

Received March 29, 2022, accepted April 19, 2022, date of publication April 22, 2022, date of current version May 3, 2022.

Digital Object Identifier 10.1109/ACCESS.2022.3169776

Attention-Based Neural Bag-of-Features Learning for Sequence Data

DAT THANH TRAN¹, NIKOLAOS PASSALIS², (Graduate Student Member, IEEE),
ANASTASIOS TEFAS², (Member, IEEE), MONCEF GABBOUJ¹, (Fellow, IEEE),
AND ALEXANDROS IOSIFIDIS³

¹Department of Computing Sciences, Tampere University, Tampere 33720, Finland

²Department of Informatics, Aristotle University of Thessaloniki, 541 24 Thessaloniki, Greece

³Department of Electrical and Computer Engineering, Aarhus University, 8000 Aarhus, Denmark

Corresponding author: Dat Thanh Tran (thanh.tran@tuni.fi)

This work was supported in part by the European Union's Horizon 2020 Research and Innovation Program under Grant 871449 (OpenDR).

ABSTRACT In this paper, we propose 2D-Attention (2DA), a generic attention formulation for sequence data, which acts as a complementary computation block that can detect and focus on relevant sources of information for the given learning objective. The proposed attention module is incorporated into the recently proposed Neural Bag of Feature (NBoF) model to enhance its learning capacity. Since 2DA acts as a plug-in layer, injecting it into different computation stages of the NBoF model results in different 2DA-NBoF architectures, each of which possesses a unique interpretation. We conducted extensive experiments in financial forecasting, audio analysis as well as medical diagnosis problems to benchmark the proposed formulations in comparison with existing methods, including the widely used Gated Recurrent Units. Our empirical analysis shows that the proposed attention formulations can not only improve performances of NBoF models but also make them resilient to noisy data.

INDEX TERMS Attention mechanism, neural bag-of-features, time-series analysis.

I. INTRODUCTION

Learning problems in many fields involve sequence data such as time-series forecasting [1], [2], audio analysis [3], [4] or natural language processing [5], [6], all of which have been extensively studied. In many application scenarios, the observed sequence is highly non-stationary and noisy, which makes the task of modeling the underlying generating process more difficult. For example, in sound source separation in which the objective is to recover different unknown sources by filtering the observed mixtures, the existence of environmental noise is inherent and often complicates the separation process. Several mathematical techniques have been proposed to model the underlying data and noise distributions or to extract hand-crafted features, capturing certain desirable properties. In financial time-series analysis, representative examples include autoregressive (AR) and moving average (MA) [7] features, which were later extended with a differencing step to eliminate nonstationarity, known

as autoregressive integrated moving average (ARIMA) [8]. Gaussian processes and Hidden Markov Model were popular mathematical frameworks in audio analysis. To ensure mathematical and computational tractability, these classical models are often formulated under many assumptions, which are sensitive to initialization and misaligned with real-world conditions, thus limiting their professional usage in practice.

During the last decade, thanks to the development in stochastic optimization techniques and computing hardware, as well as the declining costs of data acquisition and storage, a data-driven approach based on deep neural networks and stochastic optimization has replaced the classical model-based approach and convex optimization. Nowadays, many of the state-of-the-art solutions for learning with sequence data are developed on the basis of neural networks. Notably, a class of neural network architecture called Recurrent Neural Networks (RNN), which is specifically designed to process variable-length sequences and to capture sequential patterns, has become the main workforce in different application domains. Another dedicated neural formulation for sequence data is the bilinear structures [1], [9], [10], which

The associate editor coordinating the review of this manuscript and approving it for publication was Wei Liu.

were proposed to separately capture the dependencies along the temporal and spatial dimension in financial time-series. Even existing neural architectures, which were originally proposed for visual inputs such as Convolutional Neural Network (CNN) [11] and Neural Bag-of-Features (NBoF) [12], have shown competitive performances in tackling sequence data compared to dedicated statistical models [13], [14]. The advantage of neural formulations over statistical learning and traditional hand-crafted features lies in the fact that fewer assumptions are made, and data is leveraged to automatically identify and extract task-relevant features in an end-to-end fashion.

Bag-of-Features (BoF) model [15] was originally proposed to build histogram representations from images. Later, it was shown that BoF could be successfully applied to extract high-level representations for other data modalities such as video and audio [16]–[19]. Learning BoF representations consists of two steps: *dictionary learning* and *feature quantization and encoding*. In the dictionary learning step, each object is first represented by a set of low-level features, which could be, for example, a collection of local descriptors like SIFT [20] for image object or word-level vector-encodings for a sentence object. These features are then used to generate a compact dictionary (codebook) comprising of the most representative features, also known as *codewords*. In the second step, the histogram representation of each object is extracted by quantizing its low-level features using the codebook.

Recently, Neural Bag-of-Features (NBoF) [12], a neural network generalization of the BoF model, has been proposed. Similar to its predecessor, NBoF can generate a fixed-size histogram vector from variable-size inputs. This neural network generalization works as a feature extraction layer, which can be combined and optimized jointly with other neural network layers to tackle both unsupervised and supervised objectives via stochastic optimization. Since the dictionary learning step in NBoF is updated in conjunction with other layers towards the end goal of optimizing an objective function, histogram vectors synthesized by NBoF are more representative than those produced by BoF in different learning scenarios such as visual recognition, information retrieval, and financial forecasting [12], [14], [21].

While the NBoF model works well in different learning problems, the current formulations still possess some limitations. In the aggregation step, all of the quantized features are simply averaged to form the histogram vector. For sequence data, this implies that the model only allows equal contributions of the quantized features coming from different time steps to form the output representation. Similarly, the quantization results produced by each codeword are considered equally important for every sequence in the training set. These properties limit the dictionary learning, quantization, and encoding process to fully take advantage of the data-driven approach.

To incorporate a higher degree of flexibility into the NBoF model, a weighing mechanism on a sequence level is desirable. That is, for each individual sequence, the model has the

flexibility to perform a weighted sum of quantized outputs in the aggregation step, with the coefficients being adaptively changed with respect to the input sequence, or to select/discard irrelevant codewords, given the input sequence. In neural network literature, this is often achieved by having some attention mechanisms [10], [22], [23]. The idea of attention is inspired by the phenomenon observed in the human visual cortex that visual stimuli from multiple objects actively compete for neural encoding.

Although various attention mechanisms have been proposed for existing neural network architectures such as CNN [22], [24], LSTM [23], [25] or Bilinear structure [10], there is yet any formulation for the NBoF model when learning with sequence data. To have a generic attention mechanism that can be applied in a plug-and-play manner, in this work, we propose 2D-Attention (2DA), a neural network module that promotes competitions among different rows or columns in the input matrix and only (soft) selects those which win for attention. We will then demonstrate that by injecting 2DA into NBoF, we can overcome those limitations mentioned previously. The contributions of our work can be summarized as follows:

- We propose a new type of attention formulation for matrix data, which is dubbed as 2DA. The proposed layer acts as a complementary computation block, which is capable of identifying relevant sources of information to perform selective masking on the given input matrix.
- We incorporate 2DA into different stages of the Neural Bag-of-Features (NBoF) model, creating various 2DA-NBoF extensions that can enhance the feature quantization or histogram accumulation step in the NBoF model. Extensive experiments were conducted in three different application domains: financial forecasting, audio analysis, and medical diagnosis, which demonstrate the effectiveness of our attention module in improving the NBoF model. In cases of noisy input, a variant of 2DA-NBoF shows resilience to noises by filtering out the noisy source of information before the feature quantization step.

