Practical Machine Learning

Generative Adversarial Network

Sven Mayer

Generative Adversarial Network (GAN)

speaker → $f(x)$ →

$f(x)$

GAN

Face Generation

Source: https://thispersondoesnotexist.com/
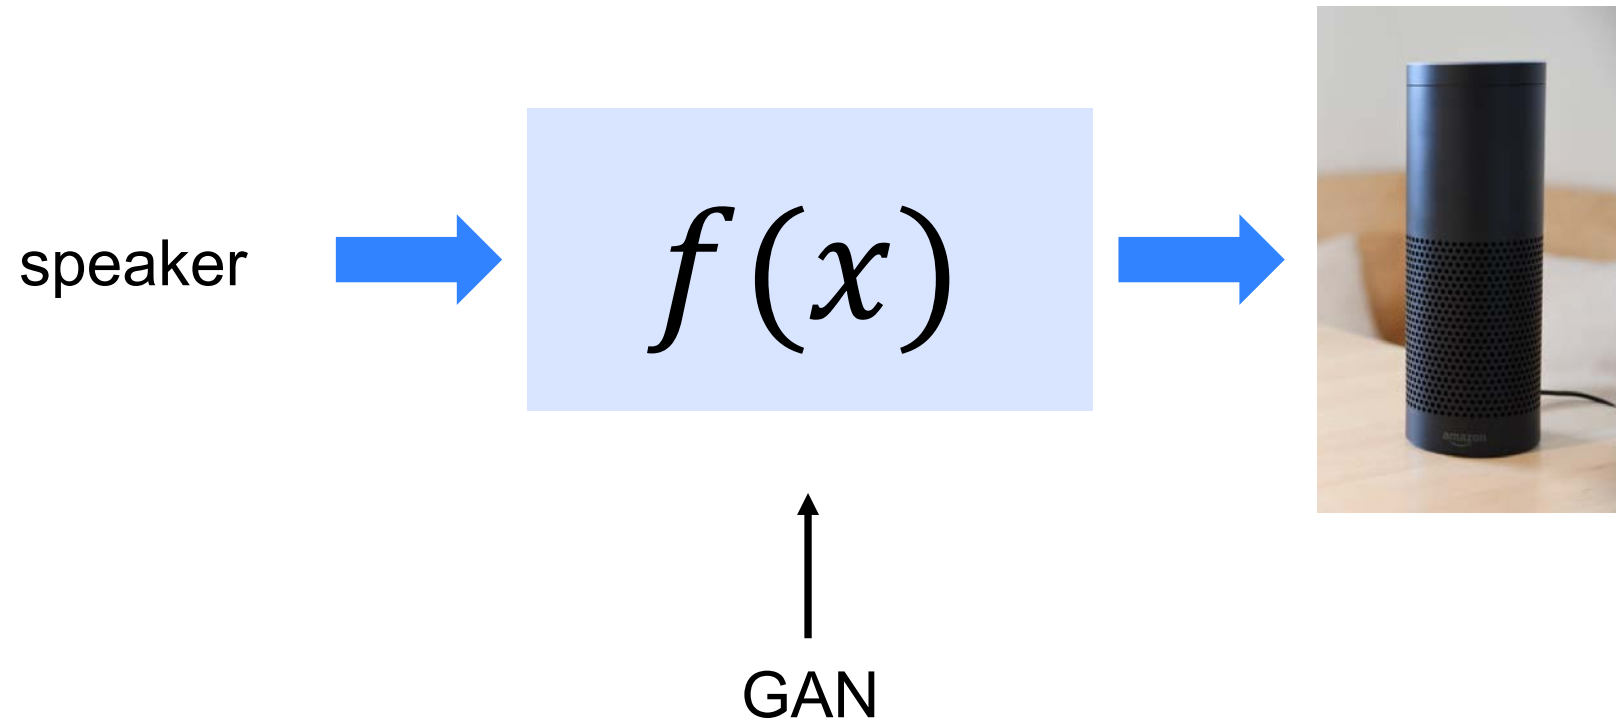
