

ABSTRACT

This thesis studies the efficacy of using progressively grown GANs for use in image inpainting through constrained image generation. This method uses the pixels in a target image to constrain a GAN. An ℓ_1 error function is constructed using these constraints, and input back-propagation is used to traverse the error manifold. A result set of inputs can be calculated in the latent space of the GAN in order to produce an image with high resemblance to the target image. It is shown that large network sizes can be beneficial to the effectiveness of inpainting with constrained image generation.