The remainder of the paper is organized as follows: in Section II, we review the NBoF model and its extensions for time-series data, as well as previously proposed attention mechanisms in the neural network literature. In Section III, we first present the proposed attention module 2DA and its interpretation. Several extensions of the NBoF model that incorporates 2DA are then presented. In Section IV, we provide details of our experiment protocols and quantitative analysis. Section V concludes our work with possible future research directions.

II. RELATED WORK

The NBoF model [12] consists of two components: a quantization layer and an accumulation layer. Each quantization neuron in the quantization layer performs like a codeword,

which can be updated via BackPropagation algorithm. In the original formulation [12], the Radial Basis Function (RBF) layer was used for feature quantization. Recently, it has been shown that the hyperbolic kernel is also effective for the feature quantization step [26]. Here we describe the original formulation with RBF layer.

Let K be the number of neurons (codewords) in the RBF layer and $\mathbf{v}_k \in \mathbb{R}^D$ be the k -th codeword. In addition, the shape of the Gaussian function modeled by each neuron can be adjusted via parameter \mathbf{w}_k . Let us denote the sequence of N features as $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_N] \in \mathbb{R}^{D \times N}$ with $\mathbf{x}_n \in \mathbb{R}^D$, $n = 1, \dots, N$. The output of the k -th RBF neuron given the input feature \mathbf{x}_n is the following:

$$\phi_{n,k} = \frac{\exp(-\|(\mathbf{x}_n - \mathbf{v}_k) \odot \mathbf{w}_k\|_2)}{\sum_{m=1}^K \exp(-\|(\mathbf{x}_n - \mathbf{v}_m) \odot \mathbf{w}_m\|_2)} \quad (1)$$

where \odot denotes element-wise product and $\mathbf{w}_k \in \mathbb{R}^D$ is the learnable weight vector that enables the shape of Gaussian kernel associated with the k -th RBF neuron to change.

As the sequence \mathbf{X} goes through the quantization layer, each feature \mathbf{x}_n is quantized as $\phi_n = [\phi_{n,1}, \dots, \phi_{n,K}]^T \in \mathbb{R}^K$, producing a sequence of quantized features $\Phi = [\phi_1, \dots, \phi_N] \in \mathbb{R}^{K \times N}$. The accumulation layer aggregates the information in Φ by calculating the averaged quantized feature:

$$\mathbf{y} = \frac{1}{N} \sum_{n=1}^N \phi_n \quad (2)$$

There have been few extensions of NBoF model for sequence data. For example, Temporal Neural Bag-of-Features (TNBoF) model with different specialized codebooks has been proposed in [27] to capture both short-term and long-term temporal information in financial time-series. In [26], the authors derived the logistic formulation of the NBoF model using the hyperbolic kernel instead of the RBF kernel for the quantization step and proposed an adaptive scaling mechanism which showed significant improvements in training stability and performance of the NBoF networks.

While the attention mechanism was biologically inspired from the perspective of visual processing, this technique has also inspired and advanced several works in sequence data analysis, notably in sequence-to-sequence learning tasks. The first attention formulation applied to sequence data was proposed in [6] for tackling machine translation tasks. In this formulation, the authors proposed to construct the context vectors in Sequence-to-sequence Recurrent Neural Network model by selectively combining some hidden states, rather than using the last hidden state as the context vector. The selection coefficients, also known as attention weights, are computed adaptively based on the given input sequence, and updated jointly with other parameters during stochastic optimization.

The successful application of attention mechanism in machine translation tasks has led to the emergence of other

attention formulations, which are designed to capture different types of salient information in sequence data. For example, in [28], the authors proposed a formulation that can detect pseudo-periods in certain types of time-series, such as energy consumption or meteorology data. To predict the future stock index, a dual-stage attention mechanism was proposed in [25] for RNN to actively select relevant exogenous series and temporal instances. Similarly, to highlight and focus on important temporal events in Limit Order Book, the authors in [10] proposed a method to calculate attention masks for bilinear networks. Although an attention formulation has been proposed for the convolutional NBoF model in [24] to estimate the true color of images capturing by different devices, this formulation only works with image data. To the best of our knowledge, there has been no attention formulation for the NBoF model to tackle sequence data.

III. PROPOSED METHODS

In this Section, we will first present 2D-Attention (2DA), our proposed attention calculation for matrix data. Then, we will show how 2DA can be used to address different limitations of the NBoF model as described in Section I. Throughout the paper, we denote scalar values by either lower-case or upper-case characters (a, b, A, B, \dots), vectors by lower-case bold-face characters ($\mathbf{x}, \mathbf{y}, \dots$), matrices by upper-case bold-face characters ($\mathbf{X}, \mathbf{Y}, \dots$), and mathematical functions by calligraphy characters $\mathcal{F}, \mathcal{G}, \dots$. In addition, we use x_{mn} to denote the element at position (m, n) in a matrix \mathbf{X} .

A. 2D-ATTENTION

A matrix $\mathbf{S} \in \mathbb{R}^{M \times N}$ is a second-order tensor which has two modes, with M and N are the dimensions of the first and second mode, respectively. The matrix representation provides a natural way to represent a signal with two different sources of information. For example, a multivariate time-series is represented by a matrix with one mode representing the temporal dimension, and the other mode represents different sources that generate individual series.

The general idea of attention mechanism is to highlight important elements in the data while discarding irrelevant ones. For data represented as a matrix \mathbf{S} , rather than considering each element in \mathbf{S} individually, we would like to actively select certain columns or rows of \mathbf{S} while discarding the others. This is because columns or rows of \mathbf{S} usually form coherent sub-groups of the data. For example, discarding some temporal events or some individual series in a multivariate series corresponds to removing some rows or columns, depending on the orientation of the matrix.

