Event details:
Thursday 23 Jul 2020, 10:30 → 11:30
https://indico.cern.ch/event/939335/
ZOOM: https://cern.zoom.us/j/96267620043
Paper: https://arxiv.org/pdf/1505.05770.pdf

Related:

Variational Inference with Normalizing Flows -->rezendee

Masked Autoregressive Flow —> Papamakario

Inverse Autoregressive Flow —> Kingma 2016
Autoregressive flows [Kingma et al., 2016] [Uria et al., 2016] [Papamakario et al., 2017]
- masked autoregressive flows [Germain et al., 2015] used in ANODE (Anomaly Detection for Density Estimation) [Nachman and Shih, 2020]

• f-VAEs: Improve VAEs with Conditional Flows [Su and Wu, 2018]

•RealNVP: Real-valued non-volume preserving transformations [Dinh et al., 2017]
- successor of NICE

•Glow: Generative Flow with Invertible 1x1 Convolutions [Kingma and Dhariwal, 2018]

• Pixel Recurrent Neural Networks s [van den Oord et al, 2016]

Questions on paper:

How is this relevant to mPP?

I guess this is something someone is already working on, but it's of course natural to try this for MC generation where you really need an explicit representation of the density. But is this too computationally expensive? → Use for Marys generator!

Put a normalizing flow to transform @Kingas latent space :)?
→ No, not necessarily better for anomaly detection

So following the paper, the key is the change of variables formula to transfer one density to another, and for that you need a bijective transformation to take the inverse, meaning if the input dimensions are huge, they are really computationally expensive to approximate well (and require a lot of data to train parameters of latent space?). What’s the ‘limit’ to problems we can solve with this?

I am interested in trying some of this out: