Quantum-Inspired Tensor Networks for Semi-Supervised Anomaly Detection in Medical Imaging

Ema Puljak,[1, 2, *] Miguel A. González Ballester,[3, 2, 4] Maurizio Pierini,[1] Artur Garcia-Saez,[2] and Michele Grossi[1]

[1] *European Organization for Nuclear Research (CERN), CH-1211 Geneva, Switzerland*
[2] *Quantic, Barcelona Supercomputing Center, 08034 Barcelona, Spain*
[3] *BCN Medtech, Dept. of Engineering, Universitat Pompeu Fabra, 08002 Barcelona, Spain*
[4] *ICREA, 08010 Barcelona, Spain*

I. MOTIVATION

In the field of medical image analysis, several challenges affect learning tasks and crucial applications. These challenges include high-dimensional data with complex correlations, the presence of noise, the importance of interpretability, and the scalability of applied algorithms. Tensor Networks, originating from the analysis of quantum many-body systems, offer a powerful and efficient framework for handling high-dimensional datasets and potentially solving these issues. Tensor networks can be applied to various medical imaging tasks, including image segmentation, classification and reconstruction. While some of these applications have been explored in recent works [1–3], this paper focuses on addressing the challenge of anomaly detection, by leveraging a combination of unsupervised and semi-supervised settings, where only a subset of anomalous samples is labeled. This approach can mitigate the problem of the typically small number of labeled anomalous data in medical imaging, and the skepticism surrounding learning solely from healthy samples without human input during the model's training.

II. METHODS

This work primarily studies one-dimensional tensor network structures, which currently offer a solid theoretical understanding of their representation capabilities, and explores the role of image preprocessing and higher-dimensional embedding as crucial steps for presenting data to the tensor network.

Aspiring to apply these models to anomaly detection in 3D images, specifically for lung cancer screening CT images, we are first conducting a baseline study on 2D slices. We use the LIDC-IDRI [4] dataset of lung cancer CT scans, and similar to recent work [5], for now, we extract 256x256 2D patches and use MaxPooling to reduce dimensionality to 64x64 due to current memory constraints. Each image is marked as benign or malignant. Before data is fed into a tensor network, two key aspects need to be considered: how to preserve spatial correlations and how to embed each pixel. To address the first concern, we are incorporating spatial information into the data representation by applying positional encoding to images. Additionally, previous research has shown that flattening 2D images to 1D vectors can result in the loss of spatial pixel correlation [1]. To address this, we follow the method that preserves local interactions by only flattening the images in small patches [6], using the *squeeze* method as described in [1, 2]. As a next step, data is mapped to high-dimensional space V by a local feature map ϕ. In this work, we use a Gaussian radial basis function (RBF) for embedding.

Our research employs a combination of unsupervised and semi-supervised learning setups, as illustrated in Fig. 1, where a small percentage of labeled anomalies (malignant nodules) are injected into the training set. With our current focus on 1D tensor network structures, in the unsupervised step, we use the Spaced Matrix Product Operator (SMPO) to be trained for anomaly detection, where the approach for achieving good performance follows that of paper [7]. The SMPO model, trained on healthy images, learns to perform the linear transformation $S : V \to \nu$ with the objective of mapping normal training instances to the surface of the hypersphere in ν:

$$L(x) = \frac{1}{|\tau|} \sum_{x \in \tau} (\log(||S\phi(x)||_2^2) - 1)^2 \tag{1}$$

* Contact author: ema.puljak@cern.ch

The data with injected anomalies is then transformed by the SMPO model, and a semi-supervised objective L_{semi} (3) is used to train a Matrix Product State to correctly classify between healthy and malignant nodules, plus penalizing small values of the function (1) for malignant samples. For testing, we use separate anomalous and normal samples, assuming a real-world scenario where labeled anomalies are typically unavailable.

$$\mathrm{CE} = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{C} y_{ic}\log(p_{ic}) \tag{2}$$

$$L_{semi} = \mathrm{CE} + p_{ic}\frac{1}{L(x)} + (1-p_{ic})L(x), \tag{3}$$

where CE is the cross-entropy loss.

