

CPSC 540: Machine Learning

VAEs and GANs

Winter 2018

Density Estimation Strikes Back

- One of the hottest topic in machine learning: **density estimation**?
 - In particular, deep learning for density estimation.
- Very fast-moving, but two most-popular methods are:
 - **Variational autoencoders (VAEs)**.
 - **Generative adversarial networks (GANs)**.
- We previously focused a lot on density estimation for digits:



Figure 3: Digits obtained by linearly interpolating between coordinates in z space of the full model.

Density Estimation Strikes Back

- These models are showing promising results going **beyond digits**:

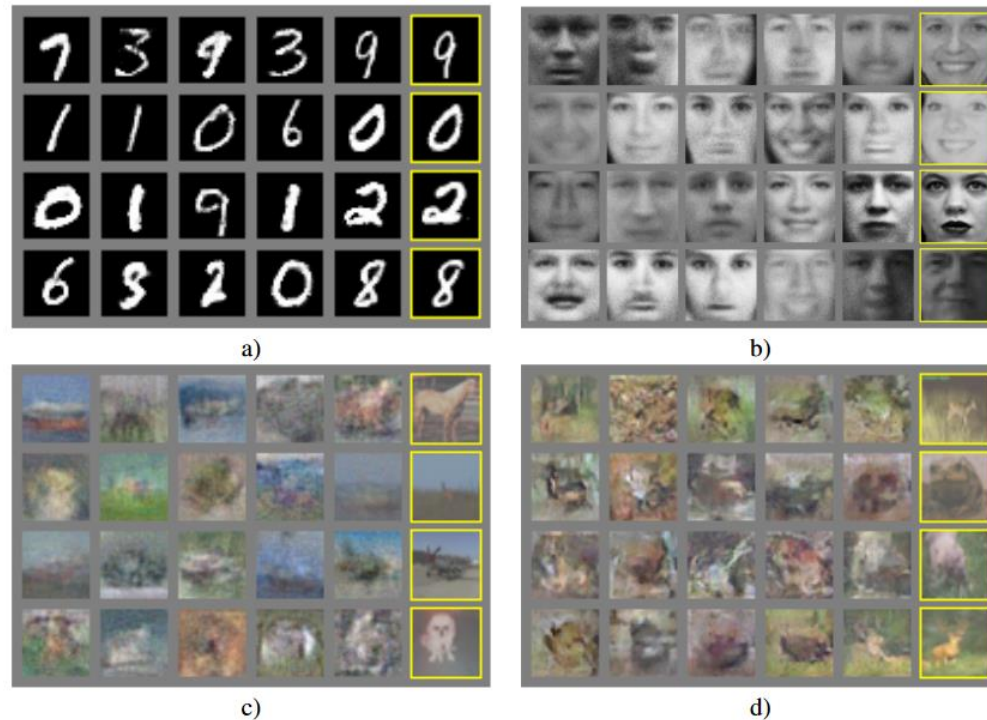




Figure 2: Visualization of samples from the model. Rightmost column shows the nearest training example of the neighboring sample, in order to demonstrate that the model has not memorized the training set. Samples are fair random draws, not cherry-picked. Unlike most other visualizations of deep generative models, these images show actual samples from the model distributions, not conditional means given samples of hidden units. Moreover, these samples are uncorrelated because the sampling process does not depend on Markov chain mixing. a) MNIST b) TFD c) CIFAR-10 (fully connected model) d) CIFAR-10 (convolutional discriminator and “deconvolutional” generator)

Autoencoders

- Autoencoders are an **unsupervised deep learning** model:
 - Use the **inputs as the output** of the neural network.





- Middle layer could be latent features in **non-linear latent-factor** model.
 - Can do outlier detection, data compression, visualization, etc.
- A non-linear generalization of PCA (old idea, never really popular).

Autoencoders for Visualization

PCA


Autoencoder

B



Denoising Autoencoder


- Denoising autoencoders add noise to the input:



- Learns a model that can remove the noise.
 - Denoising, filling in parts of the image, etc.


Autoencoders as a Generative Model

- Good autoencoder would encode any image to latent space 'z'.
 - **Encoder** converts image to a continuous space.



- **Decoder** converts from **any continuous 'z'** to images.
- We can **view the decoder as a generative model**:
 - If we sample a 'z', decoder should turn this into a realistic sample.


