Interpretable time series foundation models for anomaly prediction in sPHENIX

Tengfei Ma (tengfei.ma@stonybrook.edu)

Stony Brook University

6/18/2025 @AD4HEP Workshop

Thanks to my collaborators from BNL, RPI.

Monitoring the Detector

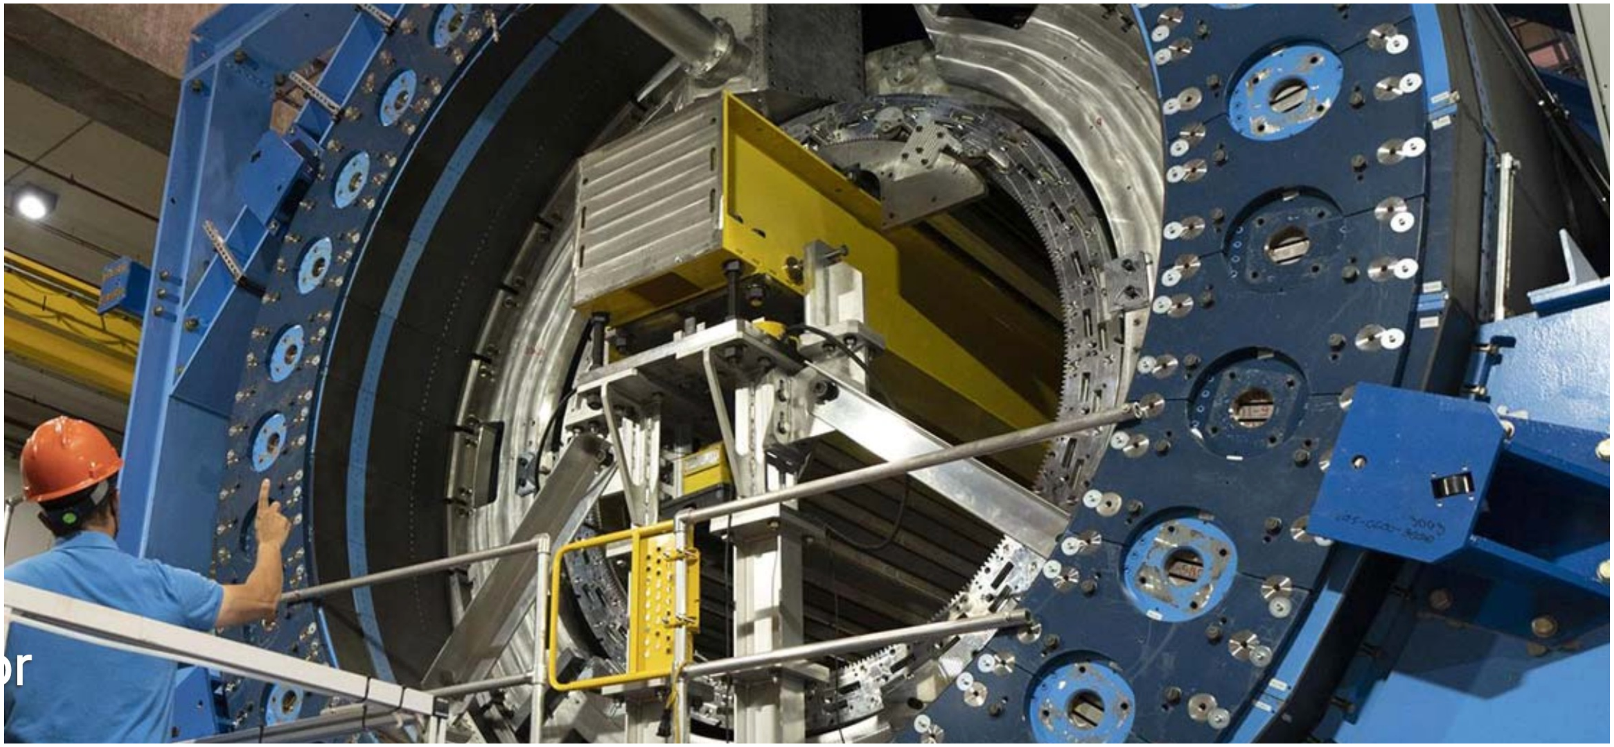

Time Series Data and Anomalies

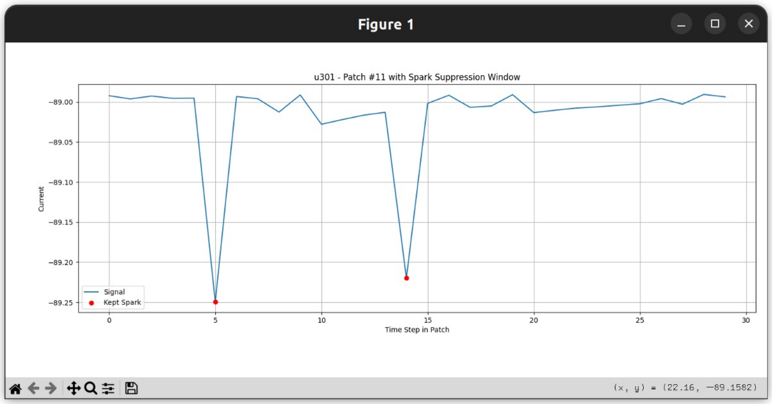

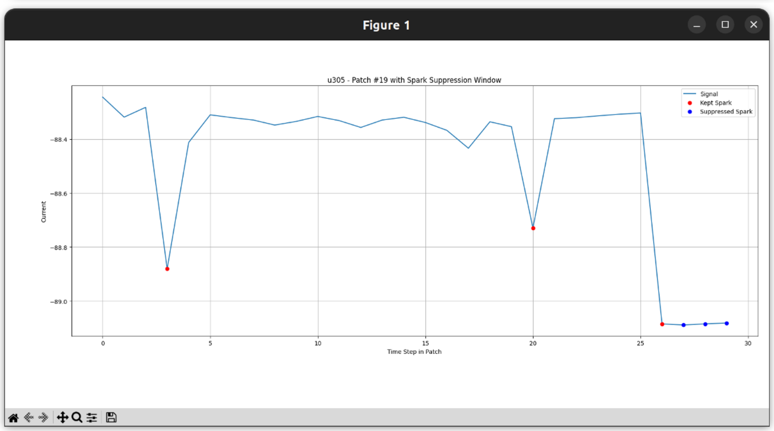

Sensory Data (multivariate time series):

1. TPC central membrane current and voltage.
2. Current and voltage measurements for each Gas Electron Multiplier (GEM) modules
3. 15 other beam conditions
4. Collected every 4 seconds

Anomalies (sparks):

1. Definition: significant drop of current/voltage compared to previous history

Anomaly Prediction

- Assume there is an anomaly event at time window *t*

- Anomaly Detection
  - Given the data S_t at t, is S_t anomaly?
- Anomaly Prediction
  - Given the data S_{t-1}, is S_t anomaly?
  - i.e. S_{t-1} is a precursor of anomaly (PoA)

