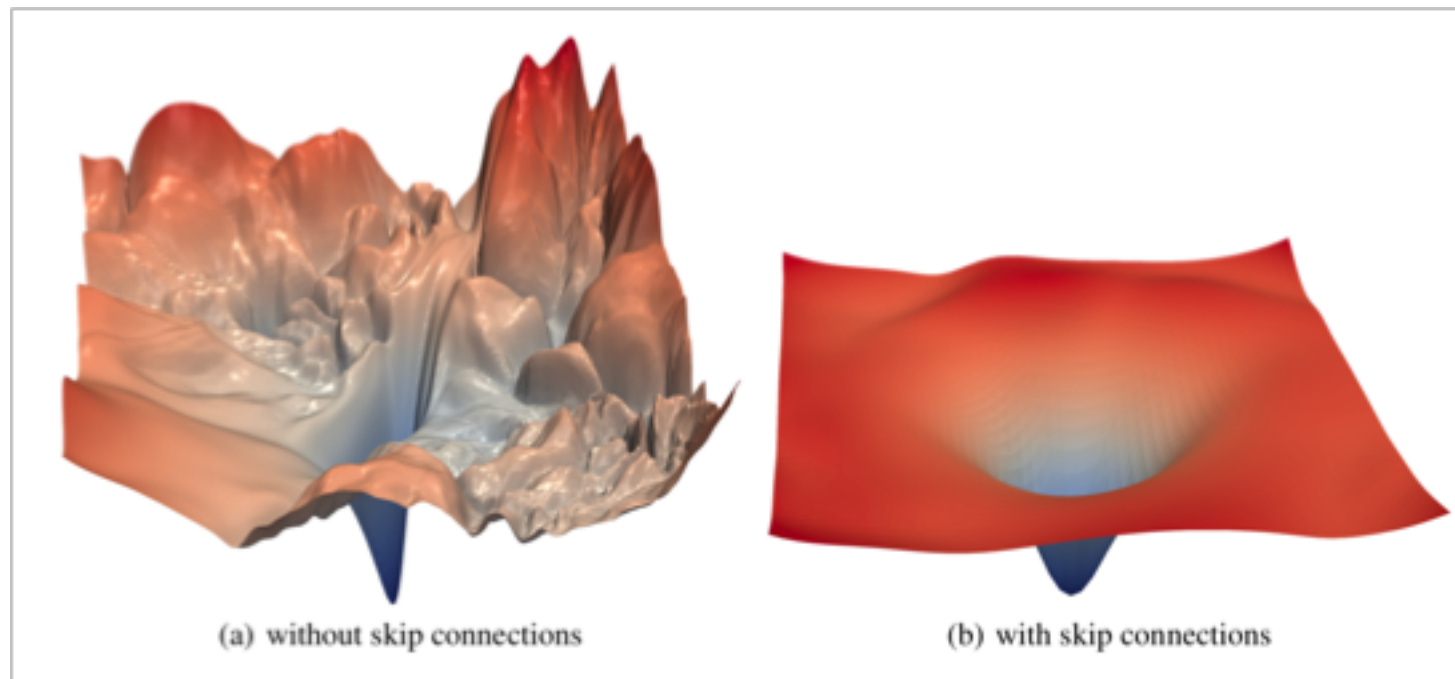


Visualizing the Loss Landscape of Neural Nets

- Analyzes the influences of architectural choices on the loss landscape and thus explains why certain „tricks“ and structures are easier to optimize



Visualizing the Loss Landscape of Neural Nets

- Summarize paper and explain background
- Code available
- Bonus: Visualize landscape for speech models

Noise2Noise: Learning Image Restoration without Clean Data

- Trains a denoiser with just noisy training data; no clean data needed!