To adaptively determine and focus on different columns or rows of a matrix, we propose 2D-Attention (2DA), with the functional form denoted by \mathcal{F}_{2DA} . This function takes a matrix $\mathbf{S} \in \mathbb{R}^{M \times N}$ as the input, and returns $\tilde{\mathbf{S}} \in \mathbb{R}^{M \times N}$ as the output. That is:

$$\tilde{\mathbf{S}} = \mathcal{F}_{2DA}(\mathbf{S}) \quad (3)$$

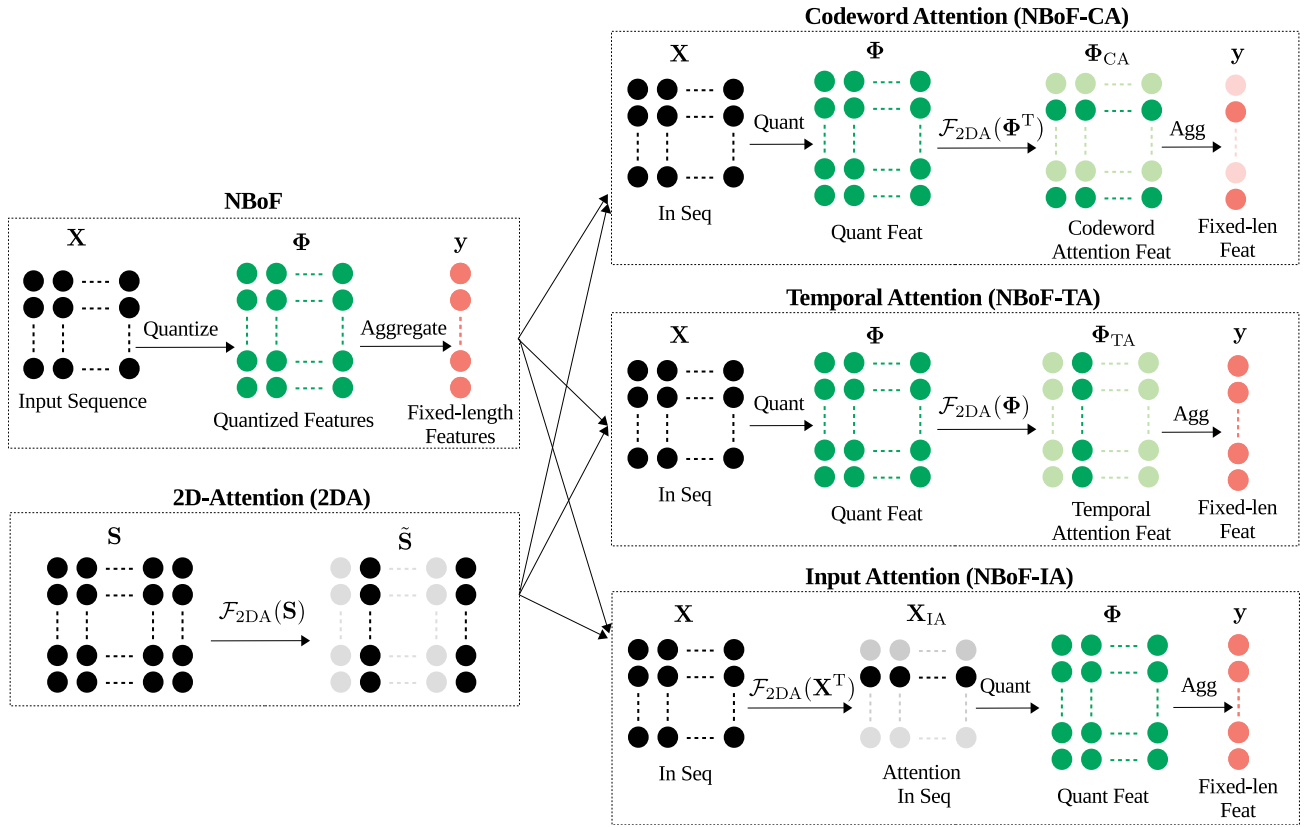


FIGURE 1. Illustration of the proposed attention formulation (2DA) and different attention-based NBoF models.

\tilde{S} can be considered as a filtered version of S , where irrelevant columns of S with respect to the learning problem are zeroed out. Here we should note that \mathcal{F}_{2DA} performs adaptive attention with respect to columns of S . To focus on different rows of S , we can simply apply \mathcal{F}_{2DA} to the transpose of S .

The selection or rejection of the columns of S is conducted via element-wise matrix multiplications as follows:

$$\tilde{S} = \tau(S \odot A) + (1 - \tau)S \quad (4)$$

where $A \in \mathbb{R}^{M \times N}$ denotes the attention mask with values in the range $[0, 1]$. Each column in A encodes the importance of the corresponding column in S . That is, the attention mask contains values that are close to 1 corresponding to those columns in S that contain important information for the downstream learning task and vice versa. In Eq.(4), parameter $\tau \in \mathbb{R}$, which is jointly optimized with other parameters, is used to allow flexible control of the attention mechanism: when S contains redundant or noisy information in its columns, the effect of attention mask A is enabled by pushing τ close to 1; on the other hand, when every column of S is necessary, i.e., there is no need for attention, pushing τ close to 0 will disable the effect of A . The necessity of attention is thus automatically determined by optimizing τ with respect to a given problem.

To calculate the attention mask A , the proposed 2DA method learns to measure the relative importance between columns of S via a specially designed weight matrix $W \in \mathbb{R}^{N \times N}$: all elements of W are learnable, i.e., they are updated during stochastic optimization, except the diagonal elements, which are fixed to $1/N$. The attention mask is calculated as follows:

$$\begin{aligned} A &= \mathcal{G}(Z) \\ Z &= SW \end{aligned} \quad (5)$$

where $\mathcal{G}(Z)$ denotes the soft-max function that is applied to every row of Z . That is, every element of A is non-negative, and each row of A sums up to 1. Similar to other attention formulations [6], [22], [25], we use soft-max normalization to promote competitions between different columns of Z .

As mentioned previously, the weight matrix W is used to measure the relative importance between columns of S , which is encoded in Z , and thus A . In order to see this, let us denote by $s_n \in \mathbb{R}^M$ and $z_n \in \mathbb{R}^M$ the n -th column of S and Z , respectively. Since $Z = SW$, the n -th column of Z , i.e., z_n , is calculated as the weighted combination of N columns of S , with the weight of the n -th column always equal to $1/N$ since the diagonal elements of W are fixed to $1/N$. In this way, element z_{mn} (in Z) encodes the relative importance of s_{mn} (in S) with respect to other s_{mk} , for $k \neq n$.

B. ATTENTION-BASED NEURAL BAG-OF-FEATURES

In this subsection, we will show how the proposed attention module 2DA can be used to address different limitations of the NBoF model described in Section I.

Codeword Attention: in the NBoF model, quantization results produced by each quantization neuron (codeword) are considered equally important for every input sequence. This property limits the feature quantization step to fully take advantage of the data-driven approach. In order to overcome this limitation, the proposed 2DA block can be applied to the quantized features to highlight or discard the outputs of certain quantization neurons. By doing so, the NBoF model is explicitly encouraged to learn a subset of specialized codewords for a given input pattern.

Particularly, given the quantized features denoted by $\Phi \in \mathbb{R}^{K \times N}$ as described in Section II, we propose to apply attention to the rows of Φ because the first mode of Φ with dimension K denotes the number of quantization neurons or codewords. Since 2DA operates on the columns of the input matrix, the attention-based quantized features is calculated as follows:

$$\Phi_{CA} = \mathcal{F}_{2DA}(\Phi^T) \quad (6)$$

where Φ^T denotes the transpose of Φ .

