WatChMaL : ν-VAE

Abhishek Abhishek, Olivia Di Matteo

TRIUMF Hyper-K Group
akajal@triumf.ca

TRIUMF-Helmholtz Meeting

November 13, 2019

WatChMaL

Water Cherenkov Machine Learning

https://www.watchmal.org

Working group established in April 2019

Goal - Apply machine learning techniques to Water Cherenkov (WC) detectors in order to :

1. Mitigate systematic and experimental uncertainties
2. Improve signal vs background discrimination
3. Utilize improvements in detector design

In this talk

1. IWCD and Hyper-K : Overview and Challenges
2. Variational Autoencoders (VAEs)
3. Semi-supervised learning
4. ν-VAE motivation
5. Results : Interpolation and Semi-supervised learning
6. Current and future work
7. Beyond VAEs
8. Constructing a discrete latent space
9. DVAE++ and QVAEs

IWCD and Hyper-K : Overview

Intermediate Water Cherenkov Detector (IWCD) : Next generation near detector with $\mathcal{O}(500)$ mPMTs and variable height for constraining the neutrino flux and cross-section uncertainties. **(Current dataset)**.

Hyper-Kamiokande (Hyper-K) : Next generation far detector with $\mathcal{O}(50k)$ PMTs and 10x fiducial volume of Super-Kamiokande (Super-K).

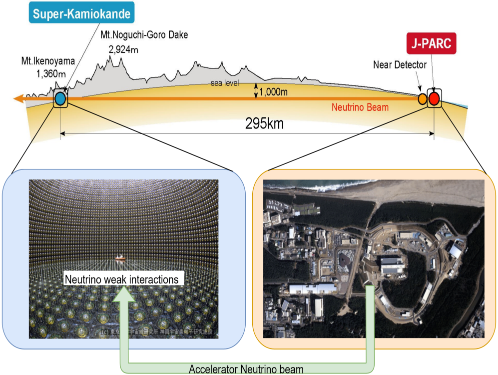

IWCD and Hyper-K : Challenges

Key challenges that we aim to address using **deep learning** :

1. High number of **background** γ **events** generating a similar signal to e^- or e^+ events due to pair production $\gamma \to e^+e^-$.
2. **Imperfect modelling** of the **water Cherenkov detectors** by simulations used for data analysis and hypothesis testing.
3. Limitation of labelled **calibration** datasets to only certain regions of the detector phase space. E.g. :
   - Michel (decay) electrons
   - Crossing and stopping muons (not produced inside the tank volume)

Variational Autoencoders (VAEs)

Class of generative models which learns a **constrained low-dimensional representation** of the input :

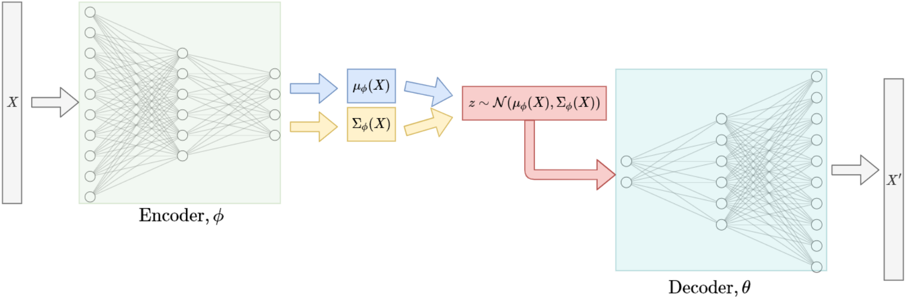

The training loss comprises of :

1. **reconstruction loss** - X vs. X'
2. **divergence loss** - $\mathcal{N}(\mu_\phi(X), \Sigma_\phi(X))$ vs. $\mathcal{N}(0, I)$

Semi-supervised learning

Many different semi-supervised models inspired by the **VAE** framework.

Simplest :

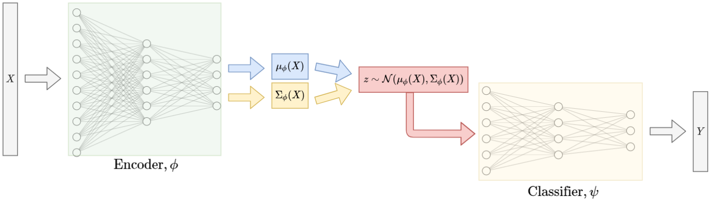

Complex :

(a) Generative model p_θ **(b)** Inference model q_ϕ

ν-VAE motivation

ν-VAE project aims to address keys challenges of IWCD and Hyper-K.