Problem with Basic Encoders as Generative Models



- Unfortunately, there is a problem with training this model.
 - It could “overfit” by mapping each image to a different point in ‘z’ space.
- Variational autoencoders:
 - Consider marginal likelihood over probabilistic decoding.
 - Add ‘z’ distribution regularizer, usually encouraging closeness to Gaussian.

Variational Autoencoder (VAE)

- Variational autoencoders (VAEs) have the same structure:
 - Encoder network $q(z | x)$, outputting parameters of a distribution.
 - Usually the mean and variance of a Gaussian, so takes 'x' and gives a Gaussian.
 - Decoder network $p(x | z)$, same as before (takes a 'z' and gives an 'x').
 - Prior distribution $p(z)$, usually a $N(0, I)$ distribution.



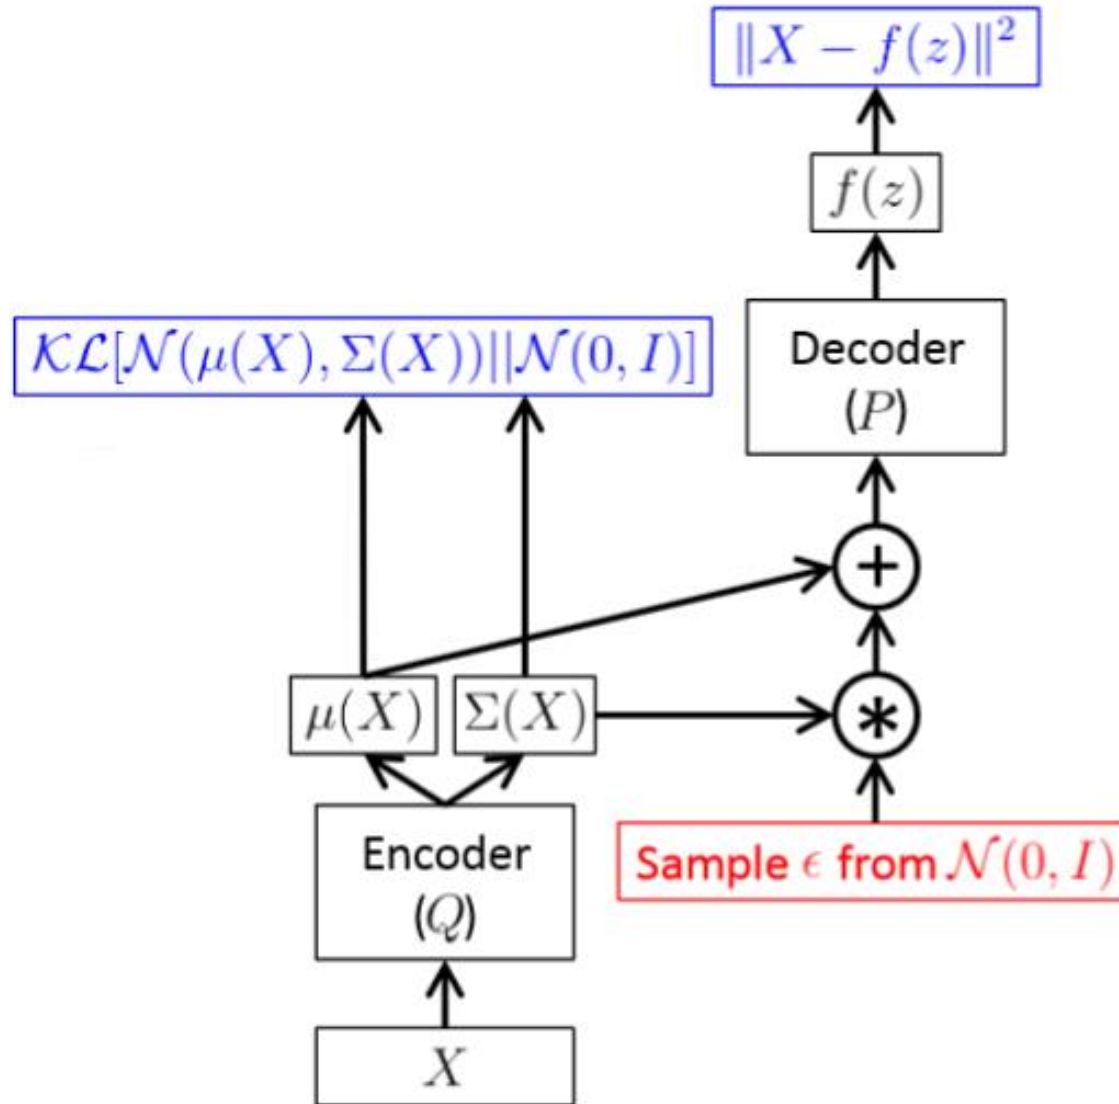
Training Variational Autoencoders

- Training: minimize **marginal decoder NLL**, regularize by prior:

$$-E_{z \sim q_{\theta}(z|x_i)} [\log p_{\phi}(x_i|z)] + KL(q_{\theta}(z|x_i) || p(z))$$

- Trained using **stochastic gradient**:
 - “Stochastic” because you choose a training example and **sample ‘z’**.
 - Sampling from encoder network is easy (Gaussian sampling).
 - Using affine property is renamed “reparameterization trick”.
- Notice again that it’s the reverse KL for tractability.
 - Equivalent to **variational inference**:
 - Using **q(z | x)** as approximation of posterior p(z | x).

Training Variational Autoencoders





Variational Autoencoder Example: MNIST

- Samples from model applied to MNIST:



Variational Autoencoder Example: MNIST

- Visualizations of latent space:



- **Non-linear** unlike PCA, but visualization is not as nice as t-SNE.
- However, goal was to produce a **generative model**:
 - Moving through latent space generates realistic digits ([video](#)).

DRAW: VAE+RNN+Attention

- Put VAE inside RNN, add attention to “draw” images:
 - <https://www.youtube.com/watch?v=Zt-7MI9eKEo>





Figure 9. Generated SVHN images. The rightmost column shows the training images closest (in L^2 distance) to the generated images beside them. Note that the two columns are visually similar, but the numbers are generally different.

(pause)

Neural Network Generative Model


- Recall the structure of a **deep belief network** and **decoder network**:



- Notice that the **edges are backwards** compared to neural networks.
 - We “generate” the features based on the latent ‘z’ variables.
- Inference is a nightmare**: observing ‘x’ makes everything dependent.

Neural Network Generative Model

- Inference is easier if we make everything **deterministic**.



- But we **need randomization** since otherwise you generate same 'x'.
- Usual assumption: **top layer comes from multivariate Gaussian**:
 - So you sample a Gaussian, and **neural network tries to convert to image**.

Generative Adversarial Network (GAN)


- So **ancestral sampling** is really easy:
 - Sample from a Gaussian, pass the sample through the network.
- But inference is still hard under the “convert Gaussian to sample”.
 - We **can't compute the likelihood** needed for training.
 - In VAEs we used a variational approximation.
- Seemingly unrelated: we've become **really good at image classification**.
- Key ideas of **generative adversarial networks (GANs)**:
 - Use ancestral sampling in this “**generator**” network.
 - Use a second “**discriminator**” network to **decide if samples look real**.
- **Discriminator “teaches” generator** to make real-looking samples.

Generative Adversarial Networks

- The generator and discriminator networks **compete**:
 - **Discriminator network** trains to **classify real vs. generated images**.
 - Tries to maximize probability of real images, minimize probability of sampled images.
 - A standard supervised learning problem.
 - **Generator network** adjusts parameters so **samples fool the discriminator**.
 - It never sees real data.
 - Trains using the **gradient of the discriminator** network.
 - Backpropagated through the network so samples look more like real images.
- Can be written as a **saddle-point** problem:

$$\min_G \max_D V(D, G) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} [\log D(\mathbf{x})] + \mathbb{E}_{\mathbf{z} \sim p_{\mathbf{z}}(\mathbf{z})} [\log(1 - D(G(\mathbf{z})))]$$

Generative Adversarial Network (GAN)




Beyond Initial GAN Model

- Improving GANs is an active research area...




Real images (CIFAR-10)



Generated images

Beyond Initial GAN Model

- Generating album covers with convolutional GANs:
 - Used uniform rather than Gaussian.



