



## Survey paper

## A survey on state-of-the-art deep learning applications and challenges

Mohd Halim Mohd Noor<sup>a</sup>, Ayokunle Olalekan Ige<sup>b</sup><sup>a</sup> School of Computer Sciences, Universiti Sains Malaysia, Jalan Universiti, Gelugor, 11800, Pulau Pinang, Malaysia<sup>b</sup> Department of Computer Science, Adekunle Ajasin University, Ikare Akoko-Ipeme Rd, Akungba-Akoko, P.M.B 001, Ondo State, Nigeria

## ARTICLE INFO

## Keywords:

Deep learning  
Fundamentals  
Computer vision  
Natural language processing  
Time series analysis  
Robotics  
Challenges

## ABSTRACT

Deep learning, a branch of artificial intelligence, is a data-driven method that uses multiple layers of interconnected units or neurons to learn intricate patterns and representations directly from raw input data. Empowered by this learning capability, it has become a powerful tool for solving complex problems and is the core driver of many groundbreaking technologies and innovations. Building a deep learning model is challenging due to the algorithm's complexity and the dynamic nature of real-world problems. Several studies have reviewed deep learning concepts and applications. However, the studies mostly focused on the types of deep learning models and convolutional neural network architectures, offering limited coverage of the state-of-the-art deep learning models and their applications in solving complex problems across different domains. Therefore, motivated by the limitations, this study aims to comprehensively review the state-of-the-art deep learning models in computer vision, natural language processing, time series analysis and pervasive computing, and robotics. We highlight the key features of the models and their effectiveness in solving the problems within each domain. Furthermore, this study presents the fundamentals of deep learning, various deep learning model types and prominent convolutional neural network architectures. Finally, challenges and future directions in deep learning research are discussed to offer a broader perspective for future researchers.

## 1. Introduction

Deep learning has revolutionized many applications across a variety of industries and research. The application of deep learning can be found in healthcare (Shamshirband et al., 2021), smart manufacturing (Wang et al., 2018), robotics (Pierson and Gashler, 2017) and cybersecurity (Dixit and Silakari, 2021), solving challenging and complex problems such as disease diagnosis, anomaly detection, object detection and malware attack detection. Deep learning is a subset of machine learning that focuses on learning from data using artificial neural networks with many layers, known as deep neural networks. An artificial neural network is a computational model that imitates the working principles of a human brain. The computational models are composed of an input layer which receives the input data, multiple processing layers that learn the representation of data and the output layer which produces the output of the model.

Prior to the reintroduction of deep learning (DL) into the research trend, pattern recognition tasks involved a transformation of the raw input data such as pixel values of an image into a feature vector that represents the internal representation of the data. The feature vector can be used by a machine learning model to detect or classify patterns in the data. This process requires feature engineering and considerable domain knowledge to design a suitable feature representation. With

deep learning, this cumbersome process can be performed automatically whereby at each processing layer known as hidden layers, the internal representation of the input data is learned or extracted in a hierarchical manner. The first layer learns the presence of basic primitive features such as edges, dots, and lines. The second layer learns patterns or motifs by recognizing the combinations of the edges, dots and lines, and the subsequent layers combine the motifs to produce more sophisticated features that correspond to the input data. This feature learning process takes place in the sequence of hidden layers until the prediction is finally produced.

Deep learning has seen numerous breakthroughs in various industries, transforming how complex problems are solved and how businesses operate. For instance, protein folding is a complex process, influenced by the intricate sequence of amino acids. Using traditional methods to determine how a protein folds would take an immense amount of computational power and time due to its complex 3D structure and vast number of possible configurations. AlphaFold (AlphaFold, 2025) has made significant progress in this area, using deep learning to predict protein structures with remarkable accuracy, helping researchers to comprehend the complete structure of a key protein associated with diseases like malaria and Parkinson's disease. In natural

\* Corresponding author.

E-mail addresses: [halimnoor@usm.my](mailto:halimnoor@usm.my) (M.H. Mohd Noor), [ayo.ige@aaua.edu.ng](mailto:ayo.ige@aaua.edu.ng) (A.O. Ige).<https://doi.org/10.1016/j.engappai.2025.111225>

Received 20 November 2024; Received in revised form 15 March 2025; Accepted 17 May 2025

Available online 16 July 2025

0952-1976/© 2025 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

**Table 1**  
Summary of related works.

Ref.	Focus	Concepts not covered
Dong et al. (2021)	A short review of the fundamentals of DL networks and discusses different types of neural networks, DL applications and challenges.	Lack of analysis of CNN architectures and limited coverage of deep learning fundamentals.
Talaei Khoei et al. (2023)	Discusses the learning approaches (supervised, unsupervised and reinforcement learnings), learning strategies, and DL challenges	Lack of fundamentals of deep learning, CNN architectures and DL applications.
Alzubaidi et al. (2021)	Discusses different types of DL networks, CNN fundamentals and architectures, DL challenges and medical imaging applications	Limited discussion on DL applications such as natural language processing and time series analysis.
Alom et al. (2019)	A short review of the fundamentals of neural networks and discusses different types of DL networks, CNN architectures and applications.	Limited discussion on DL applications and no discussion of DL challenges.
Pouyanfar et al. (2018)	Discusses different types of DL networks and DL applications and challenges	Lack of analysis of CNN architectures and limited coverage of deep learning fundamentals.
Sarker (2021)	Discusses different types of DL networks and provides a summary of DL applications	Lack of fundamentals of deep learning, analysis of CNN architectures and limited discussion on DL applications.

language processing, systems like OpenAI GPT, Google Gemini and IBM Watsonx have revolutionized chatbots and virtual assistants, enabling them to understand and respond to human language with remarkable accuracy and contextual awareness. These systems have overcome the challenge of processing vast amounts of unstructured text, allowing them to engage in more natural conversations and handle a wide range of topics. Businesses can use these systems to automate customer service, provide real-time support and improve user experience.

Several studies have been conducted to discuss the concept and application of deep learning in the last few years, as listed in Table 1. The studies addressed or focused on several aspects of deep learning, such as types of deep learning models, learning approaches and strategies, convolutional neural network (CNN) architectures, deep learning applications and challenges. In Dong et al. (2021), the authors provided fundamentals of deep learning and highlighted different types of deep learning models, such as convolutional neural networks, autoencoder and generative adversarial networks. Then, the applications of deep learning in various domains are discussed, and some challenges associated with deep learning applications are presented. Another survey (Talaei Khoei et al., 2023) provided a comprehensive analysis of supervised, unsupervised and reinforcement learning approaches and compared the different learning strategies such as online, federated and transfer learning. Finally, the current challenges of deep learning and future direction are discussed.

In Alzubaidi et al. (2021), the authors provided a comprehensive review of the popular CNN architectures used in computer vision tasks, highlighting their key features and advantages. Then, the applications of deep learning in medical imaging and the challenges are discussed. A similar survey is reported in Alom et al. (2019), where the different supervised and unsupervised deep learning models are highlighted, and the popular CNN architectures are compared and discussed. In another survey (Pouyanfar et al., 2018), the authors focused on the applications of deep learning in computer vision, natural language processing and speech and audio processing. The different types of deep learning models are also discussed. In Sarker (2021), the authors focused on the different types of deep learning models and provided a summary of deep learning applications in various domains.

Despite the existing surveys on deep learning that offer valuable insights, the increasing amount of deep learning applications and the existing limitations in the current studies motivated us to explore this topic in depth. In general, to the best of our knowledge, no survey paper focuses on the emerging trends in state-of-the-art applications and the current challenges associated with deep learning. Furthermore, the surveys do not discuss the issues and how deep learning addresses them by highlighting the key features and components in the models. Furthermore, most surveys either ignore or provide minimal coverage of the fundamentals of deep learning, which is crucial for understanding the state-of-the-art models. The main objective of this paper is to present the most important aspects of deep learning, making it accessible to a wide audience and facilitating researchers and practitioners

in advancing and leveraging its capabilities to solve complex problems across diverse domains. Specifically, we present the fundamentals of deep learning and the various types of deep learning models, including popular deep learning architectures. Then, we discuss the progress of deep learning in state-of-the-art applications, highlighting the key features of the models and their problem-solving approaches. Finally, we discuss the challenges faced by deep learning and the future research directions.

To this end, we conducted automatic search strategies using “deep learning” keyword with keywords related to deep learning applications such as “image classification”, “neural machine translation”, “text generation”, “human activity recognition” and “robotics”. Then, a manual search was carried out by scanning the references produced by the automatic search, selecting relevant works and discarding the irrelevant ones. This survey paper collected primary studies from journals, conference proceedings, and books in English only. We use scientific databases such as IEEE Explore, ScienceDirect, SpringerLink, ACM Digital Library, Scopus and ArXiv. The ArXiv repository is used because it hosts numerous manuscripts that are highly relevant to the topics discussed in this survey. These preprints often contain novel findings and methodologies that may not be available in other repositories at the time of writing. Nevertheless, we ensure that the selected preprints are of high quality, authored by researchers from reputable institutions.

The remainder of this paper is organized as follows: Section 2 describes the fundamentals of deep learning which includes layers and attention mechanisms, activation functions, model optimization and loss functions, and regularization methods. Section 3 presents the types of deep learning models, including the CNN architectures. Section 4 discusses the state-of-the-art applications of deep learning. Section 5 discusses the challenges and future directions in the field of deep learning. The conclusion is given in Section 6.

## 2. Fundamentals of deep learning

This section describes the fundamental concepts such as layer types, activation functions, training algorithms and regularization methods to provide a comprehensive understanding of the underlying principles in advancing the field of deep learning.

### 2.1. Layers

A deep learning model is characterized by having numerous hidden layers. The hidden layers are responsible for learning and extracting complex features from the input data. A hidden layer is composed of an arbitrary number of neurons which serves as the fundamental building block of a neural network as shown in Fig. 1. A neuron consists of an arbitrary number of inputs, each associated with a weight, which controls the flow of information into the neuron during the forward pass. The flow of information, or forward pass, involves the computation

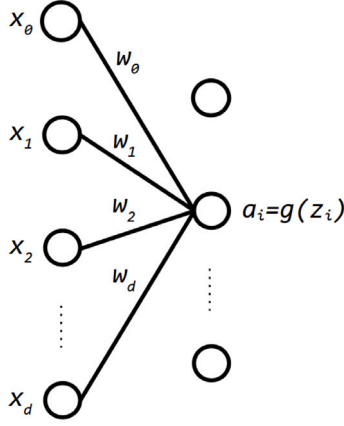


Fig. 1. A graphical representation of a neuron.

of summation of the weighted input, followed by the application of a transformation function to the weighted sum. Let  $\mathbf{x} = x_1, x_2, \dots, x_d, j = 0, 1, \dots, d$  be the input vector with  $d$  dimensions and  $\mathbf{w}_i^l$  denotes the weights that are connected to neuron  $i$  in layer  $l$ . The forward pass to neuron  $i$  in layer  $l$  is defined as

$$z_i^l = \mathbf{w}_i^l \cdot \mathbf{x} = \sum_{j=0}^d w_{i,j}^l \cdot x_j \quad (1)$$

$$a_i^l = g(z_i^l) \quad (2)$$

where  $x_j$  is the input vector of size  $d$ ,  $w_{i,j}^l$  is the weight associated with  $x_j$ , connecting the input to neuron  $i$  at layer  $l$ , and  $g$  is the transformation function also known as activation function. It is worth noting that  $w_0$  is called bias and  $x_0 = 1$ . A hidden layer wherein each neuron is connected to all neurons of the previous layer is known as the fully connected layer. The forward pass computation of a neuron  $a_i^l$  in layer  $l$  receiving a set of input from layer  $l-1$  can be generalized as follows:

$$a_i^l = g\left(\sum_{j=0}^d w_{i,j}^l \cdot a_j^{l-1}\right) = \mathbf{w}_i^l \cdot \mathbf{a}^{l-1} \quad (3)$$

where  $\mathbf{a}^{l-1}$  is the input vector from layer  $l-1$ .

Another crucial layer type is the convolutional layer, which is primarily used for processing data that has underlying structures, such as spatial patterns in image data or temporal patterns in time series data. Unlike a fully connected layer, each neuron in a convolutional layer is connected only to a subset of neurons in the previous layer, as shown in Fig. 2. As shown in the figure, each neuron in the hidden layer is connected only to a local region, a subset of nine input neurons, and the weights are shared across the input data. The weight-sharing not only significantly reduces the number of parameters of the neural network, but also allows the network to learn the same features across different spatial locations in the input (LeCun et al., 2015). Let  $\mathbf{x}$  be a two-dimensional input data with  $H \times W \times C$  such as a grayscale image, where  $H$  is the height,  $W$  is the width and  $C$  is the channel of the image data. The computation of a neuron in the convolutional layer  $l$  is defined as

$$z_{i,j,d}^l = \sum_{m=0}^{k_1} \sum_{n=0}^{k_2} w_{m,n,d}^l \cdot x_{i-s+m,j-s+n,d} \quad (4)$$

$$a_{i,j,d}^l = g(z_{i,j,d}^l) \quad (5)$$

where  $w_{m,n,d}$  is the weight connecting the input data  $x_{i,j,d}$  within the window to the neuron in layer  $l$  and  $s$  is the stride or the step size of the window as it moves across the input data. This convolution operation where the filter is represented by  $w_{m,n,d}$  slides over the input data

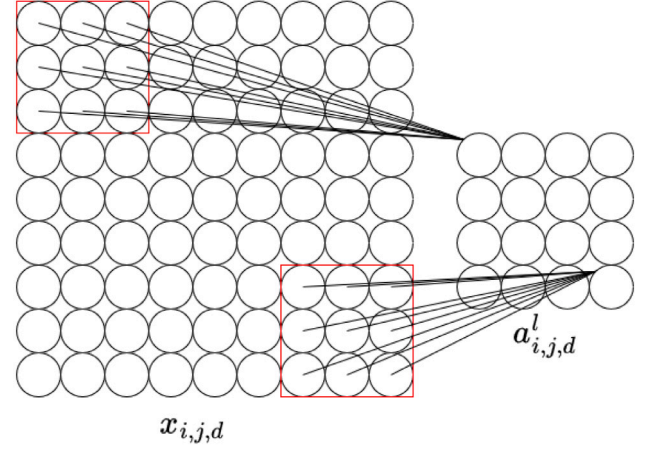


Fig. 2. A neuron is connected to a local region of the input data.

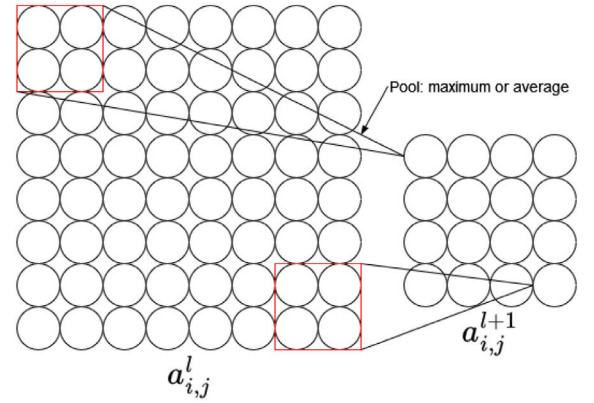


Fig. 3. Summaries of the feature map using maximum or average pooling to produce a reduced feature map.

$x_{i,j,d}$ , producing a set of output values called feature map. For a more hands-on understanding of this convolutional operation, readers are encouraged to explore CNN Explainer (poloclub/cnn-explainer, 2025), an interactive tool that demonstrates how convolutional layers process input data and generate feature maps.

Pooling layers are commonly applied after successive convolutional layers to progressively reduce the spatial dimensions of the feature maps. The spatial reduction is performed by computing the summary of a subset of the feature values in the feature maps as shown in Fig. 3. The pooling operation, as shown in the figure, can use either the maximum or average method with a pooling size of  $2 \times 2$  applied across the entire feature map, thereby reducing the size of the feature map. In addition to spatial reduction, pooling layers decrease the number of parameters and provide translation-invariant features (LeCun et al., 2015). Let  $\mathbf{a}^l$  denotes feature map (hidden layer)  $l$ , a pooling operation with a pooling size of  $m \times n$  is defined as

$$\mathbf{a}^{l+1} = \text{pool}(\mathbf{a}_{i:s+m,j:s+n}^l) \quad (6)$$

where pool is either maximum or average function and  $s$  is the stride or step size of the window as it moves across the feature map.

## 2.2. Attention mechanisms

One of the important concepts in pattern recognition is the ability to attend to and neglect certain parts of the input data based on their relevance. This is because not all parts of the input hold equal importance for making predictions. Certain features exhibit a stronger correlation

with the output while others are less relevant. To provide a simple analogy, consider a set of sensors to measure room temperature. The sensors are deployed across different corners of the room. Each sensor measures temperature in its own area, and these local measurements are used to estimate the room temperature. However, not all sensors are equally important due to various factors. For instance, a sensor near an air-conditioning or a heater might give inaccurate readings due to the external temperature fluctuation caused by these systems. A sensor near the center of the room, away from the heating or cooling sources might provide more reliable readings. Therefore, when estimating the room temperature, we give more attention to the sensors that are less affected by the localized factors. In convolutional layers, all extracted features are treated uniformly, without consideration of the varying degree of the importance of the different parts of the input data. This limitation is addressed by the introduction of attention mechanism, which can dynamically assign varying levels of significance (weights) to the different features. This flexibility enables the deep learning models to prioritize the more relevant aspects of the input data, enhancing its ability to capture the intricate dependencies for accurate prediction. Given an input data  $x$ , the process of attending to the important components of the input is given as

$$A = f(g(x), x) \quad (7)$$

where  $g$  is a composite function that performs a sequence of operations to generate the attention or the weights and  $f$  applies the generated attention  $g(x)$  on the input  $x$ .

For instance, the squeeze-and-excitation (SE) attention generates the attention through five consecutive operations (Hu et al., 2019). First, the input is vectorized using global average pooling. Then the vector is passed to two fully connected layers, where the first one with ReLU activation and the second one with sigmoid activation. SE attention was a pioneer in channel attention. The attention module assigns varying weights to the channels of the feature maps. SE attention suffers from computational cost and the use of global average, which may cause information loss at the spatial level. Several efforts have been made to improve SE attention. Global Second-order Pooling (GSoP) attention performs  $1 \times 1$  convolution on the feature maps to reduce the number of channels, and then computes the pairwise channel correlation, which is used to generate the weights (Gao et al., 2018). Efficient Channel Attention (ECA) replaced the fully connected layers with 1D convolution to reduce the number of parameters and the computational cost (Wang et al., 2020).

Temporal attention is an attention module that focuses on specific time steps in a sequence of data such as time series and video (sequence of images). In video processing such as recognizing human actions, temporal attention is used to focus on key frames at different point in time that contain crucial information for predicting the ongoing activity. Temporal adaptive module (TAM) is a temporal attention that can focus on short-term (local) information and global context information of the data (Liu et al., 2021). The composite function consists of a local branch for generating attention weights and a global branch for generating a channel-wise adaptive kernel. First, the input feature map is squeezed using global average pooling to reduce the computational cost. Subsequently, the local branch executes two 1D convolution operations, with the first convolution using ReLU, and sigmoid activation for the second to generate the local weights. The local weights are then multiplied with the feature map. Meanwhile, the global branch is composed of two fully connected layers, with the first layer using ReLU and the second layer employing softmax function to generate the adaptive kernel (weights). Self-attention is a form of temporal attention, initially proposed for machine translation to enable the deep learning models to attend different words in a sequence relative to other words (Vaswani et al., 2023). The attention module has become a fundamental building block in various natural language processing applications. To generate the attention weights, the input (word embeddings) is transformed by linear projection to

compute query, key, and value. Then, the dot product between query and key is computed, and the resultant is normalized by the square root of the size of the key. Finally, the attention weights are obtained by applying the softmax function. Self-attention is the fundamental building block of the Transformer architecture, a key deep learning model in natural language processing. For a more hand-on understanding of self-attention and Transformer, readers are encouraged to explore Transformer Explainer (poloclub/transformer-explainer, 2025), an interactive tool that demonstrates how self-attention mechanisms work in Transformer by visualizing the attention scores and how different inputs interact with each other.

Spatial attention focuses on specific regions or spatial location of the input data, enabling the deep learning models to selectively emphasize and ignore certain features. In the context of computer vision, spatial attention is crucial in capturing the spatial relationships and context within an image for accurate prediction. Attention gate is a spatial attention that can identify and focus the salient regions and suppress feature responses of the insignificant ones. The composite function consists of ReLU activation followed by  $1 \times 1$  convolution to reduce channel dimension of the feature maps to a singular feature map. Finally, sigmoid is applied to the feature map to generate attention weights (Oktay et al., 2018). The self-attention in the standard Transformer is not effective in handling image data due to its inherent sequential processing nature and lacks the ability to capture spatial dependencies and local patterns. To address this limitation, the Vision Transformer (ViT) treats images as a sequence of non-overlapping patches. A similar computational pipeline is used to generate the attention weights, the sequence of patches is transformed by linear projection to compute the query, key, and value (Dosovitskiy et al., 2021). The same operations are employed to generate the attention weights. Self-attention is computationally costly due to its quadratic complexity, especially when dealing with image data. To reduce the complexity, two learnable linear layers, independent of the input data, are adopted as the key and value vectors (Guo et al., 2023).

### 2.3. Activation functions

The role of the activation function is to transform the weighted sum into a more classifiable form. This is crucial to the learning behavior of the deep learning model, generating non-linear relationships between the input and the output of the model. The activation function, combined with many hidden layers, allows the neural network to approximate highly complex, non-linear functions. Many activation functions are available for use in neural networks, and some of the functions are shown in Figs. 4–6. The figures show the plots of the three popular activation functions. The sigmoid is a classic example of an activation function used in logistic regression. It maps the weighted sum to a value in the range of 0 to 1 which can be used for classification. The hyperbolic tangent (tanh) is another popular choice of bounded activation function, which produces an output between  $-1$  and  $1$ . Since it has a stronger gradient, neural network training often converges faster than with the sigmoid function (LeCun et al., 2012). For many years, sigmoid and hyperbolic tangent functions were the commonly used activation functions. However, both suffer from the vanishing gradient problem, which hinders the efficient training of deep neural networks with many layers (Hochreiter et al., 2001).

It was shown that neural networks with unbounded activation functions have the universal approximation property and reduce the problem of vanishing gradients. In recent years, numerous unbounded activation functions have been proposed for neural networks, with the softplus (Dugas et al., 2000) and the rectified linear unit (ReLU) (Glorot et al., 2011) activation function being notable examples. These activation functions, especially ReLU has a pivotal role in improving the training and performance of deep learning models. ReLU has been a cornerstone of deep learning models due to its computational efficiency and effectiveness in addressing the vanishing gradient problem.

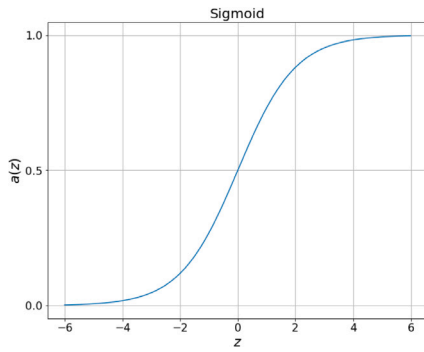


Fig. 4. Sigmoid activation function.

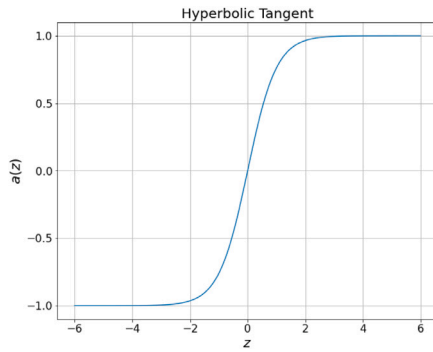


Fig. 5. Hyperbolic tangent activation function.

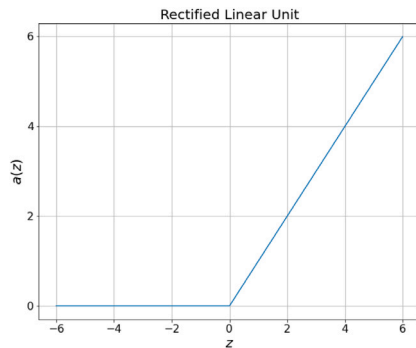


Fig. 6. Rectified linear unit activation function.

Since then, several variants of ReLU have been proposed including Leaky ReLU (Maas et al., 2013), sigmoid linear unit (Misra, 2020) and exponential linear unit (Clevert et al., 2016), each offering unique advantages for building deep learning-based applications.

