

On this \*really\* informal talk, I would like to make the group aware of our recent efforts in applying modern Deep Generative models to retinal images. We will start by giving a description on pix2pix, a model that works by translating a representation of an image into a different one through the minimization of an adversarial loss. In essence, pix2pix is a wonderful texture filler, provided that we have the geometry of the images we want to generate. I will then describe how we applied pix2pix to generate retinal images out of binary vessel trees, and explain a method that Pedro created to avoid dependence on the previous availability of such vessel trees. This extension allows us to produce a potentially infinite amount of retinal images by simply sampling a multivariate normal distribution.

The talk will be almost slides-free. The goal is not to defend our work, but rather to discuss and explore potential applications of this knowledge on different image modalities and other problems from areas of the group that do not deal with standard color images of the eye fundus.

Deep generative models are now at a peak of popularity similar to the one we experienced five years ago with Convolutional Neural Networks. It is much likely that in a matter of several months/years, all the knowledge developed in the general computer vision community will massively translate to the medical image field (as has already happened with CNNs). The work presented in this talk has been submitted a few months ago to a top-journal, with a positive feedback from the reviewers. It may be of great importance that we build on this experience, and try to create a transversal knowledge on this kind of techniques among all of us, since this could help us to better position ourselves as a group.