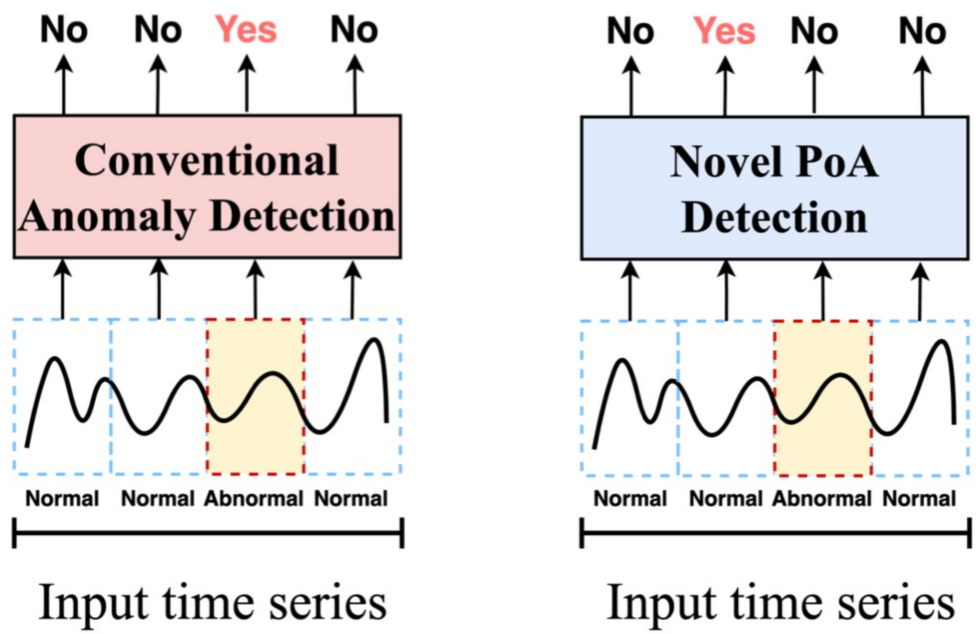

(a) Anomaly detection

(b) PoA detection

Fig from (Jhin et al. 2023).

Foundation Models

- (Generalizable) Foundation Models are reshaping the AI Research
  - Language (LLMs) and Vision
    - GPT-4, Claude3.5, Gemini 2.0, LLAMA3, …
    - DINO, SAM, CLIP, LLAVA…
  - Other domains/modalities
    - Code, Medical domain, Protein, DNA…
- Challenges on Time Series
  - Too diverse across domains
  - What are the tokens?
  - Backbone Model
    - How to encode structure information?
    - Is it scalable, expressive, interpretable?
  - What are the pretraining objectives?

Schemes of Time Series Foundation Models

- Backbones
  - Transformer (PatchTST, Autoformer), Convolution (TimesNet), MLP (Dlinear, TSmixer)
- Pretrain time series foundation model
  - Across different tasks
  - Across different domains
- Inference on a new domain/task
  - Linear probing
  - Fine tuning
- However, interpretability?
  - Continuous representation of variables
  - No meaningful tokens

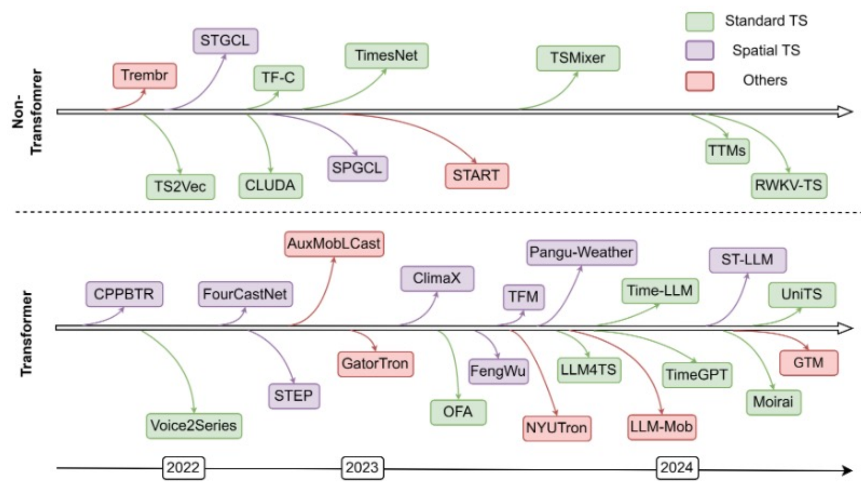

Figure 1: Roadmaps of representative TSFMs.

