



HAL
open science

Deep Learning for Time Series Classification and Extrinsic Regression: A Current Survey

Navid Mohammadi Foumani, Lynn Miller, Chang Wei Tan, Geoffrey I Webb,
Germain Forestier, Mahsa Salehi

► **To cite this version:**

Navid Mohammadi Foumani, Lynn Miller, Chang Wei Tan, Geoffrey I Webb, Germain Forestier, et al.. Deep Learning for Time Series Classification and Extrinsic Regression: A Current Survey. *ACM Computing Surveys*, 2024, 56 (9), pp.1 - 45. <10.1145/3649448>. <hal-05093756>

HAL Id: hal-05093756

<https://hal.science/hal-05093756v1>

Submitted on 2 Jun 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License



Deep Learning for Time Series Classification and Extrinsic Regression: A Current Survey

NAVID MOHAMMADI FOUMANI, Monash University, Melbourne, Australia

LYNN MILLER, Monash University, Melbourne, Australia

CHANG WEI TAN, Monash University, Melbourne, Australia

GEOFFREY I. WEBB, Monash University, Melbourne, Australia

GERMAIN FORESTIER, Monash University, Melbourne, Australia and IRIMAS, University of Haute-Alsace, Mulhouse, France

MAHSA SALEHI, Monash University, Melbourne, Australia

Time Series Classification and Extrinsic Regression are important and challenging machine learning tasks. Deep learning has revolutionized natural language processing and computer vision and holds great promise in other fields such as time series analysis where the relevant features must often be abstracted from the raw data but are not known a priori. This article surveys the current state of the art in the fast-moving field of deep learning for time series classification and extrinsic regression. We review different network architectures and training methods used for these tasks and discuss the challenges and opportunities when applying deep learning to time series data. We also summarize two critical applications of time series classification and extrinsic regression, human activity recognition and satellite earth observation.

CCS Concepts: • **Computing methodologies** → **Machine learning approaches; Supervised learning;**

Additional Key Words and Phrases: Deep learning, time series, classification, extrinsic regression, review

ACM Reference Format:

Navid Mohammadi Foumani, Lynn Miller, Chang Wei Tan, Geoffrey I. Webb, Germain Forestier, and Mahsa Salehi. 2024. Deep Learning for Time Series Classification and Extrinsic Regression: A Current Survey. *ACM Comput. Surv.* 56, 9, Article 217 (April 2024), 45 pages. <https://doi.org/10.1145/3649448>

1 INTRODUCTION

Time series analysis has been identified as one of the 10 most challenging research issues in the field of data mining in the 21st century [1]. **Time series classification (TSC)** is a key time series analysis task [2]. TSC builds a machine learning model to predict categorical class labels for data consisting of ordered sets of real-valued attributes. The many applications of time series analysis

This work was supported by an Australian Government Research Training Program (RTP) scholarship.

Authors' addresses: N. Mohammadi Foumani, L. Miller, C. W. Tan, G. I. Webb, and M. Salehi, Monash University, Melbourne, Australia; e-mails: navid.foumani@monash.edu, lynn.miller1@monash.edu, chang.tan@monash.edu, geoff.webb@monash.edu, mahsa.salehi@monash.edu; G. Forestier, Monash University, Melbourne, Australia and IRIMAS, University of Haute-Alsace, Mulhouse, France; e-mail: germain.forestier@uha.fr.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2024 Copyright held by the owner/author(s).

ACM 0360-0300/2024/04-ART217

<https://doi.org/10.1145/3649448>

include human activity recognition [3–5], diagnosis based on electronic health records [6, 7], and systems monitoring problems [8]. The wide variety of dataset types in the **University of California, Riverside (UCR)** [9] and **University of East Anglia (UEA)** [8] benchmark archive further illustrates the breadth of TSC applications. **Time series extrinsic regression (TSER)** [10] is the counterpart of TSC for which the output is numeric rather than categorical. It should be noted that the TSER is not a forecasting method but rather a method for understanding the relationship between the time series and the extrinsic variable. TSER is an emerging field with great potential to be used in a wide range of applications.

Deep learning has been very successful, especially in computer vision and natural language processing. Many modern applications integrate deep learning. Deep learning can autonomously learn informative features from raw data, eliminating the need for manual feature engineering. Consequently, there has been much interest in developing deep TSC and TSER due to their ability to learn relevant latent feature representations. It is worth noting that the majority of TSC and TSER research has focused on non-deep-learning approaches. A recent benchmark [11] shows that the deep learning method (InceptionTime [12]) is competitive but did not outperform the state of the art on benchmarking archives. One reason is that the popular UCR and UEA benchmarking archives were not designed for deep learning models. In particular, they are relatively small, while deep learning often excels when data quantities are large. Deep learning can also benefit from heightened compatibility with current hardware, particularly GPUs, leading to fast and efficient execution. Their exceptional scalability further allows seamless handling of growing data volumes and computational complexity, reinforcing their versatility in processing large datasets. Indeed, ConvTran [13], a recent deep architecture for TSC, outperforms one of the fastest conventional models, ROCKET [14], in terms of both speed and accuracy when there are more than 10k training samples.

A highly influential review paper on deep-learning-based TSC [15] was published in 2019. However, the field of research is very fast moving, and that prior survey does not cover the current state of the art. For example, it does not include InceptionTime [12], a system that consistently outperforms ResNet [16], the best performing system from the prior survey. Nor does it cover attention models, which have received huge interest in recent years and have shown excellent capacity to model long-range dependencies in sequential data and are well suited for time series modeling [17]. Many attention variants have been proposed to address particular challenges in time series modeling and have been successfully applied to TSC [13, 18, 19]. Moreover, the previous survey does not include self-supervised learning, which is emerging as a new paradigm [20]. Self-supervised learning induces supervision by designing pretext tasks instead of relying on pre-defined prior knowledge and has shown very promising results, especially in datasets with a low label regime [21–24].

In light of the emergence of attention mechanisms, self-supervised learning, and various new network configurations for TSC, a systematic and comprehensive survey on deep learning in TSC would greatly benefit the time series community. This article aims to fill that gap by summarizing recent developments in deep-learning-based time series analytics, specifically TSC and TSER. Following definitions and a brief introduction to the time series classification and extrinsic regression tasks, we propose a new taxonomy based on various methodological perspectives. Diverse architectures, including **multilayer perceptrons (MLPs)**, **convolutional neural networks (CNNs)**, **recurrent neural networks (RNNs)**, **Graph Neural Networks (GNNs)**, and attention-based models, are discussed, along with refinements made to improve performance. Additionally, various types of self-supervised learning pretexts, such as contrastive learning and self-prediction, are explored. We also conduct a review of useful data augmentation and transfer learning strategies

for time series data. Furthermore, we provide a summary of two key applications of TSC and TSER, namely Human Activity Recognition and Earth Observation.

2 BACKGROUND AND DEFINITIONS

This section begins by providing the necessary definitions and background information to understand the topic of training **deep neural networks (DNNs)** for TSC and TSER tasks. We begin by defining key terms and concepts, such as time series data and time series supervised learning. Finally, we present our proposed taxonomy of the different deep learning methods that have been used for TSC and TSER tasks.

2.1 Time Series

Time series data are sequences of data points indexed by time.

Definition 2.1. A time series X is an ordered collection of T pairs of measurements and timestamps,

$X = \{(x_1, t_1), (x_2, t_2), \dots, (x_T, t_T)\}$, where $x_i \in \mathbb{R}^D$ and t_1 to t_T are the timestamps for some measurements x_1 to x_T .

Each x_i is a D -dimensional vector of values, one for each feature captured in the series. When $D = 1$, the series is called *univariate*. When $D > 1$, the series is called *multivariate*.

2.2 Time Series Supervised Learning Tasks

This article focuses on two time series learning tasks: time series extrinsic regression and time series classification. Classification and regression are both supervised learning tasks that learn the relationship between a target variable and a set of time series. We consider learning from a dataset $D = \{(X_1, Y_1), (X_2, Y_2), \dots, (X_N, Y_N)\}$ of N time series where Y_i denotes the target variable for each X_i . It is important to note that for ease of exposition, we assume in our discussion that the series are of the same length, but most methods extend trivially to the case of unequal-length series. The main difference between TSER and TSC is that TSC predicts a categorical value for a time series from a set of finite categories, while TSER predicts a continuous value for a variable external to the input time series. Typically Y_i is a one-hot encoded vector for TSC or a numeric value for TSER.

In the context of deep learning, a supervised learning model is a neural network that executes the following functions to map the input time series to a target variable:

$$f_L(\theta_L, X) = f_{L-1}(\theta_{L-1}, f_{L-2}(\theta_{L-2}, \dots, f_1(\theta_1, X))), \quad (1)$$

where f_i represents the non-linear function and θ_i denotes the parameters at layer i . For TSC the neural network model is trained to map a time series dataset D to a set of class labels Y with C class labels. After training, the neural network outputs a vector of C values that estimates the probability of a series X belonging to each class. This is typically achieved using the softmax activation function in the final layer of the neural network. The softmax function estimates probabilities for all of the dependent classes such that they always sum to 1 across all classes. The cross-entropy loss is commonly used for training neural networks with softmax outputs or classification-type neural networks.

On the other hand, TSER trains the neural network model to map a time series dataset D to a set of numeric values Y . Instead of outputting probabilities, a regression neural network outputs a numerical value for the time series. It is typically used with a linear activation function in the final layer of the neural network. However, any non-linear functions with a single-value output such as sigmoid or ReLU can also be used. A regression neural network typically trains using the

mean square error or mean absolute error loss function. However, depending on the distribution of the target variable and the choice of final activation functions, other loss functions can be used.

2.3 TSC and TSER

TSC is a fast-growing field, with hundreds of papers being published every year [8, 9, 15, 25, 26]. The majority of works in TSC are non-deep learning based. In this survey, we focus on deep learning approaches and refer interested readers to Appendix A and benchmark papers [11, 25, 26] for more details on non-deep learning approaches. Most deep learning approaches to TSC have real-valued outputs that are mapped to a class label. TSER [10, 27] is a less widely studied task in which the predicted values are numeric, rather than categorical. While the majority of the architectures covered in this survey were designed for TSC, it is important to note that it is trivial to adapt most of them for TSER.

Deep-learning-based TSC methods can be classified into two main types: generative and discriminative [28]. In the TSC community, generative methods are often considered model based [25], aiming to understand and model the joint probability distribution of input series X and output labels Y , denoted as $p(X, Y)$. On the other hand, discriminative models focus on modeling the conditional probability of output labels Y given input series X , expressed as $p(Y|X)$.

Generative models, such as the **Stacked Denoising Auto-encoders (SDAEs)**, have been proposed by Bengio et al. [29] to identify the salient structure of input data distributions, and Hu et al. [30] used the same model for the pre-training phase before training a classifier for time series tasks. A universal neural network encoder has been developed to convert variable-length time series to a fixed-length representation [31]. Also, a **Deep Belief Network (DBN)** combined with a transfer learning method was used in an unsupervised manner to model the latent features of time series [32]. An **Echo State Network (ESN)** has been used to learn the appropriate time series representation by reconstructing the original raw time series prior to training the classifier [33]. **Generative Adversarial Networks (GANs)** are one of the popular generative models that generate new examples by learning to discriminate between real and synthetic examples. Various GANs have been developed for time series and have been reviewed in a recent survey [34]. Often, implementing generative methods is more complex due to an additional step of training. Furthermore, generative methods are typically less efficient than discriminative methods, which directly map raw time series to class probability distributions. Due to these barriers, researchers tend to focus on discriminative methods. Therefore, this survey mainly focuses on the end-to-end discriminative approaches.

2.4 Taxonomy of Deep Learning in TSC and TSER

To provide an organized summary of the existing deep learning models for TSC, we propose a taxonomy that categorizes these models based on deep learning methods and application domains. This taxonomy is illustrated in Figure 1. In Section 3, we review various network architectures used for TSC, including MLPs, CNNs, RNNs, GNNs, and attention-based models. We also discuss refinements made to these models to improve their performance on time series tasks. Additionally, various types of self-supervised learning pretexts, such as contrastive learning and self-prediction, are explored in Section 4. We also conduct a review of useful data augmentation and transfer learning strategies for time series data in Sections 5 and 6. In addition to methods, we summarize key applications of TSC and TSER in Section 7 of this article. These applications include human activity recognition and satellite earth observation, which are important and challenging tasks that can benefit from the use of deep learning models. Overall, our proposed taxonomy and the

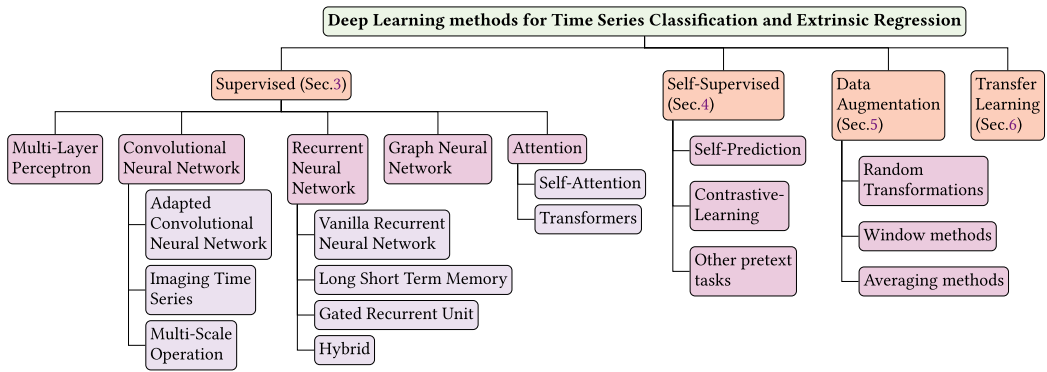


Fig. 1. Taxonomy of deep learning for TSC/TSER from the perspectives of network configuration and application domains.

discussions in these sections provide a comprehensive overview of the current state of the art in deep learning for time series analysis and outline future research directions.

3 SUPERVISED MODELS

This section reviews the deep-learning-based models for TSC and discusses their architectures by highlighting their strengths as well as limitations. More details on deep model architectures and their adaptations to time series data are available in Appendix B.

3.1 Multi-Layer Perceptron (MLP)

The most straightforward neural network architecture is a fully connected network, also called an MLP. The number of layers and neurons are defined as hyperparameters in MLP models. However, studies such as auto-adaptive MLP [35] have attempted to determine the number of neurons in the hidden layers automatically, based on the nature of the training time series data. This allows the network to adapt to the training data’s characteristics and optimize its performance on the task at hand.

One of the main limitations of using MLPs for time series data is that they are not well suited to capturing the temporal dependencies in this type of data. MLPs are feedforward networks that process input data in a fixed and predetermined order without considering the temporal relationships between the input values. Various studies used MLPs alongside other feature extractors like **Dynamic Time Warping (DTW)** to address this problem [36, 37]. DTW-NN is a feedforward neural network that exploits DTW’s elastic matching ability to dynamically align a layer’s inputs to the weights instead of using a fixed and predetermined input-to-weight mapping. This weight alignment replaces the standard dot product within a neuron with DTW. In this way, the DTW-NN is able to tackle difficulties with time series recognition, such as temporal distortions and variable pattern length within a feedforward architecture [37]. Similarly, **Symbolic Aggregate Approximation (SAX)** is used to transform time series into a symbolic representation and produce sequences of words based on the symbolic representation [38]. The symbolic time-series-based words are later used as input for training a two-layer MLP for classification.

Although the models mentioned above attempt to resolve the shortage of capturing temporal dependencies in MLP models, they still have other limitations on capturing time-invariant features [16]. Additionally, MLP models do not have the ability to process input data in a hierarchical or multi-scale manner. Time series data often exhibits patterns and structures at different scales, such as long-term trends and short-term fluctuations. MLP models fail to capture these patterns, as

Table 1. Summary of CNN Models for Time Series Classification and Extrinsic Regression

Model	Year	Baseline Architecture	Other Features
Adapted			
MC-DCNN [43]	2014	2-Stage Conv	Independent convolutions per channel
MC-CNN [44]	2015	3-Stage Conv	1D convolutions on all channels
Zhao et al. [45]	2015	2-Stage Conv	Similar architecture to MC-CNN
FCN [16]	2017	FCN	Using GAP instead of FC layer
ResNet [16]	2017	ResNet 9	Using 3-residual block
Res-CNN [49]	2019	RezNet+FCN	Using 1-residual block + FCN
DCNNs [51]	2019	4-Stage Conv	Using dilated convolutions
Disjoint-CNN [16]	2021	4-Stage Conv	Disjoint temporal and spatial convolution
Series To Image			
Wang and Oates [54]	2015	Tiled CNN	GAF, MT
Hatami et al. [55]	2017	2-Stage Conv	Recurrence plots
Karimi-Bidhendi et al. [56]	2018	Inception V3	GADF
Zhao and Cai [57]	2019	ResNet18, ShuffleNet V2	Bilinear interpolation
RPMCNN [61]	2019	VGGNet, 2-Stage Conv	Relative position matrix
Yang et al. [58]	2019	VGGNet	GASF, GADF, MTF
Multi-Scale Operation			
MCNN [62]	2016	2-Stage Conv	Identity mapping, smoothing, down-sampling
t-LeNet [63]	2016	2-Stage Conv	Squeeze and dilation
MVCNN [64]	2019	4-stage Conv	Inception V1 based
Brunel et al. [65]	2019	Inception V1	
InceptionTime [12]	2019	Inception V4	Ensemble
EEG-inception [66]	2021	InceptionTime	
Inception-FCN [67]	2021	InceptionTime + FCN	
KDCTime [68]	2022	InceptionTime	Knowledge distillation, label smoothing
LITE [69]	2023	InceptionTime	Multiplexing, dilated, and custom filters

they are only able to process input data in a single, fixed-length representation. In addition, MLPs may encounter difficulties when confronted with irregularly sampled time series data, where observations are not uniformly recorded in time. Many other deep learning models are better suited to handle time series data, such as RNNs, CNNs, and transformers, specifically designed to capture the temporal dependencies and patterns in time series data.

3.2 CNN-based Models

Several improvements have been made to CNN since the success of AlexNet in 2012 [39], such as using deeper networks, applying smaller and more efficient convolutional filters, adding pooling layers to reduce the dimensionality of the feature maps, and utilizing batch normalization to improve the stability of training [40]. They have been demonstrated to be very successful in many domains, such as computer vision, speech recognition, and natural language processing problems [40–42]. As a result of the success of CNN architectures in these various domains, researchers have also started adopting them for TSC. See Table 1 for a list of reviewed CNN models in this article.

3.2.1 Adapted CNNs for TSC and TSER. This section presents the first category, which we refer to as Adapted CNNs for TSC and TSER. The papers discussed here are mostly adaptations without any particular preprocessing or mathematical characteristics, such as transforming the series to an image or using multi-scale convolution, and therefore do not fit into one of the other categories.

The first CNN for TSC was the **Multi-Channel Deep Convolutional Neural Network (MC-DCNN)** [43]. It handles multivariate data by independently applying convolutions to each input

channel. Each input dimension undergoes two convolutional stages with ReLU activation, followed by max pooling. The output from each dimension is concatenated and passed to a fully connected layer, which is then fed to a final softmax classifier for classification. Similar to MC-DCNN, a three-layer convolution neural network was proposed for human activity recognition (MC-CNN) [44]. Unlike the MC-DCNN, this model applies 1D convolutions to all input channels simultaneously to capture the temporal and spatial relationships in the early stages. The two-stage version of the MC-CNN architecture was used by Zhao et al. [45] on the earliest version of the UCR Time Series Data Mining Archive. The authors also conducted an ablation study to evaluate the performance of the CNN models with differing numbers of convolution filters and pooling types.

Fully Convolutional Networks (FCN) [46] and **Residual Network (ResNet)** [47] are two deep neural networks that are commonly used for image and video recognition tasks and have been adapted for end-to-end TSC [16]. FCNs are a variant of CNNs designed to operate on inputs of arbitrary size rather than being constrained to fixed-size inputs like traditional CNNs. This is achieved by replacing the fully connected layers in a traditional CNN with a **Global Average Pooling (GAP)** [46]. FCN was adapted for univariate TSC [16], and similar to the original model, it contains three convolution blocks where each block contains a convolution layer followed by batch normalization and ReLU activation. Each block uses 128, 256, and 128 filters with 8, 5, and 3 filter lengths, respectively. The output from the last convolution block is averaged with a GAP layer and passed to a final softmax classifier. The GAP layer has the property of reducing the spatial dimensions of the input while retaining the channel-wise information, which allows it to be used in conjunction with a **class activation map (CAM)** [48] to highlight the regions in the input that are most important for the predicted class. This can provide useful insights into how the network is making its predictions and help identify potential improvement areas. Similar to FCN, the ResNet was also proposed in [16] for univariate TSC. ResNet is a deep architecture containing three residual blocks followed by a GAP layer and a softmax classifier. It uses residual connections between blocks to reduce the vanishing gradient effect that affects deep learning models. The structure of each residual block is similar to the FCN architecture, containing three convolution layers followed by batch normalization and ReLU activation. Each convolution layer uses 64 filters with 8, 5, and 3 filter lengths, respectively. ResNet was found to be one of the most accurate deep learning TSC architectures on 85 univariate TSC datasets [15, 25]. Additionally, integration of ResNet and FCN has been proposed to combine the strength of both networks [49].

In addition to adapting the network architecture, some research has focused on modifying the convolution kernel to suit TSC tasks better. **Dilated convolution neural networks (DCNNs)** [50] are a type of CNN that uses dilated convolutions to increase the receptive field of the network without increasing the number of parameters. Dilated convolutions create gaps between elements of the kernel and perform convolution, thereby covering a larger area of the input. This allows the network to capture long-range dependencies in the data, making it well suited to TSC tasks [51]. Recently, Disjoint-CNN [52] showed that factorization of 1D convolution kernels into disjoint temporal and spatial components yields accuracy improvements with almost no additional computational cost. Applying disjoint temporal convolution and then spatial convolution behaves similarly to the Inverted Bottleneck [53]. Like the Inverted Bottleneck, the temporal convolutions expand the number of input channels, and spatial convolutions later project the expanded hidden state back to the original size to capture the temporal and spatial interaction.