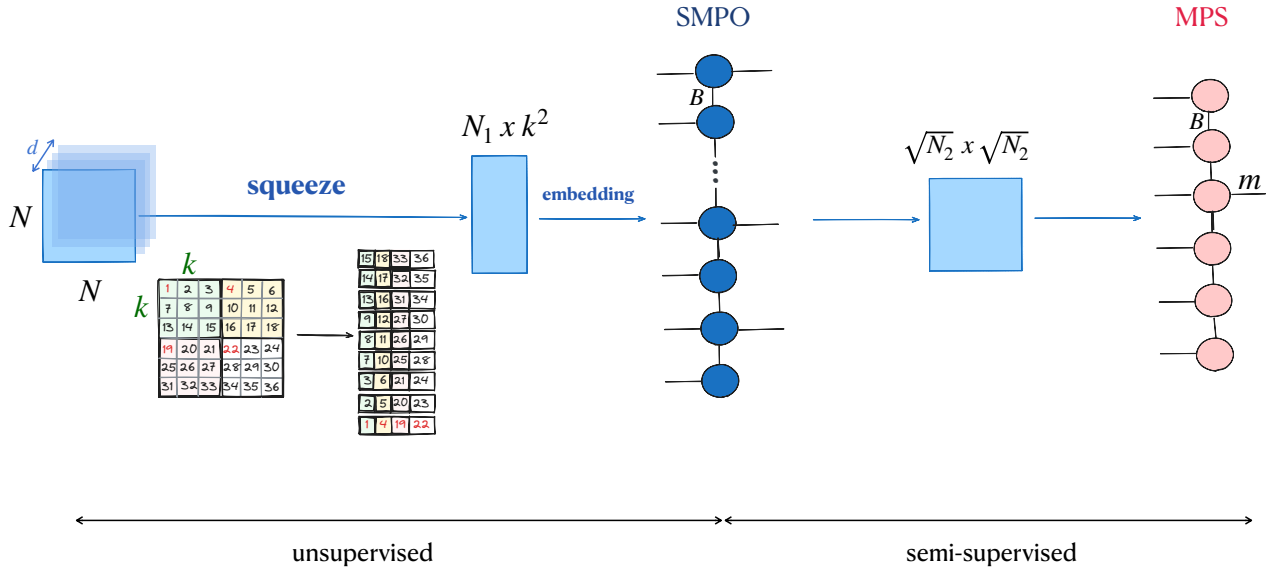

FIG. 1. Our research pipeline goes as follows: 1) data preprocessing, where d is the dimension of the positional embedding, k is the dimension of the kernel matrix for the *squeeze* function, $N_1 = N/k^2$, $N_2 = N_1/k^2$; 2) unsupervised training of the SMPO model on healthy samples; 3) semi-supervised training of the MPS model on transformed images, where m is the number of classes. B is the bond dimension.

III. RESULTS

After the preprocessing and embedding steps, the SMPO model is trained in an unsupervised manner on healthy images, achieving an Area Under the ROC Curve (AUC) score of 60.2% (first ROC curve in Figure 2). Next, the SMPO model when applied to a new set of healthy and malignant samples creates a first separation between samples, where the output is reshaped to the chosen image size (N_2xN_2). In the semi-supervised step, the MPS model, trained on these new healthy instances with 5% injected anomalies, shows improved performance, achieving an AUC score of 83.6%. For each step, the squared norm of the transformed images ($||S\phi(x)||_2^2$) is used as the anomaly score, where a smaller value indicates an anomaly. Here, S denotes the SMPO model. The distinction between malignant and healthy samples is more evident in the upper right plot of Figure 2, which is used to create the final ROC curve for the MPS model.

Additionally, we are conducting a study to understand the impact of different embedding functions. For each model, a grid search is performed to find the best hyperparameters. This includes optimizing for the amount of expressivity and correlation described by the bond dimension, as well as the output dimension for the SMPO model.