Karras, Tero, et al. "Analyzing and improving the image quality of stylegan." *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020. DOI: https://arxiv.org/abs/1912.04958
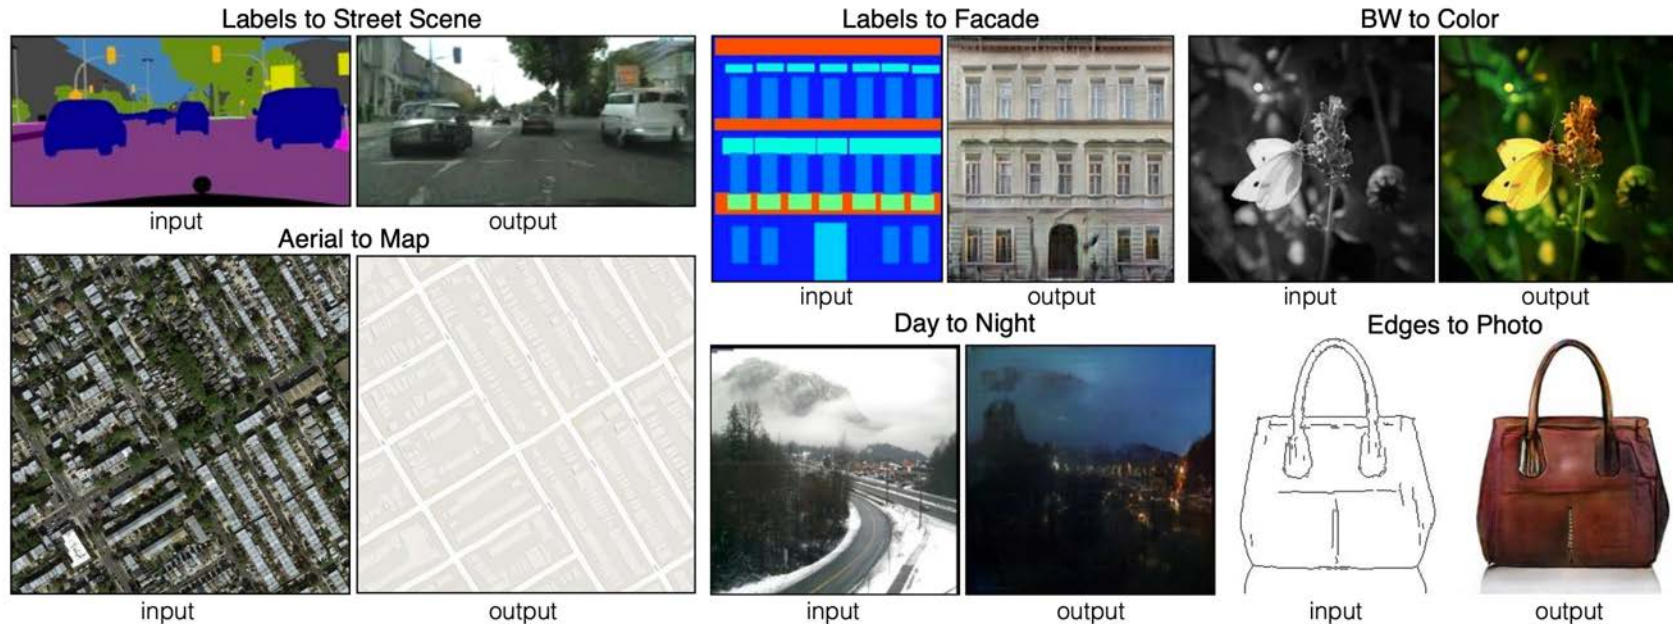
Sven Mayer