3.2.2 Imaging Time Series. In TSC, a common approach is to convert the time series data into a fixed-length representation, such as a vector or matrix, which can then be input to a deep learning model. However, this can be challenging for time series data that vary in length or have complex temporal dependencies. One solution to this problem is to represent the time series data in an

image-like format, where each time step is treated as a separate channel in the image. This allows the model to learn from the spatial relationships within the data rather than just the temporal relationships. In this context, the term *spatial* refers to the relationships between different variables or features within a single time step of the time series.

As an alternative to using raw time series data as input, Wang and Oates encoded univariate time series data into different types of images that were then processed by a regular CNN [54]. This image-based framework initiated a new branch of deep learning approaches for time series, which consider image transformation as one of the feature engineering techniques. Wang and Oates presented two approaches for transforming a time series into an image. The first generates a **Gramian Angular Field (GAF)**, while the second generates a **Markov Transition Field (MTF)**. GAF represents time series data in a polar coordinate and uses various operations to convert these angles into a symmetry matrix, and MTF encodes the matrix entries using the transition probability of a data point from one time step to another time step [54]. In both cases, the image generation increases the time series size, making the images potentially prohibitively large. Therefore, they propose strategies to reduce their size without losing too much information. Afterward, the two types of images are combined in a two-channel image that is then used to produce better results than those achieved when using each image separately. Finally, a Tiled CNN model is applied to classify the time-series images. In other studies, a variety of transformation methods, including **Recurrence Plots (RPs)** [55], **Gramian Angular Difference Field (GADF)** [56], bilinear interpolation [57], and **Gramian Angular Summation Field (GASF)** [58] have been proposed to transfer time series to input images, expecting that the 2D images could reveal features and patterns not found in the 1D sequence of the original time series.

Hatami et al. [55] propose a representation method based on RP [59] to convert the time series to 2D images with a CNN model for TSC. In their study, time series are regarded as distinct recurrent behaviors such as periodicities and irregular cyclicities, which are the typical phenomena of dynamic systems. The main idea of using the RP method is to reveal at which points some trajectories return to a previous state. Finally, two-stage convolution and two fully connected layers are applied to classify the images generated by RP. Subsequently, pre-trained Inception v3 [60] was used to map the GADF images into a 2,048-dimensional vector space. The final stage used an MLP with three hidden layers, followed by a softmax activation function [56]. Following the same framework, Chen and Shi [61] adopted the **Relative Position Matrix (RPMCNN)** and VGGNet to classify time series data using transform 2D images. Their results showed promising performances by converting univariate time series data to 2D images using relative positions between two timestamps. Following the convention, three image encoding methods, GASF, GADF, and MTF, were used to encode MTS data into 2D images [58]. They showed that the simple structure of ConvNet is sufficient for classification as it performed equally well with the complex structure of VGGNet.

Overall, representing time series data as 2D images can be difficult because preserving the temporal relationships and patterns in the data can be challenging. This transformation can also result in a loss of information, making it difficult for the model to classify the data accurately. Chen and Shi [61] have also shown that the specific transformation methods like GASF, GADF, and MTF used in this process do not significantly improve the prediction outcome.

3.2.3 Multi-scale Operation. The papers discussed here apply a multi-scale convolutional kernel to the input series or apply regular convolutions on the input series at different scales. **Multi-scale CNNs (MCNN)** [62] and **Time LeNet (t-LeNet)** [63] were considered the first models that preprocess the input series to apply convolution on multi-scale series rather than raw series. The designs of both MCNNs and t-LeNet were inspired by computer vision models, which means that they were adapted from models originally developed for image recognition tasks. These models

may not be well suited to TSC tasks and may not perform as well as models specifically designed for this purpose. One potential reason for this is the use of progressive pooling layers in these models, commonly used in computer vision models, to reduce the input data size and make it easier to process. However, these pooling layers may not be as effective when applied to time series data and may limit the performance of the model.

MCNN has a simple architecture and comprises two convolutions and a pooling layer, followed by a fully connected and softmax layer. However, this approach involves heavy data preprocessing. Specifically, before any training, they use a sliding window to extract a time series subsequence, and later, the subsequence will undergo three transformations: (1) identity mapping, (2) down-sampling, and (3) smoothing, which results in the transformation of a univariate input time series into a multivariate one. Finally, the transformed output is fed to the CNN model to train a classifier [62]. t-LeNet uses two data augmentation techniques: **window slicing (WS)** and **window warping (WW)**, to prevent overfitting [63]. The WS method is identical to MCNN's data augmentation. The second data augmentation technique, WW, employs a warping technique that squeezes or dilates the time series. WS is also adopted to ensure that subsequences of the same length are extracted for training the network to deal with multi-length time series. Therefore, a given input time series of length L is first dilated ($\times 2$) and then squeezed ($\times 1/2$) using WW, resulting in three time series of length L , $2L$, $1/2L$ that are fed to WS to extract equal-length subsequences for training. Finally, as both MCNN and t-LeNet predict a class for each extracted subsequence, majority voting is applied to obtain the class prediction for the full time series.

Inception was first proposed by Szegedy et al. [70] for end-to-end image classification. Now the network has evolved to become Inception-v4, where Inception was coupled with residual connections to improve further the performance [71]. Inspired by inception architecture, a **multivariate convolutional neural network (MVCNN)** is designed using multi-scale convolution kernels to find the optimal local construction [64]. MVCNN uses three scales of filters, 2×2 , 3×3 , and 5×5 , to extract features of the interaction between sensors. A 1D Inception model was used for Supernovae classification using the light flux of a region in space as an input MTS for the network [65]. However, the authors limited the conception of their Inception architecture to the first version of this model [70]. The Inception-ResNet [72] architecture includes convolutional layers, followed by Inception modules and residual blocks. The Inception modules are used to learn multiple scales and aspects of the data, allowing the network to capture more complex patterns. The residual blocks are then used to learn the residuals, or differences, between the input and output of the network, improving its performance.

InceptionTime [12] explores much larger filters than any previously proposed network for TSC to reach state-of-the-art performance on the UCR benchmark. InceptionTime is an ensemble of five randomly initialized inception network models, each of which consists of two blocks of inception modules. Each inception module first reduces the dimensionality of a multivariate time series using a bottleneck layer with a length and stride of 1 while maintaining the same length. Then, 1D convolutions of different lengths are applied to the output of the bottleneck layer to extract patterns at different sizes. In parallel, a max pooling layer followed by a bottleneck layer are also applied to the original time series to increase the robustness of the model to small perturbations. The outputs from the convolution and max pooling layers are stacked to form a new multivariate time series, which is then passed to the next layer. Residual connections are used between each inception block to reduce the vanishing gradient effect. The output of the second inception block is passed to a GAP layer before feeding into a softmax classifier.

The strong performance of InceptionTime has inspired a number of extensions. Like InceptionTime, EEG-inception [66] uses several inception layers and residual connections as its backbone. Additionally, noise-addition-based data augmentation of **electroencephalogram (EEG)** signals

is proposed, which increases the average accuracy. InceptionFCN [67] focuses on combining two well-known deep learning techniques, namely the Inception module and the Fully Convolutional Network [67]. In KDCTime [68], **label smoothing (LSTime)** and **knowledge distillation (KDTime)** were introduced for InceptionTime, automatically generated while compressing the inference model. Additionally, **knowledge distillation with calibration (KDC)** in KDCTime offers two calibrating strategies: **KDC by translating (KDCT)** and **KDC by reordering (KDCR)**. LITE [69] addresses InceptionTime's complexity while preserving its TSC performance. Utilizing DepthWise Separable Convolutions, LITE incorporates multiplexing, dilated convolution, and custom filters [73] to enhance efficiency.

3.3 Recurrent Neural Network

RNNs are types of neural networks built with internal memory to work with time series and sequential data. Conceptually similar to **feed-forward neural networks (FFNs)**, RNNs differ in their ability to handle variable-length inputs and produce variable-length outputs.

3.3.1 Vanilla Recurrent Neural Networks (Vanilla RNNs). RNNs for TSC have been proposed in [74]. Using RNNs, the input series have been classified based on their dynamic behavior. They used sequence-to-sequence architecture in which each sub-series of input series is classified in the first step. Then the argmax function is applied to the entire output, and finally, the neuron with the highest rate specifies the classification result. In order to improve the model parallelization and capacity, [75] proposed a two-layer RNN. In the first layer, the input sequence is split into several independent RNNs to improve parallelization, followed by a second layer that utilizes the first layer's output to capture long-term dependencies [75]. Further, RNNs have been used in some hierarchical architectures [76, 77]. Hermans and Schrauwen showed that a deeper version of RNNs could perform hierarchical processing on complex temporal tasks and capture the time series structure more naturally than a shallow version [77]. RNNs are usually trained iteratively using a procedure known as **backpropagation through time (BPTT)**. When unfolded in time, RNNs look like very deep networks with shared parameters. With deeper neural layers in RNN and sharing weights across different RNN cells, the gradients are summed up at each time step to train the model. Thus, gradients undergo continuous matrix multiplication due to the chain rule and either shrink exponentially and have small values called vanishing gradients or blow up to a very large value, referred to as exploding gradients [78]. These problems motivated the development of second-order methods for deep architectures named **long short-term memory (LSTM)** [79] and **Gated Recurrent Unit (GRU)** [80].

3.3.2 Long Short-Term Memory (LSTM). LSTM addresses the common vanishing/exploding gradient issue in vanilla RNNs by integrating memory cells with gate control into their state dynamics [79]. Due to its design nature, LSTM is suited to problems involving sequence data, such as language translation [81], video representation learning [82], and image caption generation [83]. The TSC problem is not an exception and mainly adopts a similar model to the language translation [81]. **Sequence-to-Sequence with Attention (S2SwA)** [84] incorporates two LSTMs, one encoder and one decoder, in a sequence-to-sequence fashion for TSC. In this model, the encoder LSTM accepts input time series of arbitrary lengths and extracts information from the raw data based on which the decoder LSTM constructs fixed-length sequences that can be regarded as automatically extracted features for classification.

3.3.3 Gated Recurrent Unit (GRU). GRU, another widely used variant of RNNs, shares similarities with LSTM in its ability to control information flow and memorize context across multiple time steps [80]. Similar to S2SwA [84], a **sequence auto-encoder (SAE)** based on GRU has been

defined to deal with TSC problem [85]. A fixed-size output is produced by processing the various input lengths using GRU as the encoder and decoder. The model's accuracy was also improved by pre-training the parameters on massive unlabeled data.

3.3.4 Hybrid Models. CNNs and RNNs are often combined for TSC because they have complementary strengths. As mentioned previously, CNNs are well suited for learning from spatial relationships in data, such as the patterns and correlations between the channels of different time steps in a time series. This allows them to learn useful features from the time series data that can help improve the classification performance. RNNs, on the other hand, are well suited for learning from temporal dependencies in data, such as the past values of a time series that can help predict its future values. This allows them to capture the dynamic nature of time series data and make more accurate predictions. Combining the strengths of CNNs and RNNs makes it possible to learn spatial and temporal features from the time series data, improving the model's performance for TSC. Additionally, the two models can be trained together, allowing them to learn from each other and improve the model's overall performance.

Various extensions like MLSTM-FCN [86], TapNet [87], and SMATE [88] were proposed later to deal with time series data. MLSTM-FCN extends the univariate LSTM-FCN model [89] to the multivariate case. Like the LSTM-FCN, the multivariate version comprises LSTM blocks and fully convolutional blocks for extracting features from input series. A squeeze and excite block is also added to the FCN block and can execute a form of self-attention on the output feature maps of previous layers [86]. Two further proposals for multivariate TSC are the **Time series attentional prototype Network (TapNet)** and **Semi-Supervised Spatio-Temporal (SMATE)** [87, 88]. These methods combine and seek to leverage the relative strengths of both traditional distance-based and deep learning approaches.

MLSTM-FCN, TapNet, and SMATE were designed in dual-network architectures. The input is separately fed into the CNN and RNN models, and their output is concentrated before the fully connected layer for the final task. However, one branch cannot fully use the hidden states of the other during feature extraction since the final classification results are generated by concatenating the outputs of the two branches. That motivates different types of architecture that try layer-wise integration of CNN and RNN models. This motivates different architectures, such as GCRNN [90] and CNN-LSTM [91], which aim to integrate CNNs and RNNs in a layer-wise fashion.

While RNNs are commonly used for time series forecasting, only a few studies have applied them to TSC, mainly due to four reasons: (1) RNNs typically struggle with the gradient vanishing and exploding problem due to training on long time series [92]; (2) RNNs are considered difficult to train and parallelize, so researchers are less likely to use them as they are computationally expensive [78]; (3) recurrent architectures are designed mainly to learn from the previous data to make predictions about the future [28]; and (4) RNN models can fail to effectively capture and utilize long-range dependencies in long sequences [84].

3.4 Attention-based Model

Despite the excellent performance of CNN models for capturing local temporal/spatial correlations, these models cannot effectively capture and utilize long-range dependencies. Additionally, they only consider the local order of data points rather than the overall order of all data points. Therefore, many recent studies have embedded RNNs such as LSTMs alongside the CNNs to capture this information [86, 87, 89]. The disadvantage of RNN-based models is that they are computationally expensive, and their capability to capture long-range dependencies is limited [18, 93]. On the other hand, *attention models* can capture long-range dependencies, and their broader receptive fields provide more contextual information, which can improve the models' learning capacity. The

Table 2. Summary of Attention-based Models for Time Series Classification and Extrinsic Regression

Model	Year	Embedding	Attention
MuVAN [99]	2018	Bi-GRU	Self-attention
ChannelAtt [102]	2018	RNN	Self-attention
GeoMAN [103]	2018	LSTM	Self-attention
Multi-Stage-Att [104]	2020	LSTM	Self-attention
CT_CAM [105]	2020	FCN + Bi-GRU	Self-attention
CA-SFCN [18]	2020	FCN	Self-attention
RTFN [106]	2021	CNN + LSTM	Self-attention
LAXCAT [100]	2021	CNN	Self-attention
MACNN [101]	2021	Multi-scale CNN	Squeeze-and-excitation
WHEN [107]	2023	CNN + BiLSTM	Self-attention
SAnD [112]	2018	Linear Embedding	Multi-head
T2 [115]	2021	Gaussian Process Regression + 1D Conv	Multi-head
GTN [116]	2021	Linear Embedding	Multi-head
TRANS_tf [113]	2021	Time-Frequency Features	Multi-head
FMLA [117]	2022	Deformable CNN	Multi-head
AutoTransformer [118]	2022	Multi-scale CNN + NAS	Multi-head
ConvTran [13]	2023	Disjoint-CNN	Multi-head

attention mechanism aims to enhance a network's representation ability by focusing on essential features and suppressing unnecessary ones. Not surprisingly, with the success of *attention models* in natural language processing [93, 94], many previous studies have attempted to bring the power of attention models into various domains such as computer vision [95] and time series analysis [18, 19, 96–98]. Table 2 presents a list of the attention-based models reviewed in this article.

3.4.1 Self-attention. Self-attention has been demonstrated to be effective in various natural language processing tasks due to its ability to capture long-term dependencies in text [93]. Recently, it has also been shown to be effective for TSC tasks [18, 99–101]. As we mentioned, the self-attention module is embedded in the encoder-decoder models to improve the model performance. However, only the encoder and the self-attention module have been used for TSC. Early models of TSC follow the same backbone of natural language processing models and use the recurrent-based models such as RNN [102], GRU [99] and LSTM [103, 104] for encoding the input series. For example, the **Multi-View Attention Network (MuVAN)** applies bidirectional GRUs independently to each input dimension as the encoder and then feeds all the representations into a self-attention block [99].

As a result of the excellent performance of the CNN models, many studies have attempted to encode the time series using CNNs before applying attention [18, 100, 105, 106]. **Cross-Attention Stabilized Fully Convolutional Neural Network (CA-SFCN)** [18] and **Locality-Aware eXplainable Convolutional ATtention network (LAXCAT)** [100] applied the self-attention mechanism to leverage the long-term dependencies for the MTSC task. CA-SFCN combines FCN and two types of self-attention, **temporal attention (TA)** and **variable attention (VA)**, which interact to capture the long-range dependencies and variables' interactions. LAXCAT also used temporal and variable attention to identify informative variables and the time intervals where they have informative patterns for classification. **WaveletDTW Hybrid attEntion Networks (WHENs)** [107] integrate two attention mechanisms, namely wavelet attention and DTW attention, into the BiLSTM to enhance model performance. In wavelet attention, they leverage wavelets to compute attention scores, specifically targeting the analysis of dynamic frequency components in nonstationary time series. Simultaneously, DTW attention employs the DTW distance to calculate attention scores, addressing the challenge of time distortion in multiple time series.

Several self-attention models have been developed to improve network performance [108, 109], including **Squeeze-and-Excitation (SE)** [110], which focuses on channel attention and is often used to classify time series data [86, 101, 111]. The SE block allows the whole network to use global information to selectively focus on the informative feature maps and suppress less important ones [110]. More importantly, the SE block can increase the quality of the shared lower-level representations in the early layers and becomes increasingly specialized when responding to different inputs in later layers. The weight of each feature map is automatically learned at each layer of the network, and the SE block can boost feature discrimination throughout the whole network. **Multi-scale Attention Convolutional Neural Network (MACNN)** [101] applies the different kernel size convolutions to capture different scales of information along the time axis by generating feature maps at differing scales. Then an SE block is used to enhance useful feature maps and suppress less useful ones by automatically learning each feature map's importance.

3.4.2 Transformers. The impressive performance of multi-headed attention has led to numerous attempts to adapt multi-headed attention to the TSC domain. Transformers for classification usually employ a simple encoder structure consisting of attention and feed-forward layers. **Simply Attend and Diagnose (SAnD)** [112] architecture adopted a multi-head attention mechanism similar to a vanilla transformer [93] to classify clinical time series for the first time. The model uses both positional encoding and a dense interpolation embedding technique to incorporate temporal order into representation learning. In another study that classified vibration signals [113], time-frequency features such as Frequency Coefficients and **Short Time Fourier Transformation (STFT)** spectrums are used as input embeddings to the transformers. A multi-head attention-based model was applied to raw optical satellite TSC using Gaussian Process Interpolation [114] embedding and outperformed CNNs and RNNs [115].

Gated Transformer Networks (GTNs) [116] use two-tower multi-headed attention to capture the discriminative information from the input series. Also, they merged the output of two towers using a learnable matrix named gating. To enhance locality awareness of transformers for TSC, **flexible multi-head linear attention (FMLA)** [117] integrates deformable convolutional blocks and online knowledge distillation, as well as a random mask to reduce noise. For each TSC dataset, AutoTransformer searches for the suitable network architecture using the **neural architecture search (NAS)** algorithm before feeding the output to the multi-headed attention blocks. ConvTran [13] currently stands as the state of the art in multivariate TSC. They conducted a review of existing absolute and relative position encoding methods in TSC. Based on the limitations of the current position encodings for time series, they introduced two novel ones named tAPE and eRPE for absolute and relative positions, respectively. Integrating these proposed position encodings into a transformer block and combining them with a convolution layer, they presented a novel deep learning framework for multivariate time series classification—ConvTran.

3.5 Graph Neural Networks

While both CNNs and RNNs perform well on Euclidean data, many time series problems have data that are more naturally represented as graphs [119]. For example, in a network of sensors, the sensors may be irregularly spaced, instead of the sensors forming a regular grid. A graph representation of data collected by this network can model this irregular layout more accurately than can be done using a Euclidean space. However, using standard deep learning algorithms to learn from graph structures is challenging [120]. For example, nodes may have a varying number of neighboring nodes, making it difficult to apply a convolution operation.

GNNs [121] are methods that adapt deep learning techniques to the graph domain. Much of the early research using GNNs for time series analysis concentrated on forecasting tasks [119].

Table 3. Summary of Graph Neural Network Models for Time Series Classification and Extrinsic Regression

Model	Year	GNN Type	Other Components
TGCN [126]	2019	Graph convolutional network	1D-CNN
DGCNN [127]	2020	Graph convolutional network	1x1 CNN
GraphSleepNet [128]	2020	Graph convolutional network	Temporal attention
T-GCN [129]	2021	Graph convolutional network	GRU
MRF-GCN [130]	2021	Graph convolutional network	Fast Fourier transforms (FFT)
Nhu et al. [131]	2021	Graph convolutional network	1D-CNN
DCRNN [132]	2021	Graph convolutional network	GRU
Time2Graph+ [125]	2021	Graph attention	Shapelet transform
RAINDROP [133]	2021	Graph guided network	Temporal attention
STEGON [134]	2021	Graph attention	1D-CNN
Azevedo et al. [135]	2022	Graph network block with pooling	1D-CNN
MTPool [136]	2022	Variational graph pooling	1D-CNN
SimTSC [137]	2022	Graph convolutional network	DTW, ResNet
Tulczyjew et al. [138]	2022	Graph convolutional network	Adaptive pooling
C-DGAM [139]	2023	Graph attention	1D-CNN with attention
Dufourg et al. [140]	2023	Spatio-temporal graph	Simple linear iterative clustering
TISER-GCN [124]	2023	Graph convolutional network	1D-CNN
TodyNet [123]	2023	Dynamic graph neural network	1D-CNN
LB-SimTSC [122]	2023	Graph convolutional network	Lower-bound DTW, ResNet

However, recent works consider GNNs for TSC [122, 123] and TSER [124] tasks. A list of the GNN models reviewed in this article is provided in Table 3. Time2Graph+ [125] transforms each time series into a shapelet graph. Shapelets are extracted from the time series and form the graph nodes. The graph edges are weighted based on transition probabilities between the two shapelets. Once the input graphs have been constructed, a graph attention network is used to create a representation of the time series that is fed into a classifier. SimTSC [137] constructs a pairwise similarity graph where each time series forms a node and edge weights are computed based on the DTW distance measure. Node attributes are generated using a feature vector encoder. GNN operations are used to enhance the node features based on similarities between adjacent time series. These representations are then used for the final classification step, which produces a classification for each node. LB-SimTSC [122] replaces the expensive DTW computation with the LB-Keogh lower-bounding method [141].