#### 2.4. Parameter learning and loss functions

The weights (parameters) of deep learning models are often optimized using an optimization algorithm called gradient descent, though other optimization algorithms may also be used. However, it has to be noted that gradient descent is a generic algorithm which can be used to solve a wide range of optimization problems. In general, gradient descent finds the optimal weights by iteratively updating the weights such that the weights will result in a minimum prediction error over all instances in the training set. This process is synonymous with a person at the top of a hill who wants to climb down to the ground. Just like the person may choose a path that leads to the lowest point by taking small steps based on the steepness of the slope, gradient descent makes

small adjustments to the weights, moving them in the direction that reduces the error the most. The prediction error is quantified by a loss function. For classification problems, the commonly used loss function is the negative log-likelihood loss or cross entropy loss, while the square loss and absolute loss are used for regression problems (Wang et al., 2022). The weights of layer  $l$  are updated as

$$w_{i,j}^l = w_{i,j}^l - \alpha \nabla_{w_{i,j}^l} \quad (8)$$

where  $\alpha$  is a hyperparameter called learning rate and  $\nabla_{w_{i,j}^l}$  is the gradient or the derivative of the loss function  $J$  with respect to the weight  $\frac{\partial J}{\partial w_{i,j}^l}$ .

The gradient can be computed across all training set instances, an approach known as batch gradient descent. However, this approach does not always guarantee convergence to the optimal solution, as it may get stuck in local minima or saddle points, and the same gradient is used for every weight update. An alternative approach is to perform the weight update on the basis of a single instance, but the approach results in a noisy gradient and becomes computationally intensive due to the frequent weight update. A more commonly used approach is to perform the weight update over a set of training instances, known as mini-batch gradient descent. This approach strikes a balance, providing a less noisy gradient and a more stable training process. For a more hands-on understanding of neural network training, readers are encouraged to explore TF playground (Smilkov and Carter, n.d.) and Initializing Neural Networks (Guo, n.d.), interactive tools demonstrating how neural networks are trained and the impact of hyperparameters.

Several efforts have been made to improve the efficiency of gradient descent. One of the earlier efforts is the inclusion of past rate of change in the weight update to speed up the training of deep learning models, the algorithm is called gradient descent with momentum (Qian, 1999). Another effort is to improve the training convergence by adapting the learning rate based on the occurrence of the features (Duchi et al., 2011). A more recent work utilizes both adaptive learning rate and momentum to improve the training efficiency and convergence of deep learning models (Kingma and Ba, 2017).

#### 2.5. Regularization methods

Regularization methods are employed to prevent overfitting in deep learning models and improve their generalization performance. Early stopping is a method that can detect the onset of overfitting during training by continuously monitoring the validation error. The model is considered overfitting if the validation error starts to increase at some point of the training while the training error is decreasing. However, detecting the onset of overfitting during the training of deep learning models is challenging due to the inherent stochasticity and the presence of noisy data. Several stopping criteria can be considered, such as using a threshold to check if the decrease of (average) validation error is significant and counting the number of successive increases of validation error (Prechelt, 2012).

Dropout is a regularization method that randomly switches off some neurons in the hidden layers during training with a predefined drop probability (dropout rate) (Srivastava et al., 2014). Dropout has the effect of training and evaluating a large number of different subnetworks within the models. The dropout rate is a hyperparameter that needs to be carefully tuned to balance regularization and model capacity. Different ranges of dropout rate have been suggested. The original author suggested a dropout rate between 0.5 and 0.8 (Srivastava et al., 2014) while others recommended a lower dropout rate between 0.1 and 0.3 (Park and Kwak, 2017). Furthermore, it has been suggested a low dropout rate due to the exponential increase in the volume of training data (Liu et al., 2023).

Parameter norm penalty is a regularization method that adds a penalty term consisting of the network's weights to the loss function. During the training, the penalty term discourages large weight values

and hence, constraining the model's capacity and reducing the chance of overfitting. The common penalty terms are  $L^1$  norm penalty (Tibshirani, 1996),  $L^2$  norm penalty, also known as weight decay and a combination of  $L^1$  and  $L^2$  (Zou and Hastie, 2005). An adaptive weight decay is proposed, allowing the regularization strength for each weight to be dynamically adjusted (Nakamura and Hong, 2019).

Despite the advantages of the mini-batch gradient descent, each mini-batch may comprise data from different distributions. Furthermore, the data distribution may change after each weight update, which could slow down the training process. Batch normalization overcomes this issue by normalizing the summed input to a neuron over a mini batch of training instances (Ioffe and Szegedy, 2015). An alternative method is to perform normalization across the neurons instead of the mini batch, a method known as layer normalization (Bae et al., 2016). Layer normalization is applicable in recurrent neural networks and overcomes the dependencies on the mini batch size.

### 3. Types of deep learning

Deep learning models can be categorized into deep supervised learning and deep unsupervised learning.

#### 3.1. Deep supervised learning

Deep supervised models are trained with a labeled dataset. The learning process of these models involves calculating the prediction error through a loss function and utilizing the error to adjust the weights iteratively until the prediction error is minimized. Among the deep supervised models, three important models are identified, namely multilayer perceptron, convolutional neural network and recurrent neural network.

##### 3.1.1. Multilayer perceptron

Multilayer perceptron is a neural network model with one or more hidden fully connected layers stacked between the input and output layers as shown in Fig. 7. The width (number of neurons) of the hidden layers and the depth (number of layers) of the network influence the model's ability to learn patterns in the data. Specifically, the width affects the network's ability to capture a broader range of features, while the depth facilitates the learning of hierarchical representations. Nevertheless, studies indicated that a multilayer perceptron with a single hidden layer can approximate any continuous function (Cybenko, 1989; Hornik et al., 1989). Multilayer perceptron is effective in various industries and applications from healthcare to finance (Widrow et al., 1994). However, a multilayer perceptron requires the input data to be structured in a one-dimensional format (e.g., tabular data), making it less suitable for unstructured data such as images, text, and speech. To leverage multilayer perceptron for unstructured data, a feature extraction or transformation into structured data is necessary.

##### 3.1.2. Recurrent neural network

Recurrent neural network (RNN) is a neural network model that leverages the sequential information and memory through the use of recurrent connections, allowing it to effectively process data such as time series, text, speech and other sequential patterns. As shown in Fig. 8, a recurrent neural network is characterized by the recurrent connection, which enables the network to loop back and use internal state from the previous time step to the next time step. The internal state is parameterized by a set of weights shared across the sequence of data. The training of recurrent neural networks suffers from the issue of vanishing gradient due to the challenges of propagation of gradients over a long sequence of data. Variants of recurrent neural networks are introduced to overcome the problem of vanishing gradient, such as long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) and gated recurrent memory (GRU) (Cho et al., 2014). The improved recurrent neural networks introduce memory cell

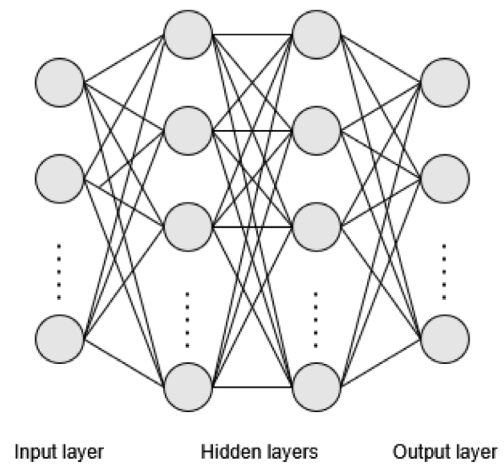


Fig. 7. A fully connected neural network.

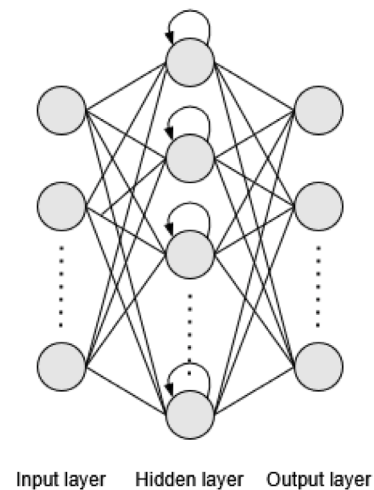


Fig. 8. A neural network with recurrent connection.

and gating mechanisms to retain and discard information in every time step, allowing for more effective learning dependencies in long sequence. The network architecture can be built using fully connected and convolutional layers (Shi et al., 2015).

LSTM is arguably the most widely used variant and has been applied in various applications such as image captioning, machine translation and sentiment analysis (Van Houdt et al., 2020). Similar to the standard RNNs, LSTMs process information steps by steps using a chain of repeating units. Each LSTM unit consists of several gates, namely the forget gate, input gate and output gate that control the flow of information from one time step to the next time step. Fig. 9 illustrates how information flows through these gates. The key component of an LSTM unit is the cell, represented by the horizontal line running through the top of the block. This cell maintains the state at each time step and is updated by the gates. Specifically, at each time step, the forget gate determines which information to be retained or discarded from the cell state based on the current input and the previous hidden state. The input gate decides what new information should be added to the cell state, while the output gate uses the updated cell state, current input and previous hidden state to produce the hidden state for the current time step.

GRU is a popular variant of RNN that is similar to LSTM but has a relatively simpler architecture. The network has two gates called reset and update to control the information flow. Unlike LSTM, GRU does

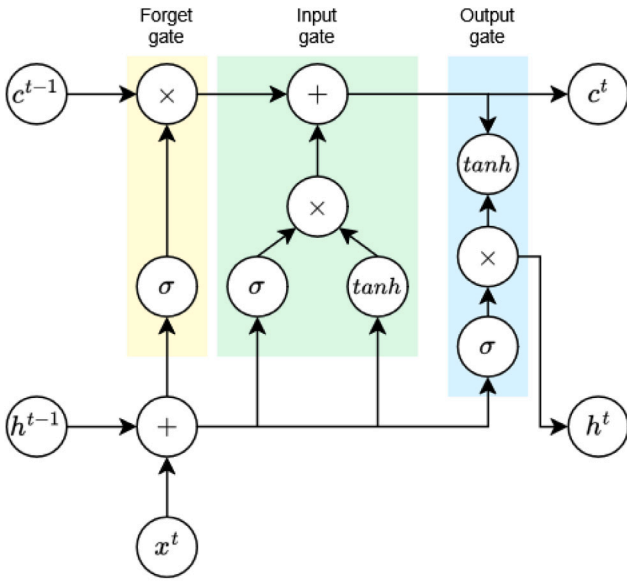


Fig. 9. The architecture of an LSTM unit.

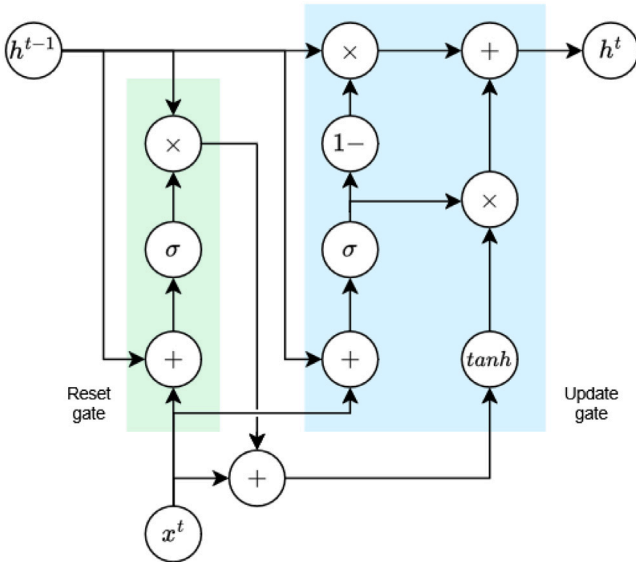


Fig. 10. The architecture of an LSTM unit.

not maintain a separate internal cell state, but uses the reset gate to determine which parts of the information to be retained and discarded, and the update gate to control how much of the previous hidden state should be passed to the current hidden state. Fig. 10 illustrates the interconnection of the gates in a GRU unit.

### 3.1.3. Convolutional neural network

Convolutional neural network (CNN) is a neural network model that preserves and leverages the spatial local information in the data through the use of convolutional layers. Fig. 11 shows a typical architecture of a convolutional neural network which consists of convolutional, pooling and fully connected layers. The convolutional and pooling layers are stacked alternately to automatically extract salient features in a hierarchical manner. The extracted features are then fed to fully connected layers to predict the outputs. The final feature maps need to be converted to a one-dimensional vector before they are fed to the fully connected layers. The conversion can be performed by

flattening the feature maps. CNN architecture is crucial in increasing the performance of the prediction, as it is designed to efficiently extract the feature representation of the input data, enabling more accurate and robust pattern recognition. Over the last decade, several CNN architectures have been proposed, whereby the focus of the improvements has been on enhancing the feature learning capabilities and addressing challenges such as vanishing gradient and diminishing feature reuse.

AlexNet is among the first CNN models that gained widespread recognition and success, marking a significant achievement in the field of deep learning for computer vision tasks (Krizhevsky et al., 2012). The model consists of five convolutional layers with maximum pooling operation performed after the first and second convolutional layers, followed by three fully connected layers. The first and second convolutional layers utilize a filter size of  $11 \times 11$  and  $5 \times 5$  respectively, and  $3 \times 3$  filter size is used for the remaining convolutional layers. ReLU activation function is used to mitigate the vanishing gradient. AlexNet is the first deep learning architecture that demonstrated CNN's capability for large-scale image recognition. While 8 layers was considered deep for its time, later deep learning architectures demonstrated that even deeper networks could achieve better performance.

ZFNet is a classic CNN model which has a similar architectural principle as AlexNet, featuring five convolutional layers with maximum pooling layers after the first and second convolution, followed by three fully connected layers (Zeiler and Fergus, 2013). The significant differences are the use of smaller filter size and stride in the convolutional layers and contrast normalization of the feature maps, which allows the model to capture better features and improve the overall performance. This configuration improves feature extraction and performance compared to AlexNet. However, like AlexNet, its limited depth restricted its ability to learn more complex hierarchical features.

Network-in-network introduces two innovative concepts to enhance the performance of the model (Lin et al., 2014). The first was introducing a block of convolutional layers consisting of  $k \times k$  convolution followed by two  $1 \times 1$  convolution operations. The pointwise convolutions are similar to applying a multilayer perceptron on the feature maps, allowing the model to approximate more abstract feature representations. In the preceding models, the final feature maps are vectorized by flattening operation for classification by the fully connected layers. Instead of flattening, network-in-network model calculates the spatial average of each feature map, and the resulting vector is fed to softmax function for classification. This approach is parameter-less, significantly reducing the number of parameters. Although this approach makes it more computationally efficient, its general performance does not always surpass models with deeper architecture.

VGGNet attempts to improve the CNN architecture by adding more convolutional layers, specifically up to 19 layers to capture more intricate feature representation from input data, followed by three fully connected layers (Simonyan and Zisserman, 2015). ReLU activation function is used to reduce vanishing gradient. Unlike AlexNet, all convolutional layers utilize a small fix filter size of  $3 \times 3$  and maximum pooling layer is added after a stack of two or three convolutional layers. This configuration allows the model to extract more discriminative features and decreases the number of parameters. The architecture provides better generalization and finer feature extraction, but it comes at the cost of a large number of parameters (138 million), making it more computationally expensive for practical applications.

GoogleNet or Inception v1 leverages the fact that visual data can be represented at different scales by incorporating a module which consists of multiple convolutional pipelines with different filter sizes (Szegedy et al., 2014). The module known as inception utilizes three kernel sizes ( $5 \times 5$ ,  $3 \times 3$ ,  $1 \times 1$ ) to capture spatial and channel information at different scales of resolution as shown in Fig. 12. This configuration enables a more effective feature extraction at both fine-grained and coarse-grained information from input data while maintaining computational efficiency. The model architecture utilizes

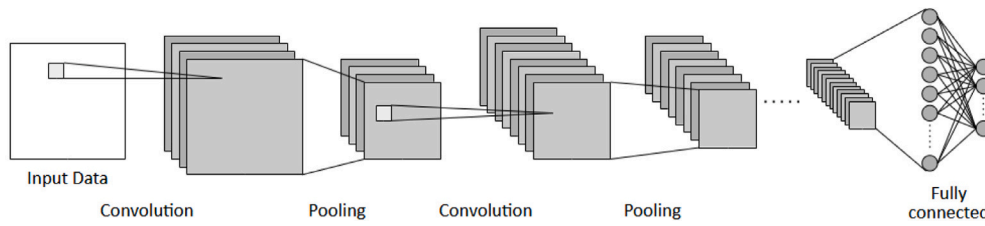


Fig. 11. A neural network with convolutional and pooling layers followed by fully connected layers.

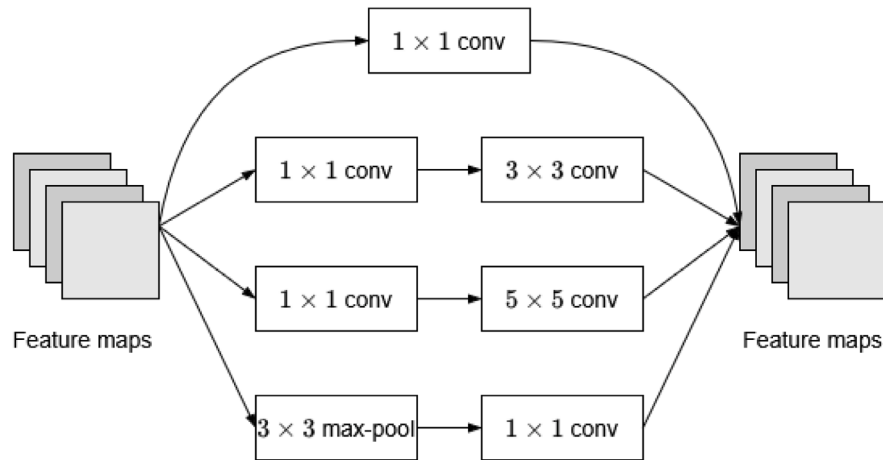


Fig. 12. The inception module.  
Source: Adapted from Szegedy et al. (2014).

the inception module at the higher layers, while the traditional convolution and maximum pooling block is used to extract primitive and basic features. The inception modules are stacked upon each other, with maximum pooling operation is performed occasionally to reduce the spatial resolution of the feature maps. GoogleNet utilizes global average pooling to vectorize the final feature maps before passing it to a fully connected layer for classification. However, the complexity of the model architecture may hinder its interpretability and tuning compared to simple models. GoogleNet has been further enhanced in later Inception models by introducing batch normalization, auxiliary classifier and deeper architecture.

Increasing the number of layers enhances the model performance, mainly for solving complex tasks. However, training a very deep neural network is challenging due to the vanishing gradient problem, where the gradients that are used to update the network become insignificant or extremely small as they are backpropagated from the output layer to the earlier layers. A model called Highway Network overcomes this issue by introducing a gating mechanism that regulates the information flow of the layers, enabling the flow of information from the earlier layers to the later layers (Srivastava et al., 2015). Consequently, this not only mitigates the vanishing gradient problem, but also renders the gradient-based training more tractable, enabling the training of very deep neural networks consisting as many as 100 layers. However, the gating mechanism increases the model complexity, making the model more resource-intensive and less suitable for real-time applications.

The gating mechanism of Highway Network increases the number of parameters for regulating the information flow. ResNet is a CNN architecture that incorporates residual (skip) connection that allows information to bypass certain layers, mitigating the vanishing gradient problem (He et al., 2015). ResNet architecture stacks residual blocks, which consists of two or three of convolutional layers with batch normalization and ReLU, and a skip connection which adds the input to the output of the final convolutional layer as shown in Fig. 13. If the input dimension does not match with the residual output

dimension, a linear projection is performed by the residual connection to match the dimensions. This concept of “feature reuse” is central to ResNet’s design, as the skip connection allows the features learned in the previous layers to be directly reused in the later layers, enhancing the model’s ability to learn hierarchical features. In comparison to the gating mechanism of Highway Network, the residual connection is parameter-free, and thus does not incur additional computational costs. Furthermore, the connections are never closed whereby all information is always passed through the layers. This innovative concept enables the training of very deep neural networks boasting as many as 152 layers. However, as the model gets deeper, the effectiveness of skip connection diminishes and the performance gain is negligible. Furthermore, ResNet can be more complex to implement and fine-tune compared to simpler architecture.

DenseNet is another CNN architecture that overcomes the vanishing gradient problem. DenseNet follows the same approach as ResNet and Highway Network, utilizing skip connection to allow information flow from the earlier layers to later layers. However, DenseNet takes this concept one step further, by introducing a dense block consisting of multiple convolution functions (layers) with each convolution function performs batch normalization followed by ReLU and  $3 \times 3$  convolution. Each convolutional layer in the dense block receives feature maps from all its preceding layers. Hence, the connection is referred to as a dense connection (Huang et al., 2017). This configuration as shown in Fig. 14 maximizes information flow, feature reuse and preserves the feed-forward nature of the network, improving the feature learning. To reduce the computational costs, a block of  $1 \times 1$  convolutional with batch normalization and maximum pooling layers known as transition block is used to reduce the spatial dimension of the feature maps. The model architecture integrates these dense and transition blocks, stacking them alternately. The network depth can reach up to 264 layers. However, Dense blocks can become computationally expensive due to the increasing number of feature maps. Furthermore, it is prone to overfitting due to the dense connection between the layers.

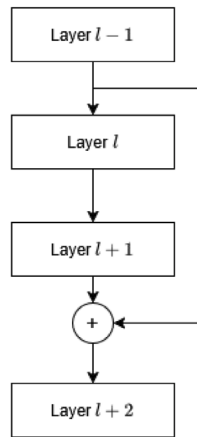


Fig. 13. A residual connection which explicitly incorporates features from previous layers into the later layers.

Source: Adapted from He et al. (2015).

Although skip connections in ResNet effectively mitigate the vanishing gradient problem, a new challenge arises in the form of diminishing feature reuse as the network becomes deeper. Diminishing feature reuse refers to the diminishing effectiveness of the previously learned feature maps in subsequent layers, impacting the final prediction. WideResNet is a CNN architecture that is based on ResNet with the aim to mitigate diminishing feature reuse problem. Instead of making the network deeper, WideResNet makes the network wider by increasing the number of channels by  $k$  factor (Zagoruyko and Komodakis, 2017). The increased width allows the model to capture more diverse features, enhancing its ability to learn complex relationships in the input data. This configuration improves feature learning efficiency while mitigating depth-related issues such as vanishing gradients and overfitting. However, wider models may increase memory usage and computational costs.

ResNeXt addresses the diminishing feature reuse by capturing more efficient and diverse features of the input data. ResNeXt introduces a concept of cardinality which is loosely based on the inception module as shown in Fig. 15. Cardinality refers to the number of independent and identical paths, where each path performs transformation of the input data, divided along the channel dimension (Xie et al., 2017). In other words, instead of solely relying on increasing the depth of the model, ResNeXt enhances the feature learning by parallelizing the feature extraction through this cardinal path. In the proposed architecture, each path configuration is similar to the residual block of ResNet. The output from each path is then aggregated to form a comprehensive and diverse representation of the input data. The skip connection is used to mitigate the vanishing gradient problem. ResNeXt improves feature learning efficiency by balancing the depth and the width of the model, but its increased architectural complexity introduces additional parameters such as the number of cardinality, and increases memory usage and computational costs.

These CNN models were trained using the ImageNet (ImageNet, n.d.) and/or CIFAR (CIFAR, n.d.) datasets. ImageNet is considered as the most influential and important dataset in deep learning for computer vision research. The dataset was used to train all popular CNN models such as VGGNet, ResNet and DenseNet due to its large number of labeled images. The dataset contains 1.2 million training images, 50,000 validation images and 100,000 test images across 1000 object classes. CIFAR is a comparatively smaller dataset and includes two variants: CIFAR-10 and CIFAR-100. Both contain 50,000 training images and 10,000 test images, with CIFAR-10 covering 10 object classes and CIFAR-100 covering 100 classes.

Figs. 16 and 17 show the top five classification error rates of the CNN models on ImageNet and CIFAR-10, respectively. Error rates

for CIFAR-100 are not reported, as it was used only by Network-in-network, DenseNet, Wide ResNet and ResNeXt. It is also noted that the error rates on ImageNet are sourced from the models' published manuscripts, which may differ from the values reported on the ILSVRC website. As shown in the figures, error rates are generally similar between the datasets, despite their differences in complexity. This could be attributed to ImageNet's large number of training images, allowing the models to generalize well across its 1000 classes. Furthermore, it is evident that the newer models, such as ResNeXt, DenseNet, Inception-v4 and ResNet significantly outperform older models due to their more advanced architectures. Among the evaluated models, DenseNet achieved the best performance on CIFAR-10 with a 3.46% error rate, while Inception-v4 achieved the lowest error rate on ImageNet at 3.70%.

### 3.1.4. Applications of deep supervised learning models

The deep supervised learning models have become integral to advancements in various applications such as medical imaging, internet of things and robotics due to their ability to extract complex hierarchical features. Typically, the studies utilize transfer learning to build the predictive models for solving the problem at hand, especially in scenarios with limited labeled data. This approach leverages pre-trained models to accelerate model development and improve accuracy. In Chougrad et al. (2018), VGGNet-16, Inception-v3 and ResNet-50 are employed as pre-trained models for breast cancer screening. Specifically, the weights of the convolutional layers are frozen, and the fully connected predicting layers are replaced with new ones. The models are then fine-tuned to adapt to the specific characteristics of the breast cancer dataset, improving their performance on the target task. In medical image segmentation, ResNet (Zhao et al., 2020) and DenseNet (Cinar et al., 2022) have been utilized to progressively capture rich spatial features from the input images, systematically reducing their dimensions while preserving essential information. These compressed feature representations are then gradually expanded and refined through upsampling operations, followed by localizing and predicting the segmentation regions.

