Questions

"Training"

1. What is an autoencoder?
2. What is the loss function used when training an autoencoder?
3. What is the difference between supervised and unsupervised learning?
4. What is an "embedding"?
5. How are word vectors (word2vec, GloVe) trained?

"Generalization"

1. What do you think would happen if the encoder output of an autoencoder is **larger** than the image size? You may assume that both the encoder and decoder have high capacity.
2. What do you think would happen if the encoder output of an autoencoder is **too small**?
3. What does **overfitting** look like in the context of an autoencoder?
4. Why do we use the sigmoid activation in an autoencoder of images?
5. Suppose that an autoencoder used only convolutional layers, with no fully-connected layers, and also no global pooling layers. Can the autoencoder be used to generate images of different sizes? If so how?
6. How can you use autoencoders for data augmentations?

(More questions will be posted soon)