Spatiotemporal GNNs model both spatial (or inter-variable) and temporal dependencies using two modules that work in tandem. The spatial module models the dependencies between the time series by applying graph convolutions over a GNN (GCNs [142]). The temporal module models the dependencies within the time series using an RNN [129, 132], 1D-CNN [134, 135], Attention [133, 139], or a combination of these [119]. The features extracted from the graph layers are then fed into the classification or regression layers to make either a single prediction [132, 133, 135, 139] or a prediction for each node [129, 134]. Spatiotemporal GCNs are often used to analyze sensor arrays, where the graph structure models the physical layout of the sensors.

A common example is EEG data, where the location of EEG electrodes is represented as a graph that is used to analyze the EEG signal. Some of these applications are epilepsy detection [131], seizure detection [126, 132], emotion recognition [127], and sleep classification [128]. Besides EEG, GCNs have also been applied to engineering applications such as machine fault diagnosis [130], slope deformation prediction [129], and seismic activity prediction [124]. MTPool [136] uses a spatiotemporal GCN for multivariate time series classification. In this study, each channel in the time series is represented by a node in the graph, and the graph edges model the correlations between the channels. The GCN is combined with temporal convolutions and a hierarchical graph pooling technique. Spatiotemporal GNNs have also been used for object-based image analysis [134] and semantic segmentation [138] of image time series. However, these assume the labels and spatial relationships are static over time. In many cases these may both change. **Spatiotemporal graphs (STGs)**, which include temporal edges as well as spatial edges, can model these dynamic relationships [140]. In STGs, each node represents an object at one timestamp. Spatial edges connect the object to adjacent objects, and temporal edges connect two objects in consecutive images if they have common pixels.

4 SELF-SUPERVISED MODELS

Obtaining labeled data for large time series datasets poses significant costs and challenges. Machine learning models trained on large labeled time series datasets often exhibit superior performance compared to models trained on sparsely labeled datasets, small datasets with limited labels, or those without supervision, leading to suboptimal performance across various time series machine learning tasks [23, 143]. As a result, rather than depending on high-quality annotations for large datasets, researchers and practitioners are increasingly shifting their focus toward self-supervised representation learning for time series.

Self-supervised representation learning, a subfield of machine learning, focuses on learning representations from data without explicit supervision [24]. In contrast to supervised learning, which relies on labeled data, self-supervised learning methods utilize the inherent structure of the data to learn valuable representations in an unsupervised manner. The learned representations can then be used for a variety of downstream tasks including classification, anomaly detection, and forecasting. This survey specifically emphasizes classification as a downstream task. We categorized self-supervised learning approaches for TSC into three groups based on the pretext. Table 4 shows a list of the self-supervised models reviewed in this article.

4.1 Contrastive Learning

Contrastive learning involves model learning to differentiate between positive and negative time series examples. **Time-Contrastive Learning (TCL)** [144], **Scalable Representation Learning (SRL or T-Loss)** [145], and **Temporal Neighborhood Coding (TNC)** [146] apply a subsequence-based sampling and assume that distant segments are negative pairs and neighbor segments are positive pairs. TNC takes advantage of the local smoothness of a signal's generative process to define neighborhoods in time with stationary properties to further improve the sampling quality for the contrastive loss function. TS2Vec [23] uses contrastive learning to obtain robust contextual representations for each timestamp hierarchically. It involves randomly sampling two overlapping subseries from input and encouraging consistency of contextual representations on the common segment. The encoder is optimized using both temporal contrastive loss and instance-wise contrastive loss.

In addition to the subsequence-based methods, other models employ instance-based sampling [21, 143, 147–150], treating each sample individually to generate positive and negative

Table 4. Summary of Self-supervised Models for Time Series Classification and Extrinsic Regression

Model	Year	Encoder Backbones	Other Features
Contrastive Learning			
TCL [144]	2016	MLP	Sequence-based contrast
T-Loss/SRL [145]	2019	Causal CNN	Sequence-based contrast
TNC [146]	2021	Bidirectional RNN	Sequence-based contrast
TS-TCC [21]	2021	CNN + Transformers	Instance/sequence-based contrast
MCL [147]	2021	FCN	Instance-based contrast
TimeCLR [148]	2021	InceptionTime	Instance-based contrast
TS2Vec [23]	2021	Dilated CNN	Sequence-based contrast
BTSF [143]	2022	Causal CNN	Instance-based contrast
TF-C [149]	2022	ResNets	Instance-based contrast
MHCCL [150]	2023	ResNet	Instance-based contrast
Self-Prediction			
BENDR [98]	2021	CNN + Transformers	Sequence masking
Voice2Series [22]	2021	CNN+Transformers	Binary masking
TST [19]	2021	Transformers	Binary masking
TARNet [151]	2022	Transformers	Binary masking
TimeMAE [152]	2023	CNN + Transformers	Sequence masking
CRT [153]	2023	Transformers	Sequence masking
Other Pretext tasks			
PHIT [154]	2023	H-InceptionTime	
Series2Vec [24]	2023	Disjoint CNN	Similarity-based representation learning

samples for contrastive loss. **Time-series Temporal and Contextual Contrasting (TS-TCC)** [21] uses weak and strong augmentations to transform the input series into two views and then uses a temporal contrasting module to learn robust temporal representations. The contrasting contextual module is then built upon the contexts from the temporal contrasting module and aims to maximize similarity among contexts of the same sample while minimizing similarity among contexts of different samples. Similarly, TimeCLR [148] introduces DTW data augmentation to enhance robustness against phase shift and amplitude change phenomena. **Bilinear Temporal-Spectral Fusion (BTSF)** [143] uses simple dropout as the augmentation method and aims to incorporate spectral information into the feature representation. Similarly, **Time-Frequency Consistency (TF-C)** [149] is a self-supervised learning method that leverages the frequency domain to achieve better representation. It proposes that the time-based and frequency-based representations, learned from the same time series sample, should be more similar to each other in the time-frequency space compared to representations of different time series samples.

4.2 Self-prediction

The primary objective of self-prediction-based self-supervised models is to reconstruct the input or representation of input data. Studies have explored using transformer-based self-supervised learning methods for TSC [19, 22, 98, 151–153], following the success of models like BERT [94]. **Bert-inspired Neural Data Representations (BENDR)** [98] uses the transformer structure to model EEG sequences and shows that it can effectively handle massive amounts of EEG data recorded with differing hardware. Another study, **Voice-to-Series with Transformer-based Attention (V2Sa)** [22], utilizes a large-scale pre-trained speech processing model for TSC.

The Transformer-based Framework (TST) [19] and TARNet [151] adapt vanilla transformers to the multivariate time series domain and use a self-prediction-based self-supervised pre-training

approach with masked data. These studies demonstrate the potential of using transformer-based self-supervised learning methods for TSC.

4.3 Other Pretext Tasks

While many pretext tasks in self-supervised learning are typically contrastive or self-predictive, specific tasks are tailored for time series data. In image-based self-supervised learning, synthetic transformations (augmentation) of an image are created, and the model learns to contrast the image and its transforms with other images in the training data, which works well for object interpretation. However, time series analysis fundamentally differs from vision or natural language processing concerning the definition of meaningful self-supervised learning tasks.

Guided by this insight, Foumani et al. [24] introduce Series2Vec, a novel self-supervised representation learning approach. Unlike other contrastive self-supervised methods in time series, which carry the risk of positive sample variants being less similar to the anchor sample than series in the negative set, Series2Vec is trained to predict the similarity between two series in both temporal and spectral domains through a self-supervised task. Series2Vec relies primarily on the consistency of the unsupervised similarity step, rather than the intrinsic quality of the similarity measurement, without the need for hand-crafted data augmentation. **Pre-trained H-InceptionTime (PHIT)** [154] is pre-trained using a novel pretext task designed to identify the originating dataset of each time series sample. The objective is to generate flexible convolution filters that can be applied across diverse datasets. Furthermore, PHIT demonstrates its capability to mitigate overfitting in small datasets.

5 DATA AUGMENTATION

In the field of deep learning, the concept of data augmentation has emerged as an important tool for enhancing performance, particularly in scenarios where the availability of training data is limited [155]. Originally proposed in computer vision, data augmentation involves a variety of transformations to images, such as cropping, rotating, flipping, and applying filters like blurring and sharpening. These transformations serve to introduce a diverse range of scenarios within the training data, thereby aiding in the development of more robust and generalizable models. However, the direct application of these image-based augmentation techniques to time series data often proves to be inadequate or inappropriate. Operations like rotation may disrupt the intrinsic temporal structure of time series data.

The challenge of overfitting is particularly pronounced in the field of deep learning models for TSC. These models are characterized by a high number of trainable parameters, which can lead to a model that performs well on training data but fails to generalize to unseen data. In such cases, data augmentation can be a valuable strategy. It offers an alternative to the costly and sometimes impractical approach of collecting additional real-world data. By generating synthetic samples from existing datasets, we can effectively augment the size and variety of our training data. The following details different investigated methods to produce synthetic time series for data augmentation.

Random Transformations. Several augmentations have been developed for the magnitude domain. Jittering, as explored by Um et al. [156], involves the addition of random noise to the time series. Another method, flipping [157], reverses the time series values. Scaling is a technique where the time series is multiplied by a factor from a Gaussian distribution. Magnitude warping, which shares similarities with scaling, distorts the series along a curve that varies smoothly. For time domain transformations, permutation algorithms play a significant role. For example, the slicing transformation involves removing sub-sequence from the series. There are also various warping methods like Random Warping [158], Time Warping [156], Time Stretching [159], and Time

Perturbation [160], each introducing different forms of distortion to the time series. Finally, in the frequency domain, transformations often utilize the Fourier transform. For example, Gao et al. [161] introduce perturbations to both the magnitude and phase spectrum following a Fourier transform.

Window methods. A primary approach in window methods is to create new time series by combining segments from various series of the same class. This technique effectively enriches the data pool with a variety of samples. Window slicing, as introduced by Cui et al. [162], involves dividing a time series into smaller segments, with each segment retaining the class label of the original series. These segments are then used to train classifiers, offering a detailed view of the data. During classification, each segment is evaluated individually, and a collective decision on the final label is reached through a voting system among the slices. Another technique is window warping, based on the DTW algorithm. This method adjusts segments of a time series along the temporal axis, either stretching or compressing them. This introduces variability in the time dimension of the data. Le Guennec et al. [163] provide examples of the application of both window slicing and window warping, showcasing their effectiveness in enhancing the diversity and representativeness of time series datasets.

Averaging methods. Averaging methods in time series data augmentation combine multiple series to form a new, unified series. This process is more difficult than it might seem, as it requires careful consideration of factors like noise and distortions in both the time and magnitude aspects of the data. In this context, weighted **DTW Barycenter Averaging (wDBA)** introduced by Forestier et al. [164] provides an averaging method by aligning time series in a way that accounts for their temporal dynamics. The practical application of wDBA is illustrated in the study by Ismail Fawaz et al. [165], where it is employed in conjunction with a ResNet classifier, demonstrating its effectiveness. Additionally, the research conducted by Terefe et al. [166] uses an auto-encoder for averaging a set of time series. This method represents a more advanced approach in time series data augmentation, exploiting the auto-encoder's capacity for learning and reconstructing data to generate averaged representations of time series.

Selection of data augmentation methods. The selection of the appropriate data augmentation technique is critical and must be adapted to the specific characteristics of the dataset and the architecture of the neural network being used. Studies like those conducted by Iwana and Uchida [167], Pialla et al. [168], and Gao et al. [169] highlight the complexity of this task. These studies demonstrate that the effectiveness of augmentation techniques can vary significantly across different datasets and neural network architectures. Consequently, a method that proves effective in one scenario may not necessarily yield similar results in another. To this end, practitioners in the field of TSC must engage in a careful and informed process of method selection and tuning. While the array of available data augmentation techniques offers a comprehensive toolkit for tackling the challenges of limited data and overfitting, their successful application depends heavily on a nuanced understanding of both the methods themselves and the specific demands of the task at hand.

6 TRANSFER LEARNING

Transfer learning, initially popularized in the field of computer vision, is increasingly becoming relevant in the domain of TSC. In computer vision, this approach involves using a pre-trained network, typically on large datasets like ImageNet [170], as a starting point rather than initiating with random network weights. This method is also related to the concept of foundation or base models, which are large-scale machine learning models trained on extensive data, often using self-supervised or semi-supervised learning. These models are adaptable to a wide array of tasks,

showcasing their versatility. The principle of transfer learning is also closely associated with domain adaptation, which focuses on applying a model trained on a source data distribution to a different, but related, target data distribution. This approach is crucial in leveraging pre-trained models for various applications, particularly in scenarios where data is scarce or specific to certain domains.

In the context of TSC, insights have been contributed by the work of Ismail Fawaz et al. [171], who conducted a study using the UCR archive. Their extensive experiments demonstrated that transfer learning could lead to positive or negative outcomes, depending on the chosen datasets for transfer. This finding underscores the importance of the relationship between source and target datasets in transfer learning efficacy. Ismail Fawaz et al. [171] also introduced an approach to predict the success of transfer learning in TSC by using DTW to measure similarities between datasets. This metric serves as a guide to select the most appropriate source dataset for a given target dataset, thereby enhancing accuracy in a majority of cases.

Other researchers have also explored transfer learning in TSC. Spiegel's [172] work on using dissimilarity spaces to enrich feature representations in TSC set a precedent for employing unconventional data sources. This approach of enhancing learning with diverse data types finds a parallel in Li et al.'s [173] method, which leverages sensor modality labels from various fields to train a deep network, emphasizing the importance of versatile data in transfer learning. Building on the concept of data diversity, Rotem et al. [174] pushed the boundaries further by generating a synthetic univariate time series dataset for transfer learning. This synthetic dataset, used for regression tasks, underscores the potential of artificial data in overcoming the limitations of real-world datasets. Furthermore, Senanayaka et al. [175] introduced the **similarity-based multi-source transfer learning (SiMuS-TL)** approach. By establishing a "mixed domain" to model similarities among various sources, Senanayaka et al. demonstrated the effectiveness of carefully selected and related data sources in transfer learning. Finally, Kashiparekh et al. [176] with their **ConvTimeNet (CTN)** focused on the adaptability of pre-trained networks across diverse time scales.

While the explored studies collectively advance our understanding of transfer learning in TSC, the field remains open for further investigation. A key challenge lies in determining the most suitable source models for transfer, a task complicated by the relative scarcity of large, curated, and annotated datasets in time series analysis compared to the field of computer vision. This restricts the utility of transfer learning in TSC, as the availability of extensive and diverse datasets is crucial for developing robust and generalizable models. Furthermore, the question of developing filters that are generic enough to be effective across a wide range of applications remains unresolved. This aspect is critical for the success of transfer learning, as the applicability of a pre-trained model to new tasks depends on the universality of its learned features. Additionally, the strategy of whether to freeze certain layers of the network during transfer or to fine-tune the entire network is another area that warrants deeper exploration.

7 APPLICATIONS: RECENT DEVELOPMENTS AND CHALLENGES

TSC and TSER techniques have been used to analyze and model time-dependent data in a wide range of applications. These include human activity recognition, Earth observation, medical diagnosis including EEG [177] and **electrocardiogram (ECG)** [178] monitoring, air quality and pollution prediction [179, 180], structural and machine health monitoring [181, 182], **Industrial Internet of Things (IIOT)** [183], energy consumption and anomaly detection [184], and bioacoustics [185].

Due to the extensive range of applications that use TSC and TSER, it is infeasible to cover them all in detail in a single review. Therefore, in this survey, we focus on just two applications: human activity recognition and satellite Earth observation. (References to recent reviews have been

provided for the other applications mentioned above.) These are two important but quite different domains and were chosen to give the reader an idea of the diverseness of time series use in deep learning. The following sections provide an overview of the use of TSC and TSER, the latest developments, and challenges in these two applications.

7.1 Human Activity Recognition

Human activity recognition (HAR) is the identification or monitoring of human activity through the analysis of data collected by sensors or other instruments [186]. The recent growth of wearable technologies and the Internet of Things has resulted in not only the collection of large volumes of activity data [187] but also easy deployment of applications utilizing this data to improve the safety and quality of human life [5, 186]. HAR is therefore an important field of research with applications including healthcare, fitness monitoring, smart homes [188], and assisted living [189].

Devices used to collect HAR data can be categorized as visual or sensor based [4, 5]. Sensor-based devices can be further categorized as object sensors (e.g., RFIDs embedded into objects), ambient sensors (motion sensors, WiFi or Bluetooth devices in fixed locations), and wearable sensors [4], including smartphones [3]. However, the majority of HAR studies use data from wearable sensors or visual devices [186]. Additionally, HAR from visual device data requires the use of computer vision techniques and is therefore out of scope for this review. Accordingly, this section reviews wearable sensor-based methods of HAR. For reviews of vision-based HAR, refer to Kong and Fu [190] or Zhang et al. [191].

The main sensors used in wearable devices are accelerometers, gyroscopes, and magnetic sensors [192], which each collect three-dimensional spatial data over time. **Inertial measurement units (IMUs)** are wearable devices that combine all three sensors in one unit [193, 194]. Wearable device studies typically collect data from multiple IMUs located on different parts of the body [195, 196]. To create a dataset suitable for HAR modeling, the sensor data is split into (usually equally sized) time windows [197]. The task is then to learn a function that maps the multi-variate sensor data for each time window to a set of activities. Thus, the data forms multi-variate time series suited to TSC.

Given the broad scope of our survey, this section necessarily only provides a brief overview of the studies using deep learning for HAR. However, there are several surveys that provide a more in-depth review of machine learning and deep learning for HAR. Lara and Labrador [197] provide a comprehensive introduction to HAR, including machine learning methods used and the principal issues and challenges. Both Nweke et al. [3] and Wang et al. [4] provide a summary of deep learning methods, highlighting their advantages and limitations. Chen et al. [5] discuss challenges in HAR and the appropriate deep learning methods for addressing each challenge. They also provide a comprehensive list of publicly available HAR datasets. Gu et al. [198] focus on deep learning methods, reviewing preprocessing and evaluation techniques as well as the deep learning models.

The deep learning methods used for HAR include both CNNs and RNNs, as well as hybrid CNN-RNN models. While some of the models include an attention module, we did not find any studies proposing a full attention or transformer model. A summary of the studies reviewed and the type of model built is provided in Table 5. Hammerla et al. [199] compared several deep learning models for HAR, including three LSTM variants, a CNN model, and a DNN model. They found that a bi-directional LSTM performed best on naturalistic datasets where long-term effects are important. However, they found that some applications need to focus on short-term movement patterns and suggested CNNs are more appropriate for these applications. Thus, research across all model types is beneficial for the ongoing development of models for HAR applications.

Many of the papers reviewed in this section used commonly available datasets to build and evaluate their models.

Table 5. Summary of HAR Deep Learning Models

Model	Year	Embedding	Other Features
Zeng et al. [200]	2014	CNN	
DCNN [201]	2015	CNN	Discrete Fourier transform
Yang et al. [202]	2015	CNN	
DeepConvLSTM [192]	2016	CNN, LSTM	
Hammerla et al. [199]	2016	CNN, LSTM	Bi-directional
Ronao and Cho [203]	2016	CNN	
Guan and Plötz [204]	2017	LSTM	Ensemble
Lee et al. [205]	2017	CNN	
Murad and Pyun [206]	2017	LSTM	Uni- and bi-directional
Ignatov [207]	2018	CNN	Statistical features
Moya Rueda et al. [208]	2018	CNN	
Yao et al. [209]	2018	CNN	Fully convolutional
Zeng et al. [210]	2018	LSTM	2 attention layers
AttnSense [211]	2019	CNN, GRU	Fast Fourier transform, 2 attention layers
InnoHAR [212]	2019	CNN, GRU	Inception
Zhang et al. [213]	2020	CNN	Attention
Challa et al. [214]	2021	CNN, LSTM	Bi-directional
CNN-biGRU [215]	2021	CNN, GRU	Bi-directional
DEBONAIR [216]	2021	ConvLSTM	
Mekruksavanich and Jitpattanakul [217]	2021	CNN, LSTM	
Mekruksavanich and Jitpattanakul [218]	2021	CNN, LSTM	Ensemble
Nafea et al. [219]	2021	CNN, LSTM	Bi-directional
Singh et al. [220]	2021	CNN, LSTM	Attention
Wang et al. [221]	2022	CNN	
Xu et al. [222]	2022	CNN, Resnet	Deformable convolutions

7.1.1 Convolutional Neural Networks. One of the most common types of convolutional kernels for HAR is the $k \times 1$ kernel. This kernel convolves k time steps together, moving along each time series in the input features in turn [221], so while weights are shared between the input features, there is no mixing between features. The outputs from the final convolutional layer are flattened and processed by fully connected layers before the final classification is made. Ronao et al. [203] performed a comprehensive evaluation of CNN models for HAR, evaluating the effect of changing the number of layers, filters, and filter sizes. The input data was collected from smartphone accelerometer and gyroscope sensors. Ignatov [207] used a one-layer CNN and augmented the extracted features with statistical features before being passed to fully connected layers. The architecture was effective with short time series (1 second) and therefore useful for real-time activity modeling. One drawback of the above method is that it forces weight sharing across all the input features. This may not be optimal, especially when using data collected from multiple devices. In this case, using a separate CNN for each device [208] allows independent weighting of the features. Similarly, as each sensor is typically tri-axial, a separate CNN can be used for each axis [200, 213]. The features extracted by each CNN are then concatenated and processed either by fully connected layers [200] or an attention head [213].

While the above two methods are the most common, other studies have proposed alternative CNNs for HAR. DCNN [201] pre-processes the sensor data using a Discrete Fourier Transform to convert IMU data to frequency signals, then uses two-dimensional convolutions to extract combined temporal and frequency features. Lee et al. [205] pre-processed the tri-axial accelerometer data to a magnitude vector, which was then processed in parallel by CNNs with varying kernel sizes, extracting features at different scales. Xu et al. [222] used deformable convolutions [223] in both a 2D-CNN and a ResNet model and found these models performed better than their non-deformable counterparts. Yao et al. [209] proposed a fully convolutional model using 2D temporal and feature convolutions. Their model has two advantages as (1) it handles arbitrary length input sequences and (2) it makes a prediction for each timestep, which avoids the need to pre-process the data into windows and can detect transitions between activities.