Temporal Attention: another limitation of the NBoF model lies in the aggregation step. In order to produce a fixed-length representation of the input sequence, the aggregation step in the NBoF model simply computes the mean of quantized features along the temporal mode. In this way, the NBoF model only allows equal contributions of all quantized features, disregarding the temporal information. In fact, the idea of giving different weights to different time instances has been adopted in previous works under different formulations [10], [25]. Using our proposed 2DA formulation, it is straightforward to enable the NBoF model to attend to salient temporal information as follows:

$$\Phi_{TA} = \mathcal{F}_{2DA}(\Phi) \quad (7)$$

Since each column of Φ contains quantized features of each time step, to obtain temporal attention-based features Φ_{TA} we simply apply \mathcal{F}_{2DA} to Φ as in Eq. (7). Φ_{TA} is then averaged along the second dimension to produce the fixed-length representation of the input sequence. Although we still perform averaging in the aggregation step, the fixed-length representation is no longer the average of the quantized features, but a weighted average. This is because each time instance (column) in Φ_{TA} has been scaled by different factors via the attention mechanism.

Input Attention: noisy data is an inherent problem in many real-world applications. Noises might surface during the data acquisition process, such as ambient noise in audio signals or power line interference and motion artifacts in Electrocardiogram signals. In other scenarios, noises are inherent in the problem formulation since the relevance between the input sources and the targets might be unclear. For example, in stock prediction, it is intuitive to use related

stocks' data, e.g., those coming from the same market sector, as the input to construct forecasting models, although some of them might be irrelevant to the movement of the target stock.

The proposed attention mechanism can also be used to filter out potential noisy series in a multivariate series as follows:

$$\mathbf{X}_{IA} = \mathcal{F}_{2DA}(\mathbf{X}^T) \quad (8)$$

where $\mathbf{X} \in \mathbb{R}^{D \times N}$ denotes an input sequence of N steps of the NBoF model as specified in Section II. Since we would like to apply attention over the individual series (rows of \mathbf{X}), \mathcal{F}_{2DA} is applied to the transpose of \mathbf{X} .

The proposed attention variants of the NBoF model are illustrated in Figure 1.

IV. EXPERIMENTS

In this section, we provide detailed descriptions and results of our empirical analysis, which demonstrate the advantages of attention-based NBoF models proposed in Section III. Experiments were conducted in different types of sequence data, namely financial time-series in stock movement prediction problem, Electrocardiogram (ECG) and Phonocardiogram (PCG) in heart anomaly detection problems, and audio recording in music genre recognition and acoustic scene classification problems.

The experiments were conducted with the recently proposed logistic formulation of the NBoF model [26], i.e., the hyperbolic kernel was used in the quantization layer. In addition, we also experimented with the temporal variant of the NBoF model as proposed in [27] with a long-term and a short-term codebook. This variant is denoted as TNBoF. The codebook attention, temporal attention, and input attention when applied to the NBoF model are denoted as NBoF-CA, NBoF-TA, and NBoF-IA, respectively. The corresponding attention variants for the TNBoF model are denoted as TNBoF-CA, TNBoF-TA, and TNBoF-IA. In addition to the NBoF and TNBoF models serving as the baseline models, we also evaluated RNN models using Gated Recurrent Units (GRU) [5].

A. FINANCIAL FORECASTING EXPERIMENTS

Although extensively studied over the last decades, financial forecasting still remains as the most challenging tasks among time-series predictions [29]. This is due to the complex dynamics of the financial markets, which make the observed data highly stationary and noisy. For this reason, we selected the stock movement prediction problem in FI2010 dataset [2] as a representative problem in time-series forecasting. FI2010 is the largest publicly available Limit Order Book (LOB) dataset, which contains approximately 4.5 million order events. The limit orders came from 5 Finnish stocks traded in Helsinki Exchange (operated by NASDAQ Nordic) over 10 business days. At each time instance, the dataset provides information (the prices and volumes) of the top 10 levels, leading to a 40-dimensional vector representation.

TABLE 1. Results on the FI2010 dataset. The second column shows performances (averaged F1 in %) of all models without using any convolution layer as preprocessing layers. The third column shows the corresponding performances (average F1 in %) when additional convolution layers were used as preprocessing layers. The best results in each column are highlighted in bold-face.

| Models | without conv | with conv |
|---|--------------|--------------------|
| <i>Prediction Horizon $H = 10$</i> | | |
| GRU [5] | 60.92±00.09 | 62.21±00.30 |
| NBoF [26] | 33.07±00.66 | 66.34±00.60 |
| NBoF-CA (our) | 40.81±00.05 | 67.56±00.02 |
| NBoF-TA (our) | 40.83±00.21 | 67.98±00.09 |
| TNBoF [27] | 36.66±00.51 | 66.74±00.36 |
| TNBoF-CA (our) | 45.61±00.16 | 67.76±00.05 |
| TNBoF-TA (our) | 45.97±00.15 | 67.88±00.13 |
| <i>Prediction Horizon $H = 20$</i> | | |
| GRU [5] | 51.61±00.25 | 53.83±00.14 |
| NBoF [26] | 38.06±00.53 | 58.85±00.05 |
| NBoF-CA (our) | 40.08±00.07 | 59.31±00.44 |
| NBoF-TA (our) | 40.34±00.06 | 60.10±00.03 |
| TNBoF [27] | 38.67±00.50 | 59.61±00.48 |
| TNBoF-CA (our) | 43.06±00.03 | 59.73±00.19 |
| TNBoF-TA (our) | 43.50±00.15 | 60.04±00.24 |
| <i>Prediction Horizon $H = 50$</i> | | |
| GRU [5] | 63.13±00.19 | 65.93±00.03 |
| NBoF [26] | 48.25±00.25 | 68.84±02.29 |
| NBoF-CA (our) | 49.34±00.17 | 73.25±00.27 |
| NBoF-TA (our) | 49.21±00.16 | 73.02±00.04 |
| TNBoF [27] | 54.06±00.14 | 69.27±01.09 |
| TNBoF-CA (our) | 57.15±00.21 | 73.77±00.37 |
| TNBoF-TA (our) | 57.41±00.06 | 73.40±00.08 |

The FI2010 dataset is used to investigate the problem of mid-price movement prediction in the next $H = \{10, 20, 50\}$ order events. The mid-price at a given time instance is the average between the best buy and best sell prices. This quantity is a virtual price since no trade can happen at this particular price at the given time instance. The movement of mid-price (stationary, increasing, decreasing) reflects the dynamic of the LOB and the market, thus plays an important role in financial analysis. The dataset provides the movement labels, given the future horizon $H = \{10, 20, 50\}$. Details regarding the FI2010 dataset and LOB can be found in [2].

We followed the same experimental setup proposed in [10], which used the first 7 days for training the models and the last 3 days to test the performances. Due to the imbalanced nature of the problem, we reported averaged F1 score per movement as the main performance metric, similar to prior experiments [1], [10]. Detailed information about the training hyper-parameters and the network architectures is provided in the Appendix.

The experiment results for FI2010 are shown in Table 1. In the second column of Table 1, we list the performances

of all models without using any convolution layers as the preprocessing layers. That is, the results in the second column of Table 1 are produced by architectures consisting of only the layer of interests (such as GRU, NBoF, and so on), plus the fully connected layers for generating predictions. In this setting, the GRU models outperform all variants of the NBoF model. This is expected since the NBoF model, by construction, is not designed to capture local features and long-term dependency in the input sequence. We can easily observe that this limitation can be partially overcome with the TNBoF variant, which uses two separate codebooks to capture the short-term and long-term dependency. By applying our proposed attention mechanism, performances of both the NBoF and TNBoF models are further boosted.

