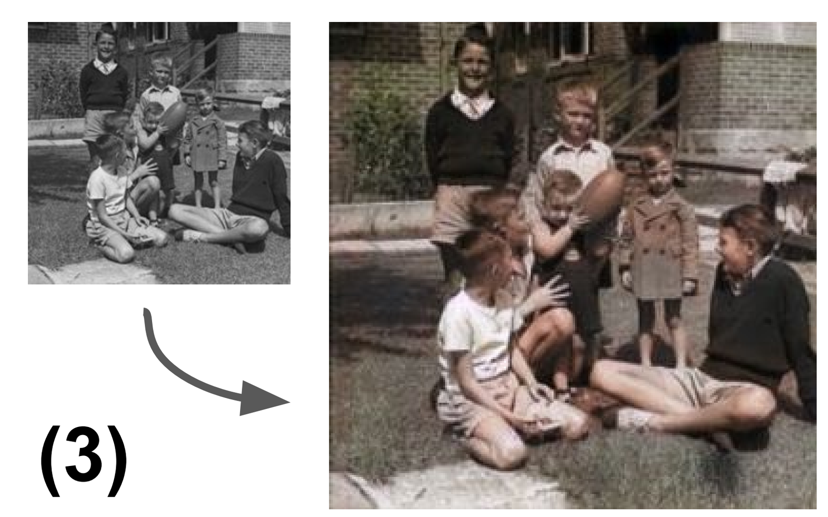

(3)

(4)

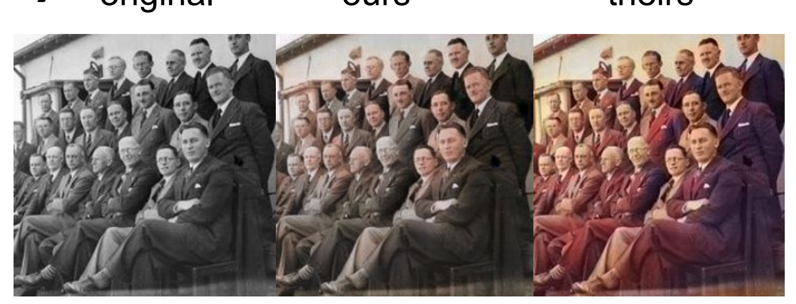

(5)

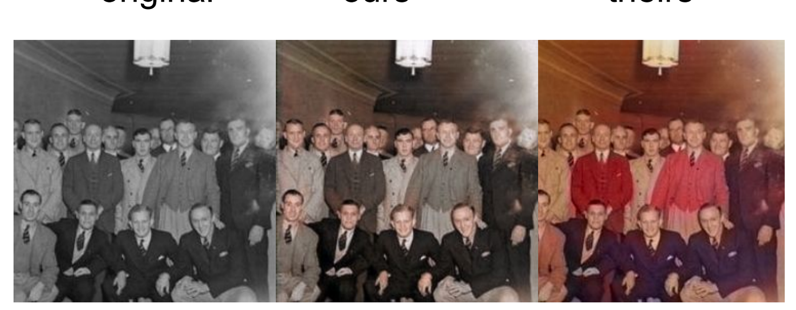

(6)

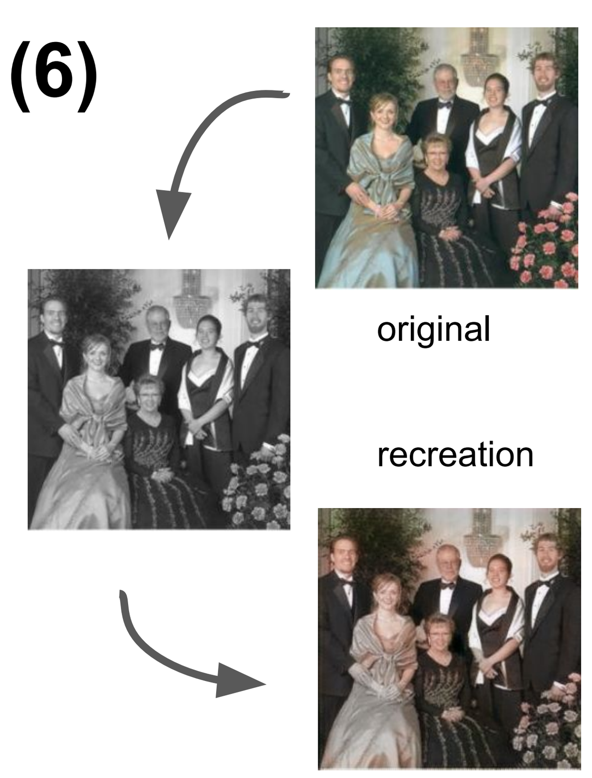

Fig. 1-3 are historical image colorizations generated by our model. Fig 4-5 compare our results with those generated by Zhang et al. (2016)'s CNN model.
Fig 6 is an example of recreating a modern image. On average, the difference for modern recreations was 0.0298 (L1) and 0.0114 (L2).

Conclusion

(7)

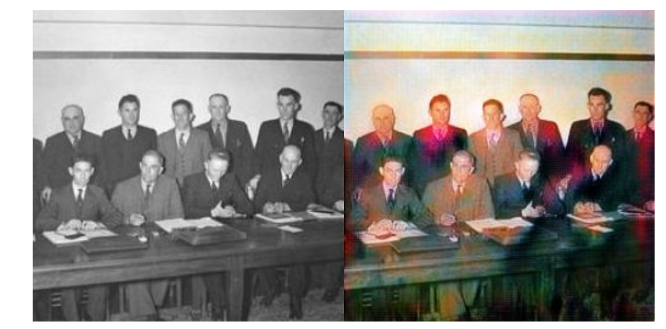

(8)

Our model was able to capture skin colors and sometimes green grass and blue skies, but it also often reduced images to sepia or hallucinated vibrant patches (Fig 7-8). Clothing, due to its diversity of possible colorings, was especially a challenge.

Future work:

- Introduce more forms of supervision in the training process. For example, we could reward brighter or rarer colors. A downside to this is that it may encourage unrealistic hallucinations.
- Experiment with different losses. Currently we use the L1 difference between images, but it may create washed out colors.
- The model may not have enough examples of the different objects in these images. Could try training on a more diverse dataset as well as a more homogeneous one (such as landscapes)

References

T. Z. P. Isola, J.-Y. Zhu and A. A. Efros. Image-to-image translation with conditional adversarial networks. arXiv preprint arXiv:1611.07004, 2016.
P. I. Jun-Yan Zhu, Taesung Park and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. arXiv preprint arXiv:1703.10593, 2017.
R. Zhang, P. Isola, and A. A. Efros. Colorful image colorization. In European Conference on Computer Vision, pages 649–666. Springer, 2016.