Noise2Noise: Learning Image Restoration without Clean Data

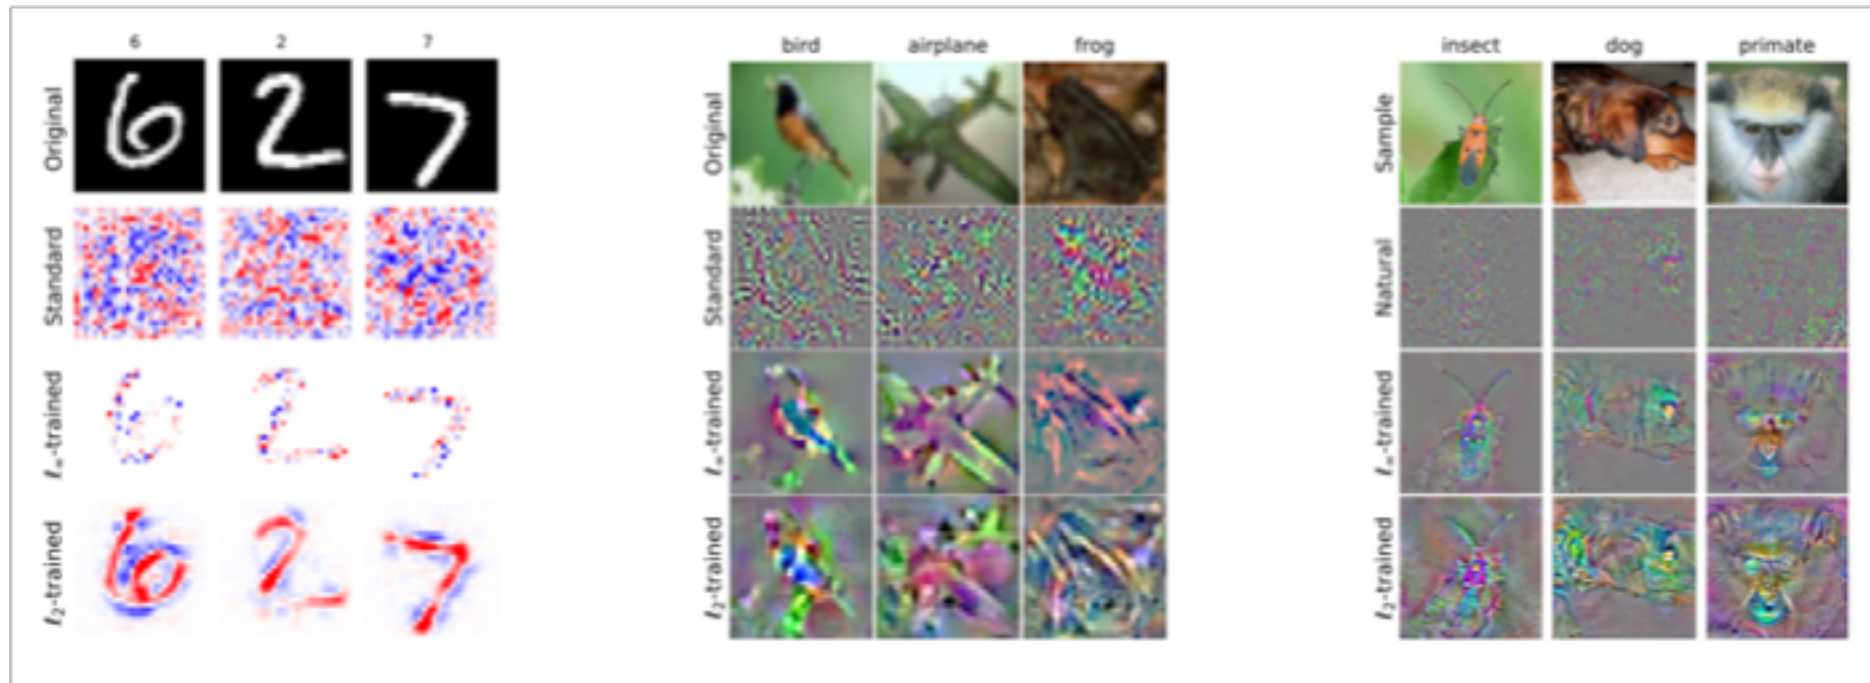
- Summarize paper and explain background
- Code available
- Bonus: Apply method to noisy speech signals

How Does Batch Normalization Help Optimization?

- BN popular technique for training NNs but poorly understood
- Shows that effectiveness of BN is not related to internal covariance shift
- Rather, BN smoothes the loss landscape and thus leads to better gradients, a more stable training and faster convergence

Robustness May Be at Odds with Accuracy

- Shows that there is a trade-off between adversarial robustness and (training) accuracy for **any** classifier
- Feature representations learned by robust classifier are fundamentally different



Spherical Latent Spaces for Stable Variational Autoencoders

- Uses a von Mises-Fisher distribution instead of a Gaussian distribution as the prior in a Variational Autoencoder
- Avoids “collapse” of latent space

