- Reminder: please fill out course survey! deadline THIS THURS

Last time
Autoencoders for unsupervised outlier detection

$\hookrightarrow L_{\text{reconstruction}}$ as anomaly score

Example: $\underbrace{\text{QCD}}_{\text{bg}}$ & $\underbrace{\text{top or gluino}}_{\text{sig}}$ jets with vanilla AE.

Using VAE for anomaly detection

$\hookrightarrow$ nice latent space (interpretable)

use it like vanilla AE w/ reconstruction error OR use clustering in latent space

---

Overdensity Detection

Goal: learn $\boxed{R(x) = \frac{p_{\text{data}}(x)}{p_{\text{bg}}(x)}}$

Claim: "ideal" or optimal anomaly score.

Pf: $p_{\text{data}}(x) = (1-\varepsilon)\, p_{\text{bg}}(x) + \varepsilon\, p_{\text{sig}}(x)$

$$\Rightarrow \underline{\underline{R(x)}} = (1-\varepsilon) + \varepsilon \frac{p_{\text{sig}}(x)}{p_{\text{bg}}(x)}$$

don't know ε or p_{sig} in advance!

R(x) linearly related to sig-bg LR!

$R(x)$ monotonic w/ sig-bg LR

$R(x) > R_c$ equivalent (sig-bg LR) > some threshold.

optimal signal-bg discrimant by NP lem

$R(x)$ is also optimal signal-bg discrimat

for any unknown signal!!

What are issues w/ learning $R(x) = \frac{p_{data}(x)}{p_{bg}(x)}$?

- need really accurate $p_{bg}(x)$ or sample from $p_{bg}(x)$

"positive-unlabeled learning"

"bg" "data"

can get $R(x)$ via binary classifier "LR trick"

- learning a classifier btw two very similar samples is difficult

- $R(x) > R_c$ → some subsample of data

still need very accurate bg estimate!

In HEP, a lot of effort → "enhancing the bump hunt"

"resonant anomaly detection"

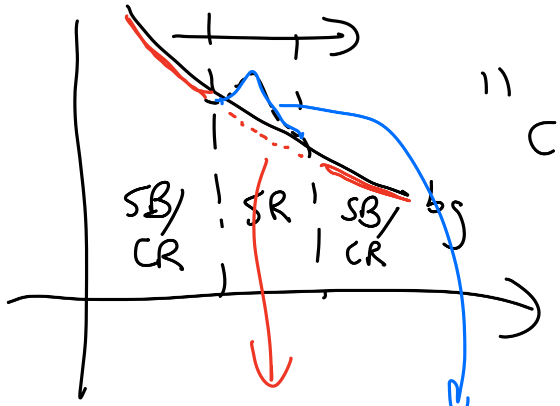

"classic bump hunt" → $Z \to \ell\ell$, $H \to \gamma\gamma$

$N_{bg} \leftrightarrow N_{data}$

$N_{sig} = N_{data} - N_{bg}$

$$\boxed{\sigma = \frac{N_{sig}}{\sqrt{N_{bg}}}}$$

"partially model agnostic"

Ex of NP search:

$pp \to Z' \to jj$ $m_{inv} = m_{jj}$ "Z' dijet resonance search"

Idea of enhancing bump hunt: maybe ∃ add'l features x where signal is localized

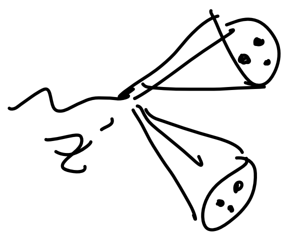

(e.g. $x = m_{j_1}, m_{j_2}, \tau_{21}^{j_1}, \tau_{21}^{j_2}, \ldots$)

→ Could learn $R(x)$ cut $R(x) > R_c$ → tpr → fpr

as long as cut doesn't sculpt bg

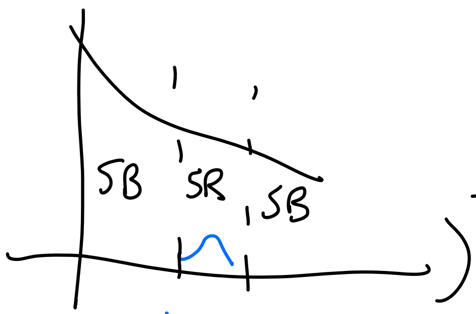

1σ

$\xrightarrow{R > R_c}$

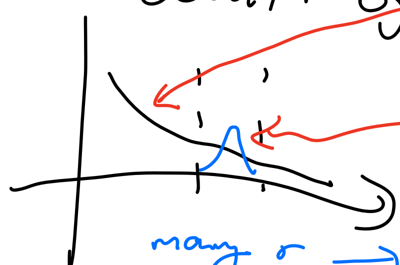

many σ → discovery!

$SIC = \text{tpr}/\sqrt{\text{fpr}}$ "significance improvement" characteristic

$\sigma = \frac{S}{\sqrt{B}} \rightarrow \frac{S \cdot \text{tpr}}{\sqrt{B}\sqrt{\text{fpr}}} = \sigma \cdot SIC$

How to learn $R(x)$ from data? Many ideas!

interpolating from SB in (x, m)

1. "CWoLa Hunting" Collins, Howe & Nachman
1805.02664, 1902.02634

→ if $p_{bg}(x)$ in SB same as $p_{bg}(x)$ in SR (so x & m are uncorrelated in bg)

then classifier btw SR & SB data gives $R(x)$.

$$\rightarrow R(x) = \frac{p_{data,SR}(x)}{p_{data,SB}(x)} = \frac{p_{data,SR}(x)}{p_{bg,SB}(x)} = \frac{p_{data,SR}(x)}{\leftrightarrow p_{bg,SR}(x)}$$

→ now actual ATLAS search

2. ANOmaly Detection through Density Estimation (ANODE)
2001.04990 Nachman & DS

- Learn conditional density estimator (normalizing flow) in SB regions $\rightarrow p_{bg}(x|m \in SB)$
- Interpolate into SR $\rightarrow p_{bg}(x|m \in SR)$

- learn second density estimator in SR

$\Rightarrow$ $\longrightarrow p_{data}(x|m)$

- construct $R(x)$ directly (take ratio)

- pros: robust to x, m correlations

- cons: DE is much harder than classification!
  so ANODE less sensitive than CWoLa!

---

3. Classifying Anomalies THrough Outer Density Estimation (CATHODE) DS + Hallin et al 2109.00546

Combine CWoLa + ANODE

- learn $\underset{= p_{bg}}{p_{data}(x|m)}$ in SB & interpolates into SR
- sample from $p_{bg}(x|m)$ in SR

  $\longrightarrow$ sample of synthetic bg events in SR!

- classifier btw data & synthetic bg events in SR

  $\longrightarrow R(x)$

  $\longrightarrow$ robust to correlations & more sensitive than either ANODE or CWoLa!