



# Qualitatively Characterizing Neural Network Optimization Problems




Ian  
Goodfellow




Oriol Vinyals



Andrew Saxe




# Traditional view of NN training




# Factored linear view

(Cartoon of  
Saxe et al 2013's  
worldview)



# Attractive saddle point view

(Cartoon of  
Dauphin et al 2014's  
worldview)




# Questions

- Does SGD get stuck in local minima?
- Does SGD get stuck on saddle points?
- Does SGD wind around numerous bumpy obstacles?
- Does SGD thread a twisting canyon?




# History written by the winners

- Visualize trajectories of (near) SOTA results
- Selection bias: looking at success
- Failure is interesting, but hard to attribute to optimization
- Careful with interpretation
  - SGD never encounters  $X$ ?
  - SGD fails if it encounters  $X$ ?

# 2-D subspace visualization

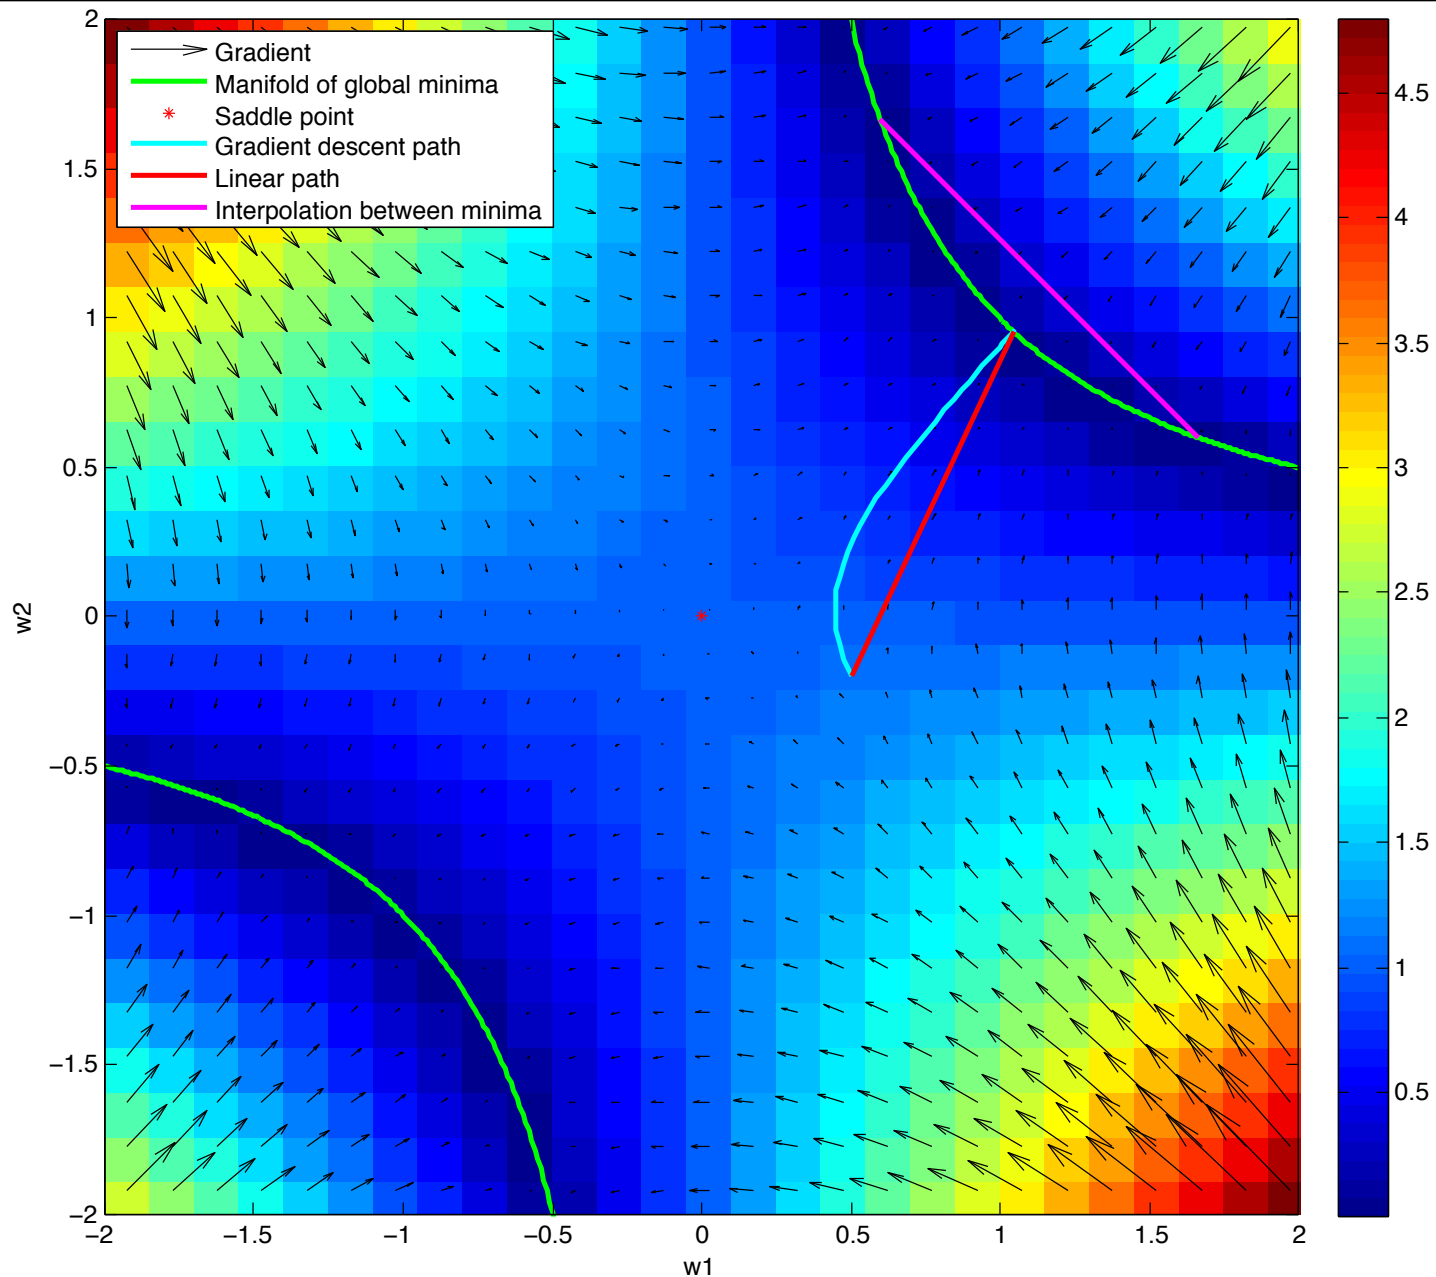


# A special 1-D subspace







# 2-parameter deep linear model

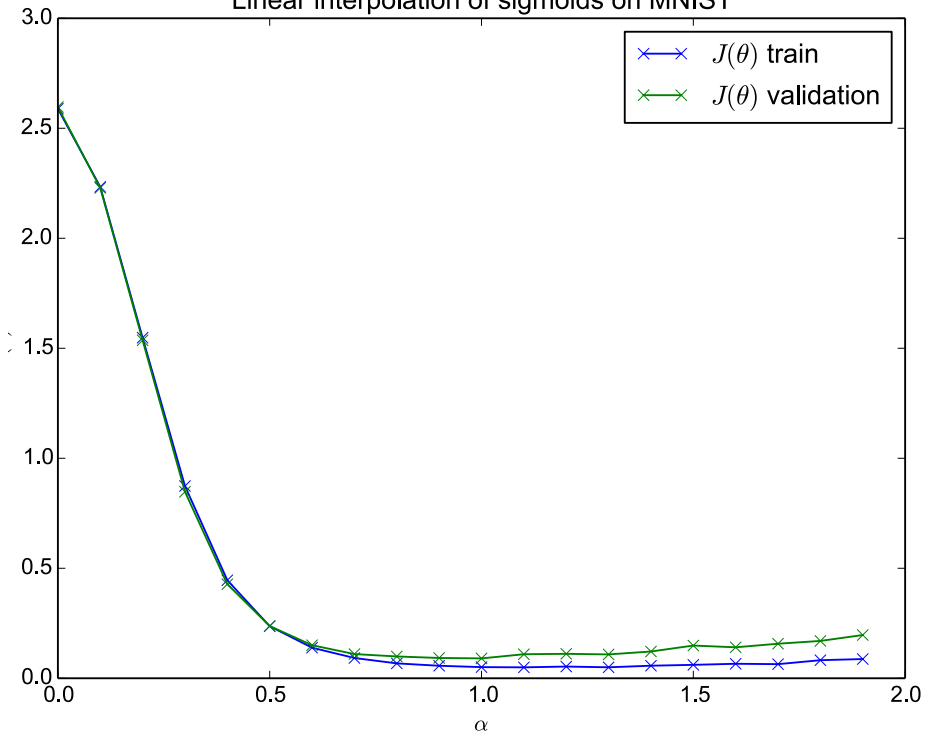