From Liang et al. 2024

Background: Discrete Neural Representation

- VQ-VAE

(Oord et al. 2017)

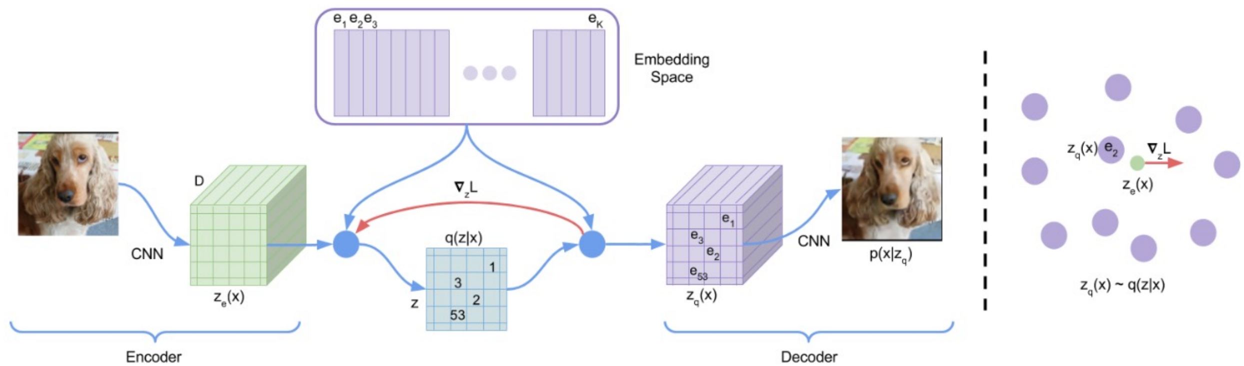

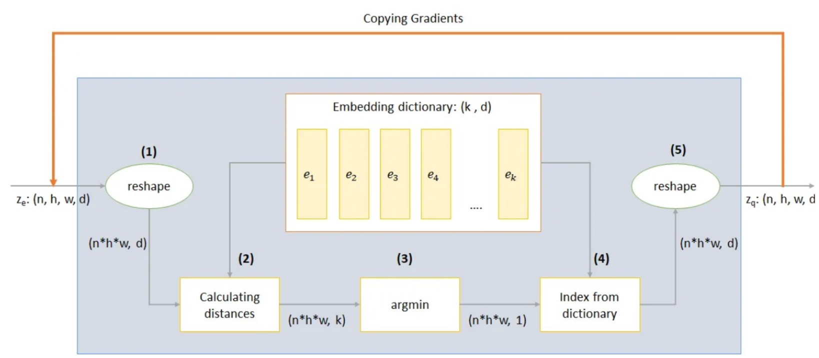

It enables better image/video generation/understanding, better global structure, more efficient generation

Shapes as Tokens

- Codebooks in VQ-VAE
  - Still not physically meaningful
  - However, possible to replace the codebook with meaningful shapes
- A relevant concept: shapelet
  - Shapelets are defined as “subsequences that are in some sense maximally representative of a class”.
  - Each time series *x* can be represented as the vector of distances to all shapelets $<d(x, s_1), d(x, s_1), \ldots>$