7.1.2 Recurrent Neural Networks. Several LSTM models have been proposed for HAR. Murad and Pyun [206] designed and compared three multi-layered LSTMs, a uni-directional LSTM, a bi-directional LSTM, and a “cascading” LSTM, which has a bi-directional first layer, followed by uni-directional layers. In each case the output from all time steps is used as input to the classification layer. Zeng et al. [210] added two attention layers to an LSTM, a sensor attention layer before the LSTM and a temporal attention layer after the LSTM. They include a regularization term they called “continuous attention” to smooth the transition between attention weights. Guan and Plötz [204] created an ensemble of LSTM models by saving the models at every training epoch, then selecting the best “M” models based on validation set results, thus aiming to reduce model variance.

7.1.3 Hybrid Models. Many recent studies have focused on hybrid models, combining both CNNs and RNNs. DeepConvLSTM [192] comprises four temporal convolutional layers followed by two LSTM layers, which the authors found to perform better than an equivalent CNN (replacing the LSTM layers with fully connected layers). As the LSTM layers have fewer parameters than fully connected layers, the DeepConvLSTM model was also much smaller. Singh et al. [220] used a CNN to encode the spatial data (i.e., the sensor readings at each timestamp) followed by a single LSTM layer to encode the temporal data, then a self-attention layer to weight the time steps. They found this model performed better than an equivalent model using temporal convolutions in the CNN layers. Challa et al. [214] proposed using three 1D-CNNs with different kernel sizes in parallel, followed by two bi-directional LSTM layers and a fully connected layer. Nafea et al. [219] also used 1D-CNNs with different kernel sizes and bi-directional LSTMs. However, they used separate branches for the CNNs and LSTMs, merging the features extracted in each branch for the final fully connected layer. Mekruksavanich and Jitpattanakul [217] compared a four-layer CNN-LSTM model with a smaller CNN-LSTM model and LSTM models, finding the extra convolutional layers improved performance over the smaller models. DEBONAIR [216] is another multi-layered model. It uses parallel 1D-CNNs, each having different kernel, filter, and pooling sizes to extract different types of features associated with different types of activity. These are followed by a combined 1D-CNN, then two LSTM layers. Mekruksavanich and Jitpattanakul [218] ensembled four different models: a CNN, an LSTM, a CNN-LSTM, and a ConvLSTM model. They aimed to produce a model for biometric user identification that could identify not only the activity being performed but also the participant performing the activity.

A few hybrid models use GRUs instead of LSTMs. InnoHAR [212] is a modified DeepConvLSTM [192], replacing the four CNN layers with inception layers and the two LSTM layers with GRU layers. The authors found this inception model performed better than both the original DeepConvLSTM model and a straight CNN model [202]. AttnSense [211] uses a Fast Fourier transform to generate frequency features, which are then convolved separately for each time step.

Attention layers are used to weight the extracted frequency features. These are then passed through a GRU with temporal attention to extract temporal features. CNN-BiGRU [215] uses a CNN layer to extract spatial features from the sensor data, then one or more GRU layers to extract temporal features. The final section of the model is a fully connected module consisting of one or more hidden layers and a softmax output layer.

7.2 Satellite Earth Observation

Ever since NASA launched the first Landsat satellite in 1972 [224], Earth-observing satellites have been recording images of the Earth's surface, providing 50 years of continuous Earth observation data that can be used to estimate environmental variables informing us about the state of the Earth. Instruments on board the satellites record reflected or emitted electromagnetic radiation from the Earth's surface and vegetation [225]. The regular, repeated observations from these instruments form **satellite image time series (SITS)** that are useful for analyzing the dynamic properties of some variables, such as plant phenology. The main modalities used for SITS analysis are multispectral spectrometers and spectroradiometers, which observe the visible and infrared frequencies and **Synthetic Aperture Radar (SAR)** systems, which emit a microwave signal and measure the backscatter.

Raw data collected by satellite instruments needs to be pre-processed before being used in machine learning. This is frequently done by the data providers to produce **analysis-ready datasets (ARDs)**. With the increasing availability of compatible ARDs from sources such as Google Earth Engine [226] and various data cubes [227, 228], models combining data from multiple data sources (multi-modal) are becoming more common. These data sources make it straightforward to obtain data that are co-registered (spatially aligned and with the same resolution and projection), thus avoiding the need for complex pre-processing.

Satellite image time series can be processed either (1) as 2D temporal and spectral data, processing each pixel independently and ignoring the spatial dimensions, or (2) as 4D data, including the two spatial dimensions, in which models thus extract spatio-temporal features. This latter method allows estimates to be made at pixel, patch, or object level; however, it requires either more complex models or spatial features to be extracted in a pre-processing step. Feature extraction can be as simple as extracting the mean value for each band. However, both clustering (TASSEL [229]) and neural-network-based methods, such as the Pixel-Set Encoder [230], have been used for more complex feature extraction. The most common use of SITS deep learning is for the classification of the Earth's surface by land cover and agricultural land by crop types. The classes used can range from very broad land cover categories (such as forest, grasslands, agriculture) through to specific crop types. Other classification tasks include identifying specific features, such as sinkholes [231], burnt areas [232], flooded areas [233], roads [234], deforestation [235], vegetation quality [236], and forest understory and litter types [237].

Extrinsic regression tasks are less common than classification tasks, but several recent studies have investigated methods of estimating water content in vegetation, as measured by the variable **Live Fuel Moisture Content (LFMC)** [238–241]. Other regression tasks include estimating the wood volume of forests [242] by using a hybrid CNN-MLP model combining a time series of Sentinel-2 images with a single LiDAR image and crop yield [243] that uses a hybrid of CNN and LSTM.

Many different approaches to learning from SITS data have been studied, with studies using all the main deep learning architectures, adapting them for multi-modal learning, and combining architectures in hybrid and ensemble models. The rest of this section reviews the architectures that have been used to model SITS data. A summary of these papers and the embedding architecture is provided in Table 6.

Table 6. Summary of SITS Deep Learning Models

Model	Year	Embedding	Other Features
Crop type classification			
TAN [244]	2019	2D-CNN and GRU	Attention—temporal
TGA [245]	2020	2D-CNN	Attention—squeeze and excitation
3D-CNN [246]	2018	3D-CNN	
DCM [247]	2020	LSTM	Self-attention
HierbiLSTM [248]	2022	LSTM	Self-attention
L-TAE [249]	2020	MLP	Attention—temporal
PSE-TAE [230, 250]	2020	MLP	Attention—temporal optionally multi-modal
SITS-BERT [251]	2021		Pre-trained transformer
Land Cover classification			
1D-CNN [252]	2017	1D-CNN and MLP	Hybrid model
1D & 2D-CNNs [253]	2017	1D-CNN; 2D-CNN	Ensemble model
TempCNN [254]	2019	1D-CNN	
TASSEL [229]	2020	1D-CNN	Self-attention
TSI [255]	2021	1D-CNN; LSTM	Ensemble model
TWINNS [256]	2019	2D-CNN and GRU	Attention—temporal; multi-modal
DuPLO [257]	2019	2D-CNN and GRU	Attention—temporal
Sequential RNN [258]	2018	2D-FCN and LSTM	Hybrid model
FG-UNET [259]	2019	UNet and 2D-CNN	Hybrid model
LSTM [260]	2017	LSTM	
HObs2sRNN [261]	2020	GRU	Attention—temporal
OD2RNN [262]	2019	GRU	Attention—temporal; multi-modal
SITS-Former [263]	2022	3D-CNN	Pre-trained transformer
Other classification tasks			
Deforestation [235]	2022	U-Net and LSTM	Hybrid model
Flood detection [233]	2020	Resnet and GRU	Hybrid model
Forest understory [237]	2022	2D-CNN and LSTM	Ensemble model
Road detection [234]	2020	U-Net and convLSTM	Hybrid model
Vegetation quality [236]	2017	LSTM; GRU	
Extrinsic regression tasks			
TempCNN-LFMC [239]	2021	1D-CNN	
Multi-tempCNN [240]	2022	1D-CNN	Multi-modal, ensemble model
LFMC estimation [238]	2020	LSTM	Multi-modal
LFMC estimation [241]	2022	1D-CNN and LSTM	Multi-modal, hybrid, ensemble
MLDL-net [243]	2020	2D-CNN and LSTM	Hybrid model
SSTNN [264]	2021	3D-CNN and LSTM	Hybrid model
MMFVE [242]	2022	2D-CNN	Hybrid model

7.2.1 Recurrent Neural Networks (RNNs). One of the first papers to use RNNs for land cover classification was Ienco et al. [260], who showed that an LSTM model out-performed non-deep-learning methods such as **Random Forest (RF)** and **Support Vector Machines (SVMs)**. However, they also showed that the performance of both RF and SVM improves if trained on features extracted by the LSTM model, and in some cases they were more accurate than the straight LSTM model. Rao et al. [238] used an extrinsic regression LSTM model to estimate LFMC in the western United States.

More commonly, however, RNNs are combined with an attention layer to allow the model to focus on the most important time steps. The OD2RNN model [262] used separate GRU layers followed by attention layers to process Sentinel-1 and Sentinel-2 data, combining the features

extracted by each source for the final fully connected layers. HOOb2sRNN [261] refined OD2RNN by using a hierarchy of land cover classifications; the model was pretrained using broad land cover classifications, then further trained using the finer-grained classifications. DCM [247] and HierbiLSTM [248] both use a bi-directional LSTM, processing the time series in both directions, followed by a self-attention transformer for a pixel-level crop-mapping model. All these studies found that adding the attention layers improved model performance over a straight GRU or LSTM model.

7.2.2 Convolutional Neural Networks (CNNs). While many authors have claimed that RNNs outperform CNNs for land cover and crop-type classification, most of these comparisons are to 2D-CNNs, which ignore the temporal ordering of SITS data [254]. However, other studies show using 1D-CNNs to extract temporal information and 3D-CNNs to extract spatio-temporal information are both effective methods of learning from SITS data. TempCNN [254] consists of three 1D convolutional layers. The output from the final convolutional layer is passed through a fully connected layer, then the final softmax classification layer. TASSEL [229] is an adaptation of TempCNN for OBIA classification, using TempCNN models to process features extracted from the objects, followed by an attention layer to weight the convolved features. TempCNN has also been adapted for extrinsic regression [239] and used for LFMC estimation [239–241].

2D-CNNs are mainly used to extract spatial or spatio-temporal features for both pixel and object classification. The model input is usually 4D and the data is convolved spatially, with two main methods used to handle the temporal dimension. In the first method, each time step is convolved separately and the extracted features are merged in later stages of the model [245]. In the second method, the time steps and channels are flattened to form a large multivariate image [242, 253]. FG-UNet [259] is a fully convolutional model that combines both the above methods, first grouping time steps by threes to produce images with 30 channels (10 spectral \times 3 temporal), which are passed through both U-Net and 2D-CNN layers. Ji et al. [246] used a 3D-CNN to convolve the spatial and temporal dimensions together, combining the strengths of 1D-CNN and 2D-CNNs. The study found that a 3D-CNN crop classification model performed significantly better than the 2D-CNN, again showing the importance of the temporal features. Another study, SSTNN [264], obtained good results for crop yield prediction by using a 3D-CNN to convolve the spatial and spectral dimensions, extracting spatio-spectral features for each time step. These features were then processed by LSTM layers to perform the temporal modeling.

7.2.3 Transformer and Attention Models. As an alternative to including attention layers with a CNN or RNN, several studies have designed models that process temporal information using only attention layers. PSE-TAE [230] used a modified transformer called a **temporal attention encoder (TAE)** for crop mapping and found the TAE performed better than either a CNN or an RNN. L-TAE [249] replaced the TAE with a lightweight transformer that is both computationally efficient and more accurate than the full TAE. Ofori-Ampofo et al. [250] adapted the TAE model for multi-modal inputs using Sentinel-1 and Sentinel-2 data for crop-type mapping. Rußwurm and Körner [265] compared a self-attention model with RNN and CNN architectures. They found that this model was more robust to noise than either RNN or CNN and suggested that self-attention is suitable for processing raw, cloud-affected satellite data.

Building on the success of pre-trained transformers for natural language processing such as BERT [94], pre-trained transformers have been proposed for Earth observation tasks [251]. Earth observation tasks are particularly suited for pre-trained models as large quantities of Earth observation data are readily available, while labeled data can be difficult to obtain [266], especially in remote locations. SITS-BERT [251] is an adaptation of BERT [94] for pixel-based SITS classification. For the pretext task, random noise is added to the pixels, and the model is trained to identify and remove this noise. The pre-trained model is then further trained for required tasks such as crop

type or land cover mapping. SITS-Former [263] modifies SITS-BERT for patch classification by using 3D-Conv layers to encode the spatial-spectral information, which is then passed through the temporal attention layers. The pretext task used for SITS-Former is to predict randomly masked pixels.

7.2.4 Hybrid Models. A common use of hybrid models is to use a CNN to extract spatial features and an RNN to extract temporal features. Garnot et al. [267] compared a straight 2D-CNN model (thus ignoring the temporal aspect), a straight GRU model (thus ignoring the spatial aspect), and a combined 2D-CNN and GRU model (thus using both spatial and temporal information) and found the combined model gave the best results, demonstrating that both the spatial and temporal dimensions provide useful information for land cover mapping and crop classification. DuPLO [257] was one of the first models to exploit this method, running a CNN and ConvGRU model in parallel, then fusing the outputs using a fully connected network for the final classifier. During training, an auxiliary classifier for each component was used to enhance the discriminative power. TWINNS [256] extended DuPLO to a multi-modal model, using time series of both Sentinel-1 (SAR) and Sentinel-2 (Optical) images. Each modality was processed by separate CNN and convGRU models, and then the output features from all four models were fused for classification.

Other hybrid models include Li et al. [244], who used a CNN for spatial and spectral unification of Landsat-8 and Sentinel-2 images, which were then processed by a GRU. MLDL-Net [243] is a 2D-CNN extrinsic regression model, using CNNs to extract time step features, which are then passed through an LSTM model to extract temporal features. Fully connected layers combine the feature sets to predict crop yield. Rußwurm and Körner [258] extracted temporal features first, using a bi-directional LSTM, then used a fully convolutional 2D-CNN to incorporate spatial information and classify each pixel in the input patch.

7.2.5 Ensemble Models. One of the easiest ways to ensemble DL models is to train multiple homogeneous models that vary only in the random weight initialization [268]. Di Mauro et al. [252] ensembled 100 LULC models with different weight initializations by averaging the softmax predictions. They found this produced a more stable and stronger classifier that outperformed the individual models. Multi-tempCNN [240], a model for LFMC estimation, is an ensemble of homogeneous models for extrinsic regression. The authors suggested that as an additional benefit, the variance of the individual model predictions can be used to obtain a measure of uncertainty of the estimates. TSI [255] also ensembles a set of homogeneous models, but instead of relying on random weight initialization to introduce model diversity, the time series are segmented and models trained on each segment.

Other methods create ensembles of heterogeneous models. Kussul et al. [253] compared ensembles of 1D-CNNs and 2D-CNNs models for land cover classification. Each model in the ensemble used a different number of filters, thus finding different feature sets useful for classification. Xie et al. [241] ensembled three heterogeneous models—a causal **temporal convolutional neural network (TCN)**, an LSTM, and a hybrid TCN-LSTM model—for an extrinsic regression model to estimate LFMC. The ensembles were created using stacking [269]. The authors compared this method to boosting their TCN-LSTM model, using Adaboost [270] to create a three-member ensemble, and found that stacking a diverse set of models out-performed boosting.

7.2.6 EO Surveys and Reviews. This survey is one of very few that include a section focusing specifically on deep learning TSC and TSER tasks using SITS data. However, there are other reviews that provide further information about related topics. Gomez et al. [271] is an older review highlighting the important role of SITS data for land cover classification. Zhu et al. [272] reviewed the advances and challenges in DL for remote sensing and the resources available that are

potentially useful to help DL address some of the major challenges facing humanity. Ma et al. [273] study the role of deep learning in Earth observation using remotely sensed data. It covers a broad range of tasks including image fusion, image segmentation, and object-based analysis, as well as classification tasks. Yuan et al. [274] provide a review of DL applications for remote sensing, comparing the role of DL versus physical modeling of environmental variables and highlighting challenges in DL for remote sensing that need to be addressed. Chaves et al. [275] reviewed recent research using Landsat 8 and/or Sentinel-2 data for land cover mapping. While not focused on SITS DL methods, the review notes the growing importance of these methods. Moskolai et al. [276] provide a review of forecasting applications using DL with SITS data that provides an analysis of the main DL architectures that are relevant for classification as well as forecasting.

8 CONCLUSION

In conclusion, this survey article has discussed a variety of deep network architectures for time series classification and extrinsic regression tasks, including multilayer perceptrons, convolutional neural networks, recurrent neural networks, and attention-based models. We have also highlighted refinements that have been made to improve the performance of these models on time series tasks. Additionally, we have discussed two critical applications of time series classification and regression, human activity recognition and satellite Earth observation. Overall, using deep network architectures and refinements has enabled significant progress in the field of time series classification and will continue to be essential for addressing a wide range of real-world problems. We hope this survey will stimulate further research using deep learning techniques for time series classification and extrinsic regression. Additionally, we provide a carefully curated collection of sources, available at https://github.com/Navidfoumani/TSC_Survey, to further support the research community.

APPENDICES

A NON-DEEP-LEARNING TIME SERIES CLASSIFICATION

In this section, we aim to give a brief introduction to the field of TSC and discuss its current status. We refer interested readers to the "bake-off" papers [11, 25, 26] that describe TSC methods in much more detail and benchmark them.

Research in TSC started with distance-based approaches that find discriminating patterns in the shape of the time series. Distance-based approaches usually consist of coupling a **1-nearest neighbor (1NN)** classifier with a time series distance measure [277, 278]. Small distortions in the time series can lead to false matches when measuring the distance between time series using standard distance measurements such as Euclidean distance [277]. A time series distance measure aims to compensate for these distortions by aligning two time series such that the alignment cost between the two are minimized. There are many time series distances proposed in the literature; among these, the *DTW* distance is one of the most popular choices for many time series tasks, due to its intuitiveness and effectiveness in aligning two time series. The 1NN-*DTW* has been the go-to method for TSC for decades. However, by comparing several time series distance measures, the work in [277] showed that as of 2015, there was no single distance that significantly outperformed *DTW* when used with a 1NN classifier. The recent Amerced *DTW* [279] distance is the first distance that is significantly more accurate than *DTW*. These individual 1NN classifiers with different distances can be ensembled together to create an ensemble, such as the **Ensemble of Elastic distances (EE)**, that significantly outperforms each of them individually [277, 278]. However, since most distances have a complexity of $O(L^2)$, where L is the length of the series, performing a nearest neighbor search becomes very costly. Hence, distance-based approaches are considered to be one of the slowest methods for TSC [280, 281].

As a result of EE, recent studies have focused mainly on developing ensembling methods that significantly outperform 1NN-DTW [278, 280–290]. These approaches use either an ensemble of tree-based approaches [289, 290] or an ensemble of different types of discriminant classifiers, such as NN with several distances and SVM on one or several feature spaces [282, 285–287]. All of these approaches share a common property—the data transformation phase where the time series is transformed into a new feature space such as the shapelets transform [286] or DTW features [285]. Taking advantage of this notion led to the development of the **Hierarchical Vote Collective of Transformation-based Ensembles (HIVE-COTE)** [281, 284]. HIVE-COTE is a meta ensemble for TSC and forms its ensemble from ensemble classifiers of multiple domains. Since its introduction in 2016 [284], HIVE-COTE has gone through a few iterations. Recently, the latest HIVE-COTE version, HIVE-COTEv2.0 (HC2), was proposed [281]. It is composed of four ensemble members, each of them being the then state of the art in their respective domains. It is currently one of the most accurate classifiers for both univariate and multivariate TSC tasks [281]. Despite being accurate on 26 multivariate and 142 univariate TSC benchmark datasets that are relatively small, HC2 scales poorly on large datasets with long time series as well as datasets with large numbers of channels.

Various work has been done on speeding up TSC methods without sacrificing accuracy [14, 278, 291–295]. A recent breakthrough is the development of Rocket [14], which was able to process 109 univariate time series datasets under 4 hours while the previous fastest one took days. Rocket leverages a large number of random convolutional filters to extract features from each series that might be relevant to classifying a series. These features are then passed to a linear model for classification. Rocket has been improved to be faster (Minirocket [291]) and more accurate (Multirocket [292] and Hydra [293]). Hydra when combined with Multirocket is now one of the fastest and most accurate methods for TSC.

B DNN ARCHITECTURES FOR TIME SERIES

In this section, we provide a descriptive overview of deep-learning-based models for TSC. The focus is on clarifying their architectures and outlining their adaptations to the specific characteristics of time series data.

B.1 Convolution Neural Networks (CNNs)

Many variants of CNN architectures have been proposed in the literature, but their primary components are very similar. Using the LeNet-5 [296] as an example, it consists of three types of layers: convolutional, pooling, and fully connected. The purpose of the convolutional layer is to learn feature representations of the inputs. Figure 2(a) shows the architecture of the t-LeNet network, which is a time-series-specific version of LeNet. This figure shows that the convolution layer is composed of several convolution kernels (or filters) used to compute different feature maps. In particular, each neuron of a feature map is connected to a region of neighboring neurons in the previous layer called the receptive field. Feature maps can be created by first convolving inputs with learned kernels and then applying an element-wise nonlinear activation function to the convolved results. It is important to note that all spatial locations of the input share the kernel for each feature map, and several kernels are used to obtain the entire feature map.

The feature value of the l th layer of the k th feature map at location (i, j) is obtained by

$$Z_{i,j,k}^l = \mathbf{W}_k^{lT} \mathbf{A}_{i,j}^{l-1} + b_k^l, \quad (2)$$

where \mathbf{W}_k^l and b_k^l are the weight vector and bias term of the k th filter of the l th layer, respectively, and $\mathbf{A}_{i,j}^{l-1}$ is the input patch centered at location (i, j) of the l th layer. Note that the kernel \mathbf{W}_k^l

that generates the feature map $Z_{i,j,k}^l$ is shared. A weight-sharing mechanism has several advantages, such as reducing model complexity and making the network easier to train. Let $f(\cdot)$ denote the nonlinear activation function. The activation value of convolutional feature $Z_{i,j,k}^l$ can be computed as

$$A_{i,j,k}^l = f(Z_{i,j,k}^l). \tag{3}$$

The most common activation functions are sigmoid, tanh, and ReLU [297]. As shown in Figure 2(a), a pooling layer is often placed between two convolution layers to reduce the resolution of the feature maps and to achieve shift invariance. Following several convolution stages—the block comprising convolution, activation, and pooling is called convolution *stage*—there may be one or more fully connected layers that aim to perform high-level reasoning. As discussed in Section 3.1, each neuron in the previous layer is connected to every neuron in the current layer to generate global semantic information. In the final layer of CNNs, there is the output layer in which the Softmax operators are commonly used for classification tasks [40].