Style Transfer

Figure 1: Many problems in image processing, graphics, and vision involve translating an input image into a corresponding output image. These problems are often treated with application-specific algorithms, even though the setting is always the same: map pixels to pixels. Conditional adversarial nets are a general-purpose solution that appears to work well on a wide variety of these problems. Here we show results of the method on several. In each case we use the same architecture and objective, and simply train on different data.
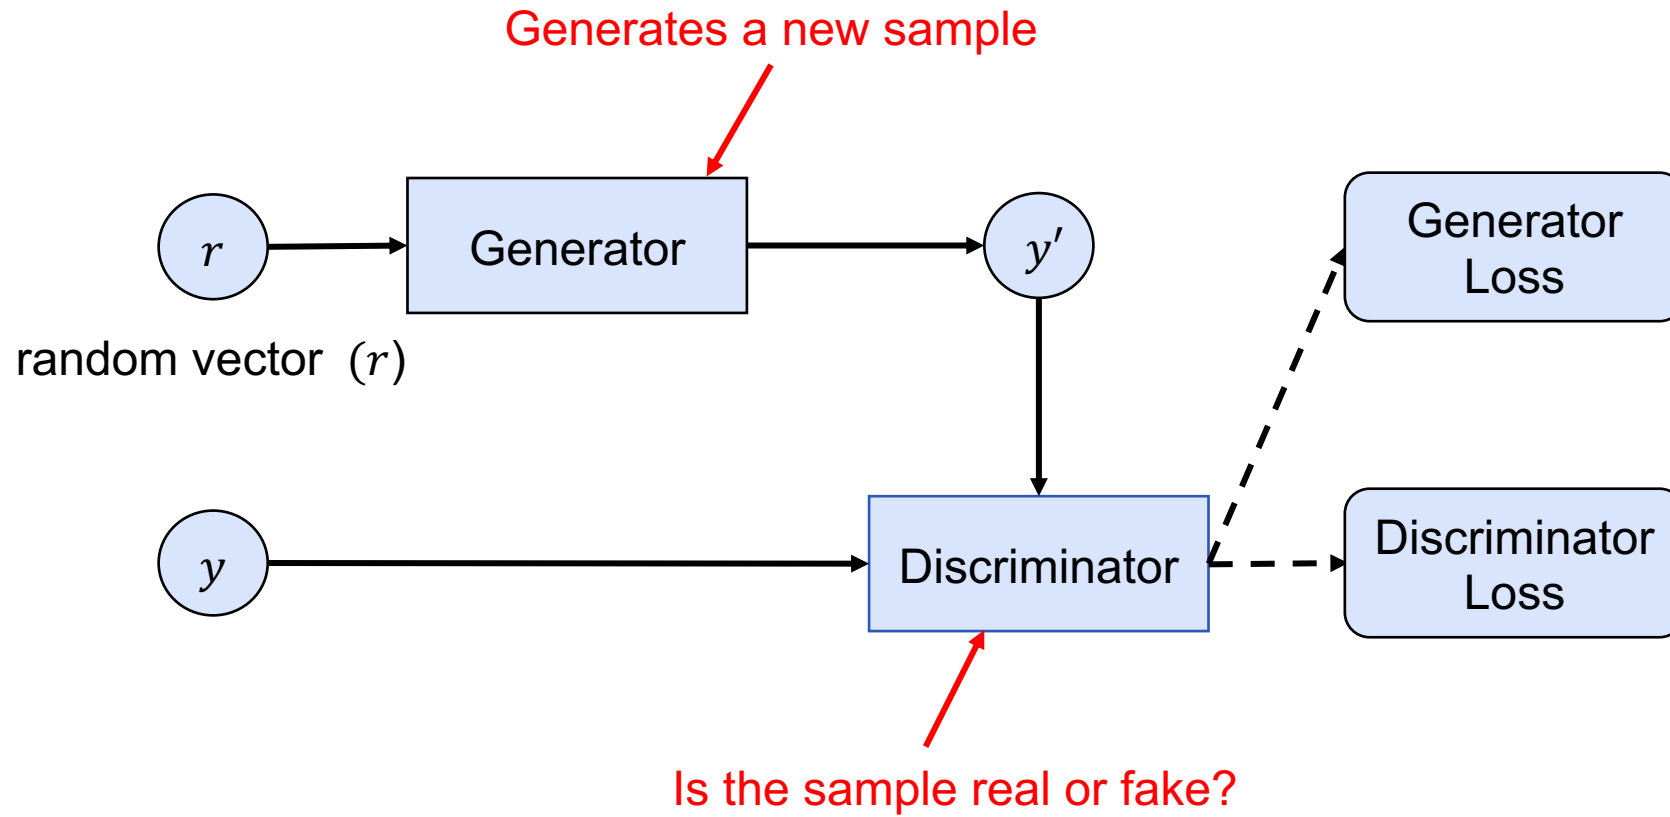
Sven Mayer

Generative Adversarial Network

- Create Data

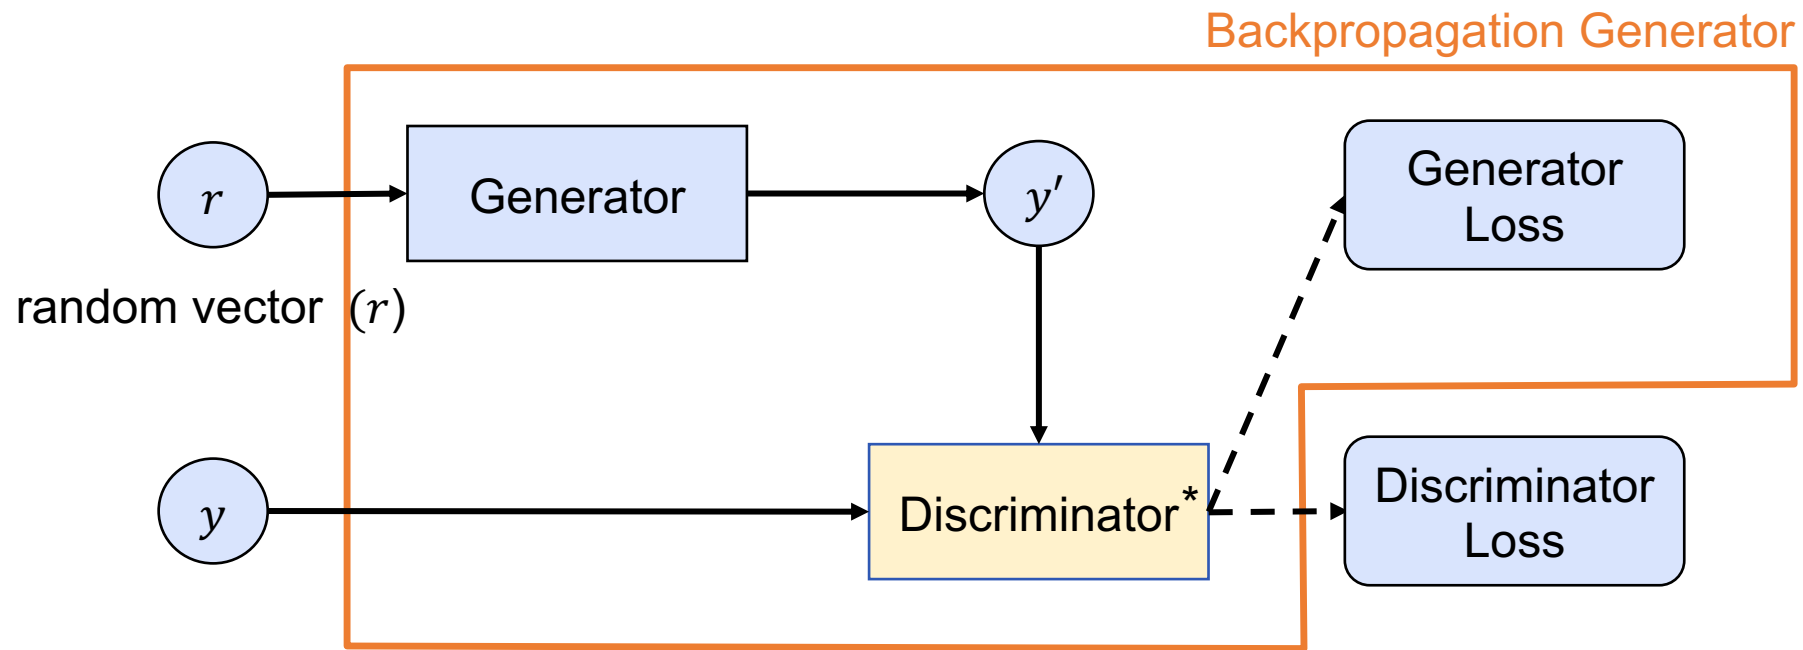- Style Transfer

- Increase Image Resolution (super-resolution)

