

- anomaly detection

- low $p(x)$ - rare?

- learn $p(x|y)$ for $y \in CR$

interpolate y into SR

set background model, fully data driven
for x

useful beyond
anomaly detection
of course!

- probably more applications out there!

4 classes of NN-based GMs:

- GANs

- VAEs

- Flows (also OE)

- Diffusion (als. AG)

Plan is to introduce class to each one, with examples

GANs: Generative Adversarial Networks

Goodfellow et al (2014)

Idea: adversarial training min-max objective

was SOTA
on images
until recently
(now diffusion)
still some advantages
- computationally
lighter & faster

saddle pt
not global min!

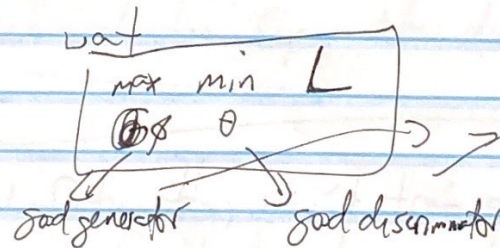
train 2 NNs
to go to defeat one another

noise data
G generator $z \rightarrow x$
D discriminator
classifies generated "fake"
vs data "real"

GAN loss:

BCE for D but also loss for G!

$$L(\theta) = - \left(\sum_{i \in \text{data}} \log D_{\theta}(x_i) + \sum_{i \in \text{gen}} \log(1 - D_{\theta}(G_{\phi}(z_i))) \right)$$



balance two objectives

Why does this work??

- formally: $\min_{\theta} L \rightarrow D_{\theta}(x) = P_{\text{data}}(x)$ (Neyman-Pearson Lemma!)

feedback into L $P_{\text{data}}(x) + P_{G_{\theta}}(x)$

$$L|_{D=D_{\theta}} = - \left(\sum_{x \in \text{data}} \log \frac{P_{\text{data}}(x)}{P_{\text{data}}(x) + P_{G_{\theta}}(x)} + \sum_{x \in \text{gen}} \log \frac{P_{G_{\theta}}(x)}{P_{\text{data}}(x) + P_{G_{\theta}}(x)} \right)$$

$$\rightarrow - \int dx \left(p_{data} \log \frac{p_{data}}{p_{data} + p_G} + p_G \log \frac{p_G}{p_{data} + p_G} \right)$$

$$= -2 \text{JSD}[p_{data}, p_G] + \text{const}$$

↓
 Jensen-Shannon Divergence — distance
 $0 \leq \text{JSD}^{(p,q)} \leq 1$ btw prob
 zero iff $p=q$ dist'n's!

$$\max_{\beta} L \Big|_{\beta=0}^{\beta=1} = \min_{\beta} \left(2 \text{JSD}[p_{data}, p_G] + \text{const} \right)$$

⇒ $p_G = p_{data}$

(can also
 prove more
 straightforwardly
 from integral
 above)

So formally, GAN objective is
 achieved when G matches data!

↓
 key pt:
 p_{data} only
 learned
 implicitly!


Unfortunately, can't train a GAN so easily,
 so formal derivation not 100% applicable...

→ go to GAN MNIST Demo

→ Training: Freeze G ←
 Train D (batch real + batch gen) repeat
 Freeze D
 Train G (batch gen)

CALOGAN example (1705.02355, 1712.10321)

- Early example of ^{deep} gen model for fast sim surrogate modeling (first in HEP?)
- Used vanilla GAN for GEANT4 e^+, γ, π^+ calor showers by ATLAS ~~GCAL~~ GCAL



3 layers of LAr + lead absorber
simplification: turn sample fraction off
(100% efficient)

"Voxelization"

$$3 \times 96 - 12 \times 12 - 12 \times 6 \\ = 504 \text{ total voxels}$$

3 GANs: e^+, γ, π^+
Data: $\left\{ \begin{array}{l} 1-100 \text{ GeV unif } E_{inc} \\ 100k \text{ showers each generated by GEANT4} \end{array} \right.$

energy in each voxel
 $\vec{E} \in \mathbb{R}^{504}$

↓
to generate from
goal: learn $p(\vec{E} | E_{inc})$

→ to overview of paper
main results: quality 50-50

speed - very fast (but CPU/GPU
comparison subtleties)

metrics?