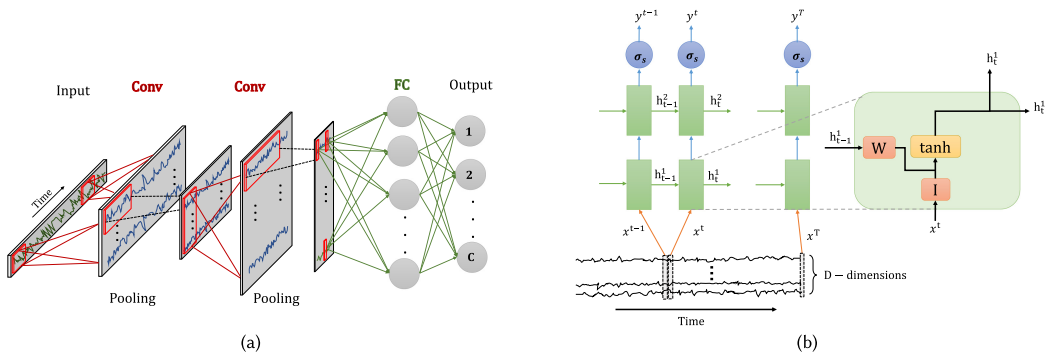


Fig. 2. (a) The architecture of the t-LeNet network. (b) The architecture of two-layer Recurrent Neural Network.

B.2 Recurrent Neural Networks (RNN)

RNNs are types of neural networks that are specifically designed to process time series and other sequential data. RNNs are conceptually similar to FFNs. While FFNs map from fixed-size inputs to fixed-size outputs, RNNs can process variable-length inputs and produce variable-length outputs. This capability is enabled by sharing parameters over time through directed connections between individual layers. RNN models for TSC can be classified as sequence to sequence or sequence to one based on their outputs. Figure 2(b) shows sequence-to-sequence architectures for RNN models, with an output for each input sub-series. On the other hand, in sequence-to-one architecture, decisions are made using only y^T and ignoring the other outputs.

At each time step t , RNNs maintain a hidden vector h , which updates as follows [298, 299]:

$$h_t = \tanh(W h_{t-1} + I x^t), \tag{4}$$

where $X = \{x^1, \dots, x^{t-1}, x^t, \dots, x^T\}$ contains all of the observation, \tanh denotes the hyperbolic tangent function, and the recurrent weight and the projection matrix are shown by W and I , respectively. The hidden-to-hidden connections also model the short-term time dependency. The hidden state h is used to make a prediction as

$$y^t = \sigma_s(W h_{t-1}), \tag{5}$$

where σ_s is a softmax function and provides a normalized probability distribution over the possible classes. As depicted in Figure 2(b), the hidden state h can be used to stack RNNs in order to build deeper networks:

$$h_t^l = \sigma(W h_{t-1}^l + I h_t^{l-1}), \quad (6)$$

where σ is the logistic sigmoid function. As an alternative to feeding each time step to the RNN, the data can be divided into time windows of ω observations, with the option for variable overlaps. Each time window is labeled with the majority response labels within the ω window.

B.3 Attention-based Model

B.3.1 Self-attention. The attention mechanism was introduced by [300] for improving the performance of encoder-decoder models [301] in neural machine translation. The encoder-decoder in neural machine translation encodes a source sentence into a vector in latent space and decodes the latent vector into a target language sentence. As shown in Figure 3(a), the attention mechanism allows the decoder to pay attention to the segments of the source for each target through a context vector c_t . For this model, a variable-length attention vector α_t , equal to the number of source time steps, is derived by comparing the current target hidden state h_t with each source hidden state \bar{h}_s as follows [302]:

$$\alpha_t(s) = \frac{\exp(\text{score}(h_t, \bar{h}_s))}{\sum_{s'} \exp(\text{score}(h_t, \bar{h}_{s'}))}. \quad (7)$$

The term *score* is referred to as an alignment model and used to compare the target hidden state h_t with each of the source hidden states \bar{h}_s , and the result is normalized to produce attention weights (a distribution over source positions). There are various choices of the scoring function:

$$\text{score}(h_t, \bar{h}_s) = \begin{cases} h_t^T W \bar{h}_s \\ v_\alpha^T \tanh(W_\alpha [h_t; \bar{h}_s]). \end{cases} \quad (8)$$

These scores influence the attention distribution, impacting how the model attends to different parts of the input sequence during predictions. As shown above, the score function is parameterized as an FFN that is jointly trained with all the other components of the model. The model directly computes soft attention, allowing the cost function's gradient to be backpropagated [300]. Given the alignment vector as weights, the context vector c_t is computed as the weighted average over all the source hidden states:

$$c_t = \sum_s \alpha_{ts} \bar{h}_s. \quad (9)$$

Accordingly, the computation path goes from $h_t \rightarrow \alpha_t \rightarrow c_t \rightarrow \tilde{h}_t$ and then makes a prediction using a *Softmax* function [302]. Note that \tilde{h}_t is a refined hidden state that incorporates both the original hidden state h_t and the context information c_t obtained through attention mechanisms.

B.3.2 Transformers. Similar to self-attention and other competitive neural sequence models, the original transformer developed for natural language processing (hereinafter the vanilla transformer) has an encoder-decoder structure that takes as input a sequence of words from the source language and then generates the translation in the target language [93]. Both the encoder and decoder are composed of multiple identical blocks. Each encoder block consists of a multi-head self-attention module and a position-wise FFN, while each decoder block inserts cross-attention models between the multi-head self-attention module and the position-wise FFN. Unlike RNNs, transformers do not use recurrence and instead model sequence information using the positional encoding in the input embeddings.

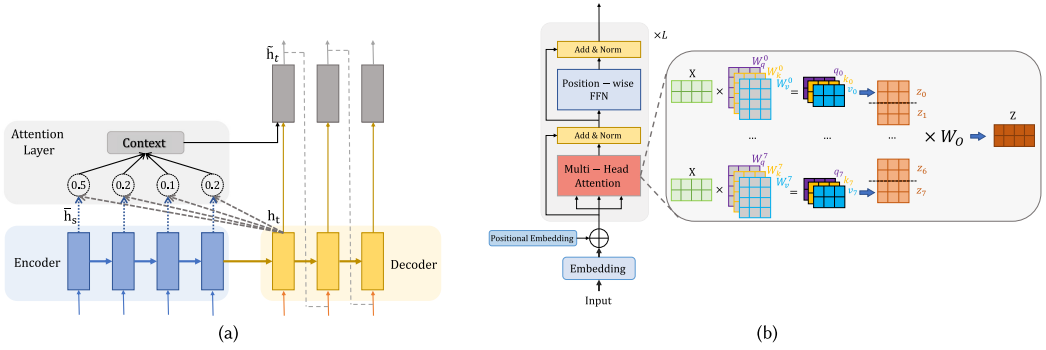


Fig. 3. (a) Self-attention mechanism. (b) Multi-head attention block.

The transformer architecture is based on finding associations or correlations between various input segments using the dot product. As shown in Figure 3(b), the attention operation in transformers starts with building three different linearly weighted vectors from the input x_i , referred to as query (q_i), key (k_i), and value (v_i):

$$q_i = W_q x_i, \quad k_i = W_k x_i, \quad v_i = W_v x_i, \tag{10}$$

where W_q , W_k , and W_v are learnable weight matrices. The output vectors z_i are given by

$$z_i = \sum_j \text{softmax} \left(\frac{q_i^T k_j}{\sqrt{d_q}} \right) v_j. \tag{11}$$

Note that the weighting of the value vector v_i depends on the mapped correlation between the query vector q_i at position i and the key vector k_j at position j . The value of the dot product tends to grow with the increasing size of the query and key vectors. As the softmax function is sensitive to large values, the attention weights are scaled by the square root of the size of the query and key vectors d_q . The input data may contain several levels of correlation information, and the learning process may benefit from processing the input data in multiple different ways. Multiple attention heads are introduced that operate on the same input in parallel and use different weight matrices W_q, W_k , and W_v to extract various levels of correlation between the input data.

REFERENCES

- [1] Qiang Yang and Xindong Wu. 2006. 10 challenging problems in data mining research. *International Journal of Information Technology & Decision Making* 5, 4 (2006), 597–604.
- [2] Philippe Esling and Carlos Agon. 2012. Time-series data mining. *ACM Computing Surveys (CSUR)* 45, 1 (2012), 1–34.
- [3] Henry Friday Nweke, Ying Wah Teh, Mohammed Ali Al-Garadi, and Uzoma Rita Alo. 2018. Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges. *Expert Systems with Applications* 105 (2018), 233–261. DOI : <http://dx.doi.org/10.1016/j.eswa.2018.03.056>
- [4] Jindong Wang, Yiqiang Chen, Shuji Hao, Xiaohui Peng, and Lisha Hu. 2019. Deep learning for sensor-based activity recognition: A survey. *Pattern Recognition Letters* 119 (2019), 3–11. DOI : <http://dx.doi.org/10.1016/j.patrec.2018.02.010>
- [5] Kaixuan Chen, Dalin Zhang, Lina Yao, Bin Guo, Zhiwen Yu, and Yunhao Liu. 2021. Deep learning for sensor-based human activity recognition: Overview, challenges, and opportunities. *ACM Computing Surveys (CSUR)* 54, 4 (2021), 1–40. DOI : <http://dx.doi.org/10.1145/3447744>
- [6] Robin Tibor Schirrmester, Jost Tobias Springenberg, Lukas Dominique Josef Fiederer, Martin Glasstetter, Katharina Eggenberger, Michael Tangermann, Frank Hutter, Wolfram Burgard, and Tonio Ball. 2017. Deep learning with convolutional neural networks for EEG decoding and visualization. *Human Brain Mapping* 38, 11 (2017), 5391–5420.
- [7] A. Rajkomar, E. Oren, K. Chen, A. M. Dai, N. Hajaj, M. Hardt, P. J. Liu, X. Liu, J. Marcus, M. Sun, P. Sundberg, H. Yee, K. Zhang, Y. Zhang, G. Flores, G. E. Duggan, J. Irvine, Q. Le, K. Litsch, A. Mossin, J. Tansuwan, D. Wang, J. Wexler, J. Wilson, D. Ludwig, S. L. Volchenboum, K. Chou, M. Pearson, S. Madabushi, N. H. Shah, A. J. Butte, M. D. Howell, C.

- Cui, G. S. Corrado, J. Dean. 2018. Scalable and accurate deep learning with electronic health records. *NPJ Digit Med.* 2018 May 8, 1 (2018), 18. DOI : [10.1038/s41746-018-0029-1](https://doi.org/10.1038/s41746-018-0029-1). PMID: 31304302; PMCID: PMC6550175.
- [8] Anthony Bagnall, Hoang Anh Dau, Jason Lines, Michael Flynn, James Large, Aaron Bostrom, Paul Southam, and Eamonn Keogh. 2018. The UEA multivariate time series classification archive, 2018. *arXiv preprint:1811.00075* (2018).
 - [9] Hoang Anh Dau, Anthony Bagnall, Kaveh Kamgar, Chin-Chia Michael Yeh, Yan Zhu, Shaghayegh Gharghabi, Chotirat Ann Ratanamahatana, and Eamonn Keogh. 2019. The UCR time series archive. *IEEE/CAA Journal of Automatica Sinica* 6, 6 (2019), 1293–1305.
 - [10] Chang Wei Tan, Christoph Bergmeir, François Petitjean, and Geoffrey I. Webb. 2021. Time series extrinsic regression. *Data Mining and Knowledge Discovery* 35, 3 (2021), 1032–1060.
 - [11] Matthew Middlehurst, Patrick Schäfer, and Anthony Bagnall. 2023. Bake off redux: A review and experimental evaluation of recent time series classification algorithms. *arXiv preprint arXiv:2304.13029* (2023).
 - [12] Hassan Ismail Fawaz, Benjamin Lucas, Germain Forestier, Charlotte Pelletier, Daniel F. Schmidt, Jonathan Weber, Geoffrey I. Webb, Lhassane Idoumghar, Pierre-Alain Muller, and François Petitjean. 2020. Inceptiontime: Finding Alexnet for time series classification. *Data Mining and Knowledge Discovery* 34, 6 (2020), 1936–1962.
 - [13] Navid Mohammadi Foumani, Chang Wei Tan, Geoffrey I. Webb, and Mahsa Salehi. 2024. Improving position encoding of transformers for multivariate time series classification. *Data Mining and Knowledge Discovery* 38, 1 (2024), 22–48.
 - [14] Angus Dempster, François Petitjean, and Geoffrey I. Webb. 2020. ROCKET: Exceptionally fast and accurate time series classification using random convolutional kernels. *Data Mining and Knowledge Discovery* 34, 5 (2020), 1454–1495.
 - [15] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery* 33, 4 (2019), 917–963.
 - [16] Zhiguang Wang, Weizhong Yan, and Tim Oates. 2017. Time series classification from scratch with deep neural networks: A strong baseline. In *2017 International Joint Conference on Neural Networks (IJCNN'17)*. IEEE, 1578–1585.
 - [17] Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. 2022. Transformers in time series: A survey. *arXiv preprint:2202.07125* (2022).
 - [18] Yifan Hao and Huiping Cao. 2020. A new attention mechanism to classify multivariate time series. In *29th International Joint Conference on Artificial Intelligence*.
 - [19] George Zerveas, Srideepika Jayaraman, Dhaval Patel, Anuradha Bhamidipaty, and Carsten Eickhoff. 2021. A transformer-based framework for multivariate time series representation learning. In *27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2114–2124.
 - [20] Xiao Liu, Fanjin Zhang, Zhenyu Hou, Li Mian, Zhaoyu Wang, Jing Zhang, and Jie Tang. 2021. Self-supervised learning: Generative or contrastive. *IEEE Transactions on Knowledge and Data Engineering* 35, 1 (2021), 857–876.
 - [21] Emadeldeen Eldele, Mohamed Ragab, Zhenghua Chen, Min Wu, Chee Keong Kwok, Xiaoli Li, and Cuntai Guan. 2021. Time-series representation learning via temporal and contextual contrasting. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI'21)*. 2352–2359.
 - [22] Chao-Han Huck Yang, Yun-Yun Tsai, and Pin-Yu Chen. 2021. Voice2series: Reprogramming acoustic models for time series classification. In *International Conference on Machine Learning*. PMLR, 11808–11819.
 - [23] Zhihan Yue, Yujing Wang, Juanyong Duan, Tianmeng Yang, Congrui Huang, Yunhai Tong, and Bixiong Xu. 2022. Ts2vec: Towards universal representation of time series. *Proceedings of the AAAI Conference on Artificial Intelligence* 36, 8 (2022), 8980–8987.
 - [24] Navid Mohammadi Foumani, Chang Wei Tan, Geoffrey I Webb, and Mahsa Salehi. 2023. Series2Vec: Similarity-based self-supervised representation learning for time series classification. *arXiv preprint arXiv:2312.03998* (2023).
 - [25] Anthony Bagnall, Jason Lines, Aaron Bostrom, James Large, and Eamonn Keogh. 2017. The great time series classification bake off: A review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery* 31, 3 (2017), 606–660.
 - [26] Alejandro Pasos Ruiz, Michael Flynn, James Large, Matthew Middlehurst, and Anthony Bagnall. 2021. The great multivariate time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery* 35, 2 (2021), 401–449.
 - [27] Chang Wei Tan, Christoph Bergmeir, Francois Petitjean, and Geoffrey I Webb. 2020. Monash university, UEA, UCR time series regression archive. *arXiv preprint:2006.10996* (2020).
 - [28] Martin Längkvist, Lars Karlsson, and Amy Loutfi. 2014. A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recognition Letters* 42 (2014), 11–24.
 - [29] Yoshua Bengio, Li Yao, Guillaume Alain, and Pascal Vincent. 2013. Generalized denoising auto-encoders as generative models. *Advances in Neural Information Processing Systems* 26 (2013), 899–907.
 - [30] Qinghua Hu, Rujia Zhang, and Yucan Zhou. 2016. Transfer learning for short-term wind speed prediction with deep neural networks. *Renewable Energy* 85 (2016), 83–95.
 - [31] Joan Serra, Santiago Pascual, and Alexandros Karatzoglou. 2018. Towards a universal neural network encoder for time series. In *CCLIA*. 120–129.

- [32] Debrup Banerjee, Kazi Islam, Keyi Xue, Gang Mei, Lemin Xiao, Guangfan Zhang, Roger Xu, Cai Lei, Shuiwang Ji, and Jiang Li. 2019. A deep transfer learning approach for improved post-traumatic stress disorder diagnosis. *Knowledge and Information Systems* 60, 3 (2019), 1693–1724.
- [33] Witali Aswolinskiy, René Felix Reinhart, and Jochen Steil. 2018. Time series classification in reservoir-and model-space. *Neural Processing Letters* 48, 2 (2018), 789–809.
- [34] Eoin Brophy, Zhengwei Wang, Qi She, and Tomás Ward. 2023. Generative adversarial networks in time series: A systematic literature review. *ACM Computer Surveys* 55, 10, Article 199 (Feb. 2023), 31 pages. DOI: <http://dx.doi.org/10.1145/3559540>
- [35] Felipe Arias Del Campo, María Cristina Guevara Neri, Osslan Osiris Vergara Villegas, Vianey Guadalupe Cruz Sánchez, Humberto de Jesús Ochoa Domínguez, and Vicente García Jiménez. 2021. Auto-adaptive multilayer perceptron for univariate time series classification. *Expert Systems with Applications* 181 (2021), 115147.
- [36] Brian Kenji Iwana, Volkmar Frinken, and Seiichi Uchida. 2016. A robust dissimilarity-based neural network for temporal pattern recognition. In *2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR'16)*. IEEE, 265–270.
- [37] Brian Kenji Iwana, Volkmar Frinken, and Seiichi Uchida. 2020. DTW-NN: A novel neural network for time series recognition using dynamic alignment between inputs and weights. *Knowledge-based Systems* 188 (2020), 104971.
- [38] Nuzhat Tabassum, Sujeendran Menon, and Agnieszka Jastrzebska. 2022. Time-series classification with SAFE: Simple and fast segmented word embedding-based neural time series classifier. *Information Processing & Management* 59, 5 (2022), 103044.
- [39] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25 (2012), 1097–1105.
- [40] Jiuxiang Gu, Zhenhua Wang, Jason Kuen, Lianyang Ma, Amir Shahroudy, Bing Shuai, Ting Liu, Xingxing Wang, Gang Wang, Jianfei Cai, et al. 2018. Recent advances in convolutional neural networks. *Pattern Recognition* 77 (2018), 354–377.
- [41] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
- [42] Seyed Navid Mohammadi Foumani and Ahmad Nickabadi. 2019. A probabilistic topic model using deep visual word representation for simultaneous image classification and annotation. *Journal of Visual Communication and Image Representation* 59 (2019), 195–203.
- [43] Yi Zheng, Qi Liu, Enhong Chen, Yong Ge, and J Leon Zhao. 2014. Time series classification using multi-channels deep convolutional neural networks. In *International Conference on Web-age Information Management*. Springer, 298–310.
- [44] Jianbo Yang, Minh Nhut Nguyen, Phyto Phyto San, Xiao Li Li, and Shonali Krishnaswamy. 2015. Deep convolutional neural networks on multichannel time series for human activity recognition. In *24th International Joint Conference on Artificial Intelligence*.
- [45] Bendong Zhao, Huanzhang Lu, Shangfeng Chen, Junliang Liu, and Dongya Wu. 2017. Convolutional neural networks for time series classification. *Journal of Systems Engineering and Electronics* 28, 1 (2017), 162–169.
- [46] Jonathan Long, Evan Shelhamer, and Trevor Darrell. 2015. Fully convolutional networks for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3431–3440.
- [47] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [48] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. 2016. Learning deep features for discriminative localization. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2921–2929.
- [49] Xiaowu Zou, Zidong Wang, Qi Li, and Weiguo Sheng. 2019. Integration of residual network and convolutional neural network along with various activation functions and global pooling for time series classification. *Neurocomputing* 367 (2019), 39–45.
- [50] Yuhong Li, Xiaofan Zhang, and Deming Chen. 2018. CSRNet: Dilated convolutional neural networks for understanding the highly congested scenes. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1091–1100.
- [51] Omolbanin Yazdanbakhsh and Scott Dick. 2019. Multivariate time series classification using dilated convolutional neural network. *arXiv preprint:1905.01697* (2019).
- [52] Seyed Navid Mohammadi Foumani, Chang Wei Tan, and Mahsa Salehi. 2021. Disjoint-CNN for multivariate time series classification. In *2021 International Conference on Data Mining Workshops (ICDMW'21)*. IEEE, 760–769.
- [53] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. MobileNet-V2: Inverted residuals and linear bottlenecks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 4510–4520.
- [54] Zhiguang Wang and Tim Oates. 2015. Encoding time series as images for visual inspection and classification using tiled convolutional neural networks. In *Workshops at the 29th AAAI Conference on Artificial Intelligence*.
- [55] Nima Hatami, Yann Gavet, and Johan Debayle. 2018. Classification of time-series images using deep convolutional neural networks. In *10th International Conference on Machine Vision (ICMV'17)*, Vol. 10696. SPIE, 242–249.

