Computer Graphics: Style transfer

The problem: style transfer

- Let us say that you have a photo of a summer landscape and you want to make it look like:
  - A painting by Van Gogh
  - It is winter instead of summer
  - It is a night instead of day
  - It is on a different planet
- The problem of **style transfer** is to take an image and change its style while preserving its **content**.
- In the case of style transfer, we assume that there is a second image that provides the style, and we want to transfer that style to the first image.

Computer graphics

- Style transfer is a problem in computer graphics.
- There are many related problems in computer graphics, such as:
  - Image editing: changing the content of an image while preserving its style
  - Image synthesis: creating a new image from scratch using a model of the world and a rendering algorithm
  - Image manipulation: changing the style of an image while preserving its content

Style transfer

- **Inputs**:
  - A content image (e.g., a photo of a summer landscape)
  - A style image (e.g., a painting by Van Gogh)
- **Outputs**:
  - A new image that combines the content of the first image with the style of the second image
- **Performance measure**:
  - User satisfaction - how much do people like the output image? This can only be assessed by showing the output image to people and asking them to rate it.
    - Expensive and time-consuming.
  - Some quantitative measure of how well the content is preserved
  - Some quantitative measure of how well the style is transferred

Let's hack together a style transfer program

- Styles can be a lot of things, but we will focus on the style of a painting, which is often characterized by:
- The colors used
  - We can modify the colors of the content image to match the colors of the style image.
  - We can create a color histogram of the style image and then modify the colors of the content image to match that histogram.
- The brush strokes
  - We can modify the texture of the content image to match the texture of the style image
  - For instance, we can take a picture of a textured surface (e.g., a piece of fabric) and add it to the content image to create a new image that has the same content

Example: color transfer

- For some painting styles the colors are so distinctive that just changing the colors of the content image can create a new image that looks like it has the same style as the style image.
    - eg. Picasso's blue period, or Van Gogh's yellows

- This would work to some degree. However, we want to have a more general definition

How it works: neural style transfer

- In 2015, a paper called "A Neural Algorithm of Artistic Style" was published by Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge.
- The paper describes a method for transferring the style of a painting to a photo using a convolutional neural network (CNN).
- The method works by optimizing a loss function that measures the difference between the content of the content image and the style of the style image.

Neural style transfer: what is style and content?

- Remember that a CNN is a hierarchical model that learns to extract features from images at different levels of abstraction.
- The lower layers of the CNN learn to extract low-level features such as edges and textures, while the higher layers learn to extract high-level features such as objects and scenes.
- The style of an image is often characterized by the low-level features, while the content of an image is often characterized by the high-level features.
- The neural style transfer method uses the activations of the CNN to measure the content and style.

Neural style transfer: how to represent style?

- The style is captured using something called a Gram matrix.
- It measures correlations between feature maps
- This captures textures, colors, and patterns, not exact shapes
- So instead of “what is where,” style is “what textures exist together.”

Process of neural style transfer

- We start with a random image (or the content image itself) and we optimize it to minimize the loss function that is a weighted sum of the content loss and style loss.
- **Content loss** → difference between generated image and content image features
- **Style loss** → difference between generated image and style image Gram matrices

Other techniques for style transfer

- There are many other techniques for style transfer, such as:
  - Using a **feedforward network** that is trained to perform style transfer in one pass (instead of optimizing an image iteratively)
  - Using a **generative adversarial network (GAN)** to learn the distribution of images in the style domain and then generating new images in that style
  - Using a **variational autoencoder (VAE)** to learn a latent representation of images and then manipulating that representation to achieve style transfer
  - Using a **transformer-based model** to learn the relationships between content and style features and then generating new images based on those relationships

Applications of style transfer: creating art

- Obvious application: creating art by transferring the style of a famous painting to a photo.
- We are relying on existing styles, so not very creative, but it can be a fun way to create new images that are inspired by existing art.
- It also raises interesting questions about creativity and originality in art.
- Obvious problems with copyright, for art that is still under copyright, but even for art that is in the public domain, there are questions about whether the new image is a derivative work or a new piece of art.

Applications of style transfer: datasets and data augmentation

- Style transfer can be used to create new images for training machine learning models.
- For instance, we have a dataset of roads in summer, but we want to train a model that can recognize roads in winter. We can use style transfer to create new images of roads in winter by transferring the style of winter images to the summer images.
- This can help to improve the performance of the model on the target domain (winter) by providing it with more training data that is similar to the target domain.

Pitfalls and dangers: A race to bottom in art

- If everyone is using the same style transfer techniques, we might end up with a lot of images that look very similar, and we might lose the diversity and creativity in art.
- If creating a new artwork is just a matter of applying a style transfer technique to an existing image, is there still a market for original art?
    - Or even just a market for craftsmanship?
- Discuss

Try it out: photos in the style of famous painters