We aim to apply :

1. **Semi-supervised learning** for training classifiers on labelled and unlabelled datasets for improved signal vs. background discrimination.
2. **VAE** framework and extensions (e.g. **DVAE/QVAE**) to approximate the complex data-generating distribution.
3. **Constrained latent representation** learned by the VAE to improve generalization to unseen regions of the detector phase space.

Results : Interpolation

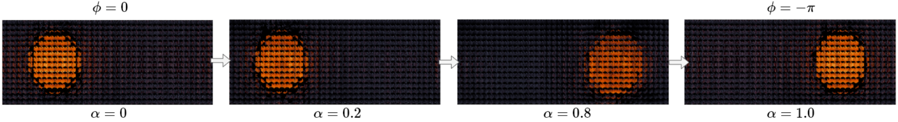

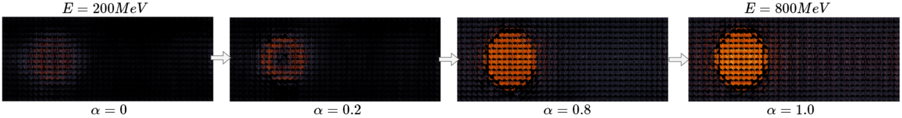

Linear interpolation $\hat{z} = \alpha z_1 + (1 - \alpha) z_2$ shows some phase space features (e.g. particle energy E) correspond linearly to directions in the latent space while other phase space features (e.g azimuthal angle ϕ) do not.

Results : Semi-supervised learning

Number of training examples	γ background rejection (%) at 50% e^- signal efficiency		γ background rejection (%) at 80% e^- signal efficiency	
	SS-CNN	**CNN**	**SS-CNN**	**CNN**
$11,250$	**77.6**	76.4	**50.7**	46.3
$22,500$	**80.4**	78.1	**54.3**	48.5
$45,000$	**80.7**	79.4	**55.9**	49.9

Initial semi-supervised models trained using both labelled and unlabelled datasets show significant improvement in signal vs. background discrimination.

Current work

Previous experiments and development conducted using events only in the barrel of the IWCD i.e. the cylinder walls without the endcaps.

However, in order to generate realistic calibration samples, we need to extend the previous work to the entire IWCD geometry.

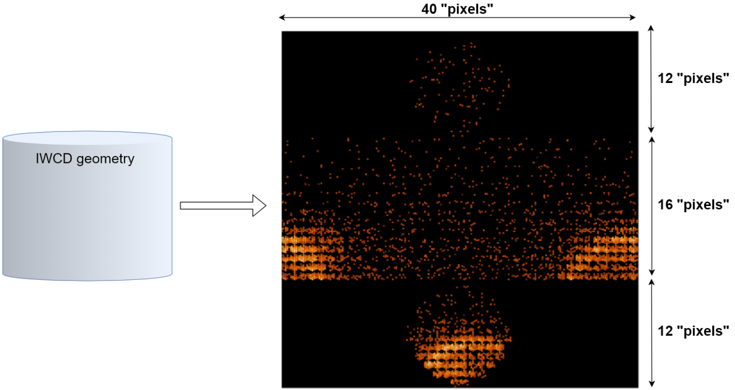

Current work - Initial results

Initial results using the same approach as the barrel-only dataset i.e. treating the detector response as a 2-d image shows promising results but with significant room for improvement.

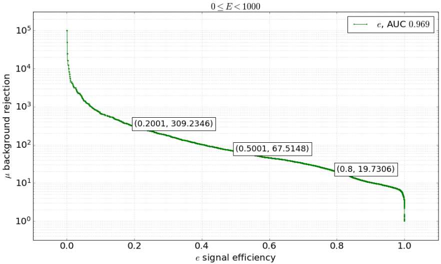

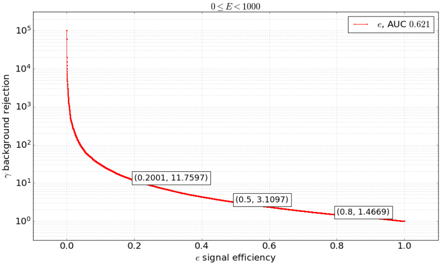

Future work

Empirically evaluate the **generalizability** of the fully-supervised and semi-supervised models to unseen regions of the detector phase space.