- [56] Saeed Karimi-Bidhendi, Faramarz Munshi, and Ashfaq Munshi. 2018. Scalable classification of univariate and multivariate time series. In *2018 IEEE International Conference on Big Data (Big Data '18)*. IEEE, 1598–1605.
- [57] Yuxuan Zhao and Zhongmin Cai. 2019. Classify multivariate time series by deep neural network image classification. In *2019 2nd China Symposium on Cognitive Computing and Hybrid Intelligence (CCHI'19)*. IEEE, 93–98.
- [58] Chao-Lung Yang, Zhi-Xuan Chen, and Chen-Yi Yang. 2019. Sensor classification using convolutional neural network by encoding multivariate time series as two-dimensional colored images. *Sensors* 20, 1 (2019), 168.
- [59] J-P. Eckmann, S. Oliffson Kamphorst, D. Ruelle, et al. 1987. Recurrence plots of dynamical systems. *Europhysics Letters* 4, 9 (1987), 17.
- [60] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2818–2826.
- [61] Wei Chen and Ke Shi. 2019. A deep learning framework for time series classification using Relative Position Matrix and Convolutional Neural Network. *Neurocomputing* 359 (2019), 384–394.
- [62] Zhicheng Cui, Wenlin Chen, and Yixin Chen. 2016. Multi-scale convolutional neural networks for time series classification. *arXiv preprint:1603.06995* (2016).
- [63] Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data augmentation for time series classification using convolutional neural networks. In *ECML/PKDD Workshop on Advanced Analytics and Learning on Temporal Data*.
- [64] Chien-Liang Liu, Wen-Hoar Hsaio, and Yao-Chung Tu. 2018. Time series classification with multivariate convolutional neural network. *IEEE Transactions on Industrial Electronics* 66, 6 (2018), 4788–4797.
- [65] Anthony Brunel, Johanna Pasquet, Jérôme Pasquet, Nancy Rodriguez, Frédéric Comby, Dominique Fouchez, and Marc Chaumont. 2019. A CNN adapted to time series for the classification of Supernovae. *Electronic Imaging* 2019, 14 (2019), 90–1.
- [66] Jingyu Sun, Susumu Takeuchi, and Ikuro Yamasaki. 2021. Prototypical inception network with cross branch attention for time series classification. In *2021 International Joint Conference on Neural Networks (IJCNN'21)*. IEEE, 1–7.
- [67] Saidrasul Usmankhujayev, Bunyodbek Ibrokhimov, Shokhrukh Baydadaev, and Jangwoo Kwon. 2021. Time series classification with InceptionFCN. *Sensors* 22, 1 (2021), 157.
- [68] Xueyuan Gong, Yin-Whar Si, Yongqi Tian, Cong Lin, Xinyuan Zhang, and Xiaoxiang Liu. 2022. KDCTime: Knowledge distillation with calibration on InceptionTime for time-series classification. *Information Science* 613 (2022), 184–203.
- [69] Ali Ismail-Fawaz, Maxime Devanne, Stefano Berretti, Jonathan Weber, and Germain Forestier. 2023. LITE: Light inception with boosting techniques for time series classification. In *2023 IEEE 10th International Conference on Data Science and Advanced Analytics (DSAA'23)*. IEEE, 1–10.
- [70] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition*. 1–9.
- [71] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A. Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *31st AAAI Conference on Artificial Intelligence*.
- [72] Mutegeki Ronald, Alwin Poulso, and Dong Seog Han. 2021. iSPLInception: An inception-ResNet deep learning architecture for human activity recognition. *IEEE Access* 9 (2021), 68985–69001.
- [73] Ali Ismail-Fawaz, Maxime Devanne, Jonathan Weber, and Germain Forestier. 2022. Deep learning for time series classification using new hand-crafted convolution filters. In *2022 IEEE International Conference on Big Data (Big Data '22)*. IEEE, 972–981.
- [74] Michael Hüskens and Peter Stagge. 2003. Recurrent neural networks for time series classification. *Neurocomputing* 50 (2003), 223–235.
- [75] Don Dennis, Durmus Alp Emre Acar, Vikram Mandikal, Vinu Sankar Sadasivan, Venkatesh Saligrama, Harsha Vardhan Simhadri, and Prateek Jain. 2019. Shallow RNN: Accurate time-series classification on resource constrained devices. *Advances in Neural Information Processing Systems* 32 (2019), 11 pages.
- [76] Santiago Fernández, Alex Graves, and Jürgen Schmidhuber. 2007. Sequence labelling in structured domains with hierarchical recurrent neural networks. In *20th International Joint Conference on Artificial Intelligence (IJCAI'07)*.
- [77] Michiel Hermans and Benjamin Schrauwen. 2013. Training and analysing deep recurrent neural networks. *Advances in Neural Information Processing Systems* 26 (2013), 190–198.
- [78] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2013. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning*. PMLR, 1310–1318.
- [79] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9, 8 (1997), 1735–1780.
- [80] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. 2014. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint:1412.3555* (2014).

- [81] Ilya Sutskever, Oriol Vinyals, and Quoc V. Le. 2014. Sequence to sequence learning with neural networks. *Advances in Neural Information Processing Systems 27* (2014), 3104–3112.
- [82] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. 2015. Long-term recurrent convolutional networks for visual recognition and description. In *IEEE Conference on Computer Vision and Pattern Recognition*. 2625–2634.
- [83] Andrej Karpathy and Li Fei-Fei. 2015. Deep visual-semantic alignments for generating image descriptions. In *IEEE Conference on Computer Vision and Pattern Recognition*. 3128–3137.
- [84] Yujin Tang, Jianfeng Xu, Kazunori Matsumoto, and Chihiro Ono. 2016. Sequence-to-sequence model with attention for time series classification. In *2016 IEEE 16th International Conference on Data Mining Workshops (ICDMW'16)*. IEEE, 503–510.
- [85] Pankaj Malhotra, Vishnu TV, Lovekesh Vig, Puneet Agarwal, and Gautam Shroff. 2017. TimeNet: Pre-trained deep recurrent neural network for time series classification. *arXiv preprint:1706.08838* (2017).
- [86] Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Samuel Harford. 2019. Multivariate LSTM-FCNs for time series classification. *Neural Networks* 116 (2019), 237–245.
- [87] Xuchao Zhang, Yifeng Gao, Jessica Lin, and Chang-Tien Lu. 2020. Tapnet: Multivariate time series classification with attentional prototypical network. In *AAAI Conference on Artificial Intelligence*, Vol. 34. 6845–6852.
- [88] Jingwei Zuo, Karine Zeitouni, and Yehia Taher. 2021. SMATE: Semi-supervised spatio-temporal representation learning on multivariate time series. In *2021 IEEE International Conference on Data Mining (ICDM'21)*. IEEE, 1565–1570.
- [89] Fazle Karim, Somshubra Majumdar, Houshang Darabi, and Shun Chen. 2017. LSTM fully convolutional networks for time series classification. *IEEE Access* 6 (2017), 1662–1669.
- [90] Sangdi Lin and George C. Runger. 2017. GCRNN: Group-constrained convolutional recurrent neural network. *IEEE Transactions on Neural Networks and Learning Systems* 29, 10 (2017), 4709–4718.
- [91] Ronald Mutegeki and Dong Seog Han. 2020. A CNN-LSTM approach to human activity recognition. In *2020 IEEE International Conference on Computational Intelligence and Communications Technologies. (ICAIIC'20)*. IEEE, 362–366.
- [92] Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. 2012. Understanding the exploding gradient problem. *CoRR, abs/1211.5063* 2, 417 (2012), 1.
- [93] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* 30 (2017).
- [94] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT 2019*, Vol. 1. Association for Computational Linguistics, Stroudsburg, PA, USA, 4171–4186. DOI : <http://dx.doi.org/10.18653/v1/N19-1423> arXiv:1810.04805
- [95] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint:2010.11929* (2020).
- [96] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyu Zhou, Wenhu Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. *Advances in Neural Information Processing Systems* 32 (2019).
- [97] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of AAAI*.
- [98] Demetres Kostas, Stephane Aroca-Ouellette, and Frank Rudzicz. 2021. BENDR: Using transformers and a contrastive self-supervised learning task to learn from massive amounts of EEG data. *Frontiers in Human Neuroscience* 15 (2021).
- [99] Ye Yuan, Guangxu Xun, Fenglong Ma, Yaqing Wang, Nan Du, Kebin Jia, Lu Su, and Aidong Zhang. 2018. Muvan: A multi-view attention network for multivariate temporal data. In *2018 IEEE International Conference on Data Mining (ICDM'18)*. IEEE, 717–726.
- [100] Tsung-Yu Hsieh, Suhang Wang, Yiwei Sun, and Vasant Honavar. 2021. Explainable multivariate time series classification: A deep neural network which learns to attend to important variables as well as time intervals. In *14th ACM International Conference on Web Search and Data Mining*. 607–615.
- [101] Wei Chen and Ke Shi. 2021. Multi-scale attention convolutional neural network for time series classification. *Neural Networks* 136 (2021), 126–140.
- [102] Ye Yuan, Guangxu Xun, Fenglong Ma, Qiuling Suo, Hongfei Xue, Kebin Jia, and Aidong Zhang. 2018. A novel channel-aware attention framework for multi-channel EEG seizure detection via multi-view deep learning. In *2018 IEEE EMBS International Conference on Biomedical & Health Informatics (BHI'18)*. IEEE, 206–209.
- [103] Yuxuan Liang, Songyu Ke, Junbo Zhang, Xiuwen Yi, and Yu Zheng. 2018. Geoman: Multi-level attention networks for geo-sensory time series prediction. In *IJCAI*, Vol. 2018. 3428–3434.
- [104] Jun Hu and Wendong Zheng. 2020. Multistage attention network for multivariate time series prediction. *Neurocomputing* 383 (2020), 122–137.

- [105] Xu Cheng, Peihua Han, Guoyuan Li, Shengyong Chen, and Houxiang Zhang. 2020. A novel channel and temporal-wise attention in convolutional networks for multivariate time series classification. *IEEE Access* 8 (2020), 212247–212257.
- [106] Zhiwen Xiao, Xin Xu, Huanlai Xing, Shouxi Luo, Penglin Dai, and Dawei Zhan. 2021. RTFN: A robust temporal feature network for time series classification. *Information Sciences* 571 (2021), 65–86.
- [107] Jingyuan Wang, Chen Yang, Xiaohan Jiang, and Junjie Wu. 2023. WHEN: A Wavelet-DTW hybrid attention network for heterogeneous time series analysis. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2361–2373.
- [108] Max Jaderberg, Karen Simonyan, Andrew Zisserman, et al. 2015. Spatial transformer networks. *Advances in Neural Information Processing Systems* 28 (2015), 2017–2025.
- [109] Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. 2018. CBAM: Convolutional block attention module. In *European Conference on Computer Vision*. 3–19.
- [110] Jie Hu, Li Shen, and Gang Sun. 2018. Squeeze-and-excitation networks. In *IEEE Conference on Computer Vision and Pattern Recognition*. 7132–7141.
- [111] Tian Wang, Zhaoying Liu, Ting Zhang, and Yujian Li. 2021. Time series classification based on multi-scale dynamic convolutional features and distance features. In *2021 2nd Asia Symposium on Signal Processing (ASSP'21)*. IEEE, 239–246.
- [112] Huan Song, Deepta Rajan, Jayaraman Thiagarajan, and Andreas Spanias. 2018. Attend and diagnose: Clinical time series analysis using attention models. In *AAAI Conference on Artificial Intelligence*, Vol. 32.
- [113] Can-can Jin and Xi Chen. 2021. An end-to-end framework combining time–frequency expert knowledge and modified Transformer networks for vibration signal classification. *Expert Systems with Applications* 171 (2021), 114570.
- [114] Carl Edward Rasmussen. 2004. *Gaussian Processes in Machine Learning*. Springer, Berlin, 63–71.
- [115] Tarek Allam Jr. and Jason D. McEwen. 2021. Paying attention to astronomical transients: Photometric classification with the time-series transformer. *arXiv preprint:2105.06178* (2021).
- [116] Minghao Liu, Shengqi Ren, Siyuan Ma, Jiahui Jiao, Yizhou Chen, Zhiguang Wang, and Wei Song. 2021. Gated transformer networks for multivariate time series classification. *arXiv preprint:2103.14438* (2021).
- [117] Bowen Zhao, Huanlai Xing, Xinhan Wang, Fuhong Song, and Zhiwen Xiao. 2022. Rethinking attention mechanism in time series classification. *arXiv preprint:2207.07564* (2022).
- [118] Yankun Ren, Longfei Li, Xinxing Yang, and Jun Zhou. 2022. AutoTransformer: Automatic transformer architecture design for time series classification. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, 143–155.
- [119] Ming Jin, Huan Yee Koh, Qingsong Wen, Daniele Zambon, Cesare Alippi, Geoffrey I. Webb, Irwin King, and Shirui Pan. 2023. A survey on graph neural networks for time series: Forecasting, classification, imputation, and anomaly detection. *arXiv* 14, 8 (July 2023), 1–27. arXiv:2307.03759 <http://arxiv.org/abs/2307.03759>
- [120] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and Philip S. Yu. 2021. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (Jan. 2021), 4–24. DOI: <http://dx.doi.org/10.1109/TNNLS.2020.2978386> arXiv:1901.00596
- [121] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2009. The graph neural network model. *IEEE Transactions on Neural Networks* 20, 1 (Jan. 2009), 61–80. DOI: <http://dx.doi.org/10.1109/TNN.2008.2005605>
- [122] Wenjie Xi, Arnav Jain, Li Zhang, and Jessica Lin. 2023. LB-SimTSC: An efficient similarity-aware graph neural network for semi-supervised time series classification. *arXiv* (Jan. 2023). arXiv:2301.04838
- [123] Huaiyuan Liu, Xianzhang Liu, Donghua Yang, Zhiyu Liang, Hongzhi Wang, Yong Cui, and Jun Gu. 2023. TodyNet: Temporal dynamic graph neural network for multivariate time series classification. *ArXiv abs/2304.05078* (2023). <https://api.semanticscholar.org/CorpusID:258059979>
- [124] Stefan Bloemheugel, Jurgen van den Hoogen, Dario Jozinović, Alberto Michelini, and Martin Atzmueller. 2023. Graph neural networks for multivariate time series regression with application to seismic data. *International Journal of Data Science and Analytics* 16, 3 (Sept. 2023), 317–332. DOI: <http://dx.doi.org/10.1007/s41060-022-00349-6> arXiv:2201.00818
- [125] Ziqiang Cheng, Yang Yang, Shuo Jiang, Wenjie Hu, Zhangchi Ying, Ziwei Chai, and Chunping Wang. 2021. Time2Graph+: Bridging time series and graph representation learning via multiple attentions. *IEEE Transactions on Knowledge and Data Engineering* 35, 2 (2021), 1–1. DOI: <http://dx.doi.org/10.1109/TKDE.2021.3094908>
- [126] Ian C. Covert, Balu Krishnan, Imad Najm, Jiening Zhan, Matthew Shore, John Hixson, and Ming Jack Po. 2019. Temporal graph convolutional networks for automatic seizure detection. In *Machine Learning for Healthcare Conference*. PMLR, 160–180.
- [127] Tengfei Song, Wenming Zheng, Peng Song, and Zhen Cui. 2020. EEG emotion recognition using dynamical graph convolutional neural networks. *IEEE Transactions on Affective Computing* 11, 3 (July 2020), 532–541. DOI: <http://dx.doi.org/10.1109/TAFFC.2018.2817622>

- [128] Ziyu Jia, Youfang Lin, Jing Wang, Ronghao Zhou, Xiaojun Ning, Yuanlai He, and Yaoshuai Zhao. 2020. GraphSleepNet: Adaptive spatial-temporal graph convolutional networks for sleep stage classification. In *IJCAI*. 1324–1330.
- [129] Zhengjing Ma, Gang Mei, Edoardo Prezioso, Zhongjian Zhang, and Nengxiong Xu. 2021. A deep learning approach using graph convolutional networks for slope deformation prediction based on time-series displacement data. *Neural Computing and Applications* 33, 21 (2021), 14441–14457.
- [130] Tianfu Li, Zhibin Zhao, Chuang Sun, Ruqiang Yan, and Xuefeng Chen. 2020. Multireceptive field graph convolutional networks for machine fault diagnosis. *IEEE Transactions on Industrial Electronics* 68, 12 (2020), 12739–12749.
- [131] D. Nhu, M. Janmohamed, P. Perucca, A. Gilligan, P. Kwan, T. O'Brien, C. W. Tan, and L. Kuhlmann. 2021. Graph convolutional network for generalized epileptiform abnormality detection on EEG. In *2021 IEEE Signal Processing in Medicine and Biology Symposium (SPMB'21)*. IEEE, 1–6.
- [132] Siyi Tang, Jared A. Dunnmon, Khaled Saab, Xuan Zhang, Qianying Huang, Florian Dubost, Daniel L. Rubin, and Christopher Lee-Messer. 2021. Self-supervised graph neural networks for improved electroencephalographic seizure analysis. In *10th International Conference on Learning Representations (ICLR'22)*, 1–23. arXiv:2104.08336
- [133] Xiang Zhang, Marko Zeman, Theodoros Tsiligkaridis, and Marinka Zitnik. 2021. Graph-guided network for irregularly sampled multivariate time series. In *10th International Conference on Learning Representations (ICLR'22)*, 1–21. arXiv:2110.05357
- [134] Alessandro Michele Censi, Dino Ienco, Yawogan Jean Eudes Gbodjo, Ruggero Gaetano Pensa, Roberto Interdonato, and Raffaele Gaetano. 2021. Attentive spatial temporal graph CNN for land cover mapping from multi temporal remote sensing data. *IEEE Access* 9 (2021), 23070–23082. DOI : <http://dx.doi.org/10.1109/ACCESS.2021.3055554>
- [135] Tiago Azevedo, Alexander Campbell, Rafael Romero-Garcia, Luca Passamonti, Richard A. I. Bethlehem, Pietro Liò, and Nicola Toschi. 2022. A deep graph neural network architecture for modelling spatio-temporal dynamics in resting-state functional MRI data. *Medical Image Analysis* 79 (July 2022), 102471. DOI : <http://dx.doi.org/10.1016/j.media.2022.102471>
- [136] Ziheng Duan, Haoyan Xu, Yueyang Wang, Yida Huang, Anni Ren, Zhongbin Xu, Yizhou Sun, and Wei Wang. 2022. Multivariate time-series classification with hierarchical variational graph pooling. *Neural Networks* 154 (Oct. 2022), 481–490. DOI : <http://dx.doi.org/10.1016/j.neunet.2022.07.032> arXiv:2010.05649
- [137] Daochen Zha, Kwei-herng Lai, Kaixiong Zhou, and Xia Hu. 2022. Towards similarity-aware time-series classification. In *Proceedings of the 2022 SIAM International Conference on Data Mining (SDM'22)*. 199–207. DOI : <http://dx.doi.org/10.1137/1.9781611977172.23>
- [138] Lukasz Tulczyjew, Michal Kawulok, Nicolas Longepe, Bertrand Le Saux, and Jakub Nalepa. 2022. Graph neural networks extract high-resolution cultivated land maps from sentinel-2 image series. *IEEE Geoscience and Remote Sensing Letters* 19 (2022), 1–5. DOI : <http://dx.doi.org/10.1109/LGRS.2022.3185407>
- [139] Le Sun, Chenyang Li, Bo Liu, and Yanchun Zhang. 2023. Class-driven graph attention network for multi-label time series classification in mobile health digital twins. *IEEE Journal on Selected Areas in Communications* 41, 10 (2023), 3267–3278. DOI : <http://dx.doi.org/10.1109/JSAC.2023.3310064>
- [140] Corentin Dufourg, Charlotte Pelletier, Stéphane May, and Sébastien Lefèvre. 2023. Graph dynamic earth net: Spatio-temporal graph benchmark for satellite image time series. In *2023 IEEE International Geoscience and Remote Sensing Symposium (IGARSS'23)*. IEEE, 7164–7167. DOI : <http://dx.doi.org/10.1109/IGARSS52108.2023.10281458>
- [141] Eamonn Keogh and Chotirat Ann Ratanamahatana. 2005. Exact indexing of dynamic time warping. *Knowledge Information Systems* 7, 3 (2005), 358–386. DOI : <http://dx.doi.org/10.1007/s10115-004-0154-9>
- [142] Thomas N. Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations (ICLR'17) - Conference Track Proceedings*. 1–14. arXiv:1609.02907
- [143] Ling Yang and Shenda Hong. 2022. Unsupervised time-series representation learning with iterative bilinear temporal-spectral fusion. In *ICML*. 25038–25054.
- [144] Aapo Hyvarinen and Hiroshi Morioka. 2016. Unsupervised feature extraction by time-contrastive learning and non-linear ICA. *Advances in Neural Information Processing Systems* 29 (2016), 3772–3780.
- [145] Jean-Yves Franceschi, Aymeric Dieuleveut, and Martin Jaggi. 2019. Unsupervised scalable representation learning for multivariate time series. *NeurIPS* 418, 32 (2019), 12 pages.
- [146] Sana Tonekaboni, Danny Eytan, and Anna Goldenberg. 2021. Unsupervised representation learning for time series with temporal neighborhood coding. *International Conference on Learning Representations*. <https://openreview.net/forum?id=8qDwejCuCN>
- [147] Kristoffer Wickstrøm, Michael Kampffmeyer, Karl Øyvind Mikalsen, and Robert Jenssen. 2022. Mixing up contrastive learning: Self-supervised representation learning for time series. *Pattern Recognition Letters* 155 (2022), 54–61.
- [148] Xinyu Yang, Zhenguo Zhang, and Rongyi Cui. 2022. TimeCLR: A self-supervised contrastive learning framework for univariate time series representation. *Knowledge-based Systems* 245 (2022), 108606.
- [149] Xiang Zhang, Ziyuan Zhao, Theodoros Tsiligkaridis, and Marinka Zitnik. 2022. Self-supervised contrastive pre-training for time series via time-frequency consistency. In *Proceedings of Neural Information Processing Systems (NeurIPS'22)*.