In other studies, the architecture of the CNN models have been enhanced to improve their feature extraction capability. In Sun et al. (2023), ResNet architecture is modified by replacing ReLU with LeakyReLU and incorporating self-attention before the predicting layers to assign weights based on the feature importance. In another study, the residual block of ResNet-18 is enhanced by integrating a channel attention before the skip connection addition (Dong et al., 2023). Additionally, a channel attention featuring two parallel processing pipelines is introduced before the predicting layers. One pipeline applies maximum pooling while the other uses average pooling, allowing the model to capture diverse features. Each pooled output is then processed by convolutional layers, and the resulting feature maps are concatenated to combine the extracted features for prediction. In Hou et al. (2024), DenseNet is modified to reduce feature reuse by controlling the number of feature maps inputs to each layer using a parameter. This parameter controls the number of feature maps based on the distance between layers, whereby if the distance between two layers is large, the input feature maps are reduced to half of the original number. Additionally, the network width is modified based on depth, leading to a gradual widening of the network as it deepens. The modification reduces the complexity of DenseNet, improving both memory usage and computational efficiency.

The LSTMs and GRUs have been exploited to capture temporal information in time series data. In Narotamo et al. (2024), AlexNet, VGG16 and ResNet50 are used to extract spatial features from ECG images, while LSTM and GRU are used to extract temporal features from the ECG signals. The features are then combined and passed to a self-attention module, followed by fully connected layers for predictions. In Zhang et al. (2024), 1D-CNN is hybridized with LSTM to exploit the spatial feature extraction capabilities of convolutional layers and the

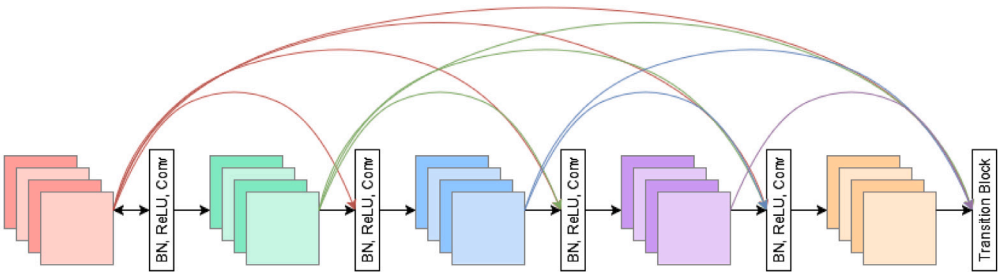


Fig. 14. A 5-layer dense block.  
Source: Adapted from Huang et al. (2017).

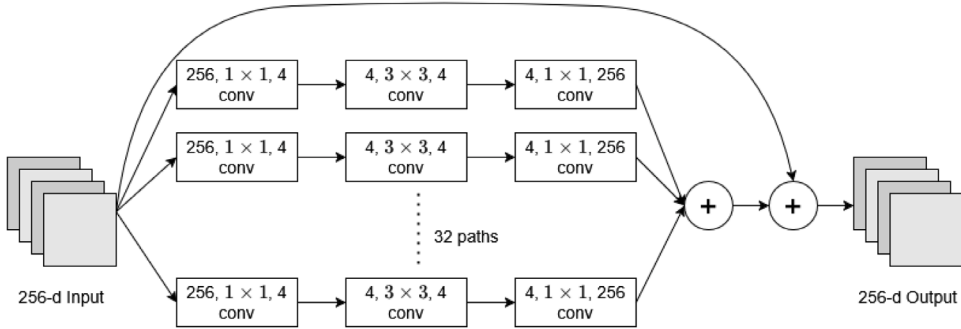


Fig. 15. A cardinal block.  
Source: Adapted from Xie et al. (2017).

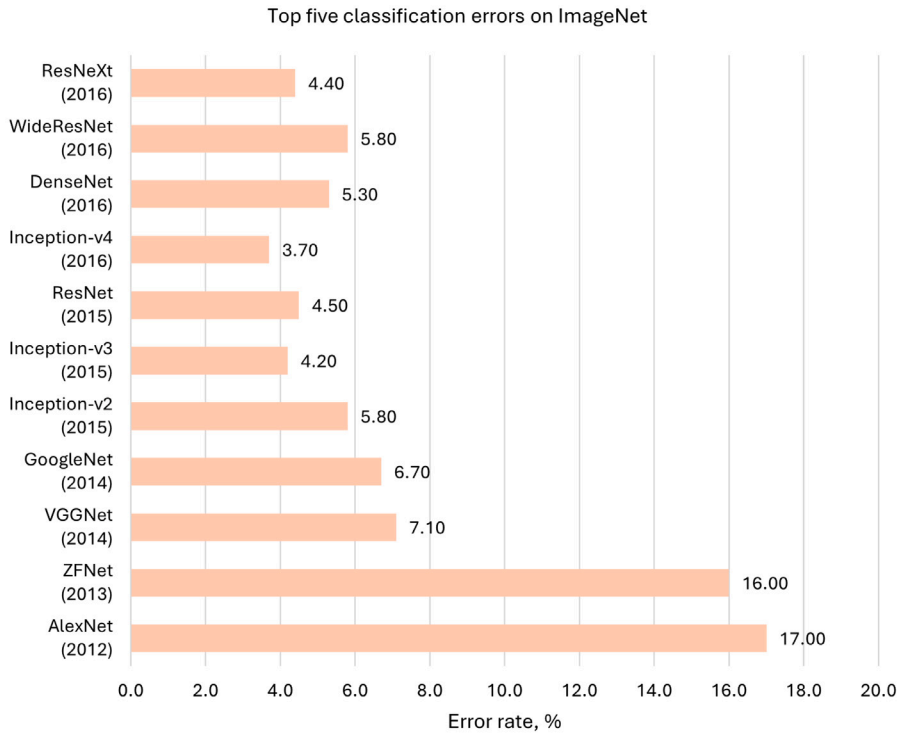


Fig. 16. Top five classification errors on ImageNet.

temporal sequence learning capabilities of LSTM. The time series data is first processed by a series of convolutional and maximum pooling layers to extract spatial features. These features are then fed into 10 layers of LSTM capture temporal dependencies and sequential patterns within the data, improving the performance of the predictions. A hybrid deep

learning model with parallel feature learning pipelines is introduced for classifying motion signals into human activities (Mohd Noor et al., 2022). The feature learning pipelines consist of convolutional and maximum pooling layers to extract local features from the sequence of segmented signals (windows). The local features are concatenated

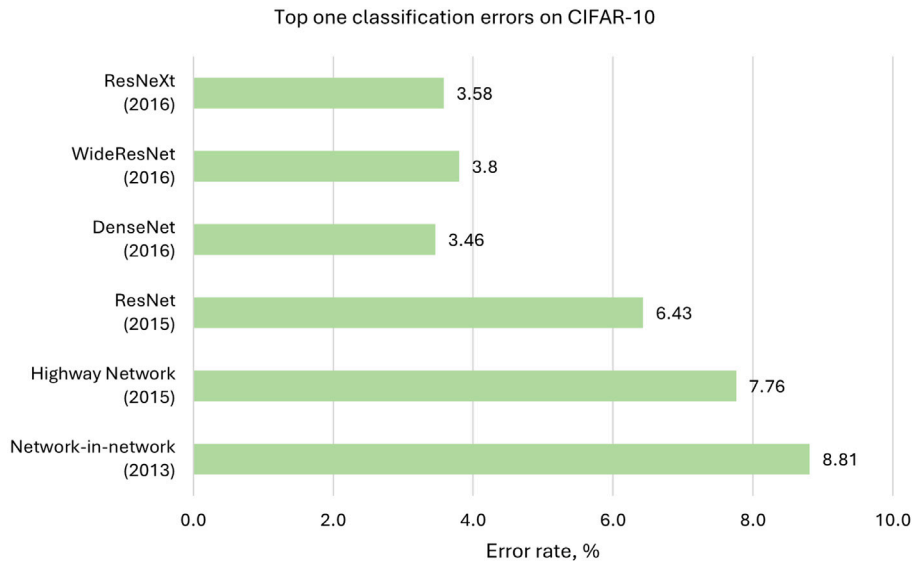


Fig. 17. Top one classification errors on CIFAR-10.

to form a sequence, which is then fed into an LSTM to capture temporal dependencies and patterns. A similar hybrid model is proposed, combining ResNet-18 with an LSTM for EEG signal classification. The feature maps produced by ResNet-18 are flattened and fed into the LSTM, followed by a fully connected layer for prediction.

### 3.2. Deep unsupervised learning

Deep unsupervised models are trained with an unlabeled dataset. The learning process of these models discovering patterns, structures, and representations within the data without relying on explicit labels or supervision. Instead, these models often learn by optimizing objective functions that capture the underlying data characteristics such as clustering, learning useful feature embeddings and reconstructing input data from compressed representations. Examples of deep unsupervised models are autoencoders, generative adversarial networks and restricted Boltzmann machines.

Restricted Boltzmann Machine is a generative neural network model that learns a probability distribution based on a set of inputs. The model consists of a visible (input) layer and a hidden layer with symmetrically weighted connections as shown in Fig. 18. The input layer represents the input data with each node corresponding to a feature or variable while the hidden layer learns the abstract representation of the input data. Restricted Boltzmann machine model is trained using contrastive divergence, an algorithm that is based on a modified form of gradient descent, utilizing a sampling-based approach to estimate the gradient (Hinton, 2012). It has found success in solving combinative problems such as dimensionality reduction, collaborative filtering and topic modeling.

Deep Belief Network can be viewed as a stack of restricted Boltzmann machines, comprising a visible layer and multiple hidden layers (Hinton et al., 2006) as shown in Fig. 19. Deep belief network has two training phases. The initial phase is known as pretraining in which the network is trained layer by layer, with each layer serves as a pre-training layer of the subsequent layers. This sequential learning allows the hidden layers learn complex hierarchical feature representation of the data. The second phase is called fine-tuning whereby the deep belief network model can be further trained with supervision to perform tasks such as classification and regression (Hinton, 2009).

An autoencoder is a generative neural network model that learns to encode the input data into a compressed representation and then reconstructs the original data from this representation. The layers that encode the input data are known as encoder while the layers that

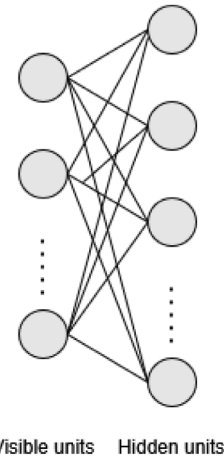


Fig. 18. A restricted Boltzmann machine.

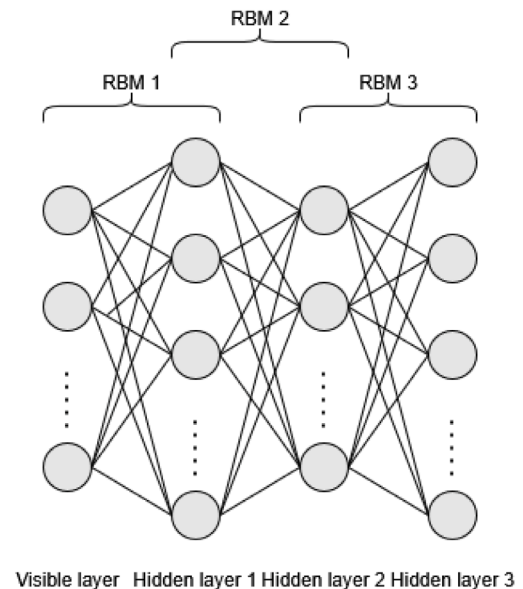


Fig. 19. A deep belief network.

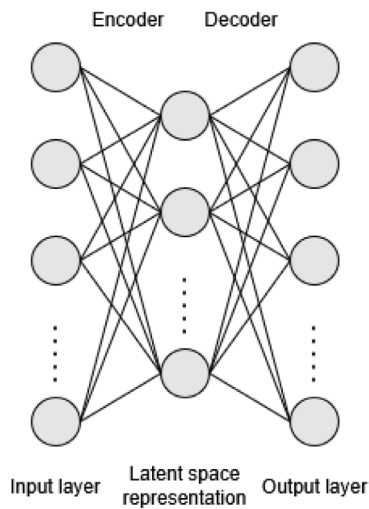


Fig. 20. An autoencoder.

responsible for the reconstruction are referred to as the decoder as shown in Fig. 20. The encoded data (hidden layer) represents the abstract features of the input data, also known as latent space or encoding. The decoder can be removed from the autoencoder, creating a standalone model that can be used for data compression and dimensionality reduction (Romero et al., 2017; Li et al., 2020). The decoder can also be replaced with predictive layers for classification task (Mohd Noor, 2021). The network architecture can be built using fully connected and convolutional layers (Li et al., 2023).

Several autoencoder variants have been introduced to improve the autoencoder's ability to capture better feature representation. Some introduced penalty terms to the loss function such as sparsity penalty (sparse autoencoder) (Ng et al., 2011) to encourage sparse representation and Jacobian Frobenius norm (contractive autoencoder) (Rifai et al., 2011) to be less sensitive to small and insignificant variations in the input data while encoding the feature representation. Others trained the autoencoder to recover original data from corrupted data with noise (Vincent et al., 2008). An improved denoising autoencoder knowns marginalized denoising autoencoder has been proposed which marginalizes the noise by adding a term that is linked to the encoding layer (Chen et al., 2012). Variational autoencoder is a variant of autoencoder that has similar architecture: encoder, latent space and decoder. Despite the similarity, instead of learning a fixed encoding, variational autoencoder learns the probability distribution of the input data in the latent space (Kingma and Welling, 2022). The model can be used to generate data by sampling from the learned probability distribution. The network architecture can be built by stacking more than one fully connected layer and convolutional layer.

Generative Adversarial Network (GAN) is another generative neural network model that is designed for generating data that adheres closely to the distribution of the original training set. The model consists of two different neural networks namely generator and discriminator as shown in Fig. 21. The generator learns to imitate the distribution of the training set given a noise vector, effectively outsmarting the discriminator. Simultaneously, during the training, the discriminator is trained to differentiate between the real data from the training set and synthetic data generated by the generator (Goodfellow et al., 2014). This intricate dynamic between the networks drives an iterative learning process whereby the generator continually refines its ability to create synthetic data that closely resembles the real data, while the discriminator enhances its ability to distinguish between authentic and fake data.

The model can be extended by providing the labels to both generator and discriminator in which the model known as conditional

GAN, capable of generating 1000 image classes (Odena et al., 2017). Conditional GANs require a labeled dataset, which might limit its application. InfoGAN is similar to conditional GAN, but the labels are substituted with latent codes, which allows the model to be trained in an unsupervised manner (Chen et al., 2016). GANs often suffer from mode collapse, where the model can only generate a single or small set of outputs. Wasserstein GAN improves the training by utilizing Wasserstein loss function, which measures the difference between the real and synthesized data distribution (Weng, 2019). ProGAN tackles the training instability of GAN by progressively growing the generator and discriminator. The idea is that the model is scaled up gradually, starting with the simplest form of the problem, and little by little the problem's complexity is increased as the training progresses (Karras et al., 2018). StyleGAN leverages the progressive GAN's approach and neural style transfer to improve the quality of the generated data (Karras et al., 2019). The model is characterized by the independent manipulation of both style and content, allowing it to generate diverse styles and high-quality data.

#### 4. State-of-the-art deep learning applications

As discussed in the previous section, the application of deep learning ranges from computer vision (Tan et al., 2020), natural language processing (Otter et al., 2021), healthcare (Esteva et al., 2019), robotics (Soori et al., 2023), education (Hernández-Blanco et al., 2019), and many others. This section presents the applications of deep learning across several areas.

##### 4.1. Computer vision

Computer vision is an essential field in artificial intelligence (AI). It is a field of study that focuses on enabling computers to acquire, analyze and interpret visual inputs to derive meaningful information. The visual inputs can take many forms such as digital images, sequence of images or video and point cloud, and the source of these inputs can be camera, LiDAR and medical scanning machine. Deep learning, specifically CNN models have been widely used in real-world computer vision applications including image classification, object detection and image segmentation. This section discusses more details about the recent advancements in deep learning models that have been achieved over the past few years.

##### 4.1.1. Image classification

Image classification is a fundamental task in computer vision, which involves categorizing an image into one of predefined classes based on the visual content. The objective of image classification is to enable computers or machines to differentiate between objects within images, in a manner similar to how humans interpret visual information. Image classification is a crucial component in various applications such as robotics, manufacturing, and healthcare. LeNet-5, introduced in 1998, is one of the earliest convolutional neural networks that was successfully trained to classify handwritten digits. The model underwent a series of improvements, including the use of tanh and average pooling, which enhanced its ability to extract hierarchical features, ultimately improving overall performance. The model architecture comprises two convolutional layers, each with an average pooling layer, followed by two fully connected layers, including the output layer (Lecun et al., 1998). Since then, numerous CNN models have been proposed based on LeNet-5 for image classification (Simard et al., 2003; Matsugu et al., 2003) but the most significant one is AlexNet in 2012, which saw a transformative breakthrough in deep learning. AlexNet is considered the first CNN model with a large number of parameters that significantly improved the performance of image classification on a very large dataset (ImageNet). The model won first place in ILSVRC 2012, improving the test error from the previous year by almost 10% (Krizhevsky et al., 2012). Numerous significant

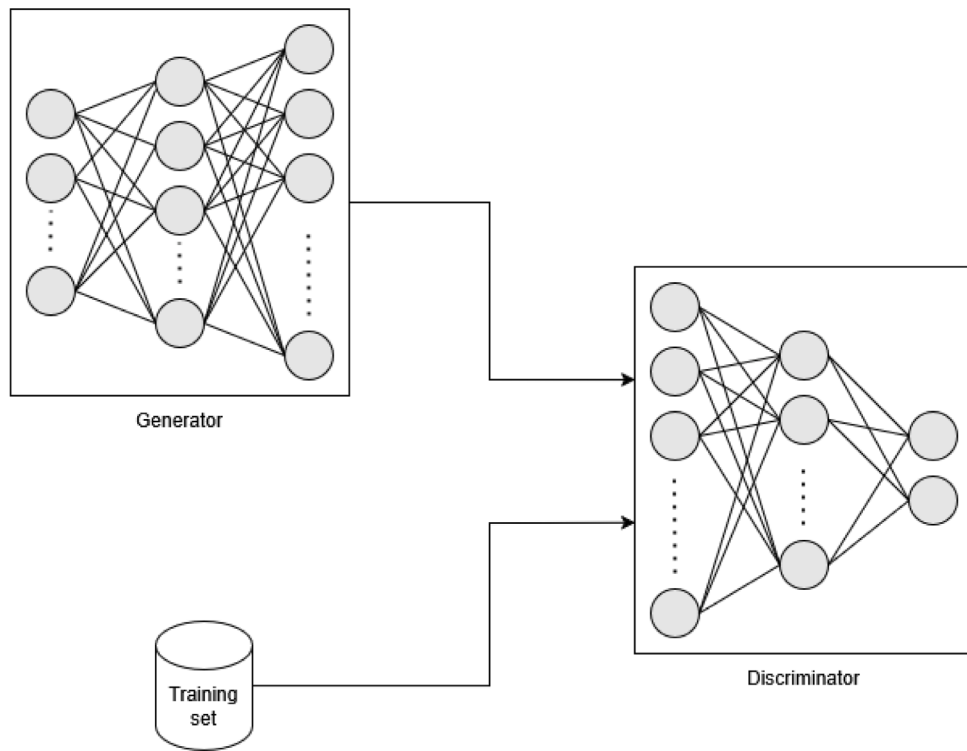


Fig. 21. A generative adversarial network.

CNN models have been introduced in subsequent ILSVRC competitions including ZFNet, VGG16, GoogleNet, ResNet and ResNext. In general, the research focused on increasing the number of layers, addressing the problem of vanishing gradient and diminishing of feature reuse.

Research in image classification continues to evolve with a focus on addressing key challenges to improve the classification performance. One notable trend is the formulation of the loss function to address problems such as neglecting well-classified instances and imbalance distribution of class labels. In a particular study, an additive term is introduced to the cross-entropy loss to reward the models for the correctly classified instances. This formulation encourages the models to also pay attention to well-classified instances while focusing on the bad-classified ones (Zhao et al., 2022). Another study proposes an asymmetric polynomial loss function using the Taylor series expansion. The loss function allows the training to selectively prioritize contributions of positive instances to mitigate the issue of imbalance between negative and positive classes (Huang et al., 2023). The asymmetric polynomial loss requires a large number of parameters to be fine-tuned and may lead to overfitting. A robust asymmetric loss is formulated by introducing a multiplicative term to control the contribution of the negative gradient and making it less sensitive to parameter optimization (Park et al., 2023). Combining multiple deep learning models improves the overall performance by leveraging the diverse strengths of individual models. However, identifying the optimal combination is non-trivial due to the large number of hyperparameters. A straightforward method is to employ the weighted sum rule (Nanni et al., 2023). To enhance the overall performance, an algorithm, named greedy soups, adds a model based on the validation accuracy (Wortsman et al., 2022). The final prediction is produced via averaging. Multi-symmetry ensembles framework improves the building of diverse deep learning models by utilizing contrastive learning (Loh et al., 2023). Then, the diverse models are sequentially combined based on their validation accuracy.

Vision transformers (ViT) offers an alternative to convolutional neural networks that have long been the dominant architecture for image classification, by leveraging self-attention mechanisms for scalable

global representation learning. Despite its effectiveness, ViT is sensitive to hyperparameter optimization and substandard performance on smaller datasets (Xiao et al., 2021). Furthermore, ViT lacks the ability to leverage local spatial features which is inherent in convolutional neural networks (Wu et al., 2021). Therefore, several studies attempt to incorporate convolutional layers into ViT architecture to improve its performance and robustness. In particular, conformer is a network architecture with two branches: CNN branch and a transformer branch to extract local and global features respectively (Peng et al., 2023). Both branches are connected by two “bridges” of  $1 \times 1$  convolution and up or down sampling operations, allowing the branches to share their features and enhance the feature representation. Both branches output predictions which are combined to produce the final prediction.

A hybrid architecture, named MaxViT, combines convolutional networks and vision transformer to address the lack of scalability issues of self-attention mechanisms when the model is trained on large input size (Tu et al., 2022). The improved vision transformer is composed of two modules whereby the first module attends local features in non-overlapping image patches and the global features are attended by processing a grid of sparse and uniform pixels. The transformer is stacked with a block of convolutional layers to extract local spatial features. The architecture of MaxViT is shown in Fig. 22. Another study proposes a convolutional transformer network, introducing the depthwise convolutional block into the ViT (Ma et al., 2024). This configuration allows the model to exploit the ability of convolutional networks to extract local spatial features while the ViT attends the extracted local features to focus on relevant information, enhancing the model’s ability to capture complex patterns and relationships. Specifically, instead of linear mapping, depth-wise convolutional mapping is used to generate query, key, and value matrices. Table 2 lists the summary of state-of-the-art image classification.

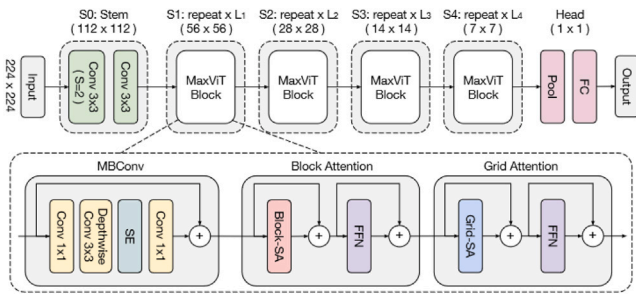
#### 4.1.2. Object detection

Deep learning plays a major role in significantly advancing the state-of-the-art in object detection performance. Region-based CNN (R-CNN) is the first breakthrough in object detection that combines CNN

**Table 2**

Summary of state-of-the-art image classification.

Ref.	Description	Results (Datasets and metrics)
Zhao et al. (2022)	The study introduced an encouraged loss that rewards well-classified examples with additive bonuses to enhance their contribution during training, addressing issues in representation learning, energy optimization, and makes the decision margin grows faster	CIFAR-10 Accuracy: 92.97% ImageNet Accuracy: 76.43%
Huang et al. (2023)	The study introduced asymmetric polynomial loss to address class imbalance by decoupling gradient contributions from positive and negative instances, increasing the impact of updates for minority instances and helping the model focus on the more challenging, less frequent cases	MS-COCO mAP: 90.08 NUS-WIDE mAP: 31.27
Park et al. (2023)	The study introduced robust asymmetric loss to address class imbalance by emphasizing minority instances through an asymmetric weighting mechanism, allowing the model to focus more on the less frequent, yet critical minority instances	ISIC2018 Accuracy: 0.852 APTOS2019 Accuracy: 0.826
Peng et al. (2023)	This study introduced dual network structure to extract local and global feature representations using convolutional layers and transformer	ImageNet Accuracy: 81.3 (Conformer-Ti) Accuracy: 83.4 (Conformer-S) Accuracy: 84.1 (Conformer-B)
Tu et al. (2022)	This study introduced MaxViT, a hybrid attention model, combining convolutional layers with transformer-based architecture whereby it uses convolutional layers in the early stages to capture local patterns efficiently, and then applies a multi-axis attention mechanism to capture both local and global dependencies more effectively	ImageNet Accuracy: 85.72 (MaxViT-T) Accuracy: 86.19 (MaxViT-S) Accuracy: 86.66 (MaxViT-B) Accuracy: 86.70 (MaxViT-L)
Ma et al. (2024)	This study introduced convolutional transformer network that integrates convolutional layers and transformer encoder blocks whereby a transformer encoder block consists of multi-head self attention and feed-forward network layers with skip connections	Diving 48 Accuracy: 77.3 (CTN-a), 78.2 (CTN-b) Epic-Kitchens Accuracy: 44.1 (CTN-a), 45.9 (CTN-b)

**Fig. 22.** The architecture of MaxViT (Tu et al., 2022).

with selective region proposals (Girshick et al., 2014). The region proposals are the candidate bounding boxes serving as the potential region of interests (objects) within the input image, and the CNN are used to extract features from the region proposals and classify the regions for object detection. An improved model, named Fast R-CNN introduces two prediction branches: object classification and bounding box regression which improves the overall performance of object detection (Girshick, 2015). However, R-CNN and Fast R-CNN models are computationally expensive and slow, thus practically infeasible for real-time applications. Addressing this issue, Fast R-CNN is integrated with a region proposal network, referred to as Faster R-CNN (Ren et al., 2015). The region proposal network (RPN) is used to efficiently generate region proposals for object detection. RPN takes an input image and output a set of rectangle object proposals, each with a confidence score to indicate the likelihood of an object's presence. To this end, RPN introduces the concept of anchor boxes, whereby multiple bounding boxes of different aspect ratios are defined over the feature maps produced by the convolutional networks. These anchor boxes are then regressed over the feature maps to localize the objects, contributing to the improved speed and effectiveness of Faster R-CNN. The training of Faster R-CNN is divided into two stages. First, the RPN is pre-trained to generate the region proposals and then, the Fast R-CNN is trained using the region proposals generated by the RPN for object detection. The backbone network responsible for extracting the features for Faster R-CNN is either ZFNet or VGG16.

In two-stage object detectors, the region proposals are generated first, and then used for object detection. The two-stage process is computationally intensive and infeasible for real-time object detection applications. You Only Look Once or YOLO proposes a one-stage detection by directly predicting bounding boxes and object's confidence score in a single forward pass through the neural network (Redmon et al., 2016). This single pass architecture significantly reduces the computational complexity, making YOLO suitable for real-time object detection applications. In YOLO, the input image is divided into  $S \times S$  grids, each grid cell is responsible for detecting the objects present in the cell. Specifically, each grid cell predicts multiple bounding boxes and associated object's confidence score, enabling simultaneous object detection across the entire image. Subsequent enhancements such as YOLOv3 (Redmon and Farhadi, 2018) and YOLOv4 (Bochkovskiy et al., 2020) are proposed, improving the model's capability and accuracy. Single Shot Multibox Detector (SSD) is another one-stage detector, which aims to address the issue of real-time object detection (Liu et al., 2016). SSD also eliminates the region proposal generation and directly predicts bounding boxes and confidence scores, reducing the computational complexity. To improve the overall performance, SSD produces the predictions from different levels of feature maps, allowing detection of objects of different sizes in the input image.

A common issue in object detection problems is the extremely imbalanced ratio of foreground to background classes. Addressing this issue, RetinaNet introduces a loss function that is based on the cross entropy called focal loss. Focal loss reduces the loss contribution of easily classified objects, allowing the model training to focus on the difficult objects (Lin et al., 2020). RetinaNet adopts the Feature Pyramid Network (FPN) (Lin et al., 2017) with ResNet as the backbone network for extracting the feature maps. FPN is a network architecture with a pyramid structure that efficiently captures multiscale feature representation, facilitating object detection across various sizes. To further improve the overall performance, EfficientDet introduces bidirectional FPN, which incorporates multi-level feature fusion to better capture multiscale feature representation (Tan et al., 2020). Also, the model utilizes EfficientNet (Tan and Le, 2019) as the backbone network to achieve a balance between computational efficiency and accuracy.

Object detection performance often relies on a post-processing step called non-maximum suppression (NMS) to eliminate duplicate detections and select the most relevant bounding boxes. Specifically, NMS

sorts all detection boxes based on their confidence scores, selects a box with the maximum score and discards the other boxes with a significant overlap with the selected box. This process is repeated on the remaining detection boxes. However, due to the inconsistency between the confidence score and the quality of object localization, NMS retains poorly localized bounding boxes with high confidence score while discarding more accurate predictions with poor confidence score. To mitigate this limitation, instead of discarding the neighboring boxes with significant overlap, soft-NMS applies Gaussian function to lower their confidence scores (Bodla et al., 2017). The idea is not to discard the neighboring bounding boxes, but gradually decline their scores based on the extent of the overlap with the selected box. This results in a smoother suppression, preserving the better-localized bounding boxes.

Adaptive NMS introduces an adaptive threshold for the suppression of bounding boxes (Liu et al., 2019). The algorithm dynamically adjusts the threshold based on the level of overlapping of the selected box with the other bounding boxes. Similar method is reported in Husham Al-Badri et al. (2023) whereby an adaptive NMS is proposed by dynamically adjusting the suppression criteria based on intersection over union (IoU) values. The IoU values are compared with a threshold and applies an additional iteration to vote for each detected proposal, ensuring better distinction between closely positioned objects. The method helps preserve multiple high-confidence bounding boxes, mitigating the issue of incorrectly merging adjacent objects into one. Redundant bounding boxes are often not filtered out due to their low IoU values with the best bounding box, leading to high false positives. In Jiang et al. (2024), a learning-based NMS is proposed to reduce false positives by integrating the NMS process into the model's learning framework. Specifically, a novel NMS-aware loss is formulated that incorporates IoU to adjust bounding box weights, boosting negative attention for low IoU bounding boxes and enhancing positive weights for high IoU bounding boxes. Furthermore, regression assisted classification branch is introduced to aid classification by leveraging regression prediction relationships between bounding boxes and their best counterparts. In Shi et al. (2024), the authors formulated similarity distance metric is proposed to evaluate the similarity between bounding boxes. The proposed method considers the location and shape of the boxes, and adapts dynamically to different datasets and object sizes. This method improves label assignment by learning hyperparameters automatically, eliminating the need for manual tuning and enhancing detection performance for tiny objects detection.

Detection Transformer (DETR) is an end-to-end trainable object detection model that leverages the transformer architecture to eliminate the need for handcrafted components such as anchor boxes and non-maximum suppression (Carion et al., 2020). The self-attention mechanism of the transformer captures the global context and relationships between different parts of the image, allowing it to localize the objects and remove duplicate predictions. The model is trained with a set of loss functions that perform bipartite matching between the predicted and ground truth objects. DETR uses ResNet as backbone network. Despite the success of DETR in simplifying and improving object detection tasks, DeTR suffers from a long training time and low performance at detecting small objects due to its reliance on the self-attention mechanism of the transformer, which lacks a multiscale feature representation. To mitigate this limitation, Deformable DETR introduces a multiscale deformable attention module which can effectively capture feature representation at different scales (Zhu et al., 2020). Furthermore, the attention module leverages deformable convolution, allowing the model to adapt to spatial variation and capture more informative features in the input data.

Dynamic DETR addresses the same issues by utilizing a deformable convolution-based FPN to learn multiscale feature representation (Dai et al., 2021). The model replaces the transformer encoder with a convolution-based encoder to attend to various spatial features and channels. Moreover, an ROI-based dynamic attention is introduced at

the transformer decoder, allowing the model to focus on the region of interests. This modification allows the model to effectively detect small objects and converge faster during training. The architecture of dynamic DETR is shown in Fig. 23. DETR utilizes bipartite matching between the ground truth and the predicted objects to assign each ground truth object to a unique prediction, limiting the localization supervision. In Jia et al. (2023), a hybrid matching scheme is introduced, combining the original one-to-one matching with an auxiliary one-to-many matching during training, improving training efficiency and detection accuracy without adding inference complexity. A training scheme known as Teach-DeTR is proposed to improve the overall performance of DeTR (Huang et al., 2023). The training scheme leverages the predicted bounding boxes by other object detection models during the training by calculating the loss of one-to-one matching between the object queries and the predicted boxes. Table 3 lists the summary of state-of-the-art object detection.

#### 4.1.3. Image segmentation

Image segmentation is another important task in which deep learning has a significant impact. One of the earliest deep learning models for image segmentation is the fully convolutional network (Long et al., 2015). A fully convolutional network consists of only convolutional layers which accepts an input of an arbitrary size and produce the predicted segmentation map of the same size. The authors adopted the AlexNet, VGG16 and GoogleNet, replace their fully connected layers with convolutional layers and append a  $1 \times 1$  convolutional layer, followed by bilinear up-sampling to match the size of the input. The model was considered a significant milestone in image segmentation, demonstrating the feasibility of deep learning for semantic segmentation trained in end-to-end manner. Deconvolution network is another popular deep learning model for semantic segmentation (Noh et al., 2015). The model architecture consists of two parts: encoder and decoder. The encoder takes an input image and uses the convolutional layers to generate the feature maps. The feature maps are fed to the decoder composed of up-sampling and deconvolutional layers to predict the segmentation map. SegNet is another encoder-decoder model for semantic segmentation (Badrinarayanan et al., 2017). The encoder is a sequence of convolutional (with ReLU) and maximum pooling blocks which is analogous to a convolutional neural network. The decoder is composed of up-sampling layers which up-samples the inputs using the memorized pooled indices generated in the encoder phase, and convolutional layers without non-linearity. The encoder progressively reduces the resolution of the input data while extracting abstract features through a series of convolutional and pooling layers. This process causes the loss of fine-grained information, degrading the overall performance of segmentation. LinkNet mitigates this limitation by passing the feature maps at several stages generated by the encoder to the decoder, hence reducing information loss (Chaurasia and Culurciello, 2017). The model architecture of LinkNet is similar to SegNet, but utilizes ResNet as the encoder.

While Faster R-CNN is a significant approach in object detection task, it has been extended to perform instance segmentation task. One such extension is Mask R-CNN which is based on Faster R-CNN, introduces an additional branch for predicting the segmentation mask (He et al., 2017). Similar to Faster R-CNN, Mask R-CNN utilizes the RPN to generate region proposals and then the region of interest alignment is applied to extract more accurate features from the proposed regions. Mask R-CNN does not leverage the multiscale feature representation which may degrade the overall performance of segmentation. To overcome this limitation, Path Aggregation Network (PANet) incorporates the FPN and introduces a bottom-up pathway to facilitate the propagation of the low-level information (Liu et al., 2018). The pathway takes the feature maps of the previous stage as input and performs  $3 \times 3$  convolution with stride 2 to reduce the spatial size of the feature maps. The generated feature maps are then fused with the feature maps from the FPN through the lateral connection. The model adopts the

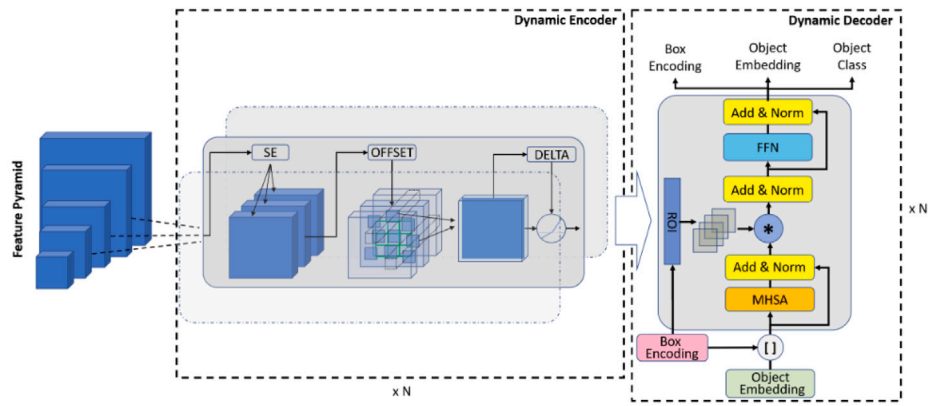


Fig. 23. The architecture of dynamic DeTR (Dai et al., 2021).

Table 3

Summary of state-of-the-art object detection.

Ref.	Description	Results (Datasets and metrics)
Husham Al-Badri et al. (2023)	This study introduced an adaptive NMS method that dynamically adjusts suppression criteria based on intersection over union (IoU) values, applying an additional iteration to improve the distinction between closely positioned objects and preserve high-confidence bounding boxes, reducing false positives	Rumex weeds IoU: 91.2, mAP: 90
Jiang et al. (2024)	This study proposed a learning-based NMS that integrates the NMS process into the model's learning framework through a novel NMS-aware loss function, which adjusts bounding box weights using IoU to reduce false positives, while introducing a regression-assisted classification branch to enhance classification by leveraging bounding box prediction relationships	CrowdHuman mAP: 90.1
Shi et al. (2024)	This study proposed a similarity distance metric to evaluate bounding box similarity by considering location and shape, dynamically adapting to different datasets and object sizes, automatically learning hyperparameters to improve label assignment, and enhancing detection performance for tiny objects	AI-TOD mAP: 26.5, AP50: 57.7, AP75: 20.5 SODA-D mAP: 32.8, AP50: 59.4, AP75: 31.3 VisDrone2019 mAP: 28.7, AP50: 50.3
Dai et al. (2021)	This study introduced Dynamic DETR, which enhances the original DETR by introducing dynamic attention mechanisms in both the encoder and decoder stages. This approach addresses limitations related to small feature resolution and slow training convergence, resulting in improved performance and efficiency in object detection tasks	COCO2017 mAP: 49.3, AP50: 68.4, AP75: 53.6
Jia et al. (2023)	This study proposed a method that combines a one-to-one matching branch with an auxiliary one-to-many matching branch during training to improve the detection accuracy of DETR	COCO2017 mAP: 59.4, AP50: 77.8, AP75: 65.4
Huang et al. (2023)	This study proposed a training scheme to enhances DETR by leveraging predicted bounding boxes from other teacher models including RCNN-based and DETR-based detectors using knowledge distillation. The method does not introduce extra parameters or computational overhead during inference	COCO2017 mAP: 58.5, AP50: 77.4, AP75: 64.8

three branches as in Mask R-CNN. MaskLab is an instance segmentation model based on Faster R-CNN, consisting of object detection, segmentation, and instance (object) center direction prediction branches (Chen et al., 2018). The direction prediction provides useful information to distinguish instances of the same semantic label, allowing the model to further refine the instance segmentation results.

Attention mechanisms have been integrated into the segmentation models to learn the weights of multiscale features at each pixel location. A multistage context refinement network introduces a context attention refinement module that is composed of two parts, context feature extraction and context feature refinement (Liu et al., 2023). The context feature extraction captures both local and global context information, fuses both contextual information and passes it to the context feature refinement while the context feature refinement removes redundant information and generates a refined feature representation, improving the utilization of contextual information. The context attention is added to the skip connection between the encoder and the decoder. Handcrafted features are often abandoned for automatic feature extraction using convolutional networks. However, it is argued that the interpretability and domain-specific knowledge embedded in handcrafted features can

provide valuable insights. To this end, an attention module based on the covariance statistic is introduced to model the dependencies between local and global context of the input image (Liu et al., 2022). Two types of attention are introduced: spatial covariance attention focuses on the spatial distribution and channel covariance attention attends to the important channels. Furthermore, the covariance attention does not require feature shape conversion, hence significantly reducing the space and time complexity of the model.

The convolutional layers use local receptive fields to process input data which can be effective for exploiting spatial patterns and hierarchical features but may find it difficult to capture global relationships across the entire image. The ViT has been leveraged to mitigate this issue in semantic segmentation (Strudel et al., 2021). Specifically, the input image is divided into patches and treated as input to the transformer to capture the long-range relationship between the patches, significantly improving the prediction of the segmentation map. Global context ViT aims to address the lack of ViT's ability to leverage local spatial features (Hatamizadeh et al., 2023). As shown in Fig. 24, the transformer consists of local and global self-attention modules. The role of global self-attention is to capture the global contextual information

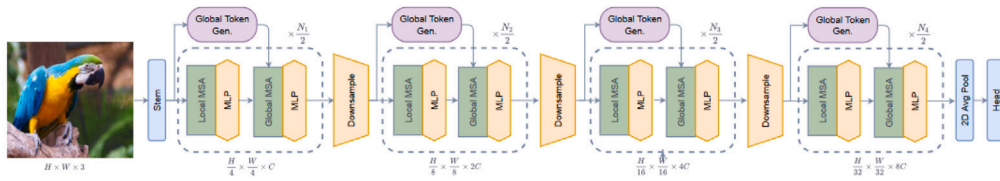


Fig. 24. The architecture of global context ViT (Hatamizadeh et al., 2023).

Table 4

Summary of state-of-the-art image segmentation.

Ref.	Description	Results (Datasets and metrics)
Liu et al. (2023)	The study proposed a multistage context refinement network with a context attention refinement module that improves feature representation by capturing local and global context, fusing information, and removing redundancies, with attention integrated into the encoder-decoder skip connections	PASCAL VOC2012 mIoU: 79.05 ADE20K MI mIoU: 41.25 Cityscapes mIoU: 79.42
Liu et al. (2022)	This study introduced a covariance attention to model the local and global dependency for the feature maps by formulating a covariance matrix, providing complementary information to improve segmentation performance	Cityscapes mIoU: 82.9 Pascal-Context mIoU: 51.7 ADE20K mIoU: 43.78
Strudel et al. (2021)	This study introduced a transformer for semantic segmentation	Cityscapes mIoU: 81.3 Pascal-Context mIoU: 59.0 ADE20K mIoU: 53.63
Hatamizadeh et al. (2023)	This study proposed an improved vision transformer architecture that employs alternating global and local self-attentions to effectively capture both local and global spatial information	ADE20K mIoU: 49.2
Shi et al. (2023)	This study introduced TSG module that uses attention maps to calculate multiscale feature weights for each image patch, enabling the model to adaptively select the optimal scale based on patch-query correlation	Cityscapes mIoU: 83.6 Pascal-Context mIoU: 64.9 ADE20K mIoU: 56.93

from different image regions while the short-range information is captured by the local self-attention. Multiscale feature representation is crucial for accurate semantic segmentation. However, the transformer often combines the features without considering their appropriate (optimal) scales, thus affecting the segmentation accuracy. Transformer scale gate is a module proposed to address the issue of selecting an appropriate scale based on the correlation between patch-query pairs (Shi et al., 2023). The transformer takes attention (correlation) maps as input and calculates the weights of the multiscale features for each image patch, allowing the model to adaptively choose the optimal scale for each patch. Table 4 lists the summary of state-of-the-art image segmentation.

#### 4.1.4. Image generation

Image generation refers to the process of creating images based on input texts. Generally, the task can be divided into three stages. The first stage is extracting features from the input text, followed by generating the image and finally controlling the image generation process to ensure the output meets specific criteria and constraints. This section focuses on the progress made in the development of deep learning models of the second stage (image generation) since it directly impacts the quality of the generated images. Variational autoencoder is one of the earliest deep learning models that is capable of generating images (Kingma and Welling, 2022). Variational autoencoder learns to generate data by capturing the underlying (Gaussian) distribution of the training data. During the generation process, the distribution parameters are sampled and passed to the decoder to generate the output image. Although the generated images are blurry and unsatisfactory, it has shown a lot of potential in image generation tasks. The

introduction of GAN significantly improved the quality of generated images. GAN consists of two connected neural networks, a generator and a discriminator that are trained simultaneously in a competitive manner (Goodfellow et al., 2014). The generator learns to generate realistic images to fool the discriminator, while the discriminator learns to distinguish between fake and real images. The generated images are less blurry and more realistic. Several enhanced models have been proposed to improve its usability and overall performance such as conditional GAN (Odena et al., 2017) which allows us to tell what image to be generated, and the deep convolutional GAN (DCGAN) (Radford et al., 2015) which provides a more stable structure for image generation. DCGAN is the basis of many subsequent improvements in GANs.

StackGAN divides the process of image generation into two stages (Zhang et al., 2017). Stage-I generates a low-resolution image by creating basic shapes and colors and the background layout using the random noise vector. Stage-II completes the details of the image and produces a high-resolution photo-realistic image. StackGAN++ is the enhanced model of StackGAN whereby it consists of multiple generators with shared parameters to generate multiscale images (Zhang et al., 2018). The generators have a progressive goal with the intermediate generators generating images of varying sizes and the deepest generator producing the photo-realistic image. HDGAN is a generative model featuring a single-stream generator with hierarchically nested discriminators at intermediate layers (Zhang et al., 2018). These layers, each connected to a discriminator, generate multiscale images. The lower resolution outputs are used to learn semantic image structures while the higher resolution outputs are used to learn fine-grained details of the image. StackGAN heavily relies on the quality of the generated

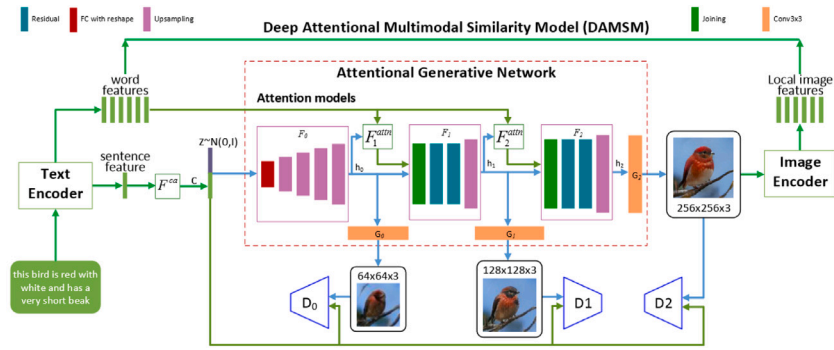


Fig. 25. The architecture of AttnGAN (Xu et al., 2018).

image in Stage-I. DM-GAN incorporates a memory network for image refinement to cope with badly generated images in Stage-I (Zhu et al., 2019). The memory network dynamically selects the words that are relevant to the generated image, and then refines the details to produce better photo-realistic images.

AttnGAN is the first to incorporate attention mechanisms into the multiple generators to focus on words that are relevant to the generated image (Xu et al., 2018). To this end, in addition to encoding the whole sentence into a global sentence vector, the text encoder encodes each word into a word vector as shown in Fig. 25. Then, the image vector is used to attend to the word vector using the attention modules at each stage of the multistage generators. Furthermore, AttnGAN introduces a loss function to compute the similarity between the generated image and the associated sentence, improving the performance of image generation. A similar work is reported whereby the model known as ResFPA-GAN, incorporates attention modules into the multiple generators (Sun et al., 2019). Specifically, a feature pyramid attention module is proposed to capture high semantic information and fuse the multiscale feature, enhancing the overall performance of the model. DualAttn-GAN improves AttnGAN by incorporating visual attention modules to focus on important features along both spatial and channel dimensions (Cai et al., 2019). This allows the model to better understand and capture both the context of the input sentence and the fine details of the image, resulting in more realistic image generation.

Although multistage generators improve image generation performance by leveraging multiscale representation, the generated images may contain fuzzy shapes with coarse features. DF-GAN replaces the multistage generators with a single-stage deep generator featuring residual connections and trained with hinge loss (Tao et al., 2022). Furthermore, DF-GAN introduces a regularization strategy on the discriminator that applies a gradient penalty on real images with matching text, allowing the model to generate more text-matching images. DMF-GAN an improved DF-GAN, incorporates three novel components designed to leverage semantic coherence between the input text and the generated image (Yang et al., 2024). The first component is the recurrent semantic fusion module, which models long-range dependencies between the fusion blocks. The second component is the multi-head attention module, which is placed towards the end of the generator to leverage the word features, forcing the generator to generate images conditioned on the relevant words. The last component is the word-level discriminator, which provides fine-grained feedback to the generator, facilitating the learning process and improving the overall quality of the generated images. Fig. 26 shows the architecture of DMF-GAN. The process of image generation involves feeding a noise vector to the generator at the very beginning of the network. However, as the generator goes deeper, the noise effect may be diminished, affecting the diversity of the image generation results. To mitigate this issue, DE-GAN incorporates a dual injection module into the single-stage generator (Jiang et al., 2024). The dual injection module consists of two text fusion layers followed by a noise broadcast operation. The text fusion layer takes

the sentence embedding and fuses it with the input feature map using the fully connected layer. Then noise is injected into the output feature map to retain the randomness in the generation process, improving diversity and generalization of the model. Table 5 lists the summary of state-of-the-art image generation.

## 4.2. Natural language processing

Natural language processing (NLP) refers to the field of AI that concerns with enabling computers to process, analyze and interpret human languages to extract useful information. Some of the common tasks in NLP are machine translation, text classification and text generation. Deep learning has been widely applied to solve real-world NLP problems. This section presents the recent advancements in deep learning models that have been designed for NLP over the past few years.

### 4.2.1. Text classification

Text classification, known as text categorization, is a task that involves assigning predefined categories or labels to a piece of text based on its content. The task is commonly used in various applications such as document classification, sentiment analysis and spam filtering. Numerous deep learning models have been proposed for text classification in the past few decades, and multilayer perceptron is one of the earliest architectures adopted to classify documents (Calvo and Ceccatto, 2000; Yu et al., 2008). The model typically has a single hidden layer with a number of units between 15 and 150. Text data is inherently sequential, as it is composed of a series of words and symbols arranged in a specific order. This property makes RNN and its variants particularly well-suited for processing and analyzing text data. In Arevian (2007), RNN with two hidden layers, each with 6 units, is used to classify news documents into eight classes. A study was conducted to investigate the variants of RNN such as LSTM and GRU for text classification (Huang and Feng, 2020). The input to the model is a sequence of words of fixed length. The input sequence is also sliced into smaller subsequences of fixed length and passed to an independent model for parallelization. A convolutional layer can extract local features, allowing the model to leverage hierarchical temporal information in textual data. A hybrid model of convolutional and LSTM architecture is proposed for text classification (Wang et al., 2019). Two parallel convolutional layers are used to extract features from word embeddings, followed by maximum pooling layers to reduce the feature dimensions. The reduced features are then concatenated and passed to LSTM for prediction.

Although CNN and RNN provide excellent results on text classification tasks, the models lack the ability to attend to specific words based on their importance and context. To address this limitation, an attention mechanism is incorporated into the model to focus on the important features, enhancing the text classification accuracy. In Liu and Guo (2019), two attention modules are introduced to capture the contextual information of the feature sequence extracted by bidirectional LSTM. The first attention module attends the sequence in

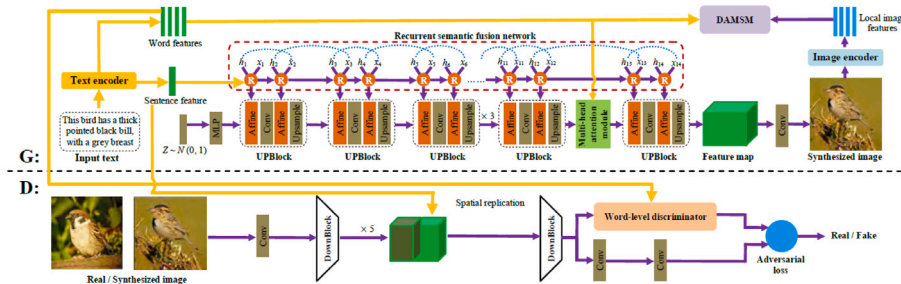


Fig. 26. The architecture of DMF-GAN (Yang et al., 2024).

**Table 5**  
Summary of state-of-the-art image generation.

Ref.	Description	Results (Datasets and metrics)
Tao et al. (2022)	This study introduced a one-stage GAN architecture, employing a Target-Aware Discriminator to enhance text-image semantic consistency and utilizes deep text-image fusion blocks for effective feature integration, resulting in improved image authenticity and relevance	CUB IS: 5.10, FID: 14.81 COCO FID: 19.32
Yang et al. (2024)	This study proposed a deep multimodal fusion generative adversarial networks that allow effective semantic interactions for fusing text information into the image synthesis process	CUB IS: 5.42, FID: 13.21 COCO IS: 36.72, FID: 15.83
Jiang et al. (2024)	This study proposed a deep learning-based method which utilizes BERT to extract sentiment information and analyze sentiment at the sentence and word level with SVM for sentiment classification	CUB IS: 4.86, FID: 18.94 COCO IS: 18.33, FID: 28.79

forward direction, while the backward sequence is attended by the second attention module. The convolutional layers are used before the bidirectional LSTM to extract features from the word embedding. The attention modules require sequential processing using RNN-based architecture such as LSTM and GRU, which may lead to information loss and distorted representations, particularly in long sequence. Furthermore, the attention modules focus on inter-sequence relationships between the input sequence and the target, ignoring the intra-sequence relationships or the dependencies between the words. In Lin et al. (2017), the deep learning model is integrated with self-attention to capture the intra-sequence relationships between the features in the sequence. A multilayer of bidirectional LSTMs is utilized to extract the feature sequence from the word embedding before the self-attention module attends the feature sequence to compute the attention weights. To further improve the overall performance, a multichannel features consisting of three input pipelines is introduced (Li et al., 2020). Each pipeline concatenates the word vector with a feature vector derived from the input sequence such as the word position, part-of-speech and word dependency parsing. The input pipeline is connected to bidirectional LSTM, followed by a self-attention module to learn the dependencies between the features in the sequence.

The transformer is a deep learning architecture that transforms sequential data using self-attention mechanisms, allowing long-range dependencies and complex patterns to be captured. The architecture is the basis of various advanced deep learning models and the Bidirectional Encoder Representations from Transformers popularly known as BERT is one of the examples that leverage transformer for pre-training on large scale textual data (Devlin et al., 2018). BERT is a bidirectional transformer encoder that is designed for various NLP tasks, capable of capturing the contextual information from both preceding and succeeding words in the input sequence. Several improvements have been made to BERT to enhance its overall performance, such as ALBERT (Lan et al., 2019), RoBERTa (Liu et al., 2019) and DeBERTa (He et al., 2020). The improvements are centered around refining the pre-training approaches such as dynamic masking of the training instances, training with a block of sentences and representing each input word using two vectors, both content and position of the word. Most of the recent works leverage BERT and its variants to capture effective feature representation of the

input sequence. In Rodrawangpai and Daungjaiboon (2022), BERT and its variants are leveraged to capture the long-range dependencies of the input tokens. The features are then passed to a layer normalization and a linear fully connected layer with dropout for classification. Similar work is reported in Murfi et al. (2024) whereby BERT is used to extract the features and the features are then passed to a hybrid of convolutional and recurrent neural networks. The traditional machine learning algorithms have been used to classify the features extracted by BERT (Hao et al., 2023). The study shows machine learning algorithms can effectively leverage the rich contextual features extracted by BERT for downstream classification tasks.

In text classification, the text labels can help in capturing the words relevant to the classification. The label-embedding attentive model is one of the earliest attempts to joint learn the label and word embeddings in the same latent space and measure the compatibility between labels and words using cosine similarity (Wang et al., 2018). The joint embedding allows the model to capture more effective text representations, increasing the overall performance of the model. LANTRN is a deep learning model that leverages label embedding extracted by BERT and entity information e.g. person name and organization name for text classification (Yan et al., 2023). The entity recognition module is based on bidirectional LSTM and conditional random field layers to calculate the probability of each word in each entity label. The model introduces a label embedding bidirectional attention to learn the attention weights of token-label and sequence-label pairs. Furthermore, a transformer consisting of RNN and a multi-head self-attention mechanism is introduced to learn local short-term dependencies of multiple short text sequences and long-term dependencies of the input sequence. Aspect refers to a specific attribute of an entity within the text, and incorporating this information enhances the model's understanding of the nuances of the text. BERT-MSL is a multi-semantic deep learning model with aspect-aware enhancement and four input pipelines: left sequence, right sequence, global sequence and aspect target (Zhu et al., 2023). The aspect-aware enhancement module takes the features extracted by BERT, and performs average pooling followed by a linear transform. Then the output is concatenated with the outputs produced by the local and global semantic learning modules. The concatenated features are then jointly attended by a multi-head attention for text classification.

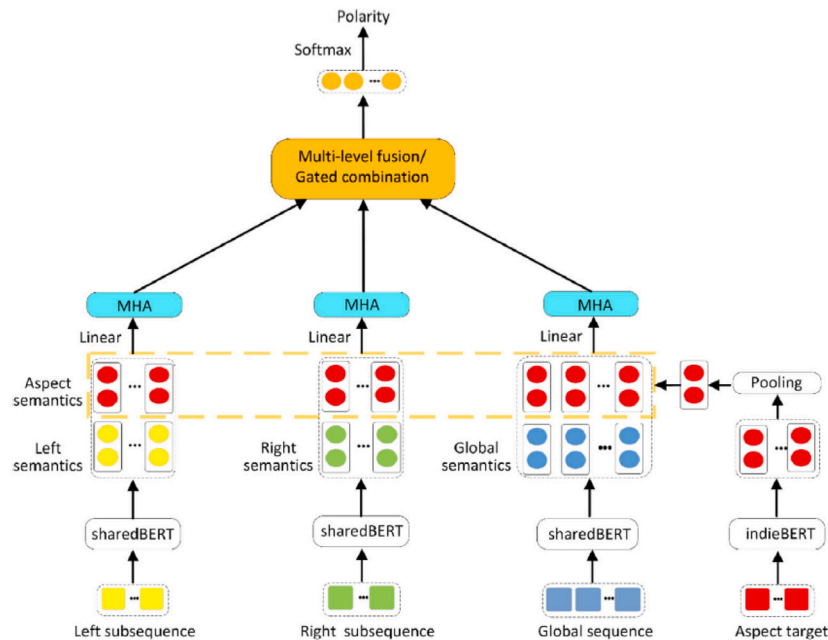


Fig. 27. The architecture of BERT-MSL (Zhu et al., 2023).

Table 6  
Summary of state-of-the-art text classification.

Ref.	Description	Results (Datasets and metrics)
Rodrawangpai and Daungjaiboon (2022)	The study proposed an enhanced RoBERTa by incorporating layer normalization and dropout to enable smoother gradients, faster training and reduced overfitting, thus, improving model generalization	NEISS Accuracy: 0.98
Murfi et al. (2024)	The study employed IndoBERT to extract features from text data which were then passed to a hybrid deep learning model for sentiment analysis	Shopee Accuracy: 0.8612 (CNN-GRU) Tokopedia Accuracy: 0.8768 (LSTM-CNN) Lazada Accuracy: 0.8710 (LSTM-CNN)
Hao et al. (2023)	This study proposed a deep learning-based method which utilizes BERT to extract sentiment information and analyze sentiment at the sentence and word level with SVM for sentiment classification	NLPCC-SCDL Accuracy: 95.12%
Yan et al. (2023)	This study proposed LANRTN model for text classification, integrating an R-Transformer with label embedding, attention mechanisms, and an entity recognition model to capture both global dependencies and label-aware contextual information	Reuters Corpus Volume I Micro F1-score: 0.893, Arxiv Academic Paper Dataset Micro F1-score: 0.718
Zhu et al. (2023)	The paper proposed a BERT-MSL model for aspect-based sentiment analysis, integrating multiple semantic learning modules and merging them using multi-head self-attention and linear transformation layers for enhanced classification	14Lap Accuracy: 82.24%, Macro-average F1-score: 78.98% 14Rest Accuracy: 89.11%, F1-score: 84.07% 15Rest Accuracy: 88.53%, Macro-average F1-score: 71.41% 16Rest Accuracy: 93.09%, Macro-average F1-score: 81.24% Twitter Accuracy: 74.89%, Macro-average F1-score: 73.43%

Fig. 27 shows the architecture of BERT-MSL. Table 6 lists the summary of state-of-the-art text classification.

4.2.2. Neural machine translation

Neural machine translation (NMT) refers to the automated process of translating text from one language to another language. Numerous deep learning models have been proposed for NMT, which can be categorized into RNN-based and CNN-based models. One of the first successful RNN-based models is the encoder-decoder (Cho et al., 2014; Sutskever et al., 2014). The model consists of two connected sub-networks (the encoder and the decoder) for modeling the translation process, as shown in Fig. 28. The encoder reads the source sentence word by word and produces a fixed-length context vector (final hidden

state). This process is known as source sentence encoding, as shown in the figure. Given the context vector, the decoder generates the target sentence (translation) word by word. This modeling of the translation can be seen as a mapping between the source sentence to the target sentence via the intermediate context vector in the semantic space. The context vector represents the summary of the input sequence's semantic meaning, providing a compressed representation that captures the essence of the source sentence. However, the compression process can sometimes result in the loss of information, especially those early in the sequence. Bidirectional RNN may mitigate the loss of information by modeling the sequence in reverse order. However, the problem can persist, particularly in cases where the input is a long sequence.

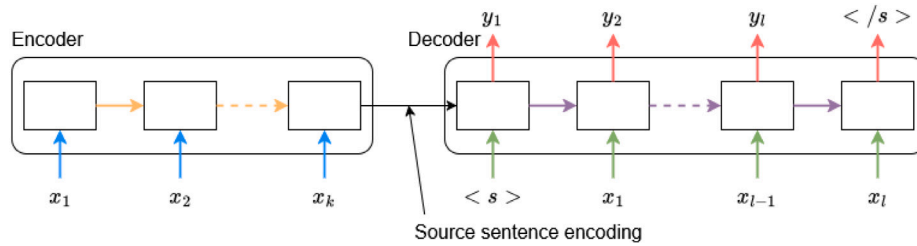


Fig. 28. The architecture of an encoder-decoder.  
Source: Adapted from [Stahlberg \(2020\)](#).

Attention mechanism was introduced to solve the problem of learning long input sequences ([Bahdanau et al., 2014](#)). Attention alleviates this issue by attending on different words of the input sequences when predicting the target sequences at each time step. Unlike the standard encoder-decoder model, attention derives the context vector from the hidden states of both the encoder and decoder, and the alignment between the source and target. This mechanism allows the model to focus on the important words, increasing the overall accuracy of the translation. Several alignment score functions have been proposed for calculating the attention weights. Some popular functions are additive ([Bahdanau et al., 2014](#)), dot-product, location-based ([Luong et al., 2015](#)), and scaled dot-product ([Vaswani et al., 2023](#)). The attention weights are calculated by attending to the entire hidden states of the encoder. This attention, also known as global attention, is computationally expensive. Instead of attending to all hidden states, local attention attends to a subset of hidden states, thus reducing the computational cost ([Luong et al., 2015](#)). Google Neural Machine Translation is a popular encoder-decoder model with an attention mechanism that significantly improves the accuracy of machine translation ([Wu et al., 2016](#)). As shown in [Fig. 29](#), the model consists of a multilayer of LSTMs with eight encoder and decoder layers and an attention connection between the bottom layer of the decoder to the top layer of the encoder. Furthermore, to deal with the challenging words to predict, a word is tokenized into subwords e.g. feud is broken down into “fe” and “ud”, allowing the model to generalize well to new and uncommon words. A year later, the self-attention mechanism was proposed, significantly improving the overall accuracy of machine translation ([Vaswani et al., 2023](#)). Self-attention, also known as intra-attention, allows the deep learning model to capture the dependencies between the input words. The self-attention mechanism is the fundamental building block of the transformer model, which has since become a cornerstone in natural language processing and other domains.

Despite the success of transformer, the model falls short in capturing nuances of human language and struggles with tasks requiring deeper understanding of context. This can be especially challenging when the tasks involve formality, colloquialism, and subtle cultural references that may not directly equivalent in the target language, resulting in inaccurate translation or losing the original meaning. One of the approaches to include context into the input sequence is concatenating the current source sentence with the previous (context) sentences and feeding the whole input to the transformer ([Lupo et al., 2023](#)). The model is trained to predict the translated sentence including the context translation. At inference time, only the translation is considered while the context translation is discarded. Furthermore, the approach encodes the sentence position and segment-shifted position to improve the distinction between current sentences and context sentences. In [Rippeth et al. \(2023\)](#), the source sentence is prefixed with the summary of the document to contextualize the input sentence. The summary is the set of salient words that represents the essence of the document, resolving ambiguity associated with the translation. A study was conducted to determine the optimal technique of aggregating contextual features ([Wu et al., 2022](#)). Three techniques were studied namely concatenation mode, flat mode and hierarchical mode, and the experimental results

indicate that concatenation mode achieved the best results. In [Kim et al. \(2023\)](#), a training method is introduced to train the deep learning machine translation model to generate translation involving honorific words. The training method indicates the honorific context in the target sentence using an honorific classifier to guide the model to attend to the related tokens. Unlike other studies where the context features are included by concatenation, the training method assigns weights to the context tokens indicated by the honorific classifier. This allows the model to generate a more accurate translation with honorifics.

The performance of transformers relies on large-scale training data. However, for the vast majority of languages, only limited amounts of training data exist. To mitigate this problem, recent studies introduce shallow transformer architectures ([Gezmu and Nürnberger, 2022](#)), explore the effect of hyperparameter finetuning ([Araabi and Monz, 2020](#)) and leveraging visual input as contextual information for the translation task ([Meetei et al., 2023](#)). The semi-supervised neural machine translation is used for translating low resource Egyptian Arabic dialects to modern standard Arabic ([Faheem et al., 2024](#)). The study utilizes three datasets: bilingual Egyptian-standard Arabic and two monolingual Egyptian Arabic and standard Arabic. First, a transformer-based model is trained in a supervised manner using the bilingual dataset. Then, the model is trained in an unsupervised manner using both monolingual datasets. The unsupervised approach employs an Encoder-Decoder model with Byte Pair Encoding for tokenization and handling unknown words. The monolingual corpora are merged and used to improve the model by iteratively generating synthetic sentence pairs between Egyptian Arabic and standard Arabic, allowing the model to learn the correspondence between them. In [Li et al. \(2024\)](#), the authors exploit monolingual corpus to enhance the bilingual dataset for model training. Furthermore, a new loss function is proposed as a replacement for traditional cross-entropy loss, allowing the model to learn with uncertainty in the presence of noise. Additionally, contrastive re-ranking is employed to refine translation results by selecting the most confident output from multiple candidates. [Table 7](#) lists the summary of state-of-the-art neural machine translation.

#### 4.2.3. Text generation

Text generation refers to the process of creating texts based on a given input, whereby the input can be in the form of texts, images, graphs, tables or even tabular data. Due to the various forms of inputs, text generation has a wide range of applications, including creative writing, image captioning and music generation. This section focuses on the progress made in text-to-text generation tasks such as question answering, dialog generation and text summarization. The recurrent neural network and its variants play an important role in text generation tasks for their strong ability to model sequential data. One of the earliest works on question answering is based on the RNN-based encoder-decoder model, whereby the encoder takes the question embedding and processes it using bidirectional LSTM, and the decoder generates the corresponding answer ([Nie et al., 2017](#)). Additionally, to prevent semantic loss and enable the model to focus on the important words in the input sequence, a convolution operation is applied to the word embedding, and an attention mechanism is then used to attend

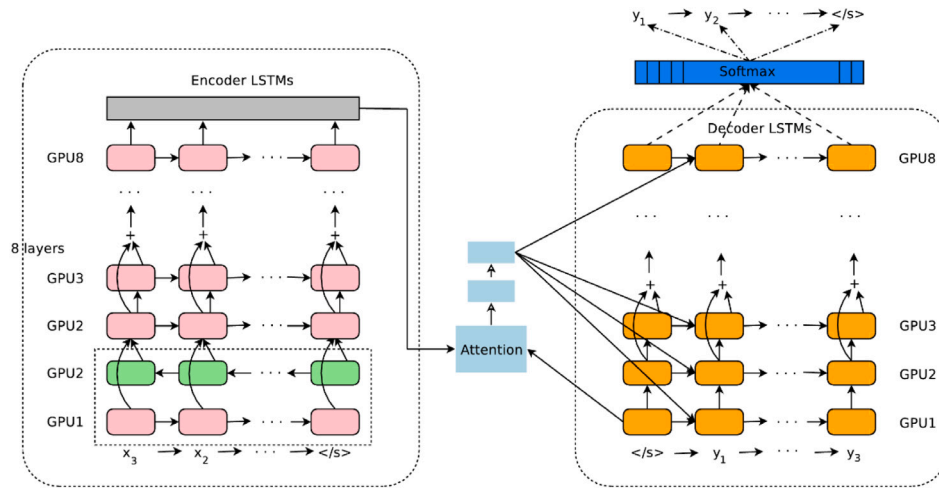


Fig. 29. The architecture of Google Neural Machine Translation (Wu et al., 2016).

Table 7

Summary of state-of-the-art neural machine translation.

Ref.	Description	Results (Datasets and metrics)
Lupo et al. (2023)	This study incorporated context into the input sequence by concatenating the current source sentence with previous context sentences, and the model is trained to predict the translation with context, while encoding sentence positions to distinguish between current and context sentences	En→De ContraPro: 82.54 En→Ru Voita: 75.94
Rippeth et al. (2023)	The paper proposed a method to enhance translation accuracy of ambiguous words by prefixing source sentences with salient words extracted from related pseudo-documents, thereby providing additional context without altering standard model architectures	En→De BLEU: 22.0, COMET: 0.785
Wu et al. (2022)	The paper investigated three most common methods to aggregate the contextual features: concatenation mode, flat mode and hierarchical mode	En→De BLEU: 30.89 De→En BLEU: 36.84 En→Zh BLEU: 20.06 Zh→En BLEU: 20.50
Kim et al. (2023)	The paper introduced formality classifier to incorporate formality-related contexts into the model training	AI-HUB (En→Ko) BLEU: 27.90 OpenSubtitles201 (En→Ko) BLEU: 20.70 BSD, AMI and OpenSubtitles2018 (En→Ja) BLEU: 13.67
Faheem et al. (2024)	This study employed a semi-supervised neural machine translation approach to translate low-resource Egyptian Arabic dialects to Modern Standard Arabic, utilizing bilingual and monolingual datasets, with a transformer-based model trained in both supervised and unsupervised manners to generate synthetic sentence pairs and improve translation accuracy	Egyptian-standard Arabic BLEU: 24 (300 word embedding) BLEU: 29.5 (512 word embedding)
Li et al. (2024)	The study proposed a new loss function to handle uncertainty in noisy data, and used contrastive re-ranking to select the most confident translation from multiple candidates	Zh→Ma BLEU: 28.12 Ma→Zh BLEU: 23.53 Zh→In BLEU: 28.91 In→Zh BLEU: 22.76

to the output of the convolution operation. Similar work is reported in Yin et al. (2016) in which a knowledge-based module is introduced to calculate the relevance score between the question and the relevant facts in the knowledge base. This improves the text (answer) generation by the decoder. Another work is described in Li et al. (2016) where an encoder–decoder with attention for dialog generation is optimized using reinforcement learning. The model is first trained in a supervised learning manner and then improved using the policy gradient method to diversify the responses. Ambiguous content in question answering sentences is a challenge in text generation and can lead to incorrect and uncertain responses. Cross-sentence context aware bidirectional model

introduces a parallel attention module to compute the co-attention weights at the sentence level, accounting for the relationships and similarities in the question and the answer (Wu et al., 2020).

The transformer has been leveraged for text generation tasks. An incremental transformer-based encoder is proposed to incrementally encode the historical sequence of conversations (Li et al., 2019). The decoder is a two-pass decoder that is based on the deliberation network, generates the next sentence. The first pass focuses on contextual coherence of the conversations, while the second pass refines the output of the first pass. BERT and ALBERT have been used as pre-trained models for question answering task (Alrowili and Vijay-Shanker, 2021). The

study found that the performance of the models is sensitive to random assignment of the initial weights especially on small datasets (Alrowili and Vijay-Shanker, 2022). T-BERTSum is a model based on BERT, designed to address the challenge of long text dependence and leveraging latent topic mapping in text summarization (Ma et al., 2021). The model integrates a neural topic module to infer topics and guide summarization, uses a transformer network to capture long-range dependencies and incorporates a multilayer of LSTM for information filtering.

The generated texts often lack diversity and may exhibit repetitive patterns. To address this issue, feature-aware conditional GAN (FAGAN) is proposed for controllable category text generation (Li et al., 2023). The generator consists of BERT, a category encoder, a relational memory core (RMC) decoder. BERT acts as a feature encoder, improves contextual representation and mitigates mode collapse while the category encoder embeds categorical information for text generation. The RMC decoder utilizes a self-attention mechanism to capture interactions between features, generating more expressive and diverse texts. The discriminator includes an additional classification head to ensure the generated texts match specified categories. Traditional text generation has largely focused on binary style transfers, but real-world applications require capturing diverse styles, which existing methods struggle to achieve. In Kwon et al. (2024), a multi-class conditioned text generation model is proposed using a transformer-based decoder with an adversarial module, a style attention module and a generation module. The adversarial module extracts style-excluded representation from the input text. The style attention mechanism then introduces the desired style into the text representation through concat attention mechanism, producing a conditioned representation. Finally, the generation module utilizes the conditioned representation to generate stylistically diverse text.

Exploiting domain knowledge is essential in reducing the semantic gap between the deep learning models and the text corpus. KeBioSum is a knowledge infusion framework to inject domain knowledge into the pre-trained BERTs for text summarization (Xie et al., 2022). In the framework, the relevant information is detected and extracted from the domain knowledge, generating label sequences of the sentences. The label data is then used to train the text summarization model using discriminative and generative training approaches, infusing the knowledge into the model. Large language models (LLMs) have been used for text generation with exceptional quality and diversity. LLMs, trained on extensive corpus data, have a deep understanding of human language, allowing them to interpret and generate texts. In Hajipoor et al. (2025), GPTGAN is introduced that leverages an LLM as a guiding mentor to a GAN-Autoencoder model for text generation. The approach involves the GPT model to generate a sequence of words given a subset of the input sequence, and the generated text is mapped into latent space by the transformer-based encoder. The latent representation is then used by the transformer-based decoder to generate the text. Furthermore, local discriminators are introduced to refine the text generation because the generated words by the GPT model may be inaccurate. Table 8 lists the summary of state-of-the-art text generation.

#### 4.3. Time series and pervasive computing

Pervasive computing, often referred to as ubiquitous computing, is the process of integrating computer technology into everyday objects and surroundings so that they become intelligent, networked, and able to communicate with one another to offer improved services and functionalities (Weiser, 1999). According to He et al. (2020), the role of pervasive computing is foremost in the field where it provides the ability to distribute computational services to the surroundings where people work, leading to trust, privacy, and identity. Examples of pervasive computing applications include smart homes with connected appliances, wearable devices that monitor health and fitness, smart cities with sensor networks for traffic management, and industrial

applications that utilize the Internet of Things (IoT) for monitoring and control. Generally, the continuous interaction of interconnected devices in pervasive computing often results in time series data, which captures the evolution of various parameters over time.

For instance, medical sensors, such as electrocardiograms (ECG) and electroencephalograms (EEG), generate time series data that contain critical diagnostic information, which deep learning can use to detect anomalies, predict diseases, and classify medical conditions with improved accuracy. Furthermore, devices such as accelerometers, magnetometers and gyroscopes, among others, can be used to capture human activity signals, which are often represented as time series of state changes (Ige and Mohd Noor, 2022). In traditional machine learning, features such as mean, variance, and others are manually extracted from times series of state changes before human activity classification. However, deep learning models automatically extract features (Mohd Noor, 2021). Furthermore, in other fields such as finance, which entail time series data, deep learning has been instrumental in stock price prediction (Singh and Srivastava, 2017), fraud detection (Zhang et al., 2021), and algorithmic trading (Lei et al., 2020), among others. Generally, deep learning networks excel at capturing intricate temporal relationships within time-series data, enabling more precise predictions and improved decision-making. Based on this, several deep learning models have been employed for feature learning across various time series and pervasive computing domains.

##### 4.3.1. Human activity recognition

Human activity recognition (HAR) finds application across various domains including intelligent video surveillance, environmental home monitoring, video storage and retrieval, intelligent human-machine interfaces, and identity recognition, among many others. It includes various research fields, including the detection of humans in video, estimating human poses, tracking humans, and analyzing and understanding time series data (Zhang et al., 2019). Despite the advancements in vision-based HAR, there exist inherent limitations. Generally, vision-based approaches heavily rely on camera systems, which may have restricted views or be affected by lighting conditions, occlusions, and complex backgrounds (Ige and Mohd Noor, 2022). Additionally, vision-based HAR struggles with identifying actions that occur beyond the range of the camera or actions that are visually similar.

Wearable sensors offer a promising alternative to overcome these limitations. By directly capturing data from the individual, wearable sensors provide more comprehensive and accurate information about human activities. The signals obtained from wearable sensors typically represent time series data reflecting state changes in activities. Deep learning models can effectively learn from these signals, allowing for robust and accurate recognition of human activities. Moreover, wearable sensors offer the advantage of mobility, enabling activity recognition in various environments and situations where vision-based systems may be impractical or ineffective (Dang et al., 2020). Generally, the time series nature of signals from wearable sensors presents an excellent opportunity for deep learning models to excel in recognizing human activities with high accuracy and reliability.

Several researchers have proposed the use of CNN, RNN, and Hybrid models for deep learning-based feature learning in wearable sensor HAR. For instance, using two-dimensional CNN (Conv2D), several researchers, as seen in Gao et al. (2021), Gupta (2021) and Erdas and Guney (2021), among others, have developed deep learning models for wearable sensor HAR, despite the time series nature of the data. This is often done by treating the time series signals from wearable sensors as 2D images by reshaping them appropriately. To achieve this, researchers often organize each time series signal into a matrix format, with time along one axis and sensor dimensions along the other, before creating a pseudo-image representation, which allows the matrix to be fed into Conv2D layers for feature extraction. Conv2D layers excel at capturing spatial patterns and relationships within images, and by treating the time series data as images, these layers can learn relevant

**Table 8**  
Summary of state-of-the-art text generation.

Ref.	Description	Results (Datasets and Metrics)
Ma et al. (2021)	The paper proposed a text summarization model that integrates neural topic model with BERT to guide the text summarization based on inferred topics. It leverages transformers for long-term dependencies and combines LSTM layers with a gated network to enhance extractive and abstractive summarization of social texts	CNN/Daily Mail Rouge-1: 43.58, Rouge-2: 20.45, Rouge-L: 34.60 Xsum Rouge-1: 39.90, Rouge-2: 17.48, Rouge-L: 29.85
Li et al. (2023)	This study proposed FA-GAN featuring BERT for contextual representation, a category encoder for embedding categorical information, and an RMC decoder for generating diverse texts, while the discriminator ensures category consistency	MR-10 BLEU-2: 0.560 MR-20 BLEU-2: 0.674 AM-30 BLEU-2: 0.748
Kwon et al. (2024)	This study proposed a multi-class conditioned text generation model, combining a transformer-based decoder with an adversarial module, style attention, and a generation module to produce stylistically diverse text	Amazon (neg→pos) self-BLEU: 0.257, S-ACC: 0.375 Amazon (pos→neg) self-BLEU: 0.096, S-ACC: 0.191 YELP (neg→pos) self-BLEU: 0.257, S-ACC: 0.198 YELP (pos→neg) self-BLEU: 0.081, S-ACC: 0.207 *S-ACC (style accuracy)
Xie et al. (2022)	The paper proposes a knowledge-based summarization model that integrates medical knowledge into the pretrained models using a lightweight knowledge adapter. It employs generative and discriminative training to predict and reconstruct PICO elements, enhancing domain-specific text summarization	CORD-19 Rouge-1: 32.04, Rouge-2: 12.61, Rouge-L: 29.10 PubMed-Long Rouge-1: 36.39, Rouge-2: 16.27, Rouge-L: 33.28
Hajipoor et al. (2025)	This study proposed GPTGAN, an approach that enhances adversarial text generation using a GPT model as a mentor to the GAN model and employing a multiscale discriminator framework to balance text quality and diversity	MSCOCO BLEU: 98.7 WMTNews BLEU: 93.4 Persian COCO BLEU: 97.9

spatial features that contribute to activity recognition. The convolution operation performed by Conv2D filters across both the time and sensor dimensions, allowing the network to identify patterns and features that may be indicative of specific activities.

Even though Conv2D can effectively capture spatial dependencies within the data, it often struggles to capture temporal dependencies inherent in time series data. Since Conv2D processes data in a grid-like fashion, it does not fully leverage the sequential nature of the time series, potentially leading to less effective feature extraction for wearable sensor HAR tasks. For this reason, recent HAR architectures have leveraged one-dimensional CNN (Conv1D) and other RNNs for automatic feature extraction. Conv1D layers are specifically designed to capture temporal dependencies within sequential data. They operate directly on the time series data without reshaping it into a 2D format, allowing them to capture temporal patterns more effectively. Conv1D layers are better suited for extracting features from time series data, making them a more natural choice for wearable sensor HAR (Mohd Noor, 2021).

For instance, Ragab et al. (2020) proposed a random search Conv1D model, and evaluated the performance of the model on UCI-HAR dataset. The result showed that the model achieved a recognition accuracy of 95.40% when classifying the six activities in the dataset. However, the model exhibited extended training times due to the dynamic nature of some activities within the dataset. To address this, Banjarey et al. (2022) proposed the use of varying kernel sizes in Conv1D layers to recognize various activities, including sitting, standing, walking, sleeping, reading, and tilting. Furthermore, a few Conv1D layers were stacked to streamline the time optimization process for training the neural network. In Baraka and Mohd Noor (2023), a signal segmentation method is proposed based on similarity between two subsequent windows. The method can not only achieve more accurate segmentation, but also distinguish between transitional and non-transitional windows. Based on the distinction, two deep learning models with convolutional and fully connected layers are trained: one for transitional activities and the other for non-transitional activities. In Baraka and Mohd Noor (2024), signal segmentation is treated as a binary

classification task, distinguishing windows as either transitional or non-transitional. To this end, a deep learning model with parallel Conv1D pipelines is proposed to capture temporal dependencies within the window sequence. Furthermore, some researchers have proposed models that combine machine learning algorithms with Conv1D in HAR, as seen in Shuvo et al. (2020). Their work presented a two-stage learning process to improve HAR by classifying activities into static and dynamic using Random Forest, before using Support Vector Machine to identify each static activity, and Conv1D to recognize dynamic activities. The result shows that the method achieved an accuracy of 97.71% on the UCI-HAR dataset.

Following these advancements, several researchers have further explored Conv1D architectures with various modifications, to enhance feature learning in activity recognition systems. For example, Han et al. (2022) developed a two-stream CNN architecture as a plug-and-play module to encode contextual information of sensor time series from different receptive field sizes. The module was integrated into existing deep models for HAR at no extra computation cost. Experiments on OPPORTUNITY, PAMAP2, UCI-HAR and USC-HAD datasets show that the module improved feature learning capabilities. A similar research is reported in Ige and Noor (2023) proposed the WSense module to address the issue of differences in the quality of features learnt, regardless of the size of the sliding window segmentation, and experimented on PAMAP2 and WISDM datasets. The results show that by plugging the WSense module into Conv1D architectures, improved activity features can be learned from wearable sensor data for human activity recognition.

**Hybrid Models:** Researchers have also proposed the use of standalone RNNs in HAR, and a hybrid of Conv1D architectures with RNNs such as LSTMs (Deep and Zheng, 2019), BiLSTMs (Luwe et al., 2022; Shi et al., 2023), GRUs (Dua et al., 2023) and BiGRUs (Imran et al., 2024) to fully harness the feature learning capabilities of both CNN and RNNs. For instance, Nafea et al. (2021), leveraged Bi-LSTM and Conv1D with increasing kernel sizes to learn features at various resolutions. Human activity features were extracted using the stacked convolutional

layers with a Bi-LSTM layer, before including a flattening layer and a fully connected layer for subsequent classification. However, the model had issues extracting quality features of dynamic activities compared to static activities. To address such issues, some research works have incorporated attention mechanisms in Conv1D-based architectures to improve feature learning of dynamic and complex activities from time series signals obtained from wearable sensors. For example, [Khan and Ahmad \(2021\)](#) designed three lightweight convolutional heads, with each specialized in feature extraction from wearable sensor data. Each head comprised stacked layers of Conv1Ds, along with embedded attention mechanisms to augment feature learning as shown in [Fig. 30](#). The results demonstrated that integrating multiple 1D-CNN heads with SE attention can enhance feature learning for Human Activity Recognition.

[Ige and Mohd Noor \(2023\)](#) designed three feature learning pipelines, each pipeline consisting of two concurrent layers of Conv1D and LSTM with maximum pooling, which are then concatenated and processed using a channel-wise attention mechanism to enhance feature learning. In [Gao et al. \(2021\)](#), a sequential channel-temporal attention is proposed for multi-modal activity recognition. The channel attention is similar to SE attention, but with average pooling and maximum pooling pipelines to squeeze to temporal dimension. The pooled features are then combined through addition operation. As for the temporal attention, average pooling and maximum pooling are applied along the channel dimension, producing two pooled features which are then combined through concatenation. The attention module is incorporated into the residual blocks of ResNet-like model to improve feature extraction. Similar work is reported in [Agac and Incel \(2024\)](#) whereby a hybrid sequential channel-spatial attention is proposed for a lightweight activity recognition model. The lightweight model consists of four convolutional blocks integrated with the attention module and two LSTM layers. Although the proposed attentions have shown a strong potential in improving model performance, their sequential processing inherently prioritize one of the feature map dimensions, impeding the model's ability to capture holistic feature representations. In [Tang et al. \(2022\)](#), a triple cross-domain attention is proposed to blend three attention branches to improve feature extraction in HAR. For each attention branch, z-pooling ([Misra et al., 2020](#)) and convolution operations are applied along the corresponding dimension to generate the attention weights. Then, the feature maps are combined via average operation. The model consists of three residual blocks where each block is integrated with the attention module. These diverse modifications and adaptations showcase the versatility and potential of deep learning models in achieving state-of-the-art in HAR systems.

**Transformers:** Transformers have been employed in human activity recognition to capture long-range dependencies in sensor data, enabling more accurate activity classification. In [Chen et al. \(2022\)](#), a transformer-based model is enhanced by integrating a bidirectional GRU and a linear fully connected layer into each encoder block. The decoder is reduced to a fully connected layer with softmax function. This hybrid model leverages the transformer's ability to capture long-range dependencies in sensor data and the bidirectional GRU's strength in modeling temporal sequences, enhancing the model performance in classifying complex and collaborative activities. Similar work is reported in [Sun et al. \(2024\)](#) whereby the encoder blocks of the transformer are integrated with 1D convolutional layers instead of fully connected layers, enabling the model to extract local features. Furthermore, self-supervised contrastive learning is employed to learn from unlabeled data before the model is fine-tuned using a smaller dataset. Although transformer-based models have shown remarkable performance, they suffer from high computational cost and memory requirements, making them less suitable for real-time or resource-constrained activity recognition. A study was conducted to determine the applicability of transformer models on a low-power ESP32 microcontroller ([Lattanzi et al., 2025](#)). The study concludes that transformer models are not suitable for tiny devices. Furthermore, the results show that a tiny transformer model with two encoder blocks achieved lower

accuracy compared to the standard LSTM model, with a difference up to 14%. However, transformer models run 3x faster than LSTM, making them a viable option provided the architecture of transformers is optimized for low-power platforms.

**Generative Models:** A reliable deep learning model requires a large amount of training data to learn the underlying patterns of the data. However, in HAR, data collection is expensive, and the available datasets are often limited in the number of samples. To overcome this limitation, researchers employ GAN models to generate synthetic data, thus augmenting the training set. In [Chan and Noor \(2021\)](#), a conditional GAN is proposed to generate realistic sensor data of different activities for human activity classification. The generator consists of four 1D convolutional and maximum pooling blocks followed by two layers of LSTM and fully connected layers, while the discriminator is a standard CNN model. An enhanced conditional GAN is proposed to improve synthetic data generation for HAR ([Jimale and Mohd Noor, 2022](#)). This architecture integrates 1D convolutional layers with multiple fully connected networks at the generator's input and discriminator's output, aiming to produce higher-quality synthetic samples compared to existing CGAN models. Similar work is reported in [Lupi3n et al. \(2024\)](#), where the conditional Wasserstein GAN is employed to generate synthetic data for HAR.

In [Kia et al. \(2024\)](#), sensor data is transformed into its frequency spectrum, forming an RGB-based feature space for human activity classification. Then, BiGAN is employed to generate synthetic frequency spectrum images to increase minority sample classes, thus diversifying the training set. The generator and discriminator of the proposed BiGAN are built using 2D convolutional and fully connected layers. In [Mohammadzadeh et al. \(2025\)](#), a conditional GAN is employed to generate synthetic sensor data for therapeutic activity recognition. The generator utilizes Inception-like modules and transposed convolutional layers, while the discriminator consists of two classification pipelines: one classifies raw signals and the other pipeline classifies the Fourier-transformed signal (frequency spectrum images). Both predictions are averaged to obtain the final prediction. [Table 9](#) lists the summary of state-of-the-art human activity recognition.

#### 4.3.2. Speech recognition

Speech, as the primary mode of human communication, has captivated researchers for over five decades, especially since the inception of AI ([Nassif et al., 2019](#)). From the earliest endeavors to understand and replicate the complexities of human speech, to contemporary advancements leveraging cutting-edge technologies, the quest for accurate and efficient speech recognition systems has been relentless. In recent years, the emergence of deep learning techniques has revolutionized the speech recognition field. Deep learning has demonstrated unparalleled success in processing and extracting intricate patterns from vast amount of data. When applied to the realm of speech recognition, deep learning have surpassed traditional approaches by learning intricate features directly from raw audio signals, circumventing the need for handcrafted features and complex preprocessing pipelines. This paradigm shift has significantly advanced the state-of-the-art in speech recognition, enabling systems to achieve unprecedented levels of accuracy and robustness across various languages, accents, and environmental conditions. Generally, deep learning has been extended to other essential applications of speech recognition, such as speaker identification ([Tirumala and Shahamiri, 2016](#); [Ye and Yang, 2021](#)), emotion recognition ([Khalil et al., 2019](#)), language identification ([Singh et al., 2021](#)), accent recognition ([Jiao et al., 2016](#)), age recognition ([S3nchez-Hevia et al., 2022](#)) and gender recognition ([Alnuaim et al., 2022](#)), among many others.

Prior to the adoption of deep learning in speech recognition, the foundation of traditional speech recognition systems was the use of Gaussian Mixture Models (GMMs), which are often combined with Hidden Markov Models (HMMs) to represent speech signals ([Srivastava and Pandey, 2022](#)). This is because a speech signal can be thought of as

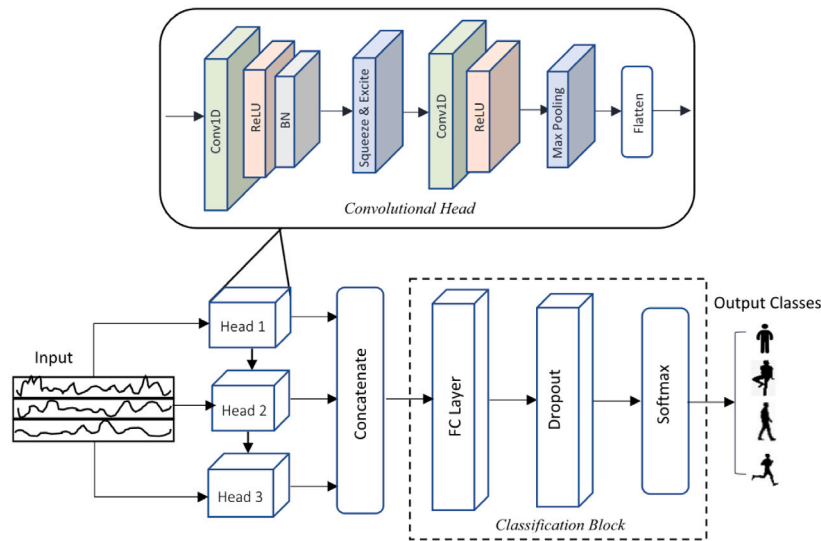


Fig. 30. The architecture of multi-head CNN model (Khan and Ahmad, 2021).

Table 9  
Summary of state-of-the-art human activity recognition.

Ref.	Description	Results (Datasets and metrics)
Dua et al. (2023)	This study proposed an Inception inspired model, combining convolutional layers with GRU. The convolutional blocks use multiple-sized filters to extract multiscale feature representations	MHEALTH Accuracy: 99.25% PAMAP2 Accuracy: 97.64%
Ige and Mohd Noor (2023)	This study proposed a model with three parallel feature learning pipelines, each pipeline has two sub-feature learning pipelines consisting of convolutional layers and bidirectional LSTM, and channel-wise attention before the classifier	PAMAP2 Accuracy: 98.52% WISDM Accuracy: 97.90%
Agac and Incel (2024)	This study proposed A lightweight model consists four convolutional blocks integrated with a sequential channel-spatial module and two LSTM layers	PAMAP2 Accuracy: 98.52% WISDM Accuracy: 97.90%
Sun et al. (2024)	A modified transformer integrated with convolutional layer instead of fully connected layer	UCI-HAR Accuracy: 95.49% Skoda Accuracy: 87.88% Mhealth Accuracy: 98.43%
Lupión et al. (2024)	This study employed a conditional Wasserstein GAN to generate synthetic human activity recognition accelerometry signals	Local Ulster University dataset Accuracy: 0.7453
Kia et al. (2024)	This study converted sensor data into its frequency spectrum to create an RGB-based feature space, employed BiGAN to generate synthetic frequency spectrum images to balance minority classes, and used 2D convolutional and fully connected layers in the generator and discriminator	Up-Fall Accuracy: 99.1% Opportunity Accuracy: 86.8% WISDM Accuracy: 99.12%
Mohammadzadeh et al. (2025)	This study employed a conditional GAN to generate synthetic sensor data for therapeutic activity recognition which has Inception-like generator and a discriminator with two classification pipelines, one for raw signals and the other for Fourier-transformed signals	Local Sharif University of Technology dataset Avg. F1-score: 0.897

a short-term stationary signal. The spectral representation of the sound wave is modeled by each HMM using a mixture of Gaussian. However, they are considered statistically inefficient for modeling non-linear or near non-linear functions (Padmanabhan and Premkumar, 2015; Nassif et al., 2019). This is because HMMs rely on a set of predefined states and transition probabilities, making assumptions about the linearity and stationarity of the underlying data. While suitable for modeling certain aspects of speech, HMMs often fall short when tasked with representing the intricate nonlinearities and variability present in speech signals. Speech, by nature, exhibits nonlinear and dynamic characteristics, with features such as intonation, rhythm, and phonetic variations challenging the simplistic assumptions of traditional statistical models

like HMMs. In other words, GMM-HMM approach had limitations in capturing complex acoustic patterns and long-term dependencies in speech (Mukhamadiyev et al., 2022).

**Sequence-to-Sequence Models:** In recent times, CNN and RNNs have been leveraged for automatic speech recognition in order to consider a longer or variable temporal window for context information extraction (Lu et al., 2020). Generally, CNNs are well-suited for capturing local patterns and hierarchical features in data, making them effective for modeling acoustic features in speech. By directly learning features from raw speech signals, CNNs bypassed the need for hand-crafted features used in traditional GMM-HMM systems. Additionally,

CNNs can capture long-range dependencies in the data, which is crucial for understanding the context of speech. Likewise, the RNNs are suitable choice for exploring extended temporal context information in one processing level for feature extraction and modeling.

Based on this, several researchers have proposed the use of both CNN and variants of RNNs for automatic speech recognition and for other speech related tasks. For instance, [Hema and Garcia Marquez \(2023\)](#), used CNN to classify speech emotions and benchmarked on a dataset consisting of seven classes (anger, disgust, fear, happiness, neutral, sadness, and surprise). However, CNNs lack the ability to model temporal dependencies explicitly. In speech recognition, understanding the temporal context of speech is essential for accurate transcription. Furthermore, speech signals are inherently sequential, and information from previous time steps is crucial for understanding the current speech segment. CNNs, by design, do not inherently capture this sequential nature. For this reason, variants of RNNs have been leveraged to collect extended contexts in speeches. This is because RNNs are designed to model sequential data by maintaining hidden states that capture information from previous time steps. This allows them to capture temporal dependencies effectively, making them well-suited for ASR tasks. In [Shewalkar et al. \(2019\)](#), the authors evaluated the performance of RNN, LSTM, and GRU on a popular benchmark speech dataset (EDLIUM). The results show that LSTM achieved the best word error rate, while the GRU optimization was faster and achieved word error rate close to that of LSTM.

However, RNN architectures process input sequences sequentially, which limits their ability to capture global context information effectively. As a result, they may struggle to understand the entire context of a spoken utterance, leading to lower transcription accuracy, particularly in tasks requiring understanding beyond local dependencies. Furthermore, most CNN and RNN automatic speech recognition systems consist of separate acoustic, pronunciation, and language modeling components that are trained independently. Usually, the acoustic model bootstraps from an existing model that is used for alignment to train it to recognize context dependent (CD) states or phonemes. The pronunciation model, curated by expert linguists, maps the sequences of phonemes produced by the acoustic model into word sequences. For this reason, Sequence-to-Sequence models are being proposed in automatic speech recognition to train the acoustic, pronunciation, and language modeling components jointly in a single system ([Prabhavalkar et al., 2017](#)). Sequence-to-Sequence models in automatic speech recognition are a class of models that aim to directly transcribe an input sequence of acoustic features such as speech spectrograms or Mel-frequency cepstral coefficients into a sequence of characters or words representing the recognized speech. There have been various sequence-to-sequence models explored in the literature, including Recurrent Neural Network Transducer (RNN-T) ([Graves, 2012](#)), Listen, Attend and Spell (LAS) ([Chan et al., 2015](#)), Neural Transducer ([Jaitly et al., 2016](#)), Monotonic Alignments ([Raffel et al., 2017](#)) and Recurrent Neural Aligner (RNA) ([Sak et al., 2017](#)).

As shown in [Fig. 31](#), the encoder component takes the input sequence of acoustic features and processes it to create a fixed-dimensional representation, often called the context vector. This representation captures the essential information from the input sequence and serves as the basis for generating the output sequence. The decoder component takes the context vector produced by the encoder and generates the output sequence. In ASR, this output sequence consists of characters or words representing the recognized speech. The decoder is typically implemented as a recurrent neural network (RNN), such as a Long Short-Term Memory (LSTM) or Gated Recurrent Unit (GRU) network, or it could be a transformer-based architecture. During training, the model learns to map input sequences to their corresponding output sequences by minimizing a suitable loss function, such as cross-entropy loss. This is typically done using techniques like backpropagation through time (BPTT) or teacher forcing, where the model is trained to predict the next token in the output sequence given

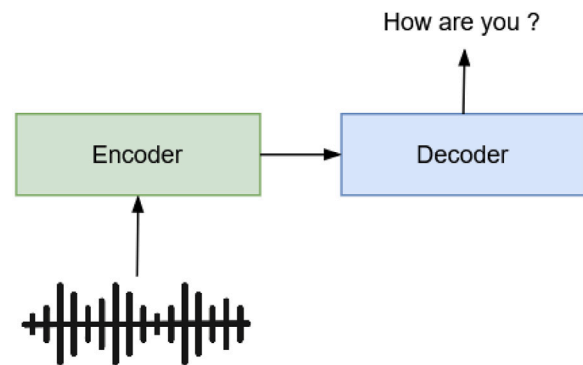


Fig. 31. Sequence-to-sequence.

the previous tokens. Thereafter, the trained model is used to transcribe unseen speech input. The encoder processes the input sequence to produce the context vector, which is then fed into the decoder to generate the output sequence. In some cases, beam search ([Szűcs and Huszti, 2019](#); [Li et al., 2018](#)) or other decoding strategies may be used to improve the quality of the generated output.

In [Chiu et al. \(2018\)](#), the authors explored various structural and optimization enhancements to their LAS Sequence to Sequence model, resulting in significant performance improvements. They introduce several structural enhancements, including the utilization of word piece models instead of graphemes and the incorporation of a multi-head attention architecture, which outperforms the commonly used single-head attention mechanism. Additionally, they investigate optimization techniques such as synchronous training, scheduled sampling, label smoothing, and minimum word error rate optimization, all of which demonstrate improvements in accuracy. The authors presented experimental results utilizing a unidirectional LSTM encoder for streaming recognition. On a 12,500-h voice search task, they observe a decrease in Word Error Rate (WER) from 9.2% to 5.6% with the proposed changes, while the best-performing conventional system achieves a WER of 6.7%. Moreover, on a dictation task, their model achieves a WER of 4.1%, compared to 5% for the conventional system. Similarly, the work of [Prabhavalkar et al. \(2017\)](#) investigated a number of sequence-to-sequence methods in automatic speech recognition. These included the RNN transducer (RNN-T), attention-based models, a new model that augments the RNN-T with attention, and a Connectionist Temporal Classification (CTC) trained system that directly outputs grapheme sequences. According to their research, sequence-to-sequence approaches can compete on dictation test sets against state-of-the-art when trained on a large volume of training data.

**Transformers:** Transformers have become the basis of state-of-the-art models in speech recognition. Wav2Vec 2.0 reads raw audio signals using multilayer convolutional encoder to generate latent speech representations ([Baevski et al., 2020](#)). The feature representation is subsequently fed to a transformer to capture the contextual information in the data. The transformer uses a convolutional layer for the positional encoding. Furthermore, the authors introduced vector quantization technique to convert the feature encoder output to discrete codes which are then used for contrastive learning during pretraining. Whisper is a robust speech recognition model proposed by OpenAI, trained on 680,000 h of labeled audio data ([Radford et al., 2022](#)). Of these, 117,000 h cover 96 non-English languages and 125,000 h consist of non-English to English translation data. The model is based on encoder-decoder transformer architecture, where the input data is first processed using convolutional layers followed by the positional encoding. Then, the input is fed into the encoder, which consists of a series of transformer blocks, each including multi-head self-attention mechanisms. The encoder's output is then fed into each decode transformer block to generate the final output. The decoder's transformer blocks

consist of multi-head self-attention and cross attention mechanisms. Conformer is a Transformer-based encoder, combining self-attention with convolutional layers to better capture both global and local features in speech signals (Gulati et al., 2020). The encoder provides a balance approach leveraging the strengths of convolution and self-attention mechanism. Conformer is often used with CTC or RNN-T, or other decoding mechanisms.

Auto-regressive decoding in transformer-based models is computationally expensive, as it requires repeated processing of the complete encoded speech, resulting in slow operation. In Stooke et al. (2025), Aligner-Encoder, a transformer-based encoder is proposed that perform internal alignment between audio and text, eliminating the need for complex decoding steps and complex dynamic programming during training. The approach simplifies the model architecture using a lightweight decoder without cross attention that processes embedding frames sequentially until an end-of-message token is generated. In Zhang et al. (2025), the Spike Window Decoding (SWD) algorithm is introduced that leverages the spike property of CTC outputs where each spike represents a strong signal indicating the presence of a specific token. By focusing on these spikes within a fixed window, SWD reduces the complexity of decoding, enabling faster and more efficient processing. Although transformer-based models are highly effective, they often capture highly entangled feature representations, leading to a lack of clear interpretability. In Wang and hamme (2024), the authors proposed the Disentangled-Transformer to enhance the interpretability of transformer-based models. The transformer disentangles the internal representations into sub-embeddings based on the various temporal characteristics of the speech signals. To this end, the loss function is improved by introducing time-invariant regularization terms for each time frame. A state-of-the-art speech recognition model called Samba-ASR is proposed utilizing Mamba architecture that is based on state-space-models (Shakhdari et al., 2025). This approach overcomes the limitations of transformer-based models such as inability to handle long-range dependencies and their quadratic scaling with input length, by efficiently capturing both local and global temporal dependencies through state-space-dynamics. Table 10 lists the summary of state-of-the-art speech recognition applications.

#### 4.3.3. Finance

Over the past few decades, computational intelligence in finance has been a hot issue in both academia and the financial sector (Ozbayoglu et al., 2020). Deep learning, especially RNN models have gained significant traction in the field of finance due to its ability to handle sequential data, since financial data often exhibit sequential dependencies, such as time series data for stock prices or historical transaction data. Within the financial industry, researchers have developed deep learning models for stock market forecasting (Singh and Srivastava, 2017), algorithmic trading (Lei et al., 2020), credit risk assessment (Shen et al., 2021), portfolio allocation (Wang et al., 2020), asset pricing (Chen et al., 2024), and derivatives markets (Ahnouch et al., 2023), among others and these models are intended to offer real-time operational solutions. In exchange rate prediction, Sun et al. (2020) developed an ensemble deep learning technique known as LSTM-B by combining a bagging ensemble learning algorithm with a long-short term memory (LSTM) neural network to increase the profitability of exchange rate trading and produce accurate exchange rate forecasting results. In comparison to previous methodologies, the authors' estimates proved to be more accurate when they looked at the potential financial profitability of exchange rates between the US dollar (USD) and four other major currencies: GBP, JPY, EUR, and CNY.

The authors in Abedin et al. (2021) proposed a Bi-LSTM-BR technique, which combined Bagging Ridge (BR) regression with Bi-LSTM as base regressors. The pre-COVID-19 and COVID-19 exchange rates of 21 currencies against the USD were predicted using the Bi-LSTM BR, and experiments showed that the proposed method outperformed ML algorithms such as DT and SVM. However, exchange rate data can

be noisy and subject to non-stationarity, which can pose challenges for predictive modeling. While bagging techniques can help mitigate the effects of noise to some extent, they may struggle to capture long-term trends or sudden shifts in the data distribution, leading to suboptimal performance. To address this, Wang et al. (2023) presented an approach for one-day ahead of time exchange rate prediction that concurrently considers both supervised and unsupervised deep representation features to enhance Random Subspace. Two crucial phases in the SUDF-RS technique are feature extraction and model building. First, LSTM and deep belief networks, respectively, extract the supervised and unsupervised deep representation features. To produce high-quality feature subsets, an enhanced random subspace approach was created that integrates a random forest-based feature weighting mechanism. Then, the matching base learner is trained using each feature subset, and the final outputs are generated by averaging the outcomes of each base learner. Experiments on EUR/USD, GBP/USD and USD/JPY showed that improved accuracy was achieved using the model.

In stock market prediction, several deep learning architectures have been proposed in the literature. For instance, Nikou et al. (2019), conducted a comparative study between the ANN, SVR, RF and an LSTM model. As compared to the other models discussed in the study, the LSTM model outperformed the others in predicting the closing prices of iShares MSCI United Kingdom. Similarly, using stock market historical data and financial news, Cai et al. (2018) used CNN and LSTM forecasting methods to generate seven prediction models. The seven models were then combined into a single ensemble model in accordance with the ensemble learning approach to create an aggregated model. However, the accuracy of all the models' predictions was low. Gudelek et al. (2017) proposed a CNN model which used a sliding window technique and created pictures by capturing daily snapshots within the window's bounds. With 72% accuracy, the model was able to forecast the prices for the following day and was able to generate 5 times the starting capital. In Eapen et al. (2019), a CNN and Bi-LSTM model with numerous pipelines was proposed, utilizing an SVM regressor model on the S&P 500 Grand Challenge dataset, and results showed enhanced prediction performance by over a factor of 6% compared to baseline models. As presented, deep learning has undeniably achieved state-of-the-art performance across various domains within finance. In Chen et al. (2024), a multi-modal deep learning model is proposed for stock market trend prediction by integrating daily stock prices, technical indicators and sentiment from daily news headlines. The model architecture consists of BERT-based branch fine-tuned on financial news to capture textual sentiment and an LSTM branch that captures temporal patterns in the time series data, including stock prices and technical indicators. Both feature representations are combined through concatenation and then passed to a fully connected layer for predictions.

A study was conducted to analyze the performance of Prophet, LSTM and Transformer. Prophet is a parametric, additive regression model based on time series decomposition (Mozaffari and Zhang, 2024). The LSTM model consists of two hidden layers, with a hidden size of 64 and a linear fully connected layer is used to generate the output. The Transformer model follows the original transformer architecture. The results show that transformer outperformed LSTM and Prophet in terms of forecasting accuracy, particularly for datasets with complex temporal dependencies. However, transformer models have high memory usage and quadratic complexity due to self-attention mechanism, which makes them inefficient for very long sequence time-series forecasting. To overcome the limitations, a StockFormer is proposed which is based on Informer (Li et al., 2025). Informer is a transformer that was designed for time-series forecasting, utilizing the ProbSparse self-attention mechanism which attends to only the most important queries for each key, reducing the computational complexity (Zhou et al., 2021). Furthermore, self-attention distilling technique is introduced to reduce redundancy and highlighting crucial information in the feature maps, helping the model to focus more on

**Table 10**  
Summary of state-of-the-art speech recognition.

Ref.	Description	Results (Datasets and metrics)
Radford et al. (2022)	This study proposed Whisper an automatic speech recognition (ASR) system that transcribes and translates spoken language using a robust deep learning model trained on diverse multilingual and multitask datasets	Various datasets (first three datasets only) LibriSpeech Clean WER: 2.7 Artie WER: 6.2 Common Voice WER: 9.0
Stooke et al. (2025)	This study introduced Aligner-Encoder, a transformer-based automatic speech recognition model that internally aligns audio-text representations before decoding, enabling a simpler and more efficient architecture trained with frame-wise cross-entropy loss	LibriSpeech-960H WER: 2.2 Voice Search WER: 3.7 YouTube videos WER: 7.3
Zhang et al. (2025)	An algorithm that leverages spike property of CTC outputs to reduce the complexity of WFST decoding process	AISHELL-1 CER: 3.89 In House CER: 2.09
Wang and hamme (2024)	A transformer-based model architecture designed to disentangle internal representations and enhance model explainability.	LibriSpeech-100H DER: 8.1 LibriMix 1 DER: 5.7 LibriMix 2 DER: 6.9 LibriMix 3 DER: 2.5 LibriMix 4 DER: 5.6
Shakhadri et al. (2025)	A speech recognition model based on state-space-models	GigaSpeech WER: 9.12 LibriSpeech Clean WER: 1.17 LS Other WER: 2.48 SPGISpeech WER: 1.84

the relevant parts of the input sequence. In [Berti and Kasneci \(2025\)](#), a transformer-based model is proposed for stock trend prediction using limit order book data (LOB). The model consists of a series of transformer LOB (TLOB) blocks which leverage dual attention mechanisms to capture both spatial and temporal dependencies, allowing the model to adaptively focus on the market microstructure of the input data. Furthermore, a multilayer of perceptron LOB (MPLOB) is introduced, comprising two blocks of two fully connected layers which operate along the feature and temporal dimensions. [Table 11](#) lists the summary of state-of-the-art finance applications.

#### 4.3.4. Electrocardiogram (ECG) classification

Disorders pertaining to the heart or blood vessels are collectively referred to as Cardiovascular Diseases (CVD) ([Liu et al., 2021](#)). According to the American Heart Association's 2023 statistics, CVD has emerged as the leading cause of death worldwide. In 2020, 19.05 million deaths were recorded from CVD globally, which signifies an increase of 18.71% from 2010, and it is believed that this number will rise to 23.6 million by 2030 ([Tsao et al., 2023](#)). Blood clots and vascular blockages caused by CVDs can cause myocardial infarction, stroke, or even death ([Liu et al., 2021](#)). Generally, early diagnosis has been shown to reduce the mortality rate of CVDs, and Electrocardiogram (ECG) signals play a crucial role in diagnosing various cardiac abnormalities and monitoring heart health. However, ECG signal has characteristics of high noise and high complexity, making it time-consuming and labor-intensive to identify certain diseases using traditional methods. The traditional approach is tedious and requires the expertise of a medical specialist. Over the past decades, the task of Long-term ECG recording classification has been significantly facilitated for cardiologists through the adoption of computerized ECG recognition practices. Throughout this period, feature extraction methods have predominantly relied on manual techniques, encompassing diverse approaches

such as wave shape functions ([Llamedo and Martínez, 2011](#)), wavelet-based features ([Mathews et al., 2018](#)), ECG morphology ([zhu wenliang wenliang et al., 2019](#)), hermite polynomials ([Desai et al., 2021](#)), and Karhunen–Loeve expansion of ECG morphology ([Crippa et al., 2015](#)), among others. These extracted features are subsequently subjected to classification using various machine learning algorithms.

More recently, the advent of deep learning has revolutionized the field by enabling automatic feature learning directly from ECG signals. This advancement holds significant promise in the realm of automated ECG classification, offering clinicians a tool for swift and accurate diagnosis. Based on this, several deep learning architectures have been proposed for feature learning of ECG signals. For instance, [Acharya et al. \(2017\)](#) developed a 9-layer CNN model to automatically identify five categories of heartbeats in ECG signals. A similar model was also developed in [Baloglu et al. \(2019\)](#). However, ECG signals often vary significantly in length, as they may contain different numbers of heartbeats. CNNs typically require fixed-length inputs, which may necessitate preprocessing steps such as padding or truncation, potentially losing important temporal information. For this reason, several architectures have leveraged RNN in ECG classification, as seen in [Singh et al. \(2018\)](#), [Prabhakararao and Dandapat \(2020\)](#) and [Wang et al. \(2023\)](#), among others. While RNNs are capable of handling sequential data, they also have limitations in capturing local patterns or short-term dependencies effectively. In ECG signals, local features such as specific waveforms or intervals can be crucial for classification. For this reason, recent works have proposed hybrid models which combine the strengths of both CNNs and RNNs to overcome some of these limitations ([Sowmya and Jose, 2022](#)).

**Hybrid Models:** The work of [Rai and Chatterjee \(2022\)](#) developed a hybrid CNN–LSTM network to evaluate the optimum performing model for myocardial infarction detection using ECG signals. The authors then experimented on 123,998 ECG beats obtained from the

**Table 11**

Summary of state-of-the-art finance applications.

Ref.	Description	Results (Datasets and metrics)
Wang et al. (2023)	This study introduces SUDF-RS technique for one-day-ahead exchange rate prediction by combining supervised and unsupervised deep representation features with an enhanced random subspace approach, leveraging LSTM and deep belief networks for feature extraction and using a random forest-based weighting mechanism to improve prediction accuracy	Data collected from investing.com RMSE: 0.448 (EUR/USD), 0.676 (GBP/USD), 0.374 (USD/JPY)
Chen et al. (2024)	A multi-modal deep learning model integrates stock prices, technical indicators, and news sentiment using a BERT-based branch for textual analysis and an LSTM branch for temporal patterns, combining their features for stock market trend prediction	Data collected using Yahoo Finance and EODHD API F1-Score: 0.45 (ATVI), 0.51 (AAPL), 0.48 (AMT), 0.53 (PLD), 0.50 (NDAQ), 0.45 (SCHW), 0.43 (BIO), 0.47 (JNJ)
Mozaffari and Zhang (2024)	This study analyzes stock closing price prediction using LSTM, Prophet, and Transformer models	American Airlines Group Inc. MSE: 0.0085 (Transformer), 0.1972 (LSTM), 16.9321 (Prophet) Atlantic American-Life Insurance MSE: 0.0118 (Transformer), 0.0095 (LSTM), 17.8601 (Prophet)
Li et al. (2025)	The study proposed a model based on Informer which utilizes ProbSparse self-attention and self-attention distilling to reduce the computation complexity	Combined data from Yahoo, alphavantage.co, alpaca.markets and polygon.io Percent profit: 1.7550
Berti and Kasneci (2025)	The study proposed a transformer-based model with dual attention mechanisms and multilayer of perceptrons for feature mixing along feature and temporal dimensions	FI-2010 F1-score: 92.81 TSLA F1-score: 60.50 INTC F1-score: 80.15

PTB diagnostic database (PTBDB) and MIT-BIH arrhythmia database (MITDB), and the result showed that by combining the capabilities of both CNN and LSTM, improved classification accuracy can be achieved. Also, in Banerjee et al. (2020), a CNN architecture was developed to extract morphological features from ECG signals. For the purpose of determining the degree of heart rate variability, another composite structure was designed using LSTM and a collection of manually created statistical features. Following that, a hybrid CNN-LSTM architecture is built using the two independent biomarkers to classify cardiovascular artery diseases, and experiments were carried out on two distinct datasets. The first is a partly noisy in-house dataset collected using an inexpensive ECG sensor, and the other is a corpus taken from the MIMIC II waveform dataset. The hybrid model proposed in the work achieved an overall classification accuracy of 88% and 93%, respectively, which surpasses the performance of standalone architectures.

An automated diagnosis method based on Deep CNN and LSTM architecture was presented in Kusuma and Jothi (2022) to identify Congestive Heart Failure (CHF) from ECG signals. Specifically, CNN was used to extract deep features, and LSTM was employed to exploit the extracted features to achieve the CHF detection goal. The model was tested using real-time ECG signal datasets, and the results show that the AUC was 99.9%, the sensitivity was 99.31%, the specificity was 99.28%, the F-Score was 98.94%, and the accuracy was 99.52%. Similar work is reported in Alamatsaz et al. (2024) whereby CNN is combined with LSTM for ECG classification. The hybrid model consists of three blocks of convolutional and maximum pooling layers with dropout followed by an LSTM with drop. Furthermore, Shapley is utilized to determine the contribution of each ECG sample on the prediction to improve model interpretability. However, since ECG signals can vary in length due to differences in recording durations or patient conditions. LSTMs are capable of handling variable-length sequences, but traditional CNNs typically require fixed-length inputs. Therefore, fusing these features effectively in a hybrid model can be challenging. Furthermore, Hybrid CNN-RNN models can be computationally intensive, especially when processing long ECG sequences or large datasets. For this reason, recent research works have proposed the use of attention mechanisms to reduce the computational burden by enabling the model to selectively attend to informative features, focusing computational resources where they are most needed. Likewise, attention mechanisms can enable the model to attend to informative

segments of the ECG signal, regardless of their length, allowing for more flexible processing of variable-length sequences.

Several researchers have leveraged attention mechanisms in standalone and hybrid architectures for improved performance. For instance, in the work of Chen et al. (2022), CNN layers were used to extract main features, while LSTM and attention were included to enhance the model's feature learning capabilities. Experiments on a 12-lead KMH ECG dataset showed that the model had high recognition rates in classifying normal and abnormal ECG signals, compared to hybrid models without attention mechanisms. Wang et al. (2021) presented a 33-layer CNN architecture with non-local convolutional block attention module (NCBAM). To extract the spatial and channel information, preprocessed ECG signals were first fed into the CNN architecture. A non-local attention further captured long-range dependencies of representative features along spatial and channel axes. Similarly, a spatio-temporal attention-based convolutional recurrent neural network (STA-CRNN) was presented in Zhang et al. (2020) with the aim of concentrating on representative features in both the spatial and temporal dimensions. The CNN subnetwork, spatiotemporal attention modules, and RNN subnetwork formed the STA-CRNN and according to findings, the STA-CRNN model was able to classify eight different forms of arrhythmias and normal rhythm with an average F1 score of 0.835. In Sun et al. (2024), the SE module is utilized in between convolutional layers and LSTM. This approach allows the model to focus on relevant channels before temporal dependencies are captured by the LSTM, improving the classification performance.

In Huang et al. (2024), a guided spatial attention mechanism is introduced to incorporate domain knowledge into the model, enhancing ECG classification performance. The guided spatial attention mechanism has an autoencoder-like structure, where the feature maps are downsampled by the encoder and upsampled by the decoder. Additionally, skip connections are employed to improve information flow between the encoder and decoder. Furthermore, an attention loss term, based on the attention weights is introduced to jointly train the guided attention mechanisms. In Aghaomidi and Wang (2024), ECG signals are classified into sleep stages using deep learning. The approach is divided into three stages. In the first stage, a deep learning model consisting of convolutional layers and liquid time-constant network (Hasani et al., 2021) is employed to compute kurtosis and skewness of the ECG signals. In the second stage, a deep learning model with convolutional

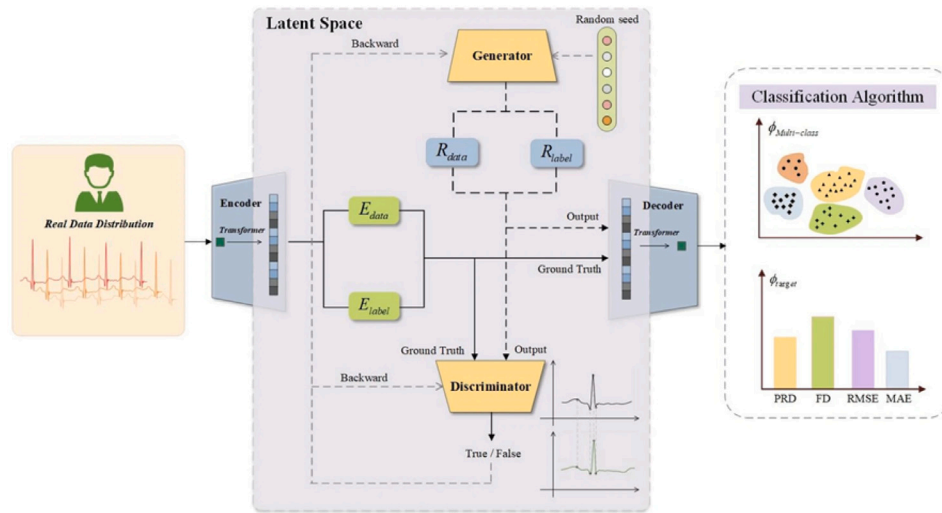


Fig. 32. The architecture of CECG-GAN (Yang et al., 2024).

layers and SE attention is used to enhance the prediction of the minority sleep stage class. In the final stage, the outputs of the models are concatenated and classified into sleep stage categories. Combining hybrid deep learning models with attention mechanisms for ECG feature learning is a promising approach that has already shown potential in ECG feature learning, according to reviewed literature. Future research can further explore semi-supervised and self-supervised learning techniques to leverage large amounts of unlabeled ECG data. This could involve pre-training models on large-scale unlabeled datasets using self-supervised learning objectives.

**Generative Models:** Deep learning models have been leveraged in the generation of synthetic ECG signals to augment real signals, as seen in Zhu et al. (2019) where a GAN model was developed to generate ECG signals that correspond with available clinical data. The GAN model used two layers of BiLSTM for the generator and CNN for the discriminator, and trained using the 48 ECG recordings of different users from the MIT-BIH dataset. The authors then compared their model with a Recurrent neural network autoencoder (RNN-AE) model and a recurrent neural network variational autoencoder (RNN-VAE) model, and the results show that their model exhibited the fastest convergence of its loss function to zero. Similar work is reported in Yang et al. (2024) in which GAN and autoencoder are employed to generate ECG signals to address the issue of imbalanced dataset as shown in Fig. 32. The generator and encoder create synthetic and real low-dimensional ECG samples, respectively. The discriminator classifies the synthetic and real low-dimensional ECG samples as real or fake, while the decoder reconstructs the ECG signals from the synthetic low-dimensional samples. This approach further refines data transformation and reconstruction quality, generating more realistic ECG signals. In Msigwa et al. (2024), cluster-based GAN is employed to improve the diversity and quality of the generated synthetic ECG signals. *K*-means algorithm is first employed to partition the ECG signals into *k* distinct groups. Then, the GAN is trained using the cluster centroids as initial reference points to guide the GAN in generating synthetic signals. Table 12 lists the summary of state-of-the-art ECG classification.

