HyperGAN:

Generating Diverse, Performant Neural Networks

Neale Ratzlaff, Fuxin Li

Oregon State University

36th ICML 2019

Oregon State University
College of Engineering

Uncertainty

- High predictive accuracy is not sufficient for many tasks

- We want to know when our models are uncertain about the data

Fixing Overconfidence

- Given many models, each model behaves differently on outlier data
- By averaging their predictions, we can detect anomalies

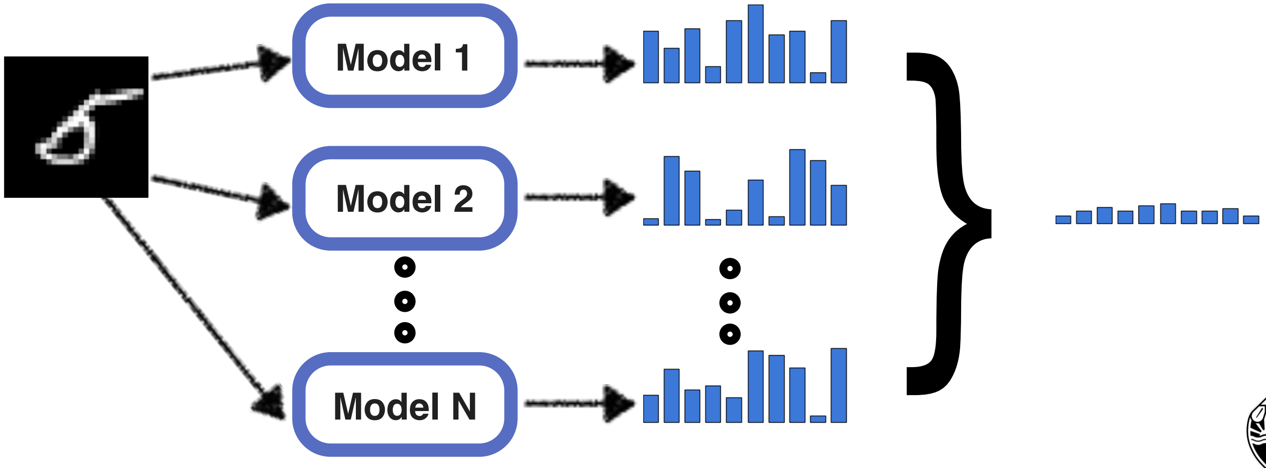

Fixing Overconfidence

- Given many models, each model behaves differently on outlier data