- [150] Qianwen Meng, Hangwei Qian, Yong Liu, Lizhen Cui, Yonghui Xu, and Zhiqi Shen. 2023. MHCCL: Masked hierarchical cluster-wise contrastive learning for multivariate time series. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 9153–9161.
- [151] Ranak Roy Chowdhury, Xiyuan Zhang, Jingbo Shang, Rajesh K. Gupta, and Dezhi Hong. 2022. TARNet: Task-aware reconstruction for time-series transformer. In *28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 14–18.
- [152] Mingyue Cheng, Qi Liu, Zhiding Liu, Hao Zhang, Rujiao Zhang, and Enhong Chen. 2023. TimeMAE: Self-supervised representations of time series with decoupled masked autoencoders. *arXiv preprint arXiv:2303.00320* (2023).
- [153] Wenrui Zhang, Ling Yang, Shijia Geng, and Shenda Hong. 2023. Self-supervised time series representation learning via cross reconstruction transformer. *IEEE Transactions on Neural Networks and Learning Systems* (2023), 1–10.
- [154] Ali Ismail-Fawaz, Maxime Devanne, Stefano Berretti, Jonathan Weber, and Germain Forestier. 2023. Finding foundation models for time series classification with a PreText task. *arXiv preprint arXiv:2311.14534* (2023).
- [155] Connor Shorten and Taghi M. Khoshgoftaar. 2019. A survey on image data augmentation for deep learning. *Journal of Big Data* 6, 1 (2019), 1–48.
- [156] Terry T. Um, Franz M. J. Pfister, Daniel Pichler, Satoshi Endo, Muriel Lang, Sandra Hirche, Urban Fietzek, and Dana Kulić. 2017. Data augmentation of wearable sensor data for Parkinson’s disease monitoring using convolutional neural networks. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction*. 216–220.
- [157] Khandakar M. Rashid and Joseph Louis. 2019. Window-warping: A time series data augmentation of IMU data for construction equipment activity identification. In *Proceedings of the International Symposium on Automation and Robotics in Construction (ISARC’19)*, Vol. 36. IAARC Publications, 651–657.
- [158] Brian Kenji Iwana and Seiichi Uchida. 2021. Time series data augmentation for neural networks by time warping with a discriminative teacher. In *2020 25th International Conference on Pattern Recognition (ICPR’21)*. IEEE, 3558–3565.
- [159] Thai-Son Nguyen, Sebastian Stueker, Jan Niehues, and Alex Waibel. 2020. Improving sequence-to-sequence speech recognition training with on-the-fly data augmentation. In *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP’20)*. IEEE, 7689–7693.
- [160] Bhavik Vachhani, Chitralekha Bhat, and Sunil Kumar Kopparapu. 2018. Data augmentation using healthy speech for dysarthric speech recognition. In *Interspeech*. 471–475.
- [161] Jingkun Gao, Xiaomin Song, Qingsong Wen, Pichao Wang, Liang Sun, and Huan Xu. 2020. RobustTAD: Robust Time Series Anomaly Detection via Decomposition and Convolutional Neural Networks. [arXiv:cs.LG/2002.09545](https://arxiv.org/abs/cs.LG/2002.09545) (2020).
- [162] Zhicheng Cui, Wenlin Chen, and Yixin Chen. 2016. Multi-Scale Convolutional Neural Networks for Time Series Classification. [arXiv:cs.CV/1603.06995](https://arxiv.org/abs/cs.CV/1603.06995) (2016).
- [163] Arthur Le Guennec, Simon Malinowski, and Romain Tavenard. 2016. Data augmentation for time series classification using convolutional neural networks. In *ECML/PKDD on Advanced Analytics and Learning on Temporal Data*.
- [164] Germain Forestier, François Petitjean, Hoang Anh Dau, Geoffrey I. Webb, and Eamonn Keogh. 2017. Generating synthetic time series to augment sparse datasets. In *2017 IEEE International Conference on Data Mining (ICDM’17)*. IEEE, 865–870.
- [165] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2018. Data Augmentation Using Synthetic Data for Time Series Classification with Deep Residual Networks. [arXiv:cs.CV/1808.02455](https://arxiv.org/abs/cs.CV/1808.02455) (2018).
- [166] Tsegamlak Terefe, Maxime Devanne, Jonathan Weber, Dereje Hailemariam, and Germain Forestier. 2020. Time series averaging using multi-tasking autoencoder. In *2020 IEEE 32nd International Conference on Tools with Artificial Intelligence (ICTAI’20)*. IEEE, 1065–1072.
- [167] Brian Kenji Iwana and Seiichi Uchida. 2021. An empirical survey of data augmentation for time series classification with neural networks. *Plos One* 16, 7 (2021), e0254841.
- [168] Gautier Piella, Maxime Devanne, Jonathan Weber, Lhassane Idoumghar, and Germain Forestier. 2022. Data augmentation for time series classification with deep learning models. In *International Workshop on Advanced Analytics and Learning on Temporal Data*. Springer, 117–132.
- [169] Zijun Gao, Lingbo Li, and Tianhua Xu. 2023. Data augmentation for time-series classification: An extensive empirical study and comprehensive survey. *arXiv preprint arXiv:2310.10060* (2023).
- [170] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 248–255.
- [171] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2018. Transfer learning for time series classification. In *2018 IEEE International Conference on Big Data (Big Data’18)*. IEEE, 1367–1376.

- [172] Stephan Spiegel. 2016. Transfer learning for time series classification in dissimilarity spaces. *Proceedings of AALTD* 78 (2016).
- [173] Frédéric Li, Kimiaki Shirahama, Muhammad Adeel Nisar, Xinyu Huang, and Marcin Grzegorzec. 2020. Deep transfer learning for time series data based on sensor modality classification. *Sensors* 20, 15 (2020), 4271.
- [174] Yarden Rotem, Nathaniel Shimoni, Lior Rokach, and Bracha Shapira. 2022. Transfer learning for time series classification using synthetic data generation. In *International Symposium on Cyber Security, Cryptology, and Machine Learning*. Springer, 232–246.
- [175] Ayantha Senanayaka, Abdullah Al Mamun, Glenn Bond, Wenmeng Tian, Haifeng Wang, Sara Fuller, T. C. Falls, Shahram Rahimi, and Linkan Bian. 2022. Similarity-based multi-source transfer learning approach for time series classification. *International Journal of Prognostics and Health Management* 13, 2 (2022).
- [176] Kathan Kashiparekh, Jyoti Narwariya, Pankaj Malhotra, Lovekesh Vig, and Gautam Shroff. 2019. ConvtimeNet: A pre-trained deep convolutional neural network for time series classification. In *2019 International Joint Conference on Neural Networks (IJCNN'19)*. IEEE, 1–8.
- [177] D. Merlin Praveena, D. Angelin Sarah, and S. Thomas George. 2022. Deep learning techniques for EEG signal applications—A review. *IETE Journal of Research* 68, 4 (2022), 3030–3037. DOI: <http://dx.doi.org/10.1080/03772063.2020.1749143>
- [178] Xinwen Liu, Huan Wang, Zongjin Li, and Lang Qin. 2021. Deep learning in ECG diagnosis: A review. *Knowledge-based Systems* 227 (2021), 107187. DOI: <http://dx.doi.org/10.1016/j.knsys.2021.107187>
- [179] Nur'atiah Zaini, Lee Woen Ean, Ali Najah Ahmed, and Marlinda Abdul Malek. 2022. A systematic literature review of deep learning neural network for time series air quality forecasting. *Environmental Science and Pollution Research* 29, 4 (Jan. 2022), 4958–4990. DOI: <http://dx.doi.org/10.1007/s11356-021-17442-1>
- [180] Bo Zhang, Yi Rong, Ruihan Yong, Dongming Qin, Maozhen Li, Guojian Zou, and Jianguo Pan. 2022. Deep learning for air pollutant concentration prediction: A review. *Atmospheric Environment* 290 (Dec. 2022), 119347. DOI: <http://dx.doi.org/10.1016/j.atmosenv.2022.119347>
- [181] Gyungmin Toh and Junhong Park. 2020. Review of vibration-based structural health monitoring using deep learning. *Applied Sciences* 10, 5 (2020), 1680. DOI: <http://dx.doi.org/10.3390/app10051680>
- [182] Nikhil M. Thoppil, V. Vasu, and C. S. P. Rao. 2021. Deep learning algorithms for machinery health prognostics using time-series data: A review. *Journal of Vibration Engineering & Technologies* 9, 6 (Sept. 2021), 1123–1145. DOI: <http://dx.doi.org/10.1007/s42417-021-00286-x>
- [183] Lei Ren, Zidi Jia, Yuanjun Laili, and Di Huang. 2023. Deep learning for time-series prediction in IIoT: Progress, challenges, and prospects. *IEEE Transactions on Neural Networks and Learning Systems* PP (2023), 1–20. DOI: <http://dx.doi.org/10.1109/TNNLS.2023.3291371>
- [184] Yassine Himeur, Khalida Ghanem, Abdullah Alsalemi, Faycal Bensaali, and Abbes Amira. 2021. Artificial intelligence based anomaly detection of energy consumption in buildings: A review, current trends and new perspectives. *Applied Energy* 287 (2021), 116601. DOI: <http://dx.doi.org/10.1016/j.apenergy.2021.116601>
- [185] Dan Stowell. 2022. Computational bioacoustics with deep learning: A review and roadmap. *PeerJ* 10 (March 2022), e13152. DOI: <http://dx.doi.org/10.7717/peerj.13152>
- [186] Neha Gupta, Suneet K. Gupta, Rajesh K. Pathak, Vanita Jain, Parisa Rashidi, and Jasjit S. Suri. 2022. Human activity recognition in artificial intelligence framework: A narrative review. *Artificial Intelligence Review* 55, 6 (Aug. 2022), 4755–4808. DOI: <http://dx.doi.org/10.1007/s10462-021-10116-x>
- [187] E. Ramanujam, Thinagaran Perumal, and S. Padmavathi. 2021. Human activity recognition with smartphone and wearable sensors using deep learning techniques: A review. *IEEE Sensors Journal* 21, 12 (June 2021), 13029–13040. DOI: <http://dx.doi.org/10.1109/JSEN.2021.3069927>
- [188] Jeffrey W. Lockhart, Tony Pulickal, and Gary M. Weiss. 2012. Applications of mobile activity recognition. In *2012 ACM Conference on Ubiquitous Computing (UbiComp'12)*. ACM Press, New York, NY, USA, 1054. DOI: <http://dx.doi.org/10.1145/2370216.2370441>
- [189] Emmanuel Munguia Tapia, Stephen S. Intille, and Kent Larson. 2004. Activity recognition in the home using simple and ubiquitous sensors. In *Lecture Notes in Computer Science*. Vol. 3001. Springer, Berlin, 158–175. DOI: http://dx.doi.org/10.1007/978-3-540-24646-6_10
- [190] Yu Kong and Yun Fu. 2022. Human action recognition and prediction: A survey. *International Journal of Computer Vision* 130, 5 (May 2022), 1366–1401. DOI: <http://dx.doi.org/10.1007/s11263-022-01594-9> arXiv:1806.11230
- [191] Hong-Bo Zhang, Yi-Xiang Zhang, Bineng Zhong, Qing Lei, Lijie Yang, Ji-Xiang Du, and Duan-Sheng Chen. 2019. A comprehensive survey of vision-based human action recognition methods. *Sensors* 19, 5 (Feb. 2019), 1005. DOI: <http://dx.doi.org/10.3390/s19051005>
- [192] Francisco Ordóñez and Daniel Roggen. 2016. Deep convolutional and LSTM recurrent neural networks for multi-modal wearable activity recognition. *Sensors* 16, 1 (Jan. 2016), 115. DOI: <http://dx.doi.org/10.3390/s16010115>

- [193] Attila Reiss and Didier Stricker. 2012. Introducing a new benchmarked dataset for activity monitoring. In *16th International Symposium on Wearable Computers*. 108–109. DOI : <http://dx.doi.org/10.1109/ISWC.2012.13>
- [194] Mi Zhang and Alexander A. Sawchuk. 2012. USC-HAD: A daily activity dataset for ubiquitous activity recognition using wearable sensors. In *2012 ACM Conference on Ubiquitous Computing (UbiComp'12)*. ACM Press, New York, NY, USA, 1036. DOI : <http://dx.doi.org/10.1145/2370216.2370438>
- [195] Daniel Roggen, Alberto Calatroni, Mirco Rossi, Thomas Holleczeck, Kilian Förster, Gerhard Tröster, Paul Lukowicz, David Bannach, Gerald Pirkel, et al. 2010. Collecting complex activity datasets in highly rich networked sensor environments. In *7th International Conference on Networked Sensing Systems*. IEEE, 233–240.
- [196] Timo Sztyler, Heiner Stuckenschmidt, and Wolfgang Petrich. 2017. Position-aware activity recognition with wearable devices. *Pervasive and Mobile Computing* 38 (July 2017), 281–295. DOI : <http://dx.doi.org/10.1016/j.pmcj.2017.01.008>
- [197] Oscar D. Lara and Miguel A. Labrador. 2013. A survey on human activity recognition using wearable sensors. *IEEE Communications Surveys & Tutorials* 15, 3 (2013), 1192–1209. DOI : <http://dx.doi.org/10.1109/SURV.2012.110112.00192>
- [198] Fuqiang Gu, Mu-Huan Chung, Mark Chignell, Shahrokh Valaee, Baoding Zhou, and Xue Liu. 2022. A survey on deep learning for human activity recognition. *Computer Surveys* 54, 8 (Nov. 2022), 1–34. DOI : <http://dx.doi.org/10.1145/3472290>
- [199] Nils Y. Hammerla, Shane Halloran, and Thomas Ploetz. 2016. Deep, convolutional, and recurrent models for human activity recognition using wearables. In *International Joint Conference on Artificial Intelligence (IJCAI'16)*. 1533–1540. arXiv:1604.08880
- [200] Ming Zeng, Le T. Nguyen, Bo Yu, Ole J. Mengshoel, Jiang Zhu, Pang Wu, and Joy Zhang. 2014. Convolutional neural networks for human activity recognition using mobile sensors. In *6th International Conference on Mobile Computing, Applications and Services*. ICST, 718–737. DOI : <http://dx.doi.org/10.4108/icst.mobica.2014.257786>
- [201] Wenchao Jiang and Zhaozheng Yin. 2015. Human activity recognition using wearable sensors by deep convolutional neural networks. In *23rd ACM International Conference on Multimedia*. ACM, New York, NY, USA, 1307–1310. DOI : <http://dx.doi.org/10.1145/2733373.2806333>
- [202] Jian Bo Yang, Minh Nhut Nguyen, Phyo Phyo San, Xiao Li Li, and Shonali Krishnaswamy. 2015. Deep convolutional neural networks on multichannel time series for human activity recognition. In *International Joint Conference on Artificial Intelligence (IJCAI '15)*, 3995–4001. DOI : <http://dx.doi.org/10.5555/2832747.2832806>
- [203] Charissa Ann Ronao and Sung-Bae Cho. 2016. Human activity recognition with smartphone sensors using deep learning neural networks. *Expert Systems with Applications* 59 (Oct. 2016), 235–244. DOI : <http://dx.doi.org/10.1016/j.eswa.2016.04.032>
- [204] Yu Guan and Thomas Plötz. 2017. Ensembles of deep LSTM learners for activity recognition using wearables. *ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 1, 2 (June 2017), 1–28. DOI : <http://dx.doi.org/10.1145/3090076> arXiv:1703.09370
- [205] Song-Mi Lee, Sang Min Yoon, and Heeryon Cho. 2017. Human activity recognition from accelerometer data using Convolutional Neural Network. In *2017 IEEE International Conference on Big Data and Smart Computing (BigComp'17)*, Vol. 83. IEEE, 131–134. DOI : <http://dx.doi.org/10.1109/BIGCOMP.2017.7881728>
- [206] Abdulmajid Murad and Jae-Young Pyun. 2017. Deep recurrent neural networks for human activity recognition. *Sensors* 17, 11 (Nov. 2017), 2556. DOI : <http://dx.doi.org/10.3390/s17112556>
- [207] Andrey Ignatov. 2018. Real-time human activity recognition from accelerometer data using Convolutional Neural Networks. *Applied Soft Computing* 62 (Jan. 2018), 915–922. DOI : <http://dx.doi.org/10.1016/j.asoc.2017.09.027>
- [208] Fernando Moya Rueda, René Grzeszick, Gernot Fink, Sascha Feldhorst, and Michael ten Hoppel. 2018. Convolutional neural networks for human activity recognition using body-worn sensors. *Informatics* 5, 2 (May 2018), 26. DOI : <http://dx.doi.org/10.3390/informatics5020026>
- [209] Rui Yao, Guosheng Lin, Qinfeng Shi, and Damith C. Ranasinghe. 2018. Efficient dense labelling of human activity sequences from wearables using fully convolutional networks. *Pattern Recognition* 78 (June 2018), 252–266. DOI : <http://dx.doi.org/10.1016/j.patcog.2017.12.024> arXiv:1702.06212
- [210] Ming Zeng, Haoxiang Gao, Tong Yu, Ole J. Mengshoel, Helge Langseth, Ian Lane, and Xiaobing Liu. 2018. Understanding and improving recurrent networks for human activity recognition by continuous attention. In *ACM International Symposium on Wearable Computers*. New York, NY, USA, 56–63. DOI : <http://dx.doi.org/10.1145/3267242.3267286> arXiv:1810.04038
- [211] Haojie Ma, Wenzhong Li, Xiao Zhang, Songcheng Gao, and Sanglu Lu. 2019. AttnSense: Multi-level attention mechanism for multimodal human activity recognition. In *28th International Joint Conference on Artificial Intelligence*. 3109–3115. DOI : <http://dx.doi.org/10.24963/ijcai.2019/431>
- [212] Cheng Xu, Duo Chai, Jie He, Xiaotong Zhang, and Shihong Duan. 2019. InnoHAR: A deep neural network for complex human activity recognition. *IEEE Access* 7 (2019), 9893–9902. DOI : <http://dx.doi.org/10.1109/ACCESS.2018.2890675>
- [213] Haoxi Zhang, Zhiwen Xiao, Juan Wang, Fei Li, and Edward Szczerbicki. 2020. A novel IoT-perceptive human activity recognition (HAR) approach using multihead convolutional attention. *IEEE Internet of Things Journal* 7, 2 (Feb. 2020), 1072–1080. DOI : <http://dx.doi.org/10.1109/JIOT.2019.2949715>

- [214] Sravan Kumar Challa, Akhilesh Kumar, and Vijay Bhaskar Semwal. 2021. A multibranch CNN-BiLSTM model for human activity recognition using wearable sensor data. *The Visual Computer* 38, 0123456789 (Aug. 2021), 4095–4109. DOI: <http://dx.doi.org/10.1007/s00371-021-02283-3>
- [215] Sakorn Mekruksavanich and Anuchit Jitpattanakul. 2021. Deep convolutional neural network with RNNs for complex activity recognition using wrist-worn wearable sensor data. *Electronics* 10, 14 (July 2021), 1685. DOI: <http://dx.doi.org/10.3390/electronics10141685>
- [216] Ling Chen, Xiaozie Liu, Liangying Peng, and Menghan Wu. 2021. Deep learning based multimodal complex human activity recognition using wearable devices. *Applied Intelligence* 51, 6 (June 2021), 4029–4042. DOI: <http://dx.doi.org/10.1007/s10489-020-02005-7>
- [217] Sakorn Mekruksavanich and Anuchit Jitpattanakul. 2021. LSTM networks using smartphone data for sensor-based human activity recognition in smart homes. *Sensors* 21, 5 (Feb. 2021), 1636. DOI: <http://dx.doi.org/10.3390/s21051636>
- [218] Sakorn Mekruksavanich and Anuchit Jitpattanakul. 2021. Biometric user identification based on human activity recognition using wearable sensors: An experiment using deep learning models. *Electronics* 10, 3 (Jan. 2021), 308. DOI: <http://dx.doi.org/10.3390/electronics10030308>
- [219] Ohoud Nafea, Wadood Abdul, Ghulam Muhammad, and Mansour Alsulaiman. 2021. Sensor-based human activity recognition with spatio-temporal deep learning. *Sensors* 21, 6 (March 2021), 2141. DOI: <http://dx.doi.org/10.3390/s21062141>
- [220] Satya P. Singh, Madan Kumar Sharma, Aime Lay-Ekuakille, Deepak Gangwar, and Sukrit Gupta. 2021. Deep ConvLSTM with self-attention for human activity decoding using wearable sensors. *IEEE Sensors Journal* 21, 6 (March 2021), 8575–8582. DOI: <http://dx.doi.org/10.1109/JSEN.2020.3045135> arXiv:2005.00698
- [221] Xing Wang, Lei Zhang, Wenbo Huang, Shuoyuan Wang, Hao Wu, Jun He, and Aiguo Song. 2022. Deep convolutional networks with tunable speed–accuracy tradeoff for human activity recognition using wearables. *IEEE Transactions on Instrumentation and Measurement* 71 (2022), 1–12. DOI: <http://dx.doi.org/10.1109/TIM.2021.3132088>
- [222] Shige Xu, Lei Zhang, Wenbo Huang, Hao Wu, and Aiguo Song. 2022. Deformable convolutional networks for multimodal human activity recognition using wearable sensors. *IEEE Transactions on Instrumentation and Measurement* 71 (2022), 1–14. DOI: <http://dx.doi.org/10.1109/TIM.2022.3158427>
- [223] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. 2017. Deformable convolutional networks. In *2017 IEEE International Conference on Computer Vision (ICCV'17)*. 764–773. DOI: <http://dx.doi.org/10.1109/ICCV.2017.89> arXiv:1703.06211
- [224] Michael A. Wulder, Joanne C. White, Samuel N. Goward, Jeffrey G. Masek, James R. Irons, Martin Herold, Warren B. Cohen, Thomas R. Loveland, and Curtis E. Woodcock. 2008. Landsat continuity: Issues and opportunities for land cover monitoring. *Remote Sensing of Environment* 112, 3 (March 2008), 955–969. DOI: <http://dx.doi.org/10.1016/j.rse.2007.07.004>
- [225] William Emery and Adriano Camps. 2017. Basic electromagnetic concepts and applications to optical sensors. In *Introduction to Satellite Remote Sensing*, William Emery and Adriano Camps (Eds.). Elsevier, Chapter 2, 43–83. DOI: <http://dx.doi.org/10.1016/B978-0-12-809254-5.00002-6>
- [226] Noel Gorelick, Matt Hancher, Mike Dixon, Simon Ilyushchenko, David Thau, and Rebecca Moore. 2017. Google earth engine: Planetary-scale geospatial analysis for everyone. *Remote Sensing of Environment* 202 (Dec. 2017), 18–27. DOI: <http://dx.doi.org/10.1016/j.rse.2017.06.031>
- [227] Grégory Giuliani, Bruno Chatenoux, Andrea De Bono, Denisa Rodila, Jean-Philippe Richard, Karin Allenbach, Hy Dao, and Pascal Peduzzi. 2017. Building an Earth observations data cube: Lessons learned from the Swiss data cube (SDC) on generating analysis ready data (ARD). *Big Earth Data* 1, 1–2 (Dec. 2017), 100–117. DOI: <http://dx.doi.org/10.1080/20964471.2017.1398903>
- [228] Adam Lewis, Simon Oliver, Leo Lymburner, Ben Evans, Lesley Wyborn, Norman Mueller, Gregory Raevksi, Jeremy Hooke, Rob Woodcock, Joshua Sixsmith, et al. 2017. The Australian geoscience data cube—Foundations and lessons learned. *Remote Sensing of Environment* 202 (2017), 276–292.
- [229] Dino Ienco, Yawogan Jean Eudes Gbodjo, Roberto Interdonato, and Raffaele Gaetano. 2020. Attentive weakly supervised land cover mapping for object-based satellite image time series data with spatial interpretation. *arXiv* (2020), 1–12. arXiv:2004.14672
- [230] Vivien Sainte Fare Garnot, Loïc Landrieu, Sebastien Giordano, and Nesrine Chehata. 2020. Satellite image time series classification with pixel-set encoders and temporal self-attention. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'20)*. IEEE, 12322–12331. DOI: <http://dx.doi.org/10.1109/CVPR42600.2020.01234>
- [231] Anurag Kulshrestha, Ling Chang, and Alfred Stein. 2022. Use of LSTM for sinkhole-related anomaly detection and classification of InSAR deformation time series. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 15 (2022), 4559–4570. DOI: <http://dx.doi.org/10.1109/JSTARS.2022.3180994>
- [232] Yifang Ban, Puzhao Zhang, Andrea Nascetti, Alexandre R. Bevington, and Michael A. Wulder. 2020. Near real-time wildfire progression monitoring with Sentinel-1 SAR time series and deep learning. *Scientific Reports* 10, 1 (Dec. 2020), 1322. DOI: <http://dx.doi.org/10.1038/s41598-019-56967-x>

