

We only want to generate the face region of an image and inpainting it to the original image background. This allows our algorithm to focus just on creating faces (and not that of background) and at the same time guarantees that we do not have background changes that could interfere with the detection or tracking algorithms. To achieve this, we provide the generative model with the masked background image together with the landmark image. The masked background image still contains the forehead region of the head. Once the generator has access to this information, it can learn to match the skin appearance of the generated face to the forehead skin color. This leads to overall better quality of visual results. In cases there are multiple faces in the same image, we detect each face on the image and sequentially apply our anonymization framework. We make sure to keep the background image because we need to generate a third face that doesn't look overly artificial. That means we will blend the face with the background.