Homework 5

Using the assigned reading listed on the course page, answer the questions below with a short response. Note that we are looking for concise statements that show understanding, not quantity. The total discussion should roughly be a page.

Generative adversarial networks

1. Goodfellow et al. state that the generator can model the true data distribution.

   (a) What are the assumptions made on the generator or the discriminator, if any? That is under what assumptions can we state that $p_G \rightarrow p_{data}$

   (b) In practice, we often see that the generator does not converge to the dataset distribution? What could be some of the reasons why this happens? (beyond the assumptions made in question 1) Any ideas on how we could fix it?

2. Suppose you train a GAN on images $\{x_i\}_{i=1}^N$. You draw a sample from the learned generator G, and it is a perfectly realistic image! However, as you draw more samples, you notice that they are all identical; you have

   $$G(z) = x_1 \;\; \forall \, z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \tag{1}$$

   This is a scenario known as "mode collapse" or "the Helvetica scenario". x_1 is now both a real image from the training data and also a fake image output by the generator. Will an optimal discriminator classify x_1 as real or fake? What score will the optimal discriminator assign to this "collapsed" generator? Explain your reasoning.

3. In CycleGAN, the goal is to map from domain A to domain B and vice versa. Hence, we need 2 generators and 2 discriminators for 2 domains. Suppose now we have n domains in which we want to map from any arbitrary domain to any other domain. How many generators and discriminators do we need now? (e.g., how does it grow with n?) Can you think of a better approach that would make it more computationally efficient?

Scene representation learning

1. Today, the majority of computer vision research is done with convolutional neural networks. What implicit assumptions does this make about our world?

2. Can you imagine a different way of processing images and reasoning about our world?

Submission: Upload a PDF of your response through Canvas by **10/19 at 1pm**.