Introduction: What problem are you trying to solve and why?

The primary objective of our project is to explore and enhance the capabilities of image style transfer, a powerful and creative application of convolutional neural networks (CNNs) . Style transfer involves applying the artistic style of one image to the content of another, allowing for the creation of unique, stylized images. In our case we chose to apply this to the movie Across the Spider-Verse. Our interest in this area was sparked by seminal works such as those by Gatys et al., which introduced the neural algorithm of artistic style transfer, and subsequent improvements by Dumoulin et al. that proposed methods for more versatile style applications.

Our project seeks to re-implement and extend these foundational approaches. The reason for choosing this particular area of research stems from its blend of art and technology, presenting a fascinating challenge in computer vision and machine learning. Specifically, our focus is on optimizing the style transfer process for better speed and accuracy, potentially expanding its application to video content. This project will primarily involve elements of structured prediction and unsupervised learning, as the task does not require labeled data in the traditional sense but rather leverages the intrinsic properties of the images themselves.

Challenges: What has been the hardest part of the project you've encountered so far?

We've had a hard time finding out to get frames from the movie to use for the transfer learning part of the project, but we have overcome those challenges. We didn't expect data collection to be a major focus of the project, but since these models perform better with more data we thought to focus more on this. For some of the universes, it was difficult to get enough data since there were very few contiguous scenes in the movie with the universe. For instance, with the universe at the very end of the movie (universe 42), the scenes would constantly change between scenes from universe 42 and Miles' universe. This made some of the data collection a little tricky and a bit tedious at times.

Insights: Are there any concrete results you can show at this point? How is your model performing compared with expectations?

Yes, we have completed our base goal of implementing a basic version of style transfer. With just testing it on Gwen's universe, we come up with some decent results but had some trouble

with Miles's universe. The model definitely has learned how to separate style from content - a key part of style transfer, but arguably still has low quality outputs.

Plan: Are you on track with your project? What do you need to dedicate more time to? What are you thinking of changing, if anything?

Yes we are on track with our project and are about to meet our target goal soon. We probably want to spend more time on the transfer learning part and maybe think of other ways to tune our model to come out stronger results. Moreover, we may want to look at other papers for more insights.