The third column of Table 1 shows the performances of all models when using two additional convolution layers as the local feature extractor, prior to applying the layer of interest. It is clear that all of the models benefit from the additional convolution layers, especially the NBoF model and its variants. In this setting, the GRU models no longer dominate the family of NBoF models. In fact, the GRU models become the worst-performing ones in the third column of Table 1. Furthermore, both codebook attention (NBoF-CA, TNBoF-CA) and temporal attention (NBoF-TA, TNBoF-TA) consistently enhance the baselines' performances, making attention-based models the best-performing ones.

Here we should note that although the baseline models (NBoF, TNBoF) use the adaptive scaling step proposed in [26] to improve training stability, we did not employ this step in attention-based models. The reason stems from the fact that adaptive scaling introduces additional degrees of freedom to the quantization step, which counteracts the constraining effects of the attention mechanism. Table 2 shows the performances of attention-based models on the FI2010 dataset, with and without the adaptive scaling step proposed in [26]. In most cases, the adaptive scaling step slightly degrades the performances of the attention-based models. As we will see in the next subsection, this effect is more noticeable in audio datasets.

B. AUDIO ANALYSIS EXPERIMENTS

One of the important types of sequence data is audio recordings. In this subsection, we present our empirical analysis using two audio datasets, representing two different applications in audio signal analysis: music genre recognition and acoustic scene classification.

In the first application, the objective is to train an acoustic system that recognizes the genre of a short musical recording. For this purpose, we conducted experiments using the *small subset* of the FMA dataset [30], which contains 8000 tracks coming from the 8 most popular genres: *pop*, *instrumental*, *experimental*, *folk*, *rock*, *international*, *electronic*, and *hip-hop*. Each audio clip is 30s long, which is transformed to Mel-spectrogram representation with 128 frequency bands using a window of 10ms with an overlap of 2.5ms.

TABLE 2. Performances (averaged F1 in %) of attention-based models on the F12010 dataset, with and without the adaptive scaling step proposed in [26]. The best results in each row are highlighted in bold-face.

| Models | adaptive scale | no adaptive scale |
|----------------------------------|---------------------|---------------------|
| <i>Prediction Horizon H = 10</i> | | |
| NBoF-CA | 66.92±00.08 | 67.56 ±00.02 |
| NBoF-TA | 67.34±00.14 | 67.98 ±00.09 |
| TNBoF-CA | 67.84 ±00.16 | 67.76±00.05 |
| TNBoF-TA | 67.16±00.32 | 67.88 ±00.13 |
| <i>Prediction Horizon H = 20</i> | | |
| NBoF-CA | 59.25±00.14 | 59.31 ±00.44 |
| NBoF-TA | 59.26±00.18 | 60.10 ±00.03 |
| TNBoF-CA | 59.75 ±00.35 | 59.73±00.19 |
| TNBoF-TA | 59.78±00.11 | 60.04 ±00.24 |
| <i>Prediction Horizon H = 30</i> | | |
| NBoF-CA | 71.93±00.14 | 73.25 ±00.27 |
| NBoF-TA | 47.89±23.87 | 73.02 ±00.04 |
| TNBoF-CA | 71.30±00.29 | 73.77 ±00.37 |
| TNBoF-TA | 72.32±00.05 | 73.40 ±00.08 |

TABLE 3. Audio analysis results of FMA and TUT-UAS2018 datasets. The second column shows performances (test accuracy in %) of all models without using any convolution layer as preprocessing layers. The third column shows the corresponding performances (test accuracy in %) when additional convolution layers were used as preprocessing layers. The best results in each column are highlighted in bold-face.

| Models | without conv | with conv |
|----------------------------|---------------------|---------------------|
| FMA Dataset | | |
| GRU [5] | 33.87±00.27 | 42.06±00.92 |
| NBoF [26] | 35.65±01.41 | 38.83±01.83 |
| NBoF-CA (our) | 38.29 ±00.95 | 41.08±01.85 |
| NBoF-TA (our) | 36.79±00.29 | 41.46±01.64 |
| TNBoF [27] | 35.25±03.50 | 39.13±00.53 |
| TNBoF-CA (our) | 37.29±00.66 | 42.67 ±01.23 |
| TNBoF-TA (our) | 36.50±01.49 | 42.58±00.91 |
| TUT-UAS2018 Dataset | | |
| GRU [5] | 56.83±00.78 | 56.89±00.93 |
| NBoF [26] | 52.02±00.18 | 55.92±01.40 |
| NBoF-CA (our) | 56.89 ±00.17 | 57.68 ±00.65 |
| NBoF-TA (our) | 56.09±00.25 | 57.63±00.30 |
| TNBoF [27] | 52.62±00.78 | 55.30±00.13 |
| TNBoF-CA (our) | 56.19±00.23 | 56.73±00.51 |
| TNBoF-TA (our) | 56.34±00.62 | 57.33±00.20 |

The preprocessing step results in the input sequence having dimensions of 128×640 .

In the second application, the objective is to train an acoustic system that can classify the type of environment based on its surrounding sounds. For this application, we used the TUT-UAS2018 dataset [31], which contains 8640 audio clips recorded from 10 urban acoustic scenes: *airport*, *shopping_mall*, *metro_station*, *street_pedestrian*,

TABLE 4. Performances (test accuracy in %) of attention-based models on FMA and TUT-UAS2018 datasets, with and without the adaptive scaling step proposed in [26]. The best results in each row are highlighted in bold-face.

| Models | adaptive scale | no adaptive scale |
|----------------------------|----------------|---------------------|
| FMA Dataset | | |
| NBoF-CA | 38.37±02.04 | 41.08 ±01.85 |
| NBoF-TA | 34.29±05.42 | 41.46 ±01.64 |
| TNBoF-CA | 39.96±00.66 | 42.67 ±01.23 |
| TNBoF-TA | 40.21±00.16 | 42.58 ±00.91 |
| TUT-UAS2018 Dataset | | |
| NBoF-CA | 40.57±14.36 | 57.68 ±00.65 |
| NBoF-TA | 56.62±00.39 | 57.63 ±00.30 |
| TNBoF-CA | 56.59±00.51 | 56.73 ±00.51 |
| TNBoF-TA | 55.05±01.25 | 57.33 ±00.20 |

public_square, *street_traffic*, *tram*, *bus*, *metro*, *park*. Similar to the FMA dataset, we also transformed each audio clip to Mel-spectrogram with 128 frequency bands using a window of 40ms with an overlap of 20ms, which results in the input sequence of size 128×500 .

For both applications, we report the test accuracy as the performance metric. Experiment results on FMA and TUT-UAS2018 dataset are shown in Table 3. Performances of all models with and without using convolution layers for feature extraction are presented in the second and third columns, respectively.