- By averaging their predictions, we can detect anomalies

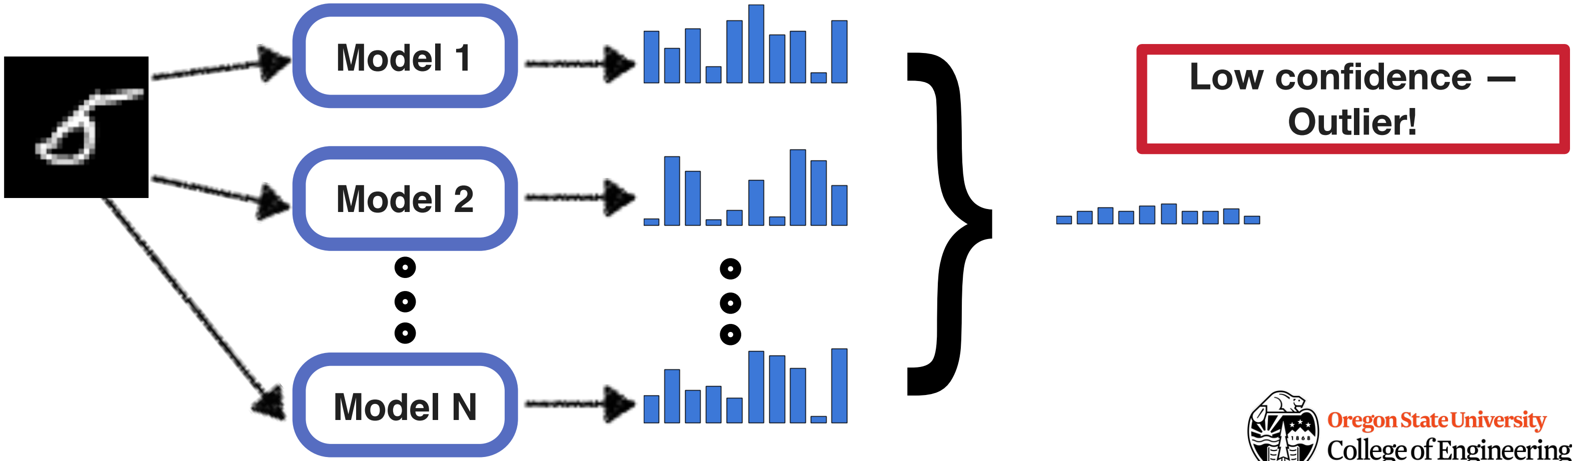

Fixing Overconfidence

- Variational inference gives a model posterior where we can sample many models
- Ensembles of models from random starts may also detect outliers

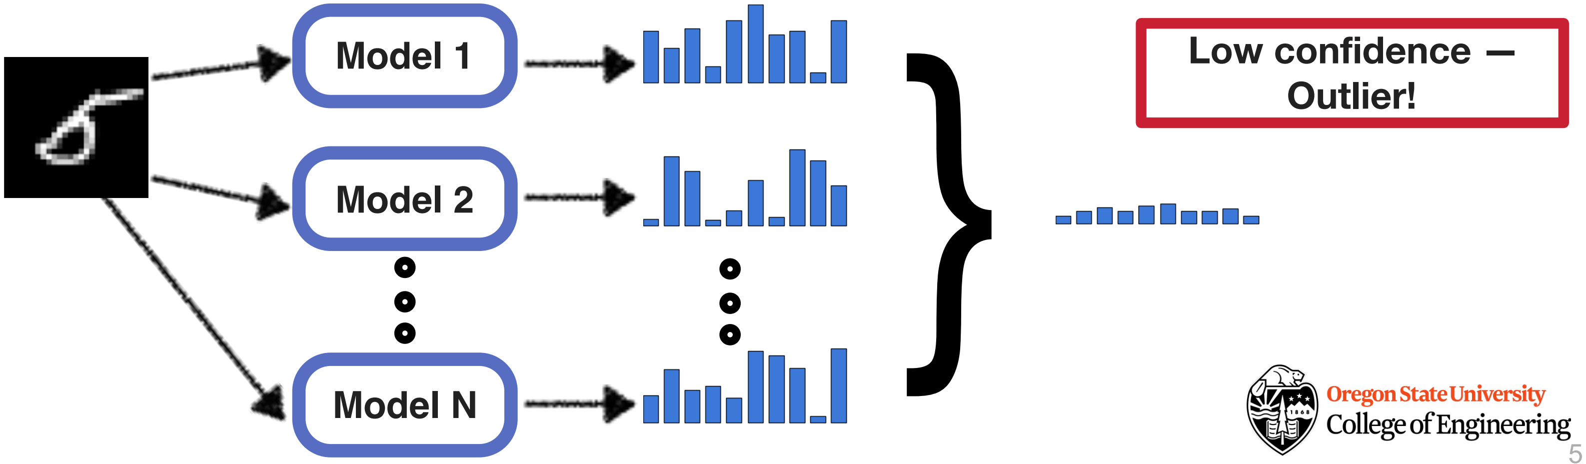

Oregon State University College of Engineering

Regularization is too Restrictive

- Learning with VI is restrictive, it cannot model the complex model posterior
- Without regularization, our outputs mode collapse, losing diversity