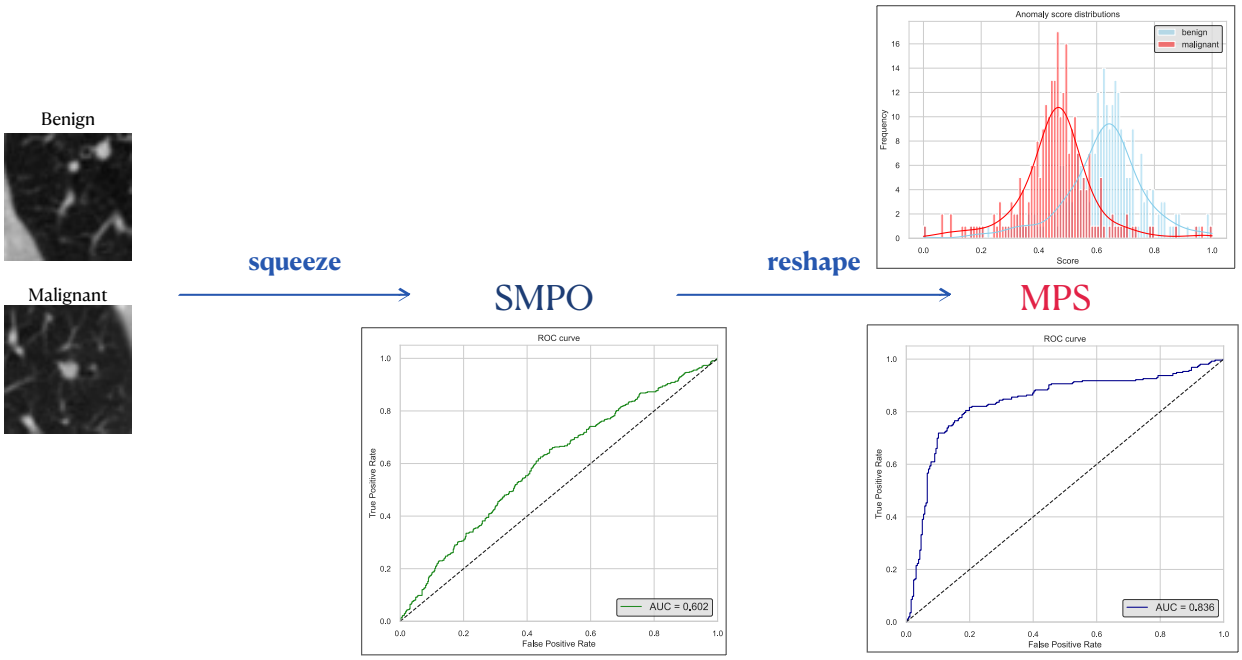

FIG. 2. Images, both malignant and benign, are preprocessed using the *squeeze* method. The SMPO model is first trained on healthy images, resulting in an AUC score of 60.2%. The data transformed by the SMPO model is then passed to the MPS model, which is trained with 5% injected anomalies. This approach improves anomaly detection performance, achieving an AUC score of 83.6%. The separation between anomalies and healthy samples is clearly visible on the anomaly score distribution plot (upper right).

IV. CONCLUSION AND FUTURE WORK

We are proposing a realistic setup based on 1D tensor network models for detecting anomalies in medical images, motivated by the limited amount of labeled samples and the need to analyze high-dimensional datasets. Preliminary results on lung nodule malignancy detection are promising. As future work, we will perform a systematic study of different data embeddings, as well as studying the expressivity and scalability of the chosen model structures. Further, we will perform exhaustive testing to compare with classical machine learning models. Finally, we will also consider the possibly of using 2D tensor network structures, towards processing 3D images.

[1] R. Selvan and E. B. Dam, Tensor networks for medical image classification (2020), arXiv:2004.10076 [cs.LG].
[2] R. Selvan, S. Ørting, and E. B. Dam, Locally orderless tensor networks for classifying two- and three-dimensional medical images (2021).
[3] R. Selvan, E. B. Dam, and J. Petersen, Segmenting two-dimensional structures with strided tensor networks, in *Information Processing in Medical Imaging* (Springer International Publishing, 2021) p. 401–414.
[4] S. G. Armato III, G. McLennan, L. Bidaut, M. McNitt-Gray, C. Meyer, *et al.*, The lung image database consortium (LIDC) and image database resource initiative (IDRI): A completed reference database of lung nodules on CT scans (2011).
[5] N. Gautam, A. Basu, and R. Sarkar, Lung cancer detection from thoracic ct scans using an ensemble of deep learning models, Neural Computing and Applications **36**, 2459 (2024).
[6] J. J. Koenderink and A. J. Van Doorn, The structure of locally orderless images, International Journal of Computer Vision (1999).
[7] J. Wang *et al.*, Anomaly detection with tensor networks, CoRR **abs/2006.02516** (2020), 2006.02516.