Interpretable time series neural representation for classification purposes

Etienne Le Naour [1,2], Ghislain Agoua [2], Nicolas Baskiotis [1], Vincent Guigue [3]

[1] Sorbonne Université, CNRS, ISIR, 75005 Paris, France
[2] EDF R&D, Palaiseau, France
[3] AgroParisTech, Palaiseau, France

DSAA 2023

Motivations

- Over the past 5 years, a lot of work on unsupervised neural representations for time series Franceschi et al. [2019], Yue et al. [2022], Zhang et al. [2022], Yang and Hong [2022], etc.

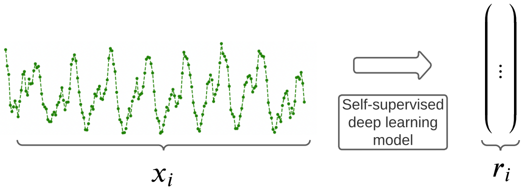

- But no works on interpretable neural representations for time series
- Although neural methods for representing time series are relatively new, non neural interpretable methods for representing time series are widely studied (e.g. Symbolic Aggregate approximation (SAX) methods)

We need to define requirements for an interpretable neural representation

Notations

- The univariate time series is denoted by the vector $\boldsymbol{x} \in \mathbb{R}^T$
- Let $\boldsymbol{r}$ be the symbolic neural representation. We denote by $\mathbb{A}$ the support (alphabet) common to all these elements: $\boldsymbol{r} = (r_1, \ldots, r_{T'}) \in \mathbb{A}^{T'}$
- ϕ_θ is the function that maps the time series into the representation
- $\psi_{\theta'}$ is the function that goes from the representation to the reconstruction space of the time series

Requirement n°1: temporal consistency

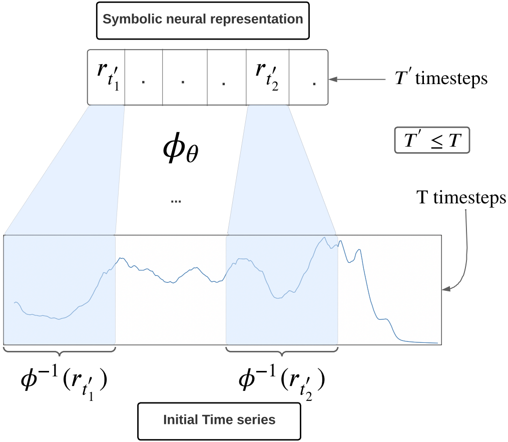

Requirement n°2: a decodable representation

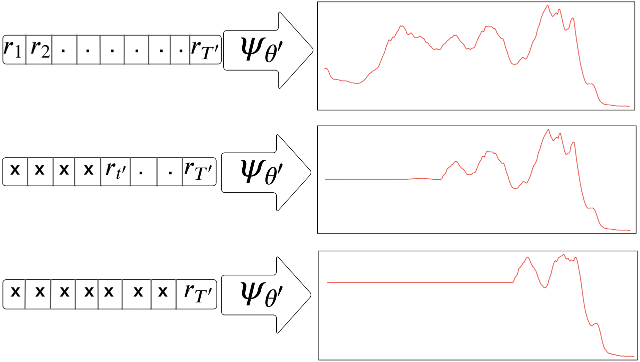

Requirement n°3: discrete symbolic representation

- The purpose of a symbolic neural representation method is to capture complex phenomena within the representation while being able to interpret and visualize the representation elements
- In addition, the support must be common to all elements of the symbolic representation
- The size of the support should be small

r_1	r_2	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$\cdot$	$r_{T'}$

$$\forall t \in \{1, \ldots, T'\}, r_t \in \mathbb{A} = \{a, b, c, d, e, f, g\}$$