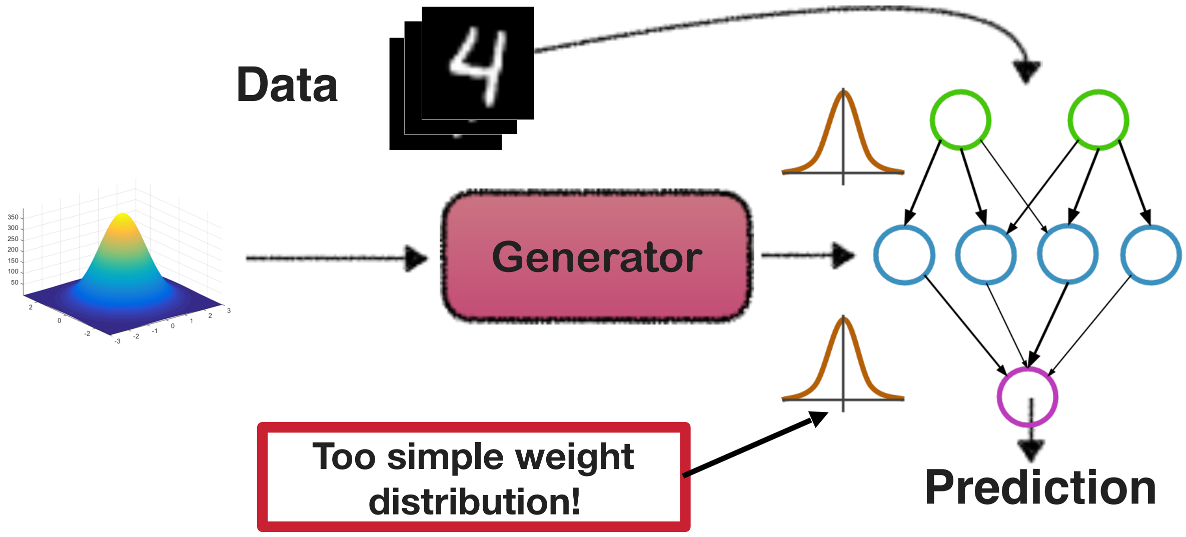

Oregon State University College of Engineering

Implicit Model Distribution

- We learn an implicit distribution over network parameters with a GAN
- We can instantly generate any number of diverse, fully trained networks