In the FMA dataset, we can easily observe significant improvements in all models when using additional convolution layers. Without any convolution layer, the NBoF and TNBoF models outperform the GRU model on average, however, with larger variances. The order reverses when additional convolution layers were used: the GRU model enjoys a huge benefit from the preprocessing layers, outperforming the NBoF and TNBoF models. In both scenarios, i.e., with or without convolution layers, the proposed attention block greatly enhances the baseline NBoF and TNBoF models, making them the best performing models in this task.

In the TUT-UAS2018 dataset, while adding convolution layers leads to noticeable improvements for the baseline NBoF and TNBoF, we observe no similar improvement for the GRU model. Similar to the FMA dataset, we observe consistent performance boost in the TUT-UAS2018 dataset by incorporating the proposed attention block to the NBoF and TNBoF models.

Similar to Section IV-A, we also conducted experiments in FMA and TUT-UAS2018 datasets to analyze the effects of the adaptive scaling step proposed in [26]. The results are shown in Table 4. The results obtained from both audio analysis tasks in Table 4 are consistent with what we observe from the stock movement prediction task in Table 2: although the adaptive scaling step can enhance NBoF and TNBoF models as

TABLE 5. Audio analysis results under noisy data setting. No preprocessing convolution layer was used in this setting.

| Models | test accuracy |
|----------------------------------|--------------------|
| Noisy FMA Dataset | |
| GRU [5] | 31.04±01.43 |
| NBoF [26] | 31.54±00.21 |
| NBoF-IA (our) | 36.21±01.93 |
| TNBoF [27] | 30.71±00.87 |
| TNBoF-IA (our) | 36.67±00.95 |
| Noisy TUT-UAS2018 Dataset | |
| GRU [5] | 56.17±01.31 |
| NBoF [26] | 41.73±12.66 |
| NBoF-IA (our) | 56.79±00.86 |
| TNBoF [27] | 51.48±00.85 |
| TNBoF-IA (our) | 56.04±00.42 |

TABLE 6. Performance (averaged F1) on AF dataset.

| Models | F1 |
|----------------|--------------------|
| GRU [5] | 76.42±00.86 |
| NBoF [26] | 78.15±00.82 |
| NBoF-CA (our) | 78.73±00.71 |
| NBoF-TA (our) | 78.55±00.90 |
| TNBoF [27] | 78.27±01.02 |
| TNBoF-CA (our) | 78.71±00.92 |
| TNBoF-TA (our) | 79.52±00.81 |

demonstrated in [26], the additional degrees of freedom introduced by this step negates the competition effects enforced by the attention mechanism, leading to performance degradation when combining both methods.

In order to evaluate how well the proposed input attention mechanism (NBoF-IA, TNBoF-IA) tackles noisy data, we simulated contaminated audio data by adding 10 synthetic frequency bands, which are generated by adding white noise to the averaged Mel coefficients. Here we should note that in this set of experiments, we did not use any convolution layers in order to gauge how well the layers of interests are resilient to noise. The results are shown in Table 5. As can be seen from Table 5, when moving from the noiseless to the noisy version of FMA and TUT-UAS2018 datasets, the accuracy of NBoF and TNBoF models dropped significantly. GRU models also exhibited similar behaviors, although the performance drops are less significant as compared to the NBoF and TNBoF. By incorporating the proposed input attention block to the NBoF and TNBoF models, we were able to achieve very similar performances compared to the noiseless scenario.

C. MEDICAL DIAGNOSIS EXPERIMENTS

Medical diagnosis, which plays a crucial role in ensuring human prosperity, is inherently an intricate process. The quality of the diagnosis is highly dependent on the expertise of

the examiner. Since it takes several years and a great amount of resources to train human experts, medical diagnosis tools have been actively developed over the past decades to assist human examiners. In our empirical analysis using medical data, we investigated the effectiveness of the proposed models in diagnosing cardiovascular diseases using publicly available Electrocardiogram (ECG) and Phonocardiogram (PCG) signals.

The AF dataset focuses on the problem of atrial fibrillation detection from ECG recordings, which are provided as the development data (training set) in the Physionet/Computing in Cardiology Challenge 2017 [32]. The dataset contains 8528 single-lead ECG recordings lasting from 9 to 60 seconds. The objective of the challenge was to classify a given recording into one of the 4 classes: normal sinus rhythm, atrial fibrillation, alternative rhythm, and noise. We followed an experimental setup similar to [33], which evaluates a given model using 5-fold cross-validation. Additionally, the recordings were clipped or padded so that they have a constant length of 30 seconds. Since a single lead ECG recording is only a univariate sequence, it is necessary to use convolution layers as preprocessing layers to extract higher-level features, before the NBoF or GRU layers. To tackle the imbalanced nature of the training set, we scaled the loss term associated with each class, with the factor inversely proportional to the number of samples in that class.

PCG signal is often used in ambulatory diagnosis in order to evaluate the heart hemodynamic status and detect potential cardiovascular problems. The data used in our experiments come from the training set provided in the Physionet/Computing in Cardiology Challenge 2016 [34]. The objective of the challenge is to develop an automatic classification method for the anomaly (normal versus abnormal) and quality (good versus bad) detection given a PCG recording.

Since the length of the recordings varies greatly, from 5 to 120 seconds, we generated 5s segments from the recordings for training the models; during the test phase, the models were used to classify 5s sub-segments (with 4s overlap) of a given recording, and the overall label is inferred from the averaged classification of the sub-segments. PCG signal captures the acoustic nature of the heart sound; thus, we extracted Mel-spectrogram with 24 frequency bands, using a window of 25ms with an overlap of 10ms to represent each segment. With a smaller size compared to the AF dataset, we only employed a 3-fold cross-validation protocol for this problem. Further details regarding our experimental setup in AF and PCG datasets are provided in the Appendix.

Table 6 shows the averaged F1 score, a metric adopted by the database [32], of all models on the AF dataset. In Table 7, we show the anomaly and quality detection performance. The performance metric used by the database [34] is calculated as the mean of sensitivity and specificity scores. In the AF dataset, the averaged F1 scores obtained from the baseline NBoF and TNBoF models are significantly higher

TABLE 7. Performance (mean of sensitivity and specificity) on PCG dataset. The higher, the better.

| Models | Anomaly Detection | Quality Detection |
|----------------|-------------------|-------------------|
| GRU [5] | 90.08±00.68 | 72.74±01.40 |
| NBoF [26] | 50.31±00.42 | 49.69±00.34 |
| NBoF-CA (our) | 88.09±00.25 | 71.98±03.00 |
| NBoF-TA (our) | 89.32±01.02 | 72.57±02.20 |
| TNBoF [27] | 54.12±05.18 | 53.40±04.21 |
| TNBoF-CA (our) | 88.68±00.95 | 72.34±01.32 |
| TNBoF-TA (our) | 88.81±01.07 | 69.45±00.89 |

than the one obtained from the recurrent model. Although the improvement margins for the NBoF model are minor in Table 6, both the NBoF and TNBoF models enjoy increases in performance when using the proposed attention blocks. The consistent performance gain produced by the attention blocks can also be observed in the PCG dataset in Table 7. In this dataset, while the NBoF and TNBoF models score far below the GRU model, the attention-based models perform nearly as well as the recurrent model.