- …

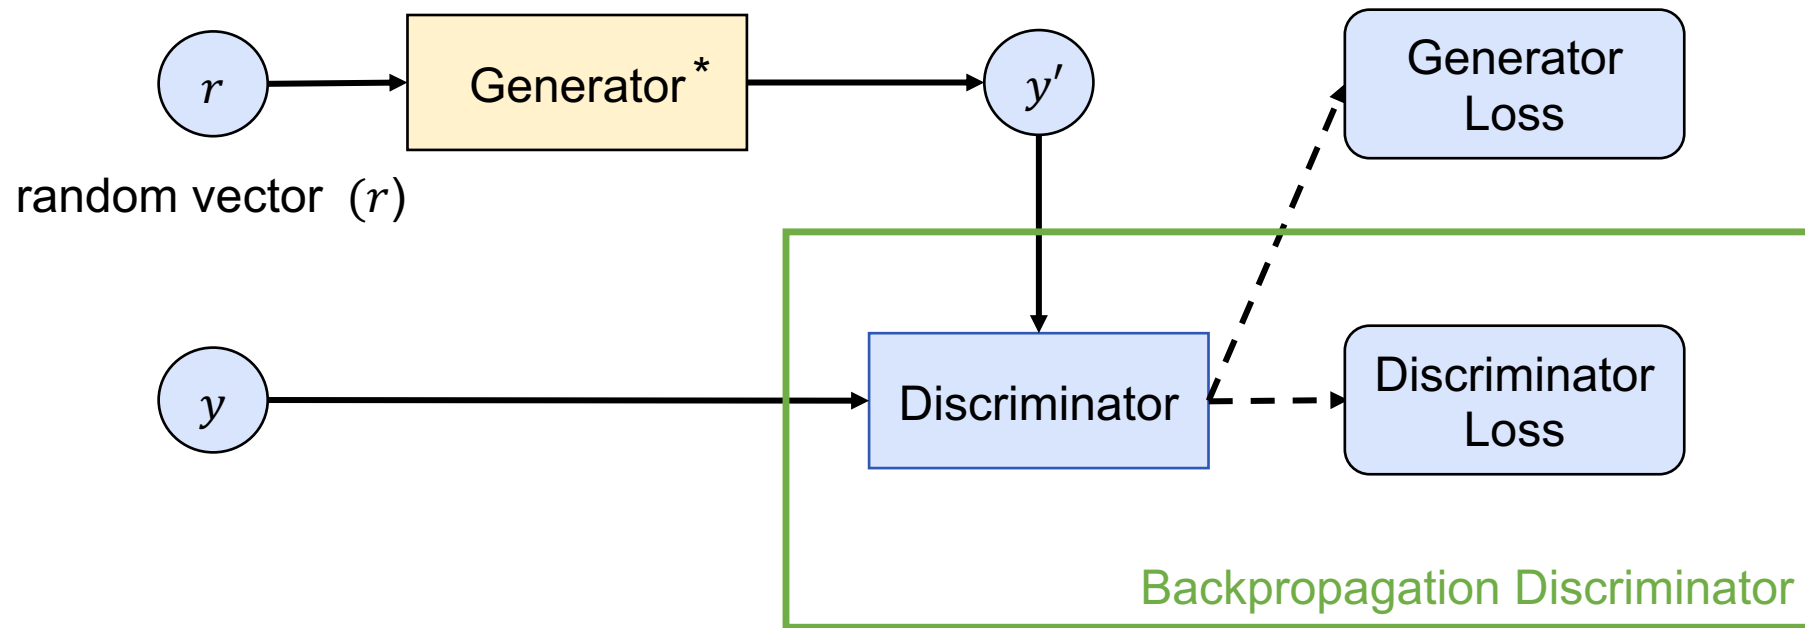# Generative Adversarial Network
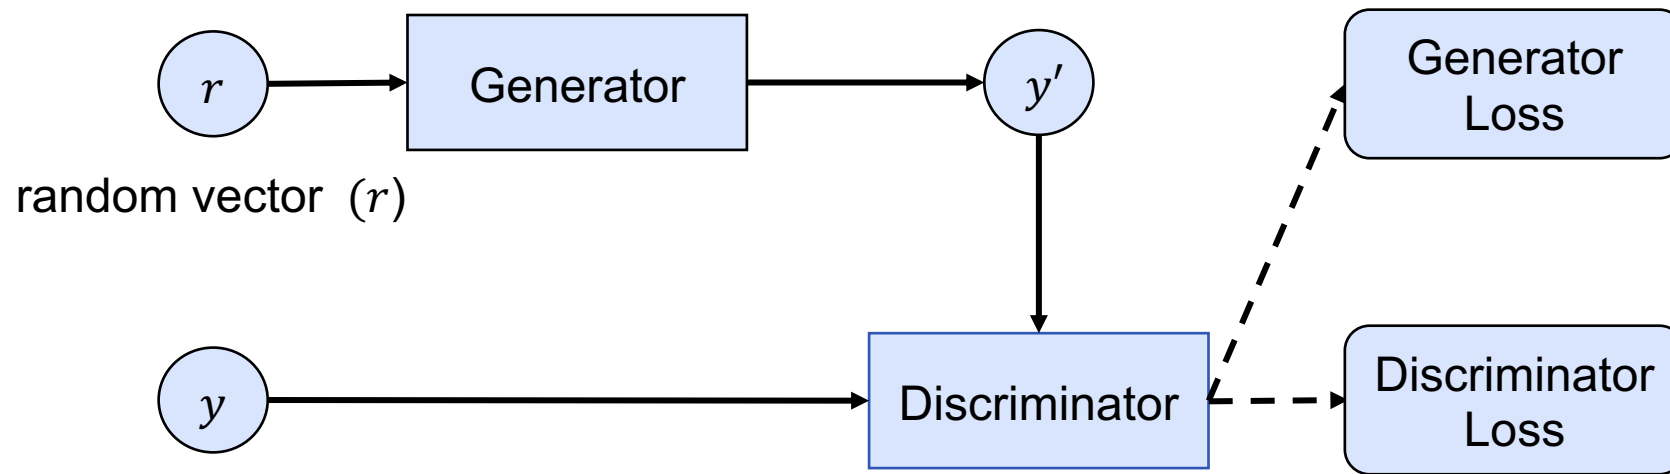
Generative Adversarial Network

Generator Training

* Frozen model

Discriminator Training

random vector (r)

* Frozen model

Training

During the training process the generator and the discriminator improve each other by both getting better each trying to outperform the other model.

Sven Mayer

What is a good Generator?

In the end the generator is should be so good to fool the discriminator. Thus, the discriminator needs to guess fake or real (accuracy is ~.5).

Sven Mayer

Conclusion

Generative Adversarial Network

- Generator

- Discriminator

- Generator tries to outperform the Discriminator to detect is if it is a real or fake image

License

This file is licensed under the Creative Commons

Attribution-Share Alike 4.0 (CC BY-SA) license:

https://creativecommons.org/licenses/by-sa/4.0

Attribution: Sven Mayer