GANs for Other Problems

- GANs for super-resolution:

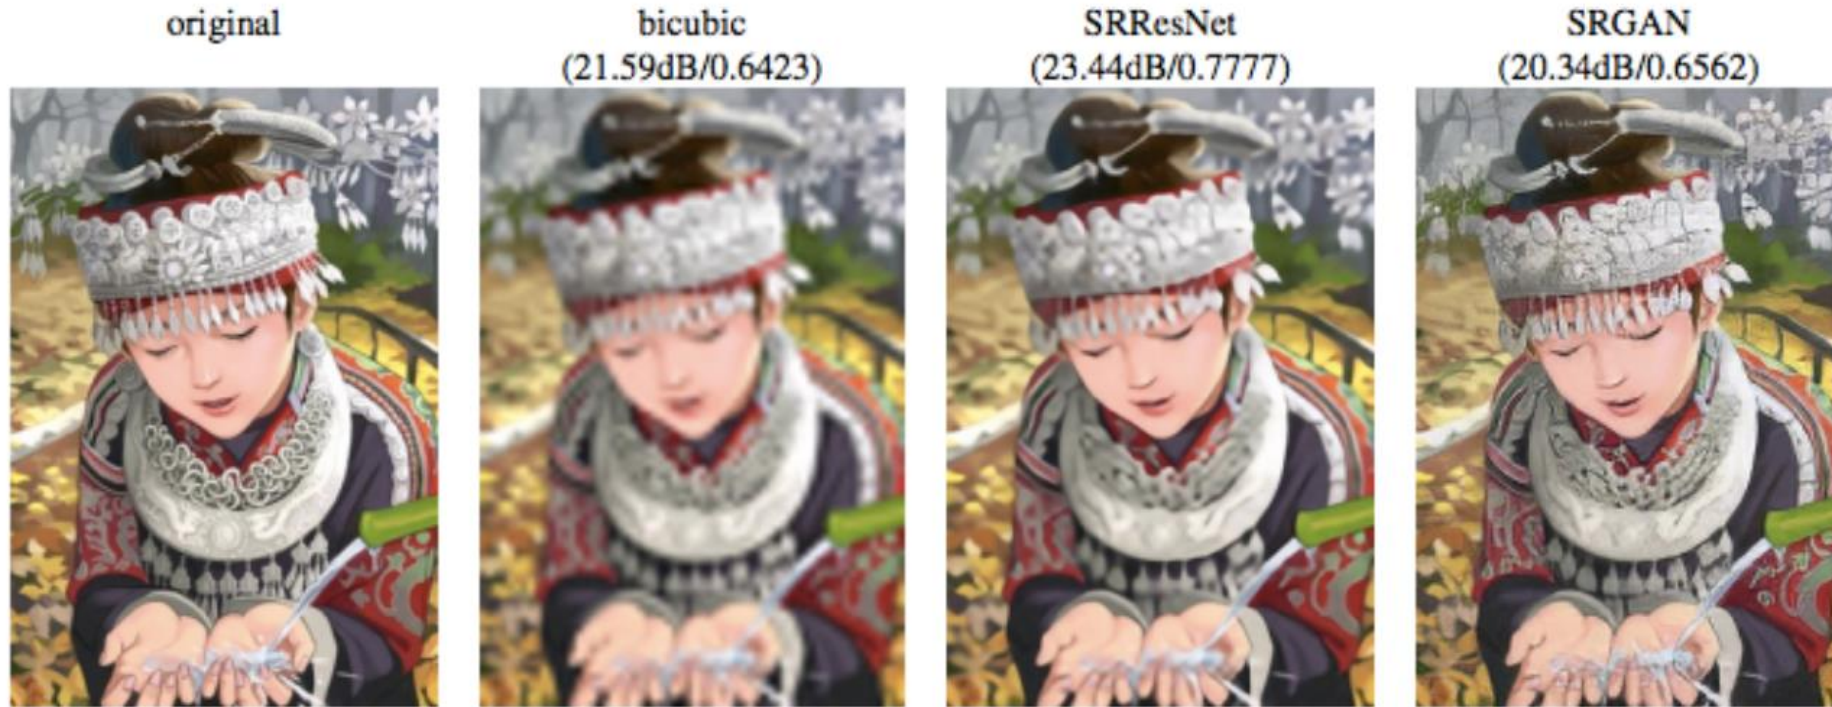


Figure 4: [Ledig et al. \(2016\)](#) demonstrate excellent single-image superresolution results that show the benefit of using a generative model trained to generate realistic samples from a multimodal distribution. The leftmost image is an original high-resolution

GANs for Other Problems

- GANs for text-to-image translation:

this small bird has a pink breast and crown, and black primaries and secondaries.



this magnificent fellow is almost all black with a red crest, and white cheek patch.



the flower has petals that are bright pinkish purple with white stigma



this white and yellow flower have thin white petals and a round yellow stamen



Figure 23: Text-to-image synthesis with GANs. Image reproduced from [Reed et al. \(2016b\)](#).