Example Shapelets

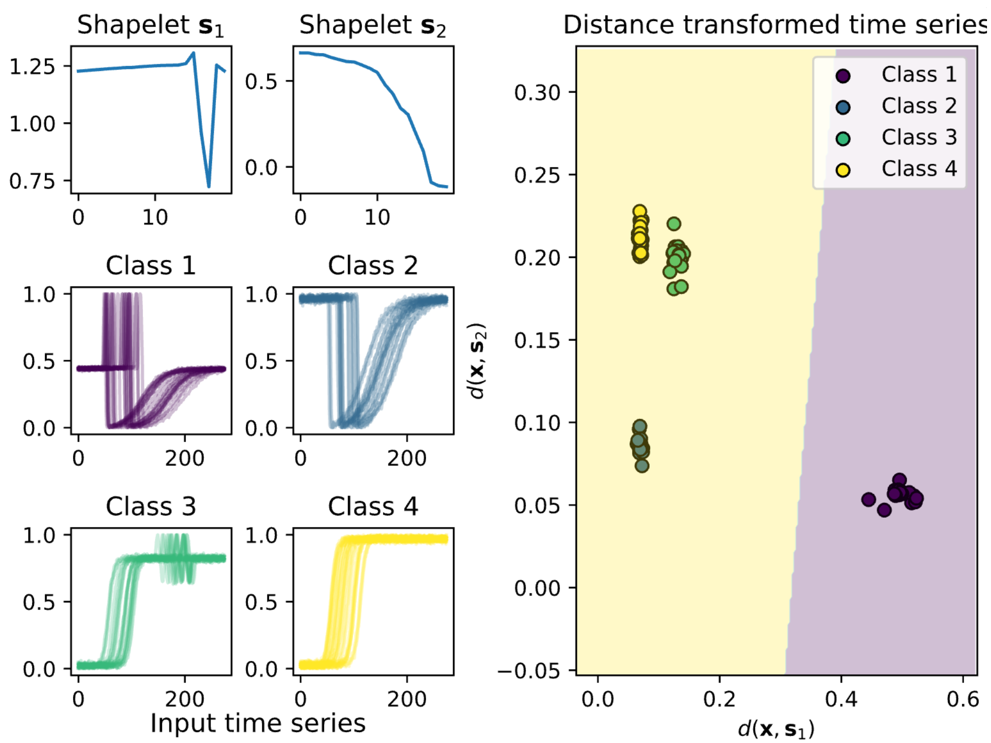

- Problems
  - Not generalizable
  - Scale/length sensitive
- Solution
  - We define a new concept "shape" which exclude scale, offset, position and information

Shape Token Representation

- An attribute tuple is a shape with meta information
  - Shapes are shared across data
  - other meta info helps define its reconstruction

For a univariate TS x, a subsequence s_k can be represented by an attribute tuple $\tau_k = (z_k, \mu_k, \sigma_k, t_k, l_k)$ where

- $z_k \in \mathbb{R}^{d_{\text{code}}}$ is the code for abstracted shape of s_k,
- $\mu_k \in \mathbb{R}^1$ is the offset of s_k,
- $\sigma_k \in \mathbb{R}^1$ is the scale (standard deviation) of s_k and $\sigma_k > 0$,
- $t_k \in \mathbb{R}^1$ is the relative starting position of s_k in x and $0 \leq t \leq T - l_{\text{min}}$,
- $l_k \in \mathbb{R}^1$ is the relative length of s_k w.r.t. to the length of x and $l_{\text{min}} \leq l \leq T - t_k$.