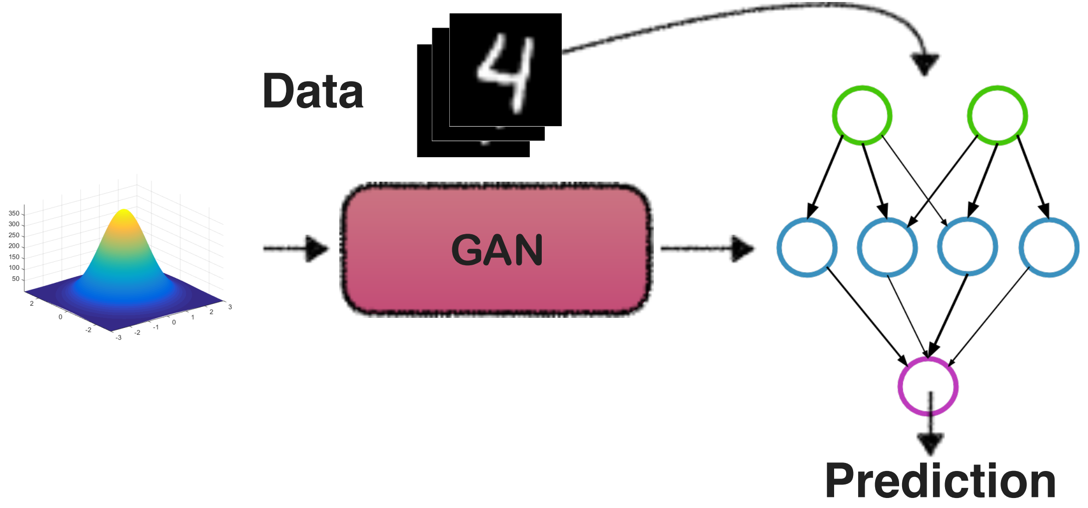

Implicit Model Distribution

- With a GAN, we can sample many networks instantly

- However, with just a Gaussian input, the generated networks tend to be similar

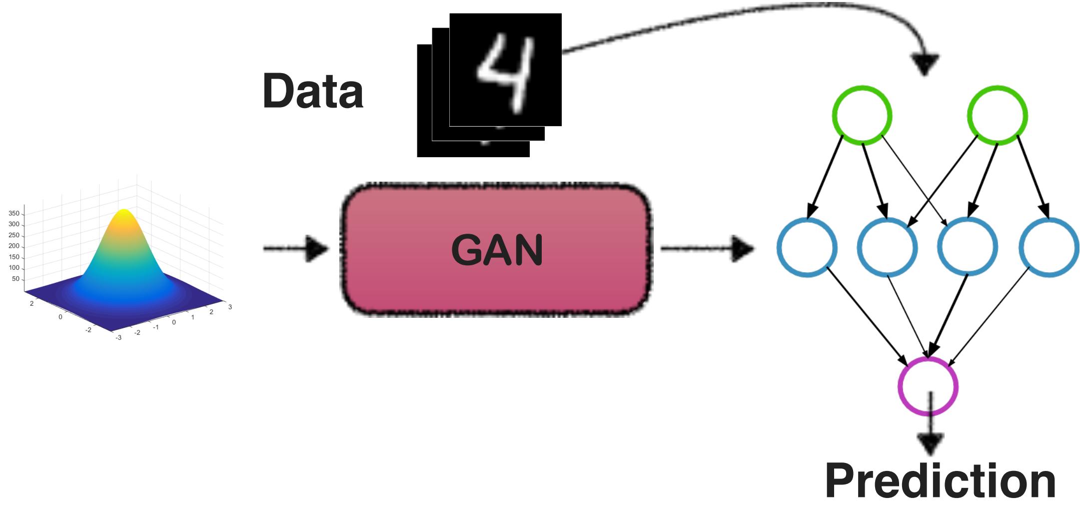

Mixer Network for Diverse Ensembles

- Want to generate *diverse* ensembles, without repeatedly training models
- Our novel Mixer, transforms the input noise to learn complex structure.
- Mixer outputs are used to generate diverse layer parameters

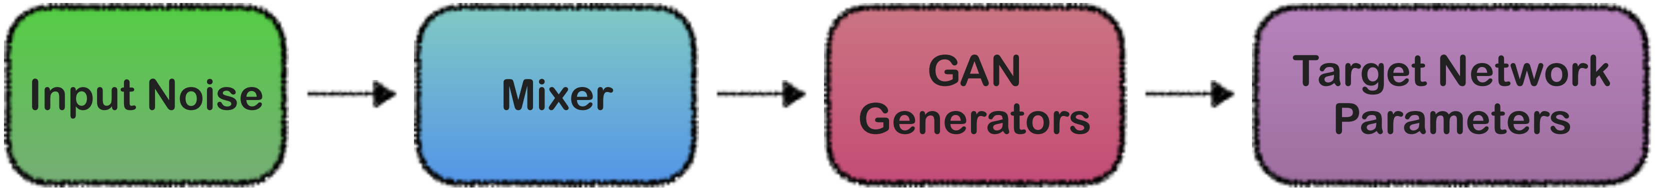

Generating Diverse Neural Networks

- Every training step we sample a new batch of networks
- The diversity given by the mixer lets us find many different models which solve the target task