- [233] C. Rambour, N. Audebert, E. Koeniguer, B. Le Saux, M. Crucianu, and M. Datcu. 2020. Flood detection in time series of optical and SAR images. *International Archives of Photogrammetry, Remote Sensing, & Spatial Information Sciences* XLIII-B2-2, B2 (Aug. 2020), 1343–1346. DOI : <http://dx.doi.org/10.5194/isprs-archives-XLIII-B2-2020-1343-2020>
- [234] G. Kamdem De Teyou, Y. Tarabalka, I. Manighetti, R. Almar, and S. Tripodi. 2020. Deep neural networks for automatic extraction of features in time series optical satellite images. *International Archives of Photogrammetry, Remote Sensing, & Spatial Information Sciences* 43 (2020), 1529–1535.
- [235] Bruno Menini Matosak, Leila Maria Garcia Fonseca, Evandro Carrijo Taquary, Raian Vargas Maretto, Hugo Do Nascimento Bendini, and Marcos Adami. 2022. Mapping deforestation in Cerrado based on hybrid deep learning architecture and medium spatial resolution satellite time series. *Remote Sensing* 14, 1 (2022), 1–22. DOI : <http://dx.doi.org/10.3390/rs14010209>
- [236] Dinh Ho Tong Minh, Dino Ienco, Raffaele Gaetano, Nathalie Lalande, Emile Ndikumana, Faycal Osman, and Pierre Maurel. 2018. Deep recurrent neural networks for winter vegetation quality mapping via multitemporal SAR Sentinel-1. *IEEE Geoscience and Remote Sensing Letters* 15, 3 (March 2018), 464–468. DOI : <http://dx.doi.org/10.1109/LGRS.2018.2794581>
- [237] Pia Labenski, Michael Ewald, Sebastian Schmidlein, and Fabian Ewald Fassnacht. 2022. Classifying surface fuel types based on forest stand photographs and satellite time series using deep learning. *International Journal of Applied Earth Observation and Geoinformation* 109 (May 2022), 102799. DOI : <http://dx.doi.org/10.1016/j.jag.2022.102799>
- [238] Krishna Rao, A. Park Williams, Jacqueline Fortin Flefil, and Alexandra G. Konings. 2020. SAR-enhanced mapping of live fuel moisture content. *Remote Sensing of Environment* 245 (2020), 111797. DOI : <http://dx.doi.org/10.1016/j.rse.2020.111797>
- [239] Liujun Zhu, Geoffrey I. Webb, Marta Yebra, Gianluca Scortechini, Lynn Miller, and François Petitjean. 2021. Live fuel moisture content estimation from MODIS: A deep learning approach. *ISPRS Journal of Photogrammetry and Remote Sensing* 179 (Sept. 2021), 81–91. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2021.07.010>
- [240] Lynn Miller, Liujun Zhu, Marta Yebra, Christoph Rüdiger, and Geoffrey I Webb. 2022. Multi-modal temporal CNNs for live fuel moisture content estimation. *Environmental Modelling & Software* 156 (Oct. 2022), 105467. DOI : <http://dx.doi.org/10.1016/j.envsoft.2022.105467>
- [241] Jiangjian Xie, Tao Qi, Wanjun Hu, Huaguo Huang, Beibei Chen, and Junguo Zhang. 2022. Retrieval of live fuel moisture content based on multi-source remote sensing data and ensemble deep learning model. *Remote Sensing* 14, 17 (Sept. 2022), 4378. DOI : <http://dx.doi.org/10.3390/rs14174378>
- [242] Kamel Lahssini, Florian Teste, Karun Reuel Dayal, Sylvie Durrieu, Dino Ienco, and Jean-Matthieu Monnet. 2022. Combining LiDAR metrics and Sentinel-2 imagery to estimate basal area and wood volume in complex forest environment via neural networks. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 15 (2022), 4337–4348. DOI : <http://dx.doi.org/10.1109/JSTARS.2022.3175609>
- [243] Jie Sun, Zulong Lai, Liping Di, Ziheng Sun, Jianbin Tao, and Yonglin Shen. 2020. Multilevel deep learning network for county-level corn yield estimation in the U.S. Corn Belt. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 13 (2020), 5048–5060. DOI : <http://dx.doi.org/10.1109/JSTARS.2020.3019046>
- [244] Zhengtao Li, Guokun Chen, and Tianxu Zhang. 2019. Temporal attention networks for multitemporal multisensor crop classification. *IEEE Access* 7 (2019), 134677–134690. DOI : <http://dx.doi.org/10.1109/ACCESS.2019.2939152>
- [245] Zhengtao Li, Gang Zhou, and Qiong Song. 2020. A temporal group attention approach for multitemporal multisensor crop classification. *Infrared Physics and Technology* 105 (2020), 103152. DOI : <http://dx.doi.org/10.1016/j.infrared.2019.103152>
- [246] Shunping Ji, Chi Zhang, Anjian Xu, Yun Shi, and Yulin Duan. 2018. 3D convolutional neural networks for crop classification with multi-temporal remote sensing images. *Remote Sensing* 10, 2 (Jan. 2018), 75. DOI : <http://dx.doi.org/10.3390/rs10010075>
- [247] Jinfan Xu, Yue Zhu, Renhai Zhong, Zhixian Lin, Jialu Xu, Hao Jiang, Jingfeng Huang, Haifeng Li, and Tao Lin. 2020. DeepCropMapping: A multi-temporal deep learning approach with improved spatial generalizability for dynamic corn and soybean mapping. *Remote Sensing of Environment* 247 (Sept. 2020), 111946. DOI : <http://dx.doi.org/10.1016/j.rse.2020.111946>
- [248] Valentin Barriere and Martin Claverie. 2022. Multimodal crop type classification fusing multi-spectral satellite time series with farmers crop rotations and local crop distribution. *arXiv preprint:2208.10838* (2022).
- [249] Vivien Sainte Fare Garnot and Loic Landrieu. 2020. Lightweight temporal self-attention for classifying satellite images time series. In *Lecture Notes in Computer Science*. Vol. 12588 LNAI. Springer International Publishing, 171–181. DOI : http://dx.doi.org/10.1007/978-3-030-65742-0_12
- [250] Stella Ofori-Ampofo, Charlotte Pelletier, and Stefan Lang. 2021. Crop type mapping from optical and radar time series using attention-based deep learning. *Remote Sensing* 13, 22 (Nov. 2021), 4668. DOI : <http://dx.doi.org/10.3390/rs13224668>

- [251] Yuan Yuan and Lei Lin. 2021. Self-Supervised pretraining of transformers for satellite image time series classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 14 (2021), 474–487. DOI : <http://dx.doi.org/10.1109/JSTARS.2020.3036602>
- [252] Nicola Di Mauro, Antonio Vergari, Teresa Maria Altomare Basile, Fabrizio G. Ventola, and Floriana Esposito. 2017. End-to-end learning of deep spatio-temporal representations for satellite image time series classification. In *DC@PKDD/ECML*.
- [253] Nataliia Kussul, Mykola Lavreniuk, Sergii Skakun, and Andrii Shelestov. 2017. Deep learning classification of land cover and crop types using remote sensing data. *IEEE Geoscience and Remote Sensing Letters* 14, 5 (May 2017), 778–782. DOI : <http://dx.doi.org/10.1109/LGRS.2017.2681128>
- [254] Charlotte Pelletier, Geoffrey Webb, and François Petitjean. 2019. Temporal convolutional neural network for the classification of satellite image time series. *Remote Sensing* 11, 5 (March 2019), 523. DOI : <http://dx.doi.org/10.3390/rs11050523>
- [255] Peng Dou, Huanfeng Shen, Zhiwei Li, and Xiaobin Guan. 2021. Time series remote sensing image classification framework using combination of deep learning and multiple classifiers system. *International Journal of Applied Earth Observation and Geoinformation* 103 (2021), 102477. DOI : <http://dx.doi.org/10.1016/j.jag.2021.102477>
- [256] Dino Ienco, Roberto Interdonato, Raffaele Gaetano, and Dinh Ho Tong Minh. 2019. Combining Sentinel-1 and Sentinel-2 satellite image time series for land cover mapping via a multi-source deep learning architecture. *ISPRS Journal of Photogrammetry and Remote Sensing* 158 (2019), 11–22. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2019.09.016>
- [257] Roberto Interdonato, Dino Ienco, Raffaele Gaetano, and Kenji Ose. 2019. DuPLO: A DUal view Point deep Learning architecture for time series classification. *ISPRS Journal of Photogrammetry and Remote Sensing* 149 (March 2019), 91–104. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2019.01.011>
- [258] Marc Rußwurm and Marco Körner. 2018. Multi-temporal land cover classification with sequential recurrent encoders. *ISPRS International Journal of Geo-Information* 7, 4 (March 2018), 129. DOI : <http://dx.doi.org/10.3390/ijgi7040129>
- [259] Andrei Stoian, Vincent Poulain, Jordi Inglada, Victor Poughon, and Dawa Derksen. 2019. Land cover maps production with high resolution satellite image time series and convolutional neural networks: Adaptations and limits for operational systems. *Remote Sensing* 11, 17 (2019), 1–26. DOI : <http://dx.doi.org/10.3390/rs11171986>
- [260] Dino Ienco, Raffaele Gaetano, Claire Dupaquier, and Pierre Maurel. 2017. Land cover classification via multitemporal spatial data by deep recurrent neural networks. *IEEE Geoscience and Remote Sensing Letters* 14, 10 (Oct. 2017), 1685–1689. DOI : <http://dx.doi.org/10.1109/LGRS.2017.2728698>
- [261] Yawogan Jean Eudes Gbodio, Dino Ienco, Louise Leroux, Roberto Interdonato, Raffaele Gaetano, and Babacar Ndao. 2020. Object-based multi-temporal and multi-source land cover mapping leveraging hierarchical class relationships. *Remote Sensing* 12, 17 (Aug. 2020), 2814. DOI : <http://dx.doi.org/10.3390/rs12172814>
- [262] Dino Ienco, Raffaele Gaetano, Roberto Interdonato, Kenji Ose, and Dinh Ho Tong Minh. 2019. Combining Sentinel-1 and Sentinel-2 time series via RNN for object-based land cover classification. In *2019 IEEE International Geoscience and Remote Sensing Symposium (IGARSS'19)*. IEEE, 4881–4884. DOI : <http://dx.doi.org/10.1109/IGARSS.2019.8898458>
- [263] Yuan Yuan, Lei Lin, Qingshan Liu, Renlong Hang, and Zeng-Guang Zhou. 2022. SITS-Former: A pre-trained spatio-spectral-temporal representation model for Sentinel-2 time series classification. *International Journal of Applied Earth Observation and Geoinformation* 106 (Feb. 2022), 102651. DOI : <http://dx.doi.org/10.1016/j.jag.2021.102651>
- [264] Mengjia Qiao, Xiaohui He, Xijie Cheng, Panle Li, Haotian Luo, Lehan Zhang, and Zhihui Tian. 2021. Crop yield prediction from multi-spectral, multi-temporal remotely sensed imagery using recurrent 3D convolutional neural networks. *International Journal of Applied Earth Observation and Geoinformation* 102 (2021), 102436.
- [265] Marc Rußwurm and Marco Körner. 2020. Self-attention for raw optical satellite time series classification. *ISPRS Journal of Photogrammetry and Remote Sensing* 169 (2020), 421–435. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2020.06.006>
- [266] Devis Tuia, Claudio Persello, and Lorenzo Bruzzone. 2016. Domain adaptation for the classification of remote sensing data: An overview of recent advances. *IEEE Geoscience and Remote Sensing Magazine* 4, 2 (2016), 41–57. DOI : <http://dx.doi.org/10.1109/MGRS.2016.2548504>
- [267] V. Sainte Fare Garnot, Loic Landrieu, Sebastien Giordano, and Nesrine Chehata. 2019. Time-space tradeoff in deep learning models for crop classification on satellite multi-spectral image time series. In *2019 IEEE International Geoscience and Remote Sensing Symposium (IGARSS'19)*. IEEE, 6247–6250.
- [268] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. 2019. Deep neural network ensembles for time series classification. In *2019 International Joint Conference on Neural Networks (IJCNN'19)*, Vol. 2019. IEEE, 1–6. DOI : <http://dx.doi.org/10.1109/IJCNN.2019.8852316>
- [269] David H. Wolpert. 1992. Stacked generalization. *Neural Networks* 5, 2 (Jan. 1992), 241–259. DOI : [http://dx.doi.org/10.1016/S0893-6080\(05\)80023-1](http://dx.doi.org/10.1016/S0893-6080(05)80023-1)
- [270] Yoav Freund and Robert E. Schapire. 1996. Experiments with a new boosting algorithm. In *13th International Conference on Machine Learning*. 148–156.

- [271] Cristina Gómez, Joanne C. White, and Michael A. Wulder. 2016. Optical remotely sensed time series data for land cover classification: A review. *ISPRS Journal of Photogrammetry and Remote Sensing* 116 (2016), 55–72. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2016.03.008>
- [272] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. 2017. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine* 5, 4 (2017), 8–36. DOI : <http://dx.doi.org/10.1109/MGRS.2017.2762307>
- [273] Lei Ma, Yu Liu, Xueliang Zhang, Yuanxin Ye, Gaofei Yin, and Brian Alan Johnson. 2019. Deep learning in remote sensing applications: A meta-analysis and review. *ISPRS Journal of Photogrammetry and Remote Sensing* 152 (June 2019), 166–177. DOI : <http://dx.doi.org/10.1016/j.isprsjprs.2019.04.015>
- [274] Qiangqiang Yuan, Huanfeng Shen, Tongwen Li, Zhiwei Li, Yun Jiang, Hongzhang Xu, Weiwei Tan, Qianqian Yang, Jiwen Wang, Jianhao Gao, and Liangpei Zhang. 2020. Deep learning in environmental remote sensing: Achievements and challenges. *Remote Sensing of Environment* 241 (May 2020), 111716. DOI : <http://dx.doi.org/10.1016/j.rse.2020.111716>
- [275] Michel E. D. Chaves, Michelle C. A. Picoli, and Ieda D. Sanches. 2020. Recent applications of Landsat 8/OLI and Sentinel-2/MSI for land use and land cover mapping: A systematic review. *Remote Sensing* 12, 18 (sep 2020), 3062. DOI : <http://dx.doi.org/10.3390/rs12183062>
- [276] Waytehad Rose Moskolai, Wahabou Abdou, Albert Dipanda, and Kolyang. 2021. Application of deep learning architectures for satellite image time series prediction: A review. *Remote Sensing* 13, 23 (Nov. 2021), 4822. DOI : <http://dx.doi.org/10.3390/rs13234822>
- [277] Jason Lines and Anthony Bagnall. 2015. Time series classification with ensembles of elastic distance measures. *Data Mining and Knowledge Discovery* 29, 3 (2015), 565–592.
- [278] Chang Wei Tan, François Petitjean, and Geoffrey I. Webb. 2020. FastEE: Fast ensembles of elastic distances for time series classification. *Data Mining and Knowledge Discovery* 34, 1 (2020), 231–272.
- [279] Matthieu Herrmann and Geoffrey I. Webb. 2021. Amercing: An intuitive, elegant and effective constraint for dynamic time warping. *arXiv preprint:2111.13314* (2021).
- [280] Anthony Bagnall, Michael Flynn, James Large, Jason Lines, and Matthew Middlehurst. 2020. On the usage and performance of the hierarchical vote collective of transformation-based ensembles version 1.0 (hive-cote v1. 0). In *International Workshop on Advanced Analytics and Learning on Temporal Data*. 3–18.
- [281] Matthew Middlehurst, James Large, Michael Flynn, Jason Lines, Aaron Bostrom, and Anthony Bagnall. 2021. HIVE-COTE 2.0: A new meta ensemble for time series classification. *Machine Learning* 110, 11 (2021), 3211–3243.
- [282] Anthony Bagnall, Jason Lines, Jon Hills, and Aaron Bostrom. 2015. Time-series classification with COTE: The collective of transformation-based ensembles. *IEEE Transactions on Knowledge and Data Engineering* 27, 9 (2015), 2522–2535.
- [283] Jason Lines, Sarah Taylor, and Anthony Bagnall. 2018. Time series classification with HIVE-COTE: The hierarchical vote collective of transformation-based ensembles. *ACM Trans. Knowl. Discov. Data* 12, 5, Article 52 (October 2018), 35 pages. <https://doi.org/10.1145/3182382>
- [284] Jason Lines, Sarah Taylor, and Anthony Bagnall. 2016. Hive-Cote: The hierarchical vote collective of transformation-based ensembles for time series classification. In *2016 IEEE 16th International Conference on Data Mining (ICDM'16)*. IEEE, 1041–1046.
- [285] Rohit J. Kate. 2016. Using dynamic time warping distances as features for improved time series classification. *Data Mining and Knowledge Discovery* 30, 2 (2016), 283–312.
- [286] Aaron Bostrom and Anthony Bagnall. 2015. Binary shapelet transform for multiclass time series classification. In *International Conference on Big Data Analytics and Knowledge Discovery*. Springer, 257–269.
- [287] Patrick Schäfer. 2015. The BOSS is concerned with time series classification in the presence of noise. *Data Mining and Knowledge Discovery* 29, 6 (2015), 1505–1530.
- [288] Jon Hills, Jason Lines, Edgaras Baranauskas, James Mapp, and Anthony Bagnall. 2014. Classification of time series by shapelet transformation. *Data Mining and Knowledge Discovery* 28, 4 (2014), 851–881.
- [289] Houtao Deng, George Runger, Eugene Tuv, and Martyanov Vladimir. 2013. A time series forest for classification and feature extraction. *Information Sciences* 239 (2013), 142–153.
- [290] Mustafa Gokce Baydogan, George Runger, and Eugene Tuv. 2013. A bag-of-features framework to classify time series. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35, 11 (2013), 2796–2802.
- [291] Angus Dempster, Daniel F. Schmidt, and Geoffrey I. Webb. 2021. Minitrocket: A very fast (almost) deterministic transform for time series classification. In *27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 248–257.
- [292] Chang Wei Tan, Angus Dempster, Christoph Bergmeir, and Geoffrey I. Webb. 2022. MultiRocket: Multiple pooling operators and transformations for fast and effective time series classification. *Data Mining and Knowledge Discovery* 36 (June 2022), 1623–1646. DOI : <http://dx.doi.org/10.1007/s10618-022-00844-1> arXiv:2102.00457

- [293] Angus Dempster, Daniel F. Schmidt, and Geoffrey I. Webb. 2023. Hydra: Competing convolutional kernels for fast and accurate time series classification. *Data Mining and Knowledge Discovery* 37 (2023), 1–27.
- [294] Benjamin Lucas, Ahmed Shifaz, Charlotte Pelletier, Lachlan O’Neill, Nayyar Zaidi, Bart Goethals, François Petitjean, and Geoffrey I. Webb. 2019. Proximity forest: An effective and scalable distance-based classifier for time series. *Data Mining and Knowledge Discovery* 33, 3 (2019), 607–635.
- [295] Matthieu Herrmann, Chang Wei Tan, Mahsa Salehi, and Geoffrey I. Webb. 2023. Proximity forest 2.0: A new effective and scalable similarity-based classifier for time series. *arXiv preprint arXiv:2304.05800* (2023).
- [296] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86, 11 (1998), 2278–2324.
- [297] Vinod Nair and Geoffrey E. Hinton. 2010. Rectified linear units improve restricted Boltzmann machines. In *ICML*.
- [298] Salah Hihi and Yoshua Bengio. 1995. Hierarchical recurrent neural networks for long-term dependencies. *Advances in Neural Information Processing Systems* 8 (1995), 493–499.
- [299] Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. 2013. How to construct deep recurrent neural networks. *arXiv preprint:1312.6026* (2013).
- [300] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint:1409.0473* (2014).
- [301] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. *arXiv preprint:1406.1078* (2014).
- [302] Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint:1508.04025* (2015).

Received 17 January 2023; revised 15 December 2023; accepted 31 January 2024