We will consider this in the context of the **IWCD calibration** datasets :

1. Michel (Decay) electrons
2. Crossing and Stopping cosmic muons

Next steps also include :

1. Development of the semi-supervised learning approaches
2. Augmentation of neural network architectures for full tank geometry

Beyond VAEs

Consider:

- What if we don't think the prior distribution is as simple as a product of Gaussians?
- What if it would make more sense to have latent variables that are *discrete*?
- What about a mix of discrete and continuous latent variables?

Discrete VAEs (DVAEs) provide us with a framework for making more expressive generative models. They also provide a bridge between purely classical VAEs and quantum VAEs (QVAEs), where we can in principle use a quantum annealer to generate our samples.

Constructing a discrete latent space

In DVAEs, the latent variables take only the values 0 and 1.

Problem: a network with discrete nodes would not be a continuous, differentiable function. We couldn't do backpropagation!

Solution: Perform a relaxation to continuous variables using a **smoothing function**.

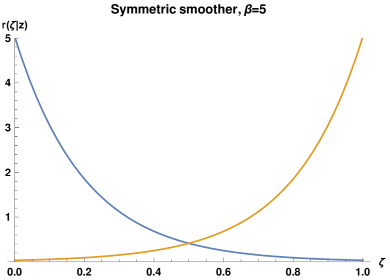

Constructing a discrete latent space

Within the network discrete variables will be smoothed, but we can still build a prior distribution over discrete variables.

Restricted Boltzmann machines (RBMs) are often chosen as the prior. They are more expressive than a product of Gaussians, and are trained along with the encoder and decoder.

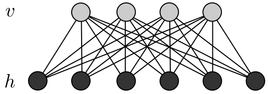

$$E(v,h) = -\sum_i a_i v_i - \sum_i b_i h_i - \sum_{ij} W_{ij} h_i v_j$$

Every state of the RBM has an associated energy; the probability of being in a state (v, h) of particular energy is given by the Boltzmann distribution:

$$\Pr(v,h) = \frac{e^{-E(v,h)}}{Z_\theta}, \quad Z_\theta = \sum_s e^{-E(v,h)} \tag{1}$$

DVAE++

We want to implement a type of DVAE with a hierarchically structured latent space with a discrete portion, and a continuous portion.

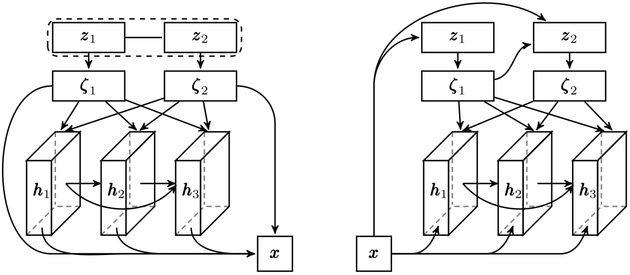

(a) generative model (b) inference model

Motivation: water Cherenkov detector outputs have both an event type (discrete variables), and physical properties that give it its shape (continuous variables).

Image: Vahdat, A., Macready, W. G., Bian, Z., Khoshaman, A., & Andriyash, E. (2018). DVAE++: Discrete variational autoencoders with overlapping transformations. 35th International Conference on Machine Learning, ICML 2018, 11, 8008–8023.

QVAEs

Training an RBM can be computationally expensive - relies on a lot of **sampling**.

Idea: source the sampling to a *quantum annealer* to speed up training.

D-Wave has developed the theoretical tools and the physical implementation.

We are about to submit a Mitacs proposal and enter into collaboration with them to explore the potential for QVAEs in particle physics.

Backup

T2K and Hyper-K : Physics goals

T2K open questions :

1. Is Charge conjugate Parity **(CP)** symmetry violated in the lepton sector ?
2. Is the neutrino mass hierarchy "normal" or "inverted" ?
3. Are there sterile neutrino states ?

Hyper-K open questions :

4. Does the proton remains stable forever ?
5. Observation of cosmic neutrinos from supernovae ?

Variational Autoencoders (VAEs)

Latent deep generative model which maximizes a lower bound to the true log-likelihood of the data under the model :

$$logp(x) \geq \text{ELBO} = \underbrace{-KL[q_\phi(z|x)||p(z)]}_{\text{- KL Loss}} + \underbrace{E_{q_\phi(z|x)}[logp_\theta(x|z)]}_{\text{- Reconstruction Loss}}$$