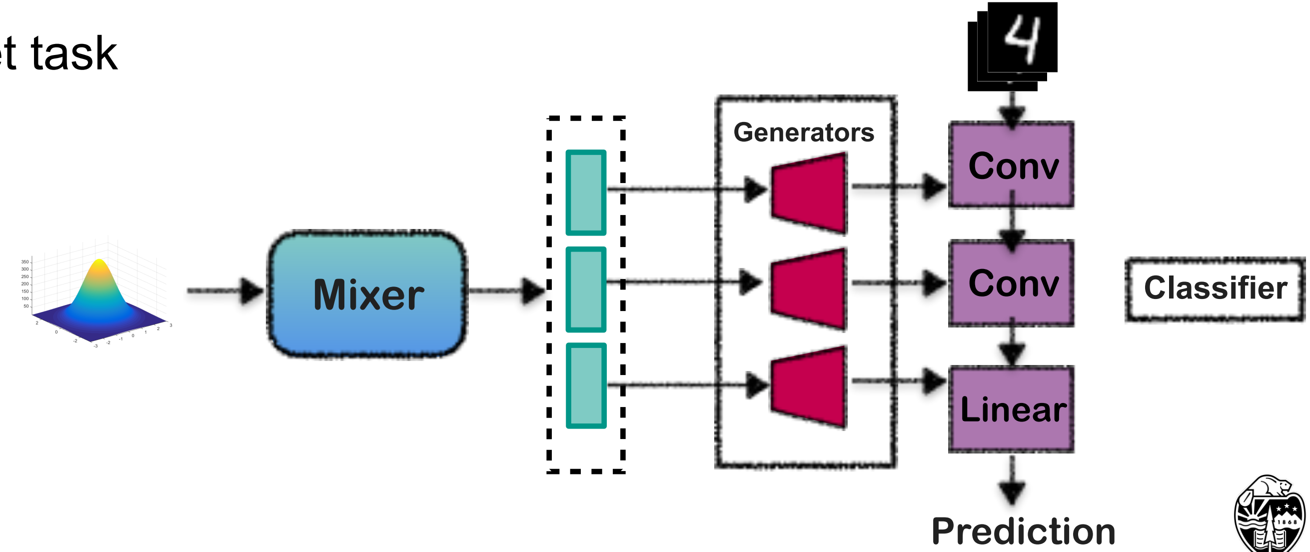

HyperGAN Training: Full Architecture

- Prevent mode collapse by regularizing the Mixer with a Discriminator
- We use the target loss to train HyperGAN

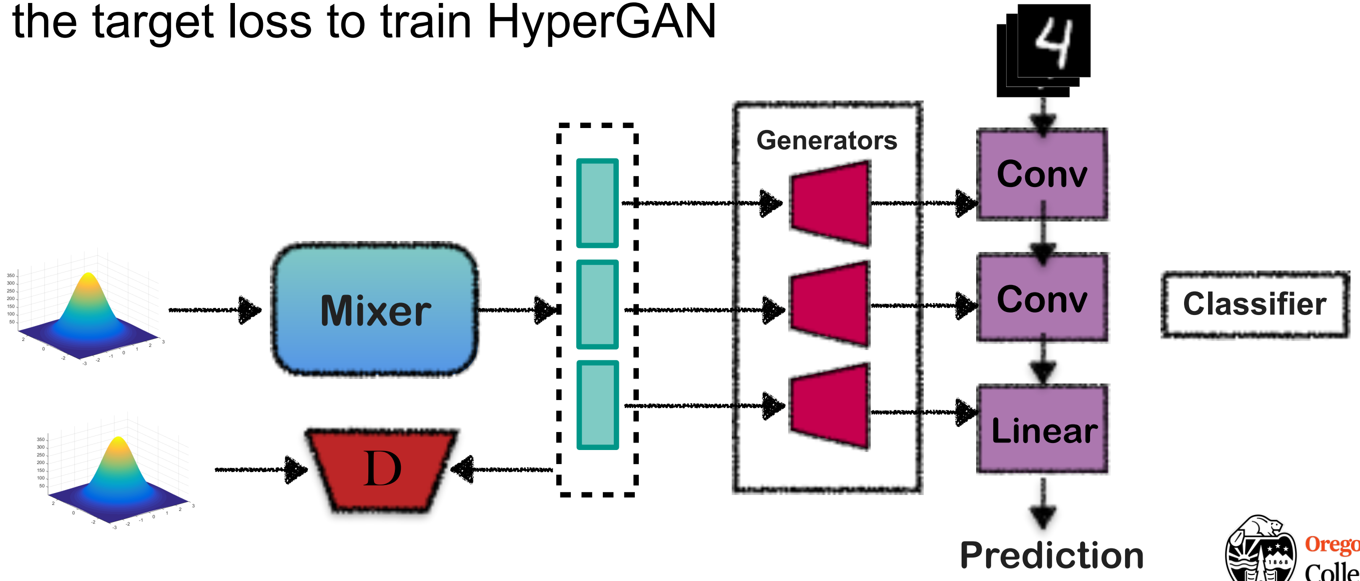

Weight Diversity

- HyperGAN learns diverse weight posteriors beyond simple Gaussians imposed by variational inference