GANs for Other Problems

- GANs for text-to-image translation:

This small blue bird has a short pointy beak and brown on its wings



This bird is completely red with black wings and pointy beak



A small sized bird that has a cream belly and a short pointed bill



A small bird with a black head and wings and features grey wings



Figure 25: StackGANs are able to achieve higher output diversity than other GAN-based text-to-image models. Image reproduced from [Zhang et al. \(2016\)](#).

GANs for Other Problems

- GANs for image manipulation:
 - <https://www.youtube.com/watch?v=9c4z6YsBGQ0>
 - <https://www.youtube.com/watch?v=FDELBFSeqQs>

GANs for Other Problems

- GANs for image-to-image translation:



- <https://affinelayer.com/pixsrv>



Figure 7: [Isola et al. \(2016\)](#) created a concept they called image to image translation, encompassing many kinds of transformations of an image: converting a satellite photo into a map, converting a sketch into a photorealistic image, etc. Because many of these conversion processes have multiple correct outputs for each input, it is necessary to use generative modeling to train the model correctly. In particular, [Isola et al. \(2016\)](#) use a GAN. Image to image translation provides many examples of how a creative algorithm designer can find several unanticipated uses for generative models. In the future, presumably many more such creative uses will be found.

GANs for Other Problems

- Recent works try to avoid needing to have image pairs:
 - Adds extra part regularizing mapping in both directions.



In Progress...



Figure 30: GANs on 128×128 ImageNet seem to have trouble with the idea of three-dimensional perspective, often generating images of objects that are too flat or highly axis-aligned. As a test of the reader's discriminator network, one of these images is actually real.

Figure 29: GANs on 128×128 ImageNet seem to have trouble with counting, often generating animals with the wrong number of body parts.




Figure 31: GANs on 128×128 ImageNet seem to have trouble coordinating global structure, for example, drawing "Fallout Cow," an animal that has both quadrupedal and bipedal structure.

May not work as well in real life as in papers:

<https://twitter.com/search?q=edges2cats>

<https://arxiv.org/pdf/1701.00160.pdf>

Improving Resolution

- New generative models are appearing at a very-fast rate:




Figure 33: PPGNs are able to generate diverse, high resolution images from ImageNet classes. Image reproduced from [Nguyen et al. \(2016\)](#).

Improving Resolution

- A lot of work on trying to **improve resolution**:
 - Fake celebrities: <https://www.youtube.com/watch?v=VrgYtFhVGmg>