V. CONCLUSION

In this paper, we proposed 2D-Attention, a generic attention mechanism for data represented in the form of matrices. The proposed attention computation can be used in a plug-and-play manner, and can be updated jointly with other components in a computation graph. Using the proposed attention block, we further proposed three variants of the Neural Bag-of-Features model when learning with sequence data. Our extensive experiments in financial forecasting, audio analysis and medical diagnosis demonstrated that the proposed attention consistently led to performance gains for the Neural Bag-of-Features models. Since 2D-Attention is a generic attention computation method for matrices, investigating its efficacy in other neural network models is an interesting research direction in the future works.

APPENDIX

In all of our experiments, we used ADAM optimizer for stochastic optimization. Weight decay (0.0001) or max-norm constraint (4.0) was used to for regularization. In addition, dropout (0.2) was applied to the output of the layer before the classification layer. In all models, before the output layer, there is a fully-connected layer with 512 neurons. For NBoF, TNBoF and the attention models, we used 256 codewords in the quantization layer. Correspondingly, the number of units in GRU model was set to 256. Details that are specific to each experiment are provided below:

- *Financial Forecasting Experiments:* All models were trained for 80 epochs, with the initial learning rate set to 0.001. The learning rate was decreased by a factor of 0.1 at epoch 11 and 51. We followed [10] and scaled the loss term associated with each class with a factor that is inversely proportional to the number of samples

of each class to counter the effect of class imbalanced. In experiments that used convolution layers as preprocessing layers, we used two 1D convolution layers, each of which has 64 filters with the filter size set to 5 and the stride set to 1. Batch normalization was used after each convolution layer, followed by the ReLU activation.

- *Audio Analysis Experiments:* The setup is similar to the financial forecasting experiments, except for the configuration of convolution layers: four 1D convolution layers with the filter size of 5 were used; the first two convolution layers have 32 filters, which are followed by a max-pooling layer to reduce the temporal dimension by half. The last two convolution layers have 64 filters. After each convolution layer, we applied batch normalization, followed by ReLU activation.
- *Medical Diagnosis Experiments:* in both AF and PCG datasets, all models were trained for 90 epochs, with the initial learning rate set to 0.001, which was decreased to 0.0001 at epoch 11, then to 0.00001 at epoch 71. For the AF dataset, we adopted the convolution architecture proposed in [33] as the first computation block in all models. For PCG dataset, we used five 1D convolution layers with the filter size set to 3 as the preprocessing layers: the first two layers have 32 filters with strides of 1; the third layer has 64 filters with strides of 2; the fourth layer has 64 filters with strides of 1; the last layer has 128 filters with strides of 2. After each convolution layer, we applied batch normalization, followed by ReLU activation.

ACKNOWLEDGMENT

This publication reflects the authors' views only. The European Commission is not responsible for any use that may be made of the information it contains.

REFERENCES

- [1] D. T. Tran, M. Magris, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Tensor representation in high-frequency financial data for price change prediction," in *Proc. IEEE Symp. Ser. Comput. Intell. (SSCI)*, Nov. 2017, pp. 1–7.
- [2] A. Ntakaris, M. Magris, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Benchmark dataset for mid-price forecasting of limit order book data with machine learning methods," *J. Forecasting*, vol. 37, no. 8, pp. 852–866, Dec. 2018.
- [3] A. Graves, A.-R. Mohamed, and G. Hinton, "Speech recognition with deep recurrent neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, May 2013, pp. 6645–6649.
- [4] D. Stowell, D. Giannoulis, E. Benetos, M. Lagrange, and M. D. Plumbley, "Detection and classification of acoustic scenes and events," *IEEE Trans. Multimedia*, vol. 17, no. 10, pp. 1733–1746, Oct. 2015.
- [5] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," 2014, *arXiv:1406.1078*.
- [6] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," 2014, *arXiv:1409.0473*.
- [7] E. Slutsky, "The summation of random causes as the source of cyclic processes," *Econometrica, J. Econ. Soc.*, vol. 5, pp. 105–146, Apr. 1937.
- [8] G. C. Tiao and G. E. Box, "Modeling multiple time series with applications," *J. Amer. Statist. Assoc.*, vol. 76, no. 376, pp. 802–816, 1981.

- [9] D. T. Tran, M. Gabbouj, and A. Iosifidis, "Multilinear class-specific discriminant analysis," *Pattern Recognit. Lett.*, vol. 100, pp. 131–136, Dec. 2017.
- [10] D. T. Tran, A. Iosifidis, J. Kannianen, and M. Gabbouj, "Temporal attention-augmented bilinear network for financial time-series data analysis," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 30, no. 5, pp. 1407–1418, May 2018.
- [11] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [12] N. Passalis and A. Tefas, "Neural bag-of-features learning," *Pattern Recognit.*, vol. 64, pp. 277–294, Apr. 2017.
- [13] B. Zhao, H. Lu, S. Chen, J. Liu, and D. Wu, "Convolutional neural networks for time series classification," *J. Syst. Eng. Electron.*, vol. 28, no. 1, pp. 162–169, Feb. 2017.
- [14] N. Passalis, A. Tsantekidis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Time-series classification using neural bag-of-features," in *Proc. 25th Eur. Signal Process. Conf. (EUSIPCO)*, Aug. 2017, pp. 301–305.
- [15] S. Lazebnik, C. Schmid, and J. Ponce, "Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. (CVPR)*, vol. 2, Jun. 2006, pp. 2169–2178.
- [16] Y.-G. Jiang, C.-W. Ngo, and J. Yang, "Towards optimal bag-of-features for object categorization and semantic video retrieval," in *Proc. 6th ACM Int. Conf. Image Video Retr. (CIVR)*, 2007, pp. 494–501.
- [17] M. Riley, E. Heinen, and J. Ghosh, "A text retrieval approach to content-based audio retrieval," in *Proc. Int. Symp. Music Inf. Retr. (ISMIR)*, 2008, pp. 295–300.
- [18] A. Iosifidis, A. Tefas, and I. Pitas, "Multidimensional sequence classification based on fuzzy distances and discriminant analysis," *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 11, pp. 2564–2575, Nov. 2012.
- [19] A. Iosifidis, A. Tefas, and I. Pitas, "Discriminant bag of words based representation for human action recognition," *Pattern Recognit. Lett.*, vol. 49, pp. 185–192, Nov. 2014.
- [20] D. G. Lowe, "Object recognition from local scale-invariant features," in *Proc. IEEE Int. Conf. Comput. Vis.*, vol. 2, Sep. 1999, pp. 1150–1157.
- [21] N. Passalis and A. Tefas, "Entropy optimized feature-based bag-of-words representation for information retrieval," *IEEE Trans. Knowl. Data Eng.*, vol. 28, no. 7, pp. 1664–1677, Jul. 2016.
- [22] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, attend and tell: Neural image caption generation with visual attention," in *Proc. Int. Conf. Mach. Learn.*, 2015, pp. 2048–2057.
- [23] V. Mnih, N. Heess, and A. Graves, "Recurrent models of visual attention," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 2204–2212.
- [24] F. Laakom, N. Passalis, J. Raitoharju, J. Nikkanen, A. Tefas, A. Iosifidis, and M. Gabbouj, "Bag of color features for color constancy," 2019, *arXiv:1906.04445*.
- [25] Y. Qin, D. Song, H. Chen, W. Cheng, G. Jiang, and G. Cottrell, "A dual-stage attention-based recurrent neural network for time series prediction," 2017, *arXiv:1704.02971*.
- [26] N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Temporal logistic neural bag-of-features for financial time series forecasting leveraging limit order book data," 2019, *arXiv:1901.08280*.
- [27] N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Temporal bag-of-features learning for predicting mid price movements using high frequency limit order book data," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 4, no. 6, pp. 774–785, Dec. 2020.
- [28] Y. G. Cinar, H. Mirisae, P. Goswami, E. Gaussier, A. Ait-Bachir, and V. Strijov, "Position-based content attention for time series forecasting with sequence-to-sequence RNNs," in *Proc. Int. Conf. Neural Inf. Process. Cham, Switzerland: Springer*, 2017, pp. 533–544.
- [29] O. B. Sezer, M. U. Gudelek, and A. M. Ozbayoglu, "Financial time series forecasting with deep learning : A systematic literature review: 2005–2019," *Appl. Soft Comput.*, vol. 90, May 2020, Art. no. 106181.
- [30] M. Defferrard, K. Benzi, P. Vandergheynst, and X. Bresson, "FMA: A dataset for music analysis," 2016, *arXiv:1612.01840*.
- [31] A. Mesaros, T. Heittola, and T. Virtanen, "TUT acoustic scenes 2017, evaluation dataset," Zenodo, Tech. Rep., Nov. 2017. [Online]. Available: <https://doi.org/10.5281/zenodo.1040168>, doi: 10.5281/zenodo.1040168.
- [32] G. Clifford, C. Liu, B. Moody, L.-W. Lehman, I. Silva, Q. Li, A. Johnson, and R. Mark, "AF classification from a short single lead ECG recording: The physionet computing in cardiology challenge 2017," in *Proc. Comput. Cardiol. Conf. (CinC)*, Sep. 2017, pp. 1–4.
- [33] F. Andreotti, O. Carr, M. A. F. Pimentel, A. Mahdi, and M. De Vos, "Comparing feature based classifiers and convolutional neural networks to detect arrhythmia from short segments of ECG," in *Proc. Comput. Cardiol. Conf. (CinC)*, Sep. 2017, pp. 1–4.
- [34] G. Clifford, C. Liu, D. Springer, B. Moody, Q. Li, R. Abad, J. Millet, I. Silva, A. Johnson, and R. Mark, "Classification of normal/abnormal heart sound recordings: The PhysioNet/computing in cardiology challenge 2016," in *Proc. Comput. Cardiol. Conf. (CinC)*, Sep. 2016, pp. 609–612.