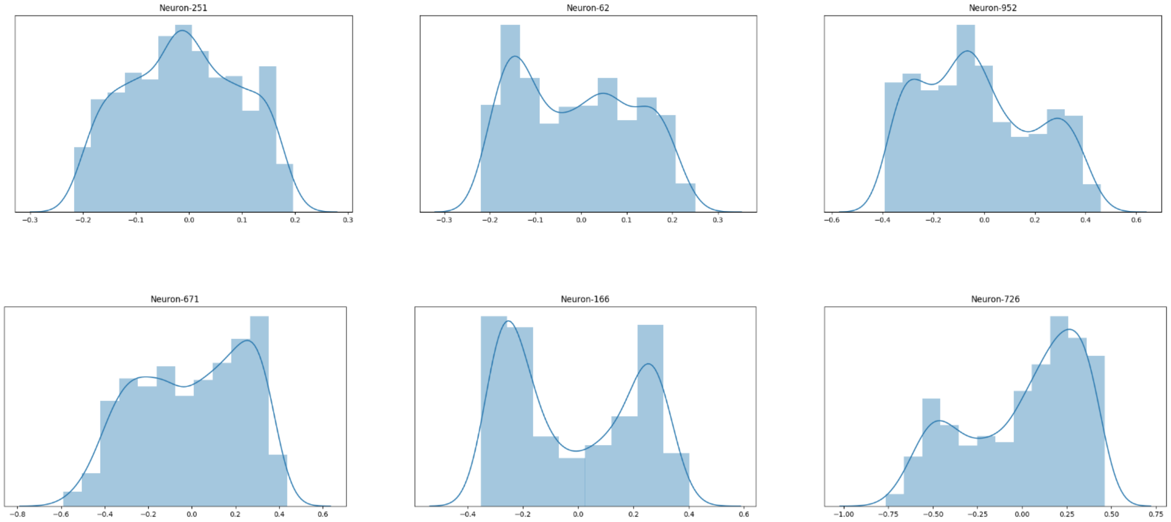

Results - Classification

- MNIST 5000: train on 5k example subset.
- CIFAR-5: Restricted subset of CIFAR-10

Method	MNIST	MNIST 5000	CIFAR-5	CIFAR-10
1 network	98.64 ±.3	96.69 ±.3	84.50 ±.6	76.32 ±.3
5 networks	98.75 ±.3	97.24 ±.14	85.51 ±.2	76.84 ±.1
10 networks	99.22 ±.09	97.33 ±.1	85.54 ±.2	77.52 ±.09
100 networks	**99.31 ±.02**	**97.71 ±.05**	**85.81 ±.02**	**77.71 ±.03**
APD	98.61	96.35	83.21	75.62
MNF	99.30	97.52	84.00	76.71
MC Dropout	98.73	95.58	84.00	72.75
Random Start	99.14	97.09	83.84	74.79

Out of Distribution Experiments

- Outlier detection on CIFAR-10 and MNIST datasets
  - MNIST ⟶ notMNIST

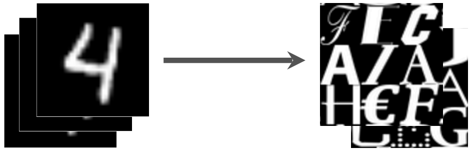

  - CIFAR (0-4) ⟶ CIFAR (5-9)

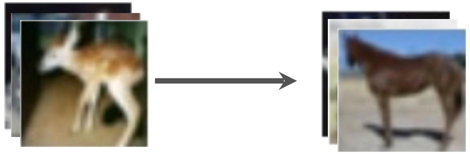

  - Adversarial Examples: FGSM and PGD

Our increased diversity allows us to outperform other methods

Conclusion

- HyperGAN generates diverse models
- Makes few assumptions about output weight distribution
- Method is straightforward and extensible

Come to our poster for more details!