Requirement n°4: shift equivariance properties

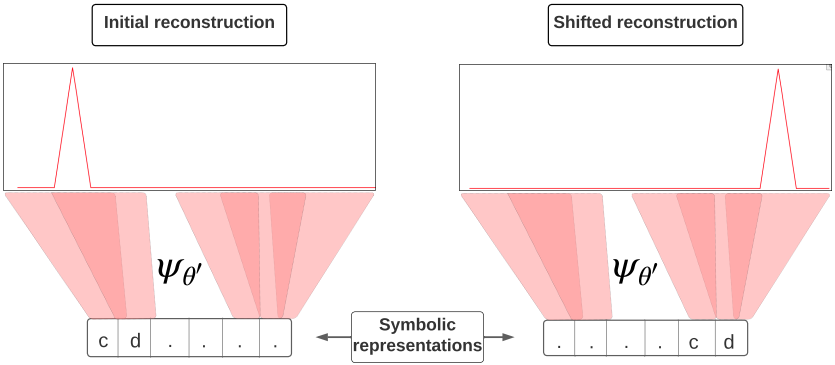

Requirement n°5: a representation adjustable to the frequency level

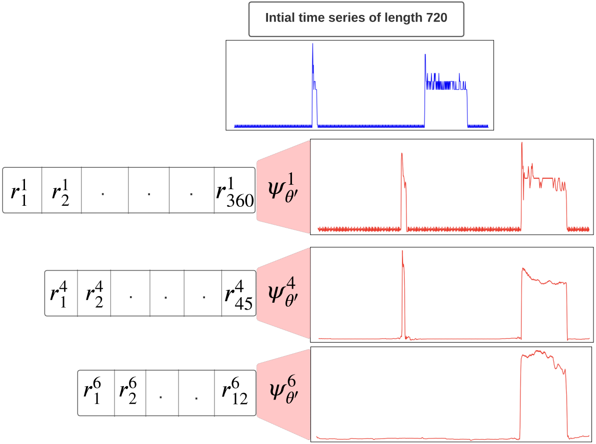

Proposed model

Global architecture

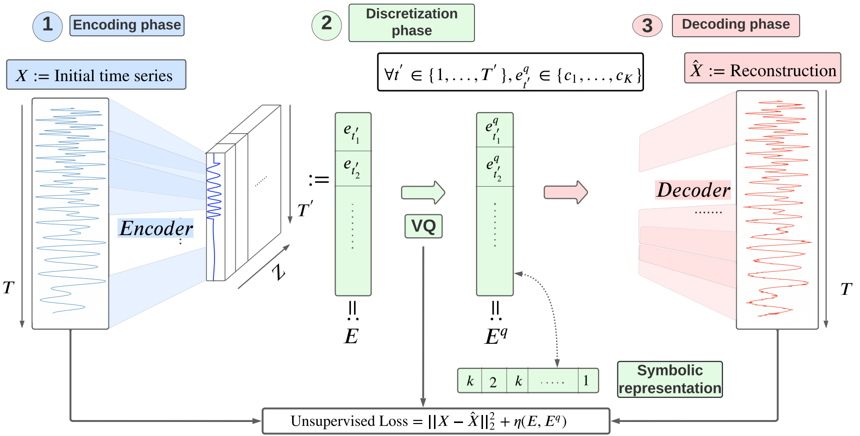

- Inspired from the Vector-Quantization Variational Auto-Encoder [van den Oord et al., 2017]

Encoder

- The encoder consists of B blocks with identical structure
- Inside one block the structure is :

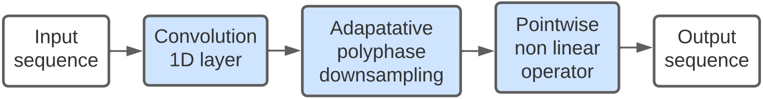

- Each block divides the length of the input sequence by two and preserves the shift equivariance property through adaptive polyphase downsampling (Chaman and Dokmanic [2021])
- The decoder blocks are define symmetrically to the encoder blocks

Discretization mechanism: vector quantization

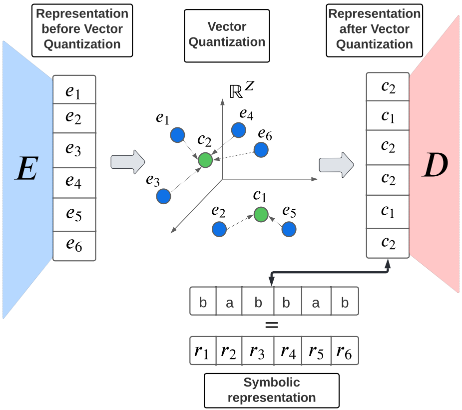

$$\boldsymbol{e}^{\boldsymbol{q}}_{\boldsymbol{t'}} \leftarrow \boldsymbol{c_k} \quad \text{where} \quad k = \arg\min_j ||\boldsymbol{e_{t'}} - \boldsymbol{c_j}||_2^2.$$

Downstream task: classification over extracted representations

Interpretable classification on a unique representation

- For a sample i, we can extract for the representation $\boldsymbol{r_i}$, the vector $\boldsymbol{h_i}$ which indicates if a symbolic subsequence is present in $\boldsymbol{r_i}$.
- d stands for the cardinal of all the subsequence space
- Thus, $\boldsymbol{h_i}$ is a vector of size d composed of 0 and 1 elements ($\boldsymbol{h_i} \in \{0,1\}^d$).
- Then we can use $\boldsymbol{h_i}$ to solve the classification problem using logistic regression

$$\arg\min_{w,b} \frac{1-\rho}{2}\boldsymbol{w}^{\boldsymbol{T}}\boldsymbol{w} + \rho\|\boldsymbol{w}\|_1 + \lambda\sum_{i=1}^{n}\log\left(\exp\left(-y_i\left(\boldsymbol{h_i}^T\boldsymbol{w}+b\right)\right)+1\right),$$

Quantitative results on UCR archive

Table: Accuracy on 25 UCR datasets compare to other interpretable methods. The best results are in bold and the second best results are underlined.