Stony Brook University

Brookhaven National Laboratory

VQShape [1] : an interpretable time series foundation model

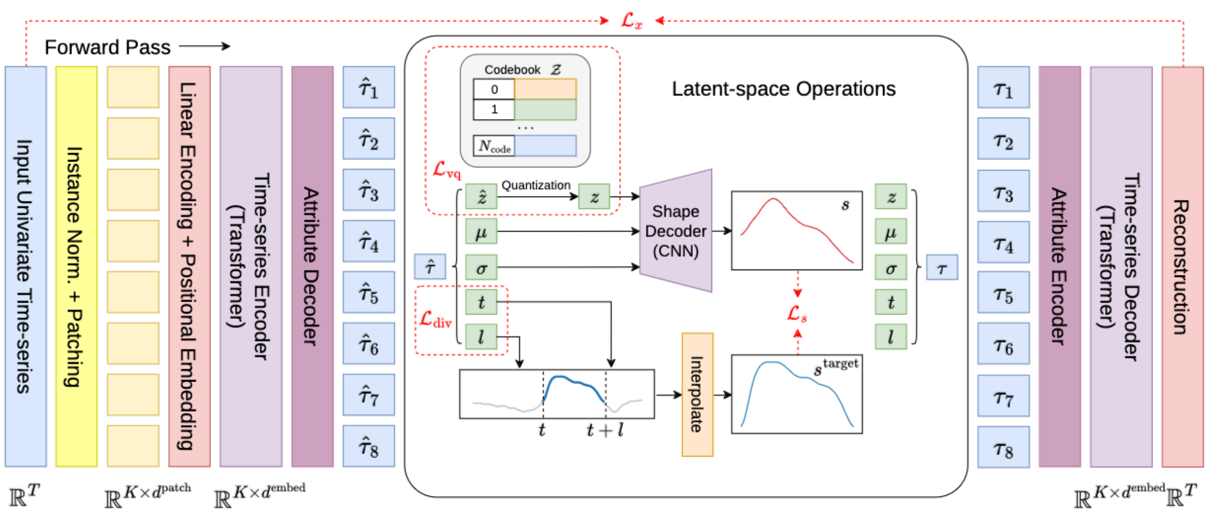

Figure 1: Overview of VQShape

Key Ideas:

- Shapes should be reused, but can be scaled in different time series
- Shape code can reconstruct a time series subsequence defined by the attribute tuple
- The combination of shapes can reconstruct original time series

[1] Wen et al. “Abstracted Shapes as Tokens -- A Generalizable and Interpretable Model for Time-series Classification”. NeurIPS 2024.

Stony Brook University

Brookhaven National Laboratory

Components

- 1. Time series encoding
  - Backbone: PatchTST
  - Output *K* embeddings
- 2. Attribute decoding
  - Take one embedding vector h_k

$$\hat{\tau}_k = (\hat{z}_k, \mu_k, \sigma_k, t_k, l_k) = \mathcal{A}_{\text{dec}}(h_k), \text{ where } \begin{cases} \hat{z}_k = f_z(h_k), \\ \mu_k = f_\mu(h_k), \\ \sigma_k = \texttt{softplus}(f_\sigma(h_k)), \\ t_k = \texttt{sigmoid}(f_t(h_k)) \cdot (1 - l_{\min}), \\ l_k = \texttt{sigmoid}(f_l(h_k)) \cdot (1 - t_k) + l_{\min}. \end{cases}$$

- 3. map shape to codebook (vector quantization)

$$z_k = \underset{z_q \in \mathcal{Z}}{\arg\min} \|\hat{z}_k - z_q\|.$$

Pretraining objective

$$\mathcal{L}_{\text{pretrain}} = \lambda_x \mathcal{L}_x + \lambda_s \mathcal{L}_s + \lambda_{\text{vq}} \mathcal{L}_{\text{vq}} + \lambda_{\text{div}} \mathcal{L}_{\text{div}}$$

- L_x is reconstruction loss of time series
- L_s is reconstruction loss of subsequences
- L_{vq} is vector quantization loss
  - To ensure codebook coverage, we add entropy regularization

$$\mathcal{L}_{\text{vq}} = \underbrace{\|\hat{z} - \mathbf{sg}(z)\|_2^2 + \lambda_{\text{commit}} \|\mathbf{sg}(\hat{z}) - z\|_2^2}_{\text{quantization}} + \underbrace{\mathbb{E}\left[H(q(\hat{z}, \mathcal{Z}))\right] - H(\mathbb{E}\left[q(\hat{z}, \mathcal{Z})\right])}_{\text{codebook usage}}$$

- L_{div} ensures attribute tuples have diverse positions and scales

$$\mathcal{L}_{\text{div}} = \frac{1}{K^2} \sum_{k_1=1}^{K} \sum_{k_2=1}^{K} \mathbb{1}(k_1 \neq k_2)\, \mathtt{relu}\left(\epsilon - \|\kappa(t_{k_1}, l_{k_1}) - \kappa(t_{k_2}, l_{k_2})\|_2^2\right),$$

$$\text{where} \quad \kappa(t_k, l_k) = \begin{bmatrix} \cos(t_k \pi) \cdot \ln(l_k) / \ln(l_{\min}) \\ \sin(t_k \pi) \cdot \ln(l_k) / \ln(l_{\min}) \end{bmatrix}.$$