DAT THANH TRAN received the M.Sc. degree (Hons.) in data engineering and machine learning and the Ph.D. degree (Hons.) in machine learning from Tampere University, Finland, in 2019 and 2022, respectively. He has published several scientific articles in a wide range of topics, ranging from computer vision, bio-signal, and time-series analysis to financial forecasting. His current research interests include high-frequency financial data analysis, model compression and acceleration, and edge AI.



NIKOLAOS PASSALIS (Graduate Student Member, IEEE) received the B.Sc. degree in informatics, the M.Sc. degree in information systems, and the Ph.D. degree in informatics from the Aristotle University of Thessaloniki, Thessaloniki, Greece, in 2013, 2015, and 2018, respectively. Since 2019, he has been a Postdoctoral Researcher with the Aristotle University of Thessaloniki, while from 2018 to 2019, he also conducted postdoctoral research with the Faculty of Information Sciences, Tampere University, Finland. He has coauthored more than 45 journal articles and 60 conference papers. Over 1600 citations have been recorded to his publications and his H-index is 21 according to Google scholar. His research interests include deep learning, information retrieval, time-series analysis, and computational intelligence.



ANASTASIOS TEFAS (Member, IEEE) received the B.Sc. degree in informatics and the Ph.D. degree in informatics from the Aristotle University of Thessaloniki, Greece, in 1997 and 2002, respectively. Since 2017, he has been an Associate Professor with the Department of Informatics, Aristotle University of Thessaloniki, where he was a Lecturer and an Assistant Professor, from 2008 to 2017. He participated in 20 research projects financed by national and European funds. He is the Coordinator of the H2020 Project OpenDR, "Open Deep Learning Toolkit for Robotics." He is an Area Editor of *Signal Processing: Image Communications* journal. He has coauthored 137 journal articles, 255 papers in international conferences, and contributed 17 chapters to edited books in his area of expertise. Over 8000 citations have been recorded to his publications and his H-index is 45 according to Google scholar. His current research interests include computational intelligence, deep learning, pattern recognition, machine learning, digital signal and image analysis and retrieval, computer vision, and robotics.



MONCEF GABBOUJ (Fellow, IEEE) received the B.S. degree in electrical engineering from Oklahoma State University, Stillwater, in 1985, and the M.S. and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, USA, in 1986 and 1989, respectively.

He is currently a Professor in information technology with the Department of Computing Sciences, Tampere University, Tampere, Finland. He was an Academy of Finland Professor, from 2011 to 2015. He held several visiting professorships at different universities. He is also the Finland Site Director of the NSF IUCRC funded the Center for Visual and Decision Informatics. He has published two books and over 850 journals and conference papers and supervised 54 doctoral and 76 master's theses. His research interests include big data analytics, multimedia content-based analysis, indexing and retrieval, artificial intelligence, machine learning, pattern recognition, nonlinear signal and image processing and analysis, voice conversion, and video processing and coding.

Dr. Gabbouj is a member of the Academia Europaea and the Finnish Academy of Science and Letters. He is the Past Chairman of the IEEE CAS TC on DSP and a Committee Member of the IEEE Fourier Award for Signal Processing. He served as a Distinguished Lecturer for the IEEE CASS. He served as an associate editor and a guest editor for many IEEE, and international journals and the General Chair for IEEE SPS Conference and CAS Flagship Conference, ICIP and ISCAS, and ICME 2021. He was a recipient of the 2017 Finnish Cultural Foundation for Art and Science Award, the 2015 TUT Foundation Grand Award, the 2012 Nokia Foundation Visiting Professor Award, the 2005 Nokia Foundation Recognition Award, and several best paper awards. He is the Finland Site Director of the NSF IUCRC funded the Center for Visual and Decision Informatics (CVDI) and led the Artificial Intelligence Research Task Force of the Ministry of Economic Affairs and Employment funded Research Alliance on Autonomous Systems (RAAS).



ALEXANDROS IOSIFIDIS is currently a Professor with Aarhus University, Denmark. He has contributed to more than 30 research and development projects financed by EU, Finnish, and Danish funding agencies and companies. He has coauthored 92 articles in international journals and 120 papers in international conferences and workshops proposing novel machine learning techniques and their application in a variety of problems. He is a member of the IEEE Technical

Committee on Machine Learning for Signal Processing. He is the Associate Editor-in-Chief of *Neurocomputing* journal covering the research area of neural networks and an Associate Editor of IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS.

...