Datasets	Ours	SAX SEQL	SAX VSM	FS	LTS	DTW CV
Coffee	0.964	**1.000**	0.929	0.929	**1.000**	**1.000**
Computers	**0.728**	<u>0.676</u>	0.620	0.500	0.584	0.620
DistalPhalanxOAG	0.755	<u>0.818</u>	**0.842**	0.655	0.779	0.626
DistalPhalanxOC	<u>0.732</u>	0.718	0.728	**0.750**	0.719	0.725
DistalPhalanxTW	<u>0.640</u>	**0.748**	0.604	0.626	0.626	0.633
Earthquakes	0.734	**0.789**	<u>0.748</u>	0.705	0.741	0.727
ECG5000	**0.932**	0.924	0.910	0.923	**0.932**	0.925
FordA	<u>0.883</u>	0.851	0.827	0.787	**0.957**	0.691
GunPoint	0.940	<u>0.987</u>	<u>0.987</u>	0.947	**1.000**	0.913
Ham	0.705	0.705	**0.810**	0.648	0.667	0.600
Herring	**0.656**	0.578	<u>0.625</u>	0.531	<u>0.625</u>	0.531
ItalyPowerDemand	0.906	0.734	0.816	0.917	**0.970**	<u>0.955</u>
LargeKitchenApp	<u>0.864</u>	0.760	**0.877**	0.560	0.701	0.795
PhalangesOC	<u>0.748</u>	0.717	0.710	0.744	**0.765**	0.761
ProximalPhalanxOC	0.818	0.818	<u>0.828</u>	0.804	**0.834**	0.790
ProximalPhalanxOAG	0.839	<u>0.844</u>	0.824	0.780	**0.849**	0.785
ProximalPhalanxTW	0.771	**0.792**	0.610	0.702	<u>0.776</u>	0.756
RefrigerationDevices	0.533	<u>0.541</u>	**0.653**	0.333	0.515	0.440
ScreenType	<u>0.499</u>	0.461	**0.512**	0.413	0.429	0.411
ShapeletSim	<u>0.994</u>	<u>0.994</u>	0.717	**1.000**	0.950	0.700
SmallKitchenApp	**0.795**	<u>0.776</u>	0.579	0.333	0.664	0.672
Strawberry	**0.962**	0.954	<u>0.957</u>	0.903	0.911	0.946
Wafer	0.975	0.993	**0.999**	<u>0.997</u>	0.996	0.995
Wine	<u>0.759</u>	0.556	**0.963**	<u>0.759</u>	0.500	0.611
Worms	**0.714**	0.536	0.558	<u>0.649</u>	0.610	0.532
Mean	**0.793**	0.770	0.769	0.715	0.764	0.725

Qualitative results on the *ShapeletSim* dataset

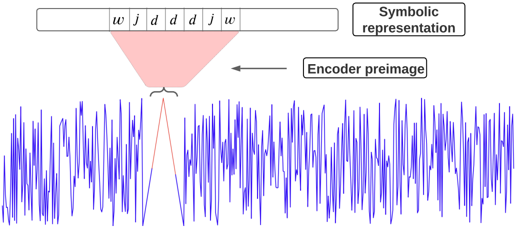

Figure: Local interpretability *ShapeletSim*

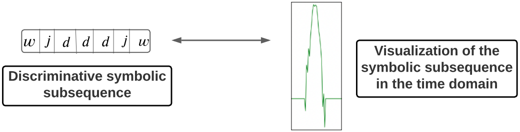

Figure: Global interpretability *ShapeletSim*

Thank you for your attention!

References I

A. Chaman and I. Dokmanic. Truly shift-equivariant convolutional neural networks with adaptive polyphase upsampling. In *55th Asilomar Conference on Signals, Systems, and Computers, ACSSC*, pages 1113–1120. IEEE, 2021. doi: 10.1109/IEEECONF53345.2021.9723377.

J. Franceschi, A. Dieuleveut, and M. Jaggi. Unsupervised scalable representation learning for multivariate time series. In *Advances in Neural Information Processing Systems 32*, pages 4652–4663, 2019.

A. van den Oord, O. Vinyals, and K. Kavukcuoglu. Neural discrete representation learning. In *Advances in Neural Information Processing Systems 30*, pages 6306–6315, 2017.

L. Yang and S. Hong. Unsupervised time-series representation learning with iterative bilinear temporal-spectral fusion. In *International Conference on Machine Learning, ICML*, volume 162 of *Proceedings of Machine Learning Research*, pages 25038–25054. PMLR, 2022.

References II

Z. Yue, Y. Wang, J. Duan, T. Yang, C. Huang, Y. Tong, and B. Xu. Ts2vec: Towards universal representation of time series. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI*, pages 8980–8987. AAAI Press, 2022.

X. Zhang, Z. Zhao, T. Tsiligkaridis, and M. Zitnik. Self-supervised contrastive pre-training for time series via time-frequency consistency. *CoRR*, abs/2206.08496, 2022. doi: 10.48550/arXiv.2206.08496.