Example Shapes

Example Shapes
and Time Series

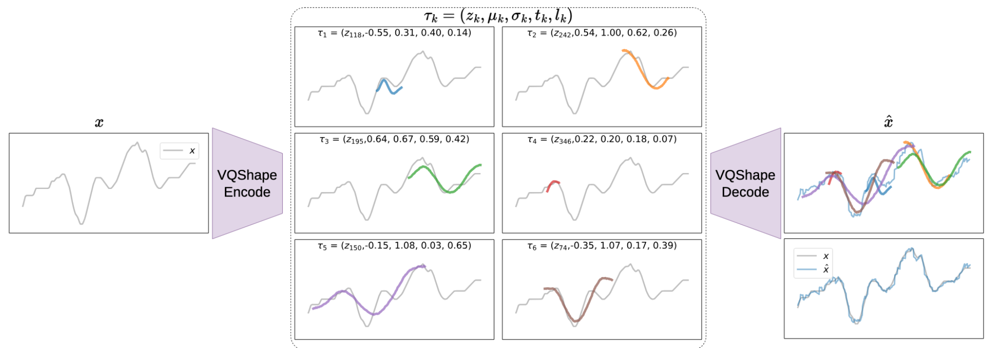

Figure 3: An example of abstracted shapes and their attributes (i.e., token representations) extracted by VQShape. For better presentation, we visualize 6 of the 64 shapes.