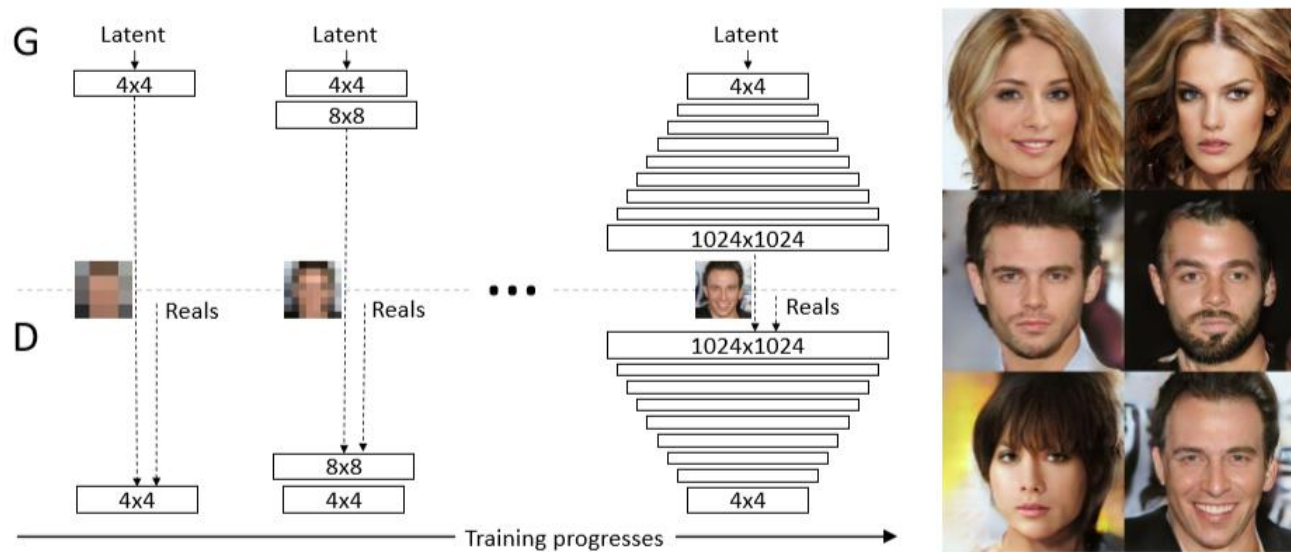


Figure 1: Our training starts with both the generator (G) and discriminator (D) having a low spatial resolution of 4×4 pixels. As the training advances, we incrementally add layers to G and D, thus increasing the spatial resolution of the generated images. All existing layers remain trainable throughout the process. Here $N \times N$ refers to convolutional layers operating on $N \times N$ spatial resolution. This allows stable synthesis in high resolutions and also speeds up training considerably. On the right we show six example images generated using progressive growing at 1024×1024 .

(pause)

Remaining Topics

- Major topics we didn't cover in 340 or 540:
 - Online learning (data coming in over time).
 - Active learning (semi-supervised where you choose examples to label).
 - Causality (distinguishing cause from effect.).
 - Learning theory (VC dimension).
 - Probabilistic context-free grammars (recursive version of Markov chains).
 - Relational models (“object oriented” graphical models).
 - Sub-modularity (discrete version of convexity).
 - Spectral methods (consistent HMM parameter estimation).
- The biggest topic we didn't cover is probably **reinforcement learning**:
 - Read Sutton and Barto's “Introduction to Reinforcement Learning”.
 - You can also take EECE 592 or Michiel van de Panne's graphics course.

A Word of Caution

- ML world is really exciting right now, but proceed with caution:
 - ML should still be combined with **rigorous testing**, **sanity checking**, and **considering misuse cases**.
 - “Microsoft deletes ‘teen girl’ AI after it became a Hitler-loving sex robot within 24 hours”:
 - <https://www.telegraph.co.uk/technology/2016/03/24/microsofts-teen-girl-ai-turns-into-a-hitler-loving-sex-robot-wit>
 - “Amazon AI Designed to Choose Phone Cases Terribly Malfunctions, Fills Store with 31,000+ Hilarious Products”:
 - <https://www.boredpanda.com/funny-amazon-ai-designed-phone-cases-fail>
 - “Uber video shows the kind of crash self-driving cars are made to avoid”:
 - <https://www.wired.com/story/uber-self-driving-crash-video-arizona/>
 - “One pixel attack for fooling deep neural networks”:
 - <https://arxiv.org/abs/1710.08864>
 - “Failures of Gradient-Based Deep Learning”:
 - <https://arxiv.org/abs/1703.07950>
 - “Meaningless Comparisons Lead to False Optimism in Medical Machine Learning”:
 - <http://www.arxiv.org/abs/1707.06289>
 - It’s important to get a sense of **what can and can’t** be done (now and in near-future).
 - Many industry people have unrealistic expectations.

What's Next?

- “Calling Bullshit in the Age of Big Data”:
 - <https://www.youtube.com/playlist?list=PLPnZfvKID1Sje5jWxt-4CSZD7bUI4gSPS>
 - There is a lot of **bullshit in the machine learning world** right now.
 - E.g., **cherry-picking of examples** in papers and **overfitting to test sets**.
 - You should try to start recognizing obvious non-sense, and not accidentally produce non-sense yourself!
- I’m putting material from all my courses (“**All of Machine Learning**”) here:
 - <https://www.cs.ubc.ca/~schmidtm/Courses/AllOfML>
 - (I’ll try to keep this up to date and exhaustive.)
- Our Machine Learning Reading Group (topic undecided for the summer):
 - <http://www.cs.ubc.ca/labs/lci/mlrg>
- Thank you for your patience (this course is not easy to organize), and good luck with the next steps!