# Maxout / MNIST experiment

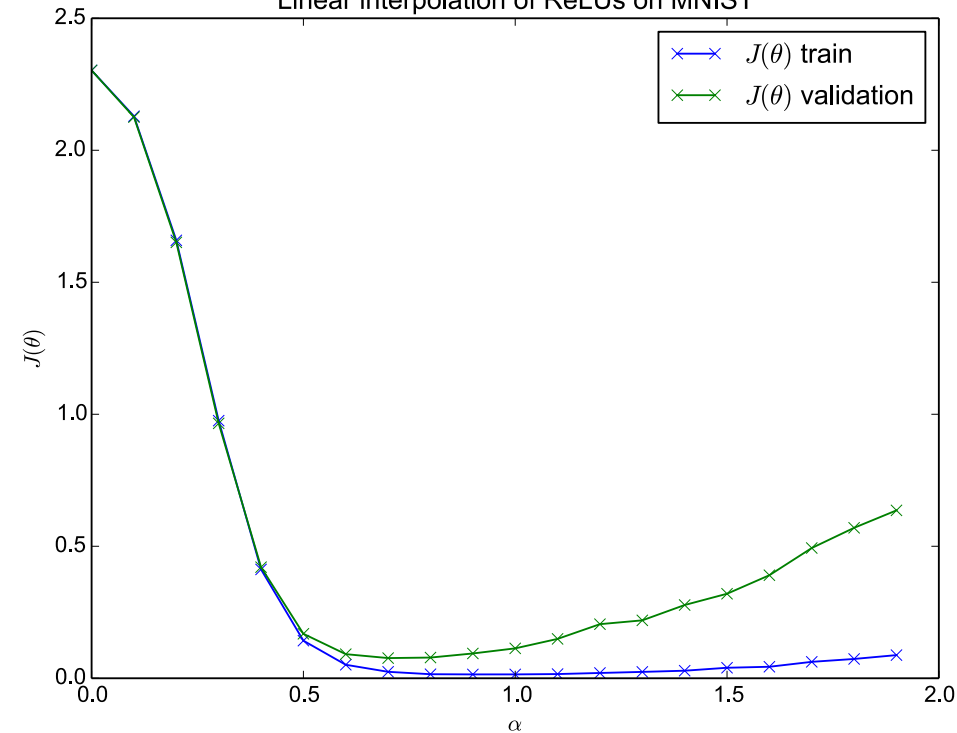


# Other activation functions



Linear interpolation of sigmoids on MNIST



Linear interpolation of ReLUs on MNIST




# Convolutional network




A small barrier


# LSTM




# MP-DBM



# 3-D Visualization

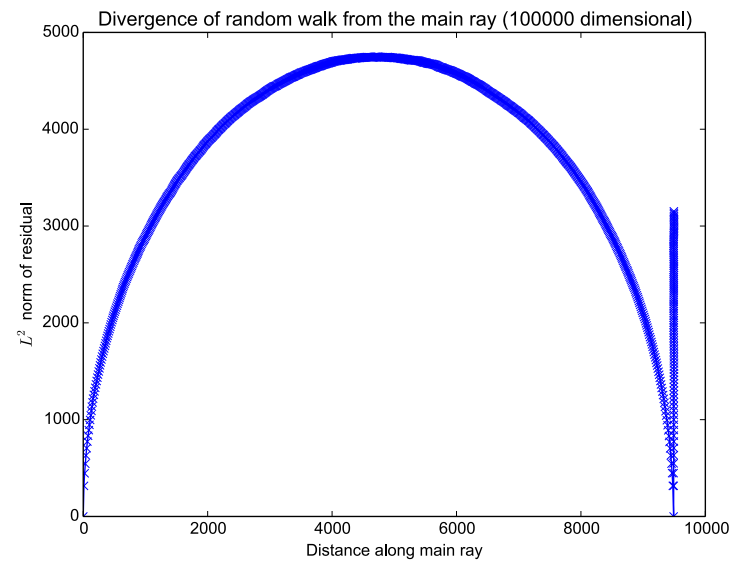

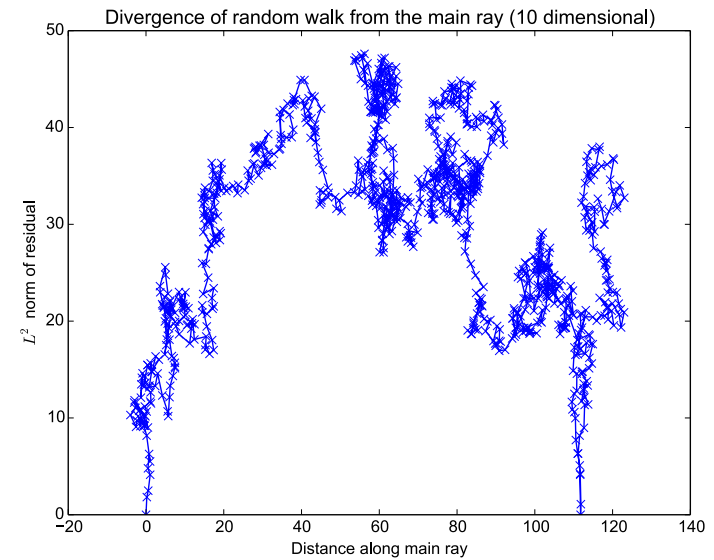
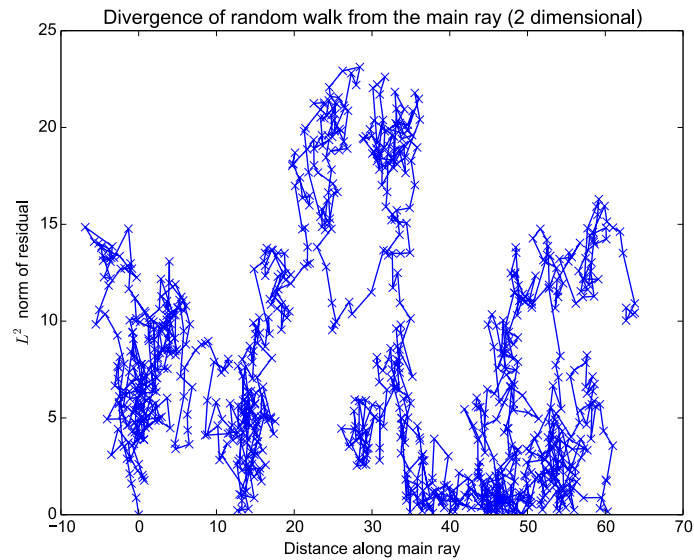


# 3-D MP-DBM visualization







# Random walk control




# 3-D Plots Without Obstacles



LSTM




Adversarial  
ReLUs




Factored Linear

# 3-D Plot of Adversarial Maxout



SGD naturally exploits  
negative curvature!



Obstacles!

# Conclusion

- For most problems, there exists a linear subspace of monotonically decreasing values
- For some problems, there are obstacles between this subspace the SGD path
- Factored linear models capture many qualitative aspects of deep network training
- See more visualizations at our poster / demo / paper