The Problem

- Discriminator dominates generator
  - Discriminator too good?
  - Generator too bad?

- Potential solutions:
  - Hyperparameter variation- adjust learning rate of D and G
  - Not always so stable- stabilize discriminator training
  - Reduce layers in D
  - Increase layers in G

- Standard angular GAN converged with:
  - Lr G = 10^{-3}
  - Lr D = 10^{-5}

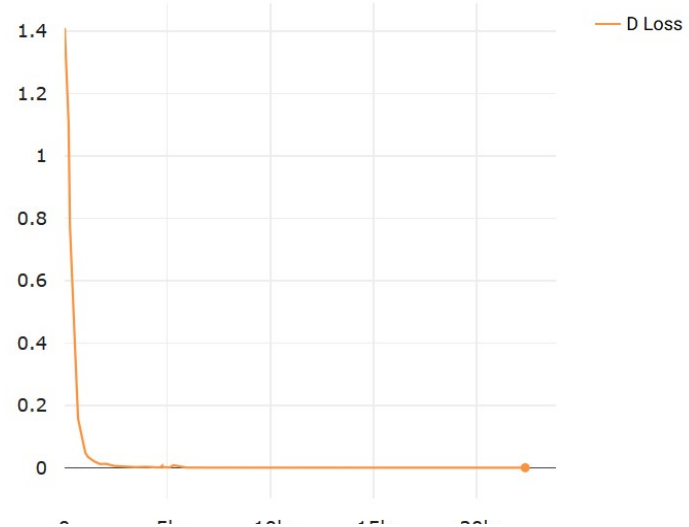

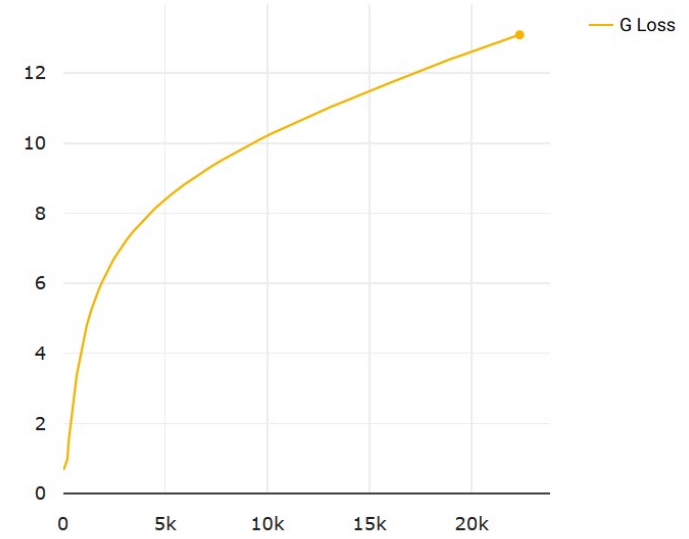

Attempting to get seedGAN to converge

- SeedGAN with same architecture as before seeding (tan theta etc.)
  - Lr G = 10^{-3} , Lr D = 10^{-5}
  - Lr G = 10^{-3} , Lr D = 10^{-6}
  - Lr G = Lr D = 10^{-5}
- SeedGAN with 1 fc layer for D:
  - Lr G = 10^{-3} , Lr D = 10^{-6} – D still won…
  - Lr G = Lr D = 10^{-5} – G beat D…
- SeedGAN with spectral norm:
  - Lr G = Lr D = 10^{-5}
  - Lr G = 10^{-5} , Lr D = 10^{-6}
  - Lr G = 10^{-6} , Lr D = 10^{-5}
  - Lr G = 10^{-5} , Lr D = 10^{-7}
  - Lr G = 10^{-4} , Lr D = 10^{-7}
  - Lr G = 10^{-3} , Lr D = 10^{-5}

Next steps

- What is going wrong?
  - Different latent space each time?
  - Noise fed to generator wrong? - currently add seed to U(-1,1)
  - Generator not deep enough?- add more layers to G
  - More hyperparameters that need scanning?

- Any other ideas?