Example Shapes
in Classification

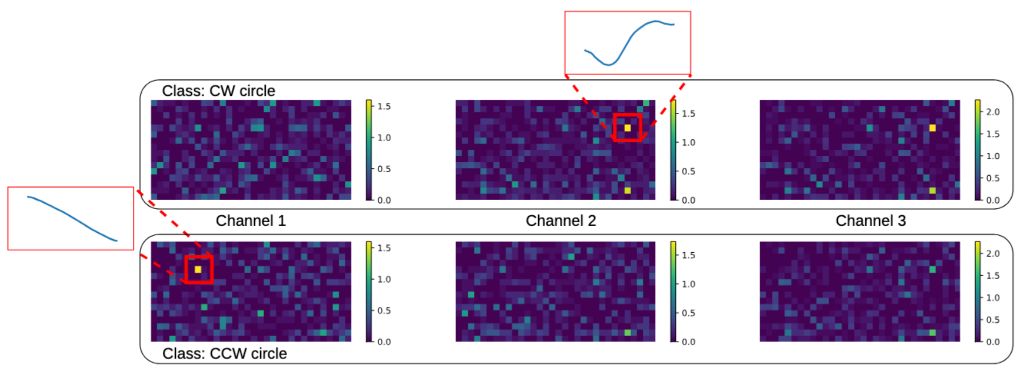

Experient Results on UEA Benchmarks

	Classical		Supervised					Unsupervised Representation				Pre-trained		
	DTW	STRF	DLinear	Autoformer	FEDformer	PatchTST	TimesNet	TS-TCC	TST	T-Rep	TS2Vec	MOMENT	UniTS	VQShape
Statistics with N/A														
Mean Accuracy	0.648	0.660	0.635	0.570	0.612	0.669	0.710	0.682	0.630	*0.719*	0.712	0.686	0.629	**0.723**
Median Accuracy	0.711	0.679	0.673	0.553	0.586	0.756	0.797	0.753	0.620	0.804	**0.812**	0.759	0.684	*0.810*
Mean Rank	7.138	7.828	8.690	9.448	7.750	8.296	*5.143*	7.172	8.448	5.207	**4.897**	5.929	9.828	5.621
Median Rank	7.0	8.0	9.0	10.0	8.0	8.0	5.0	8.0	9.0	*4.0*	**3.0**	5.5	10.0	5.0
Num. Top-1	3	2	1	0	1	0	1	3	1	4	**6**	5	0	**6**
Num. Top-3	8	5	5	0	8	2	8	5	4	*12*	**16**	11	0	9
Num. Win/Tie	14	20	22	25	19	20	14	18	20	15	13	13	25	-
Num. Lose	15	9	7	4	9	7	14	11	9	14	16	15	4	-
Wilcoxon p-value	0.206	0.023	0.002	0.000	0.051	0.000	0.898	0.156	0.022	0.536	0.576	0.733	0.000	-
Statistics without N/A														
Mean Accuracy	0.642	0.658	0.635	0.561	0.601	0.657	0.703	0.669	0.623	*0.710*	0.704	0.688	0.618	**0.720**
Median Accuracy	0.714	0.712	0.690	0.552	0.585	0.739	0.797	0.752	0.638	*0.802*	0.748	0.759	0.679	**0.812**
Mean Rank	7.231	7.692	8.923	9.462	7.731	8.346	5.308	7.538	8.462	*5.192*	**5.038**	5.885	10.115	5.538
Median Rank	6.5	7.5	9.5	9.5	8.0	8.5	5.000	8.0	9.5	*4.5*	**3.0**	5.5	10.5	*4.5*
Num. Top-1	3	2	1	0	1	0	1	3	1	3	**5**	**5**	0	**5**
Num. Top-3	7	5	5	0	8	2	7	4	4	*10*	**14**	*10*	0	8
Num. Win/Tie	13	18	20	22	17	20	13	17	19	14	12	12	23	-
Num. Lose	13	8	6	4	9	6	13	9	7	12	14	14	3	-
Wilcoxon p-value	0.187	0.036	0.001	0.001	0.111	0.001	0.803	0.094	0.036	0.653	0.696	1.000	0.000	-

Performance for Anomaly Prediction

Model	*Precision*	*Recall*	*F1-Score*
Transformer	0.53	0.51	0.52
LightGBM	0.81	0.61	0.69
VQShape	0.77	0.80	0.78

1. Transformer underfit due to small training data
2. VQShape is better due to pretraining on other datasets, it can even achieve 0.76 F1 by only adding a linear probing.

Stony Brook University

Brookhaven National Laboratory

Top Activated Codebook Tokens from VQShape

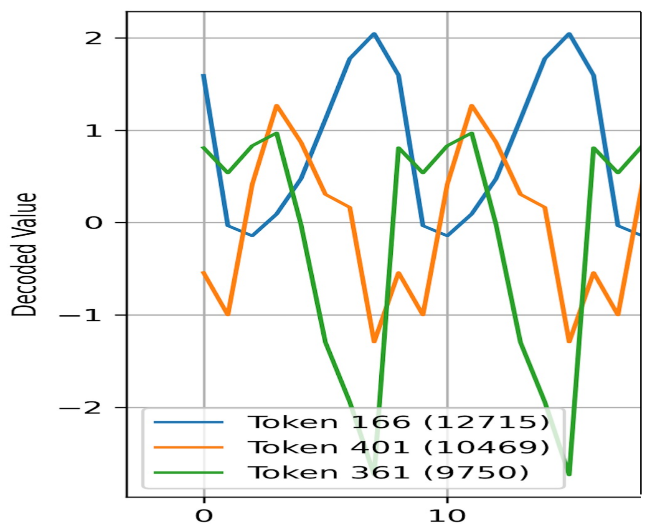

Here we visualize the reconstructed shapes of 3 frequently activated codebook tokens.
These shapes represent the most common local signal patterns captured by the VQShape.

Ongoing and Future work

- Ongoing work
  - Adding graphs to connect all sensor variables
    - Demonstrated useful in our previous works (Ma et al. UAI 2023)
  - Designing a new self-supervised learning framework which are more suitable for the anomaly detection/prediction task
  - Adding more data to capture bigger sparks and more patterns
- Future work
  - Online adaptation
  - Move to other AD tasks (e.g. JETs)

Stony Brook University

Brookhaven National Laboratory

Thank you!

Looking forward to questions and potential collaboration!