Color Me This:
Automatic Image Colorization with CNNs and Inception-ResNet-V2 combined with User-Guided Image Colorization with U-Net architecture.

Abstract
Given a grayscale image as input, we attempt to provide an estimation of the image's RGB colors, with the help of user-guided color hints. We implement automatic colorization using a CNN encoder-decoder architecture combined with a pre-trained Inception-ResNet-V2 model for feature extraction. The automatic colorization is then used to generate a palette of five suggested colors via K-means clustering. The user can then input color hints onto the grayscale image, which are fed into a U-Net architecture to generate a colorized output in real-time in a user interface that we provide. The output of the user-hint model is conditioned on an arbitrary number of hints provided.

Introduction
Color plays an important role in conveying information. As such, the aim of colorization can range anywhere from making an image more visually appealing, offering a new perspective, giving a glimpse of history, or perhaps simply allowing your grandparents to go on a vivid trip down memory lane. For this purpose, we have chosen to address a subset of image colorization—flower colorization. In flower colorization, the color channels of a grayscale image of a flower are predicted. By tackling flower colorization, we intend to deliver transferable results with the limited training resources available to us. The hope is that the colorization network can learn basic patterns: that grass is green, the sky is blue, etc. while respecting the multi-modal nature of flower color. Additionally, a model that excels at colorizing flowers may be moderately transferable to other domains by recognizing features shared across all natural images; namely object boundaries and color gradients. In order to address the above-mentioned multimodal nature of flower colors, we've chosen to integrate user hints. We condition our output images on points and corresponding colors that users provide. The hope is for the model to propagate the user hints to color entire objects, making it convenient to color general natural images. Because of our integration of user hints, images with color that's already known can still be converted to grayscale and used as an input with our model. This enables images to be recolored; a task that's normally very time-consuming. Users can provide sparse color hints that do not reflect the ground truth image and still end up with creative colorizations.

Link to GitHub repository: https://github.com/EtomicBomb/cs143-final-project-proposal