#### 4.3.5. Electroencephalography (EEG) classification

Three-dimensional scalp surface electrode readings provide a dynamic time series that is called Electroencephalogram (EEG) signal (Schirrmester et al., 2017). Brain waves obtained from an EEG can effectively depict both the psychological and pathological states of a human. The human brain is acknowledged to be a fascinating and incredibly complicated structure. Numerous brain signals, including functional magnetic resonance imaging (fMRI), near-infrared spectroscopy

(NIRS), electroencephalograms (EEGs), and functional near-infrared spectroscopy (fNIR), among others have been collected and used to study the brain (Gao et al., 2021). Due to the EEG's non-invasive, affordable, accessible, and excellent temporal resolution characteristics, it has become the most utilized approach. However, the signal-to-noise ratio of EEG signal is low, meaning that sources with no task-relevant information frequently have a stronger effect on the EEG signal than those that do. These characteristics often make end-to-end feature learning for EEG data substantially more challenging (Schirrmester et al., 2017). Based on this, several methods have been leveraged for improved feature extraction in EEG signals across several domains including Motor imagery (Ang and Guan, 2017), anxiety disorder (Shen et al., 2022), epileptic seizure detection (Boonyakitanont et al., 2020), sleep pattern analysis and disorder detection (Sharma et al., 2021; Vaquerizo-Villar et al., 2023), and Alzheimer's disease detection (Modir et al., 2023), and many others.

EEG Motor Imagery (MI) is a technique used to study brain activity associated with the imagination of movement. It involves recording electrical activity generated by the brain through electrodes placed on the scalp. MI tasks typically involve imagining performing a specific motor action, such as moving a hand or foot, without physically executing the movement, and has been leveraged in smart healthcare applications such as post-stroke rehabilitation and mobile assistive robots, among others (Altafheri et al., 2023). Prior to the advent of deep learning, motor imagery EEG data are passed through various steps before classification using traditional ML techniques. Pre-processing, feature extraction, and classification are the three primary stages that traditional approaches usually take while processing MI-EEG signals. Pre-processing includes a number of operations, including signal filtering (choosing the most valuable frequency range for MI tasks), channel selection (identifying the most valuable EEG channels for MI tasks), signal normalization (normalizing each EEG channel around the time axis), and artifact removal (removing noise from MI-EEG signals). Independent component analysis (ICA) is the most often utilized technique for removing artifacts (Brunner et al., 2007; Delorme et al., 2007; Jafarifarmand and Badamchizadeh, 2019). In contrast to the traditional approach, deep learning architectures can automatically extract complex features from raw MI-EEG data without the need for laborious feature extraction and pre-processing. Based on this, several deep learning architectures have been proposed for MI-EEG feature learning, as seen in Zhang et al. (2019), Kumar et al. (2016) and Tibrewal et al. (2022), among others. For instance Schirrmester et al. (2017), categorized MI-EEG signals using three CNNs with varying architectures, and the number of convolutional layers varied from

**Table 12**

Summary of state-of-the-art ECG classification.

Ref.	Description	Results (Datasets and metrics)
Sun et al. (2024)	A hybrid model consists of three blocks of convolutional layer, SE module and an LSTM	MIT-BIH ECG Accuracy: 98.5%
Alamatsaz et al. (2024)	A lightweight hybrid model consists of three blocks of convolutional and maximum pooling layers and an LSTM	MIT-BIH and LTAF Accuracy: 98.24%
Chen et al. (2022)	A hybrid model consists of three convolutional pipelines with different kernel size, two LSTM layers and an attention module	KMUH Accuracy: 96.02% CPSC-2018 Accuracy: 94.05%
Huang et al. (2024)	This study proposed a guided spatial attention mechanism to enhance ECG classification performance	CPSC2018 and PTB-XL Chapman Cinc2017 F1-score: 77.56% (STC), 88.53% (PC), 82.54% (WPW), 85.37% (AF)
Aghaomidi and Wang (2024)	A novel three-stage approach for sleep stage classification extracts statistical features using a Feature Imitating Network, enhances minority sleep stage detection, and integrates outputs for five-class classification	MIT-BIH Polysomnographic Accuracy: 80.79%
Yang et al. (2024)	This study proposed a generative model by combining conditional GAN and autoencoder	CSPC2020 dataset F1-score: 0.9823 (Normal), 0.9824 (Premature ventricular), 0.9825 (Premature Supraventricular)
Msigwa et al. (2024)	This study introduces a generative model by combining GAN and K-means	MIT-BIH arrhythmia Accuracy: 99.66%

two layers to a five-layer deep ConvNet to a thirty-one-layer residual network.

In Dai et al. (2019), the authors proposed an approach for classifying MI-EEG signals which blend variational autoencoder with CNN architecture. The VAE decoder was used to fit the Gaussian distribution of EEG signals, and the time, frequency, and channel information from the EEG signal were combined to create a novel representation of input, and the proposed CNN-VAE method was optimized for the input. Experiments showed that by combining both deep learning architectures, improved features were learnt, which led to a high classification performance on the BCI Competition IV dataset 2b. Li et al. (2017) employed optimal wavelet packet transform (OWPT) for the generation of feature vectors from MI-EEG signals. These vectors were then utilized to train an LSTM network, which demonstrated satisfactory performance on dataset III from the BCI Competition 2003. However, the model has an excessively intricate structure. To address this, Li et al. (2020) introduced a technique that merges continuous wavelet transform (CWT) with a simplified convolutional neural network to enhance the accuracy of recognizing MI-EEG signals. By employing CWT, MI-EEG signals were converted into time-frequency image representations. Subsequently, these image representations were fed into the SCNN for feature extraction and classification. Experiments on the BCI Competition IV Dataset 2b demonstrate that, on average, the classification accuracy across nine subjects reached 83.2%. However, the computational complexity of the model was quite high, due to the processing of time-frequency image representations. The conversion of MI-EEG signals into time-frequency images using CWT requires significant computational resources.

The authors in Hwang et al. (2023) introduced a classification framework based on Long Short-Term Memory (LSTM) to improve the accuracy of classifying four-class motor imagery signals from EEG. The authors used sliding window technique to capture time-varying EEG signal data, and employed an overlapping-band-based Filter Bank Common Spatial Patterns (FBCSP) method to extract subject-specific spatial features. Experiments on the BCI Competition IV dataset 2a, showed that their model achieved an average accuracy of 97%, compared to existing methods. Also, in the classification of Alzheimer's disease, Zhao and He (2015) employed a deep learning network to analyze EEG data. The deep learning model was evaluated on a dataset that consisted of fifteen (15) patients with clinically confirmed Alzheimer's disease and fifteen (15) healthy individuals, and results showed that improved features were learnt and compared the results to the traditional methods. This has prompted the use of deep learning in Alzheimer's disease detection, as seen in Xia et al. (2023), where the authors used CNN for

diagnosing Alzheimer's Disease. To address challenges posed by limited data and overfitting in deep learning models designed for Alzheimer's Disease detection, the authors explored the use of overlapping sliding windows to augment the EEG data collected from 100 subjects (comprising 49 AD patients, 37 mild cognitive impairment patients, and 14 healthy controls subjects). After assembling the augmented dataset, a modified Deep Pyramid Convolutional Neural Network (DPCNN) was used to classify the enhanced EEG signals.

**Hybrid Models:** In Hermawan et al. (2024), the authors developed three deep learning architectures (CNN, LSTM, and hybrid CNN-LSTM), with each model chosen for its effectiveness in handling the intricate characteristics of EEG data for epilepsy detection. Each architecture offers distinct advantages, with CNN excelling in spatial feature extraction, LSTM in capturing temporal dynamics, and the hybrid model combining these strengths. The CNN model, consisting of 31 layers, attained the highest accuracy, achieving 91% on the first benchmark dataset and 82% on the second dataset using a 30-s threshold, selected for its clinical significance. In the work of Abdulwahhab et al. (2024), EEG waves' time-frequency image and raw EEG waves served as input elements for CNN and LSTM models. Two signal processing methods, namely Short-Time Fourier Transform (STFT) and CWT, were employed to generate spectrogram and scalogram images, sized at  $77 \times 75$  and  $32 \times 32$ , respectively. The experimental findings demonstrated detection accuracies of 99.57% and 99.26% for CNN inputs using CWT Scalograms on the Bonn University dataset and 99.57% and 97.12% using STFT spectrograms on the CHB-MIT dataset. Similarly, in emotion recognition, several deep learning models have been leveraged with EEG signals. For instance, in Pandey and Seeja (2022), a subject-independent emotion recognition model was proposed, which utilizes Variational Mode Decomposition (VMD) for feature extraction and DNN as the classifier. Evaluation against the benchmark DEAP dataset demonstrates superior performance of this approach compared to other techniques in subject-independent emotion recognition from EEG signals.

Some researchers have also combined EEG signals with facial expression and speech in emotion recognition, as seen in Hassouneh et al. (2020), Pan et al. (2024), and Wang et al. (2023), among others. In Pan et al. (2024), a multi-modal emotion recognition framework is proposed, utilizing three deep learning models to extract features from facial expressions, speech and EEG signals. To process EEG signals, the authors designed a tree-like LSTM model that extracts temporal features at different stages to capture multiscale feature representations. The output of each model is then weighted to generate the final predictions. In Wang et al. (2023), the feature maps extracted by the deep learning

**Table 13**  
Summary of state-of-the-art EEG classification.

Ref.	Description	Results (Datasets and metrics)
Hermawan et al. (2024)	A hybrid CNN-LSTM model consists of a convolutional layer to extract local features and two LSTM layers for temporal modeling	UPenn and Mayo Clinic's Seizure Detection Challenge Accuracy: 89% F1-score (avg.): 87%
Pan et al. (2024)	A tree-like LSTM to capture multiscale temporal feature representations from EEG signals	MAHNOB-HCI Accuracy: 0.9450 (valence), 0.7628 (arousal)
Wang et al. (2023)	Parallel convolutional pipelines for extracting local and global features from EEG signals	DEAP Accuracy: 85.72% (valence), 87.97% (arousal) MAHNO B-HCI Accuracy: 85.98% (valence), 85.23% (arousal)
Abdulwahhab et al. (2024)	A framework comprises two separated deep learning-based models using a 2D Convolutional Neural Network and an LSTM to process EEG signal images and normalized EEG signals	Bonn Accuracy: 99.75% CHB-MIT Accuracy: 97.12%
Song et al. (2024)	A conditional GAN integrated with SE attention in the discriminator and employing a cropped training strategy for improved data generation	BCI Competition IV 2a Accuracy: 81.3% BCI Competition IV 2b Accuracy: 90.3%
Corley and Huang (2025)	A Wasserstein GAN model is used to enhance the spatial resolution of EEG signals for classification by generating high-resolution signals from low-resolution inputs while simultaneously interpolating missing channels, addressing noise and artifact contamination	Dataset V (Berlin BCI) Accuracy: 83.88% (scale 2), 82.00% (scale 4)
Cai et al. (2025)	A dual-branch hybrid model that denoises EEG signals by separately learning features from clean and noisy inputs, fusing them through fully connected layers, and leveraging convolutional layers with self-attention mechanisms	EEGdenoiseNet(EEG, EOG, EMG), MIT-BIH Arrhythmia (ECG) and semi-simulated EEG/EOG $\eta$ (artifact reduction): 82.35% (EMG), 91.80% (EOG), 88.97% (ECG), 82.07% (EOG+EMG)

models are combined through concatenation, and then they are fed into a fully connected classifier to generate the predictions. For EEG signals, two deep learning models are designed: one for extracting local features and the other for extracting global features. Both models are implemented using convolutional layers. However, EEG signals can vary significantly across individuals, making it challenging to generalize models across different subjects.

**Generative Models:** GANs have been employed to enhance EEG signal classification. In Song et al. (2024), a conditional GAN model is proposed to overcome the limited number of EEG samples. To improve data generation quality, the discriminator is integrated with SE attention and a cropped training strategy is employed to leverage the entire spectrum of the EEG signals. In Corley and Huang (2025), a Wasserstein GAN model is employed to enhance the spatial resolution of EEG signals for EEG classification. The model takes low-resolution signals and generates the corresponding high-resolution signals while simultaneously interpolating the missing channels. EEG signals are often corrupted with noise and artifacts. DHCT-GAN is a dual-branch hybrid model designed to generate denoised EEG signals by processing both clean and noisy/artifact-corrupted signals (Cai et al., 2025). The generator consists of dual-branch feature learning pipelines, where one accepts clean signals while the other handles noisy/artifact signals to learn feature representations. The feature maps are then fused using two fully connected layers with a tanh activation function to generate denoised EEG signals. Both feature learning branches are identical, consisting of convolutional layers with local and global self-attention mechanisms, while the discriminators are implemented using convolutional and fully connected layers. Table 13 lists the summary of state-of-the-art EEG classification.

#### 4.4. Robotics

Deep learning is a crucial component in modern robotics, enabling machines to perceive and interpret their surroundings. By leveraging sensors and cameras, robots can identify objects, navigate environments and interact with humans more effectively.

##### 4.4.1. Object identification

Deep learning plays a crucial role in object identification, allowing robots to recognize, classify, and differentiate between various objects in real-time. In Le and Ngo (2025), an intelligent system, integrating deep learning with robotic kinematic control is proposed for waste classification. The system consists of a robotic arm with a gripper and a camera placed in front of the robot to capture real-time video. YOLO model is trained to recognize and localize objects, allowing the system to determine the location, pick them up and classify them. A similar system is reported in Vukicevic et al. (2025), where a robotic arm is used to pick up objects on a conveyor and a camera is mounted above the conveyor to capture real-time video. The waste sorting process is divided into two tasks: localization of the waste objects and their classification. The object localization is performed using Segment Anything architecture (SAM) (Kirillov et al., 2023) which is based on vision transformer. Several lightweight SAM models are proposed to solve the localization task efficiently while minimizing computational requirements. For classification, object images are cropped using SAM's output masks and classified by a MobileNetV2.

A deep learning-based method is proposed to detect tile edges and corners for tile-paving mobile robotic systems (Liu et al., 2025). The camera is mounted on a tile-grabbing platform that is parallel to the ground and approximately 10 to 20 cm above it. The system utilizes a YOLOv8 for semantic segmentation with rule-based post-processing techniques. An unmanned ground vehicle-based is proposed for automated structural damage inspection (Ge and Sadhu, 2025). The robotic system utilizes LiDAR and a camera, integrating them via Robot Operating System for efficient control and data processing. The images are processed for the damage detection and segmentation using an improved YOLOv7 model, while the point cloud data is processed for real-time localization and 3D mapping using an enhanced Keep It Small and Simple-Iterative Closest Point algorithm. A similar system is reported in Dai et al. (2024), where object detection and tracking approach is proposed to fuse point cloud and visual data. The approach leverage YOLOv8 real-time object detection capabilities, while LiDAR provides accurate 3D spatial information and distance measurements, enhancing perception reliability. A fusion method aligns LiDAR data with camera images, enabling accurate 3D bounding boxes and object tracking. Table 14 lists the summary of state-of-the-art robot object identification.

**Table 14**

Summary of state-of-the-art robot object identification.

Ref.	Hardware platform	Model	Task
Le and Ngo (2025)	Robotic arm, AI Stereo ZED 2 camera, gripper, Intel Xeon E5 series CPU @ 2.40 GHz	YOLOv4	Waste classification
Vukicevic et al. (2025)	PC workstation, industrial camera, AMD Threadripper 3970X 32 cores CPU @ 3.79 GHz, 128 GB RAM, two Titan RTX 24 GB GPUs	Segment Anything architecture (SAM)	Waste sorting
Kirillov et al. (2023)	Mobile robot, camera, Nvidia RTX 3060 GPU	YOLOv8	Tile edges and corners detection
Ge and Sadhu (2025)	Husky Unmanned ground vehicle, Intel Realsense camera, Velodyne VLP-16 LiDAR, Intel i7 CPU	YOLOv7	Structural damage inspection
Dai et al. (2024)	Agilex Scout Mini robot, Velodyne LiDAR, Intel D435 camera	YOLOv8	Object tracking

#### 4.4.2. Path extraction and navigation

Deep learning models have been applied for robot navigation systems. In Alotaibi et al. (2024), LiDAR and a camera are to capture point cloud and visual data for a comprehensive understanding and navigation in the environments. The study experiments with Faster R-CNN, YOLOv5 and YOLOv8 for object detection. Additionally, the paper compares the performance of these systems in various real-world environments, showcasing their potential to enhance autonomous navigation. Mobile robot navigation system is proposed utilizing semantic segmentation to determine the drivable paths (Misir, 2024). First, Deeplabv3 (Chen et al., 2017) model with a ResNet-50 backbone is applied to segment scene images to extract drivable area. A perspective transformation then maps the segmented images to real-world space. Following the transformation, the image is divided into grids to determine the optimal path while avoiding obstacles. Finally, a PID controller guides the robot along the smoothed path to ensure accurate navigation in the environment.

In Cao et al. (2024), a method for orchard robot navigation utilizing a modified YOLOv8 model is proposed. The model locates vine tree trunks and identifies the center points of tree trunks at the lower end of the detection boxes. The least squares method is then applied to fit reference lines on both sides of the trunk. Furthermore, a multi-scale attention mechanism is incorporated into the model to prioritize relevant features, enhancing the model performance. In Liu et al. (2025), deep learning is employed for navigation path extraction for agricultural robots. The authors propose Single-Stage Navigation Path Extraction Network (NPENet) to predict the road centerline in a single stage, eliminating the need for complex multi-stage processes such as line detection and segmentation. The proposed model uses residual blocks for feature extraction which includes batch normalization, leaky ReLU and convolutional layers, generating three key outputs: detecting if the robot is on the road, estimating the navigation line's angle and predicting the vanishing point of the road. These outputs enable the robot to determine its trajectory and real-time navigation in orchard environments. Table 15 lists the summary of path extraction and navigation.

#### 4.4.3. Human-robot interaction

Deep learning has significantly advanced human-robot interaction, particularly through the integration of large language models (LLMs) designed for natural language processing. Despite their extensive knowledge, LLMs have limitations, as they can sometimes generate inaccurate information, a phenomenon known as hallucination (Huang et al., 2025). To address this issue, LLMs need to be connected to the physical world by providing real-world contextual information to the models, enabling them to generate accurate responses that are relevant to the situation. This process is known as LLM grounding. In Asuzu et al. (2025), the authors explore the use of LLMs in human-robot interaction, focusing on collaborative planning between humans and robots. The LLM is guided to generate outputs within the predefined set of robot skills using a few-shot prompting, while vision models such as SAM-CLiP and ViLD are used for object identification within the scene, enabling the robot to perform actions in the real-world. By leveraging

LLMs, robots can understand and process natural language instructions and execute the corresponding actions, enabling efficient collaboration with human users. LLMs need to be aligned with human intentions to ensure effective communication and interaction between humans and robots. In Chen et al. (2024), a 3D scene graph is employed to represent objects and environments captured by 3D instance segmentation and 2D image classification. Given the 3D scene graph, a complex task is decomposed using LLM through step-by-step reasoning by progressively retrieving relevant nodes from the graph. During the task solving, human intent is automatically identified, and the system generates responses accordingly.

In Allgeuer et al. (2024), a modular system architecture that can be easily extended is proposed to integrate LLM models for human-robot conversation and collaboration. Fig. 33 shows the system architecture, where the chat manager is the central to the system that coordinates the state and inference of the LLM model. The robot perception such as recognizing objects and human gesture is handled separately by different deep learning models while robot actions are operated by speech generation model and actuators for emotion expression, gaze control and arm movement. This approach allows the robot to engage in open-ended conversation and collaborate naturally with humans. Grounding LLMs can be challenging due to robot's multi-modal perception system. Furthermore, the incoming data has different sampling rate which makes data alignment difficult, causing the robotic system to miss valuable information. In Wang et al. (2024), a framework is proposed for processing multi-modal inputs to generate coherent narratives about robot experiences including its internal status, current observation, planning and critical alert. The framework is divided into three stages: key event selection, experience summarization and narration generation modules. First, multi-modal data across three categories: environment, internal and planning are aligned by sampling at a fixed rate and associating each frame at the nearest timestamp. Based on the aligned data, key events are selected by detecting changes in optical flow, RGB images and joint states using heuristic and normalization techniques to identify moment of interest. Using the detected events, the robotic data is converted into natural language summaries, and an LLM generates user-friendly explanation by filtering repetitive and irrelevant information. Table 16 lists the summary of human-robot interaction.

## 5. Research challenges

Despite the success of deep learning in various applications, there are still fundamental challenges faced by researchers and practitioners. This section discusses the challenges ethical, technical and domain-specific perspectives.

### 5.1. Ethical issues

Interpretability and explainability is crucial for building trust and understanding how predictive models make decisions, especially in high-stake applications such as healthcare and medical image analysis (Tonekaboni et al., 2019). However, as deep learning models

**Table 15**  
Summary of state-of-the-art path extraction and navigation.

Ref.	Hardware platform	Model	Task
Alotaibi et al. (2024)	Mobile robot, 2MP Logitech webcam, LiDAR, Raspberry Pi 4, 4 GB RAM	YOLOv8	Semantic navigation
Misir (2024)	Mobile robot, camera, Nvidia Jetson Nano	Deeplabv3+ (ResNet50 backbone)	Path planning
Cao et al. (2024)	Mobile robot, Intel RealSense D455 camera, Intel Core i7 10870H CPU @ 2.20 GHz, Nvidia GeForce GTX 1650 Ti GPU	YOLOv8	Navigation line extraction
Liu et al. (2025)	Mobile robot, Intel RealSense cameras, Nvidia Jetson Xavier	NPENet	Navigation path extraction

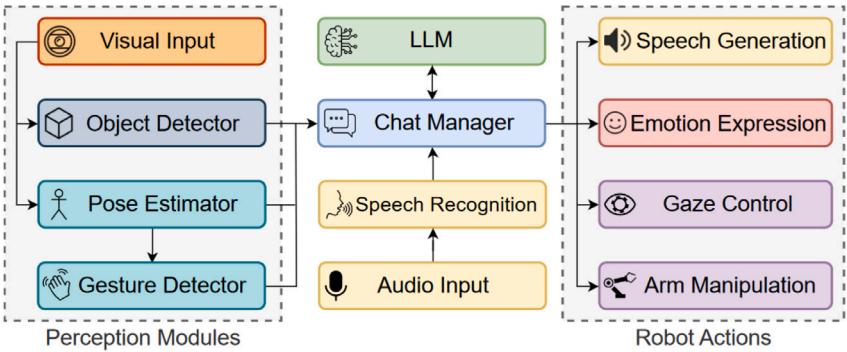


Fig. 33. A grounded chat architecture (Allgeuer et al., 2024).

**Table 16**  
Summary of state-of-the-art human–robot interaction.

Ref.	Hardware platform	Model	Task
Asuzu et al. (2025)	Niryo Ned robot and Intel RealSense D435i camera	LLM: GPT-3.5, Object detection: SAM-CLiP and ViLD	Robot planning
Chen et al. (2024)	N.A.	LLM: GPT-4 turbo	Human collaboration
Allgeuer et al. (2024)	Neuro-Inspired COLlaborator (NICOL) robot	LLM: GPT, Vicuna, Mistral Human detection: YOLOX Pose estimation: HRNet Object detection: ViLD	Conversation and collaboration
Wang et al. (2024)	MStretch SE3 robot, Intel RealSense D435i and D405	Object segmentation: YOLO World	Robot behavior announcement

become more intricate and complex with numerous layers, subnetworks and a large number of parameters, the models are often perceived as a “black box” and difficult to explain in terms of decision-making processes. Therefore, it is crucial for the researchers to focus on methods that provide insights into how a deep learning model performs the prediction and how its decisions are influenced by the input data, making it more transparent and trustworthy. Numerous methods have been proposed for interpreting and explaining the decisions of deep learning models, which can be categorized into visualization (feature attribution), model distillation and intrinsic (explainable by itself). Visualization methods involve the use of scientific visualization such as saliency maps or heatmaps to express the explanation by highlighting the degree of association between the inputs and the predictions (Tjoa et al., 2023). The heatmaps identify the saliency of the input features influencing the model’s predictions. The visualization approach is simple and intuitive and can be applied to tabular data and image data. Furthermore, it can be used to identify and debug issues in deep learning models, leading to improved performance and robustness.

Model distillation is an approach to approximating a complex model by fitting a simpler model using the training set. The simpler model is built typically using a simpler or interpretable algorithm such as linear regression, decision tree or rule-based methods (Li and Shen, 2024). In this approach, the simpler model is trained to resemble the predictive behavior of the complex model. Then, the simpler model may serve as the proxy or surrogate model for explaining the complex model. Model

distillation can be used together with visualization to further enhance the interpretability of the complex model (Termritthikun et al., 2023). Model distillation seeks explanations of the models that were never designed to be explainable. Ideally, the explanation of a deep learning model’s prediction should be included as part of the model output, or the explanation can be derived from the architecture of the model. This is because an intrinsic model can learn not only the mapping between the input and output, but also generate an explanation of the prediction that is faithful to the model’s behavior. Attention mechanisms are the key to this approach, providing a form of attention weights that can be used to explain why the model made a particular decision (Xiong et al., 2022). Another type of intrinsic approach is to train the model to simultaneously perform the prediction task and generate the explanation for its predictions (Fernandes et al., 2023). This “additional task” can be in the form of a text explanation or model prototype which embeds the semantic meaning of the prediction. However, the intrinsic approach is more difficult to apply because the user needs additional knowledge and understanding of the model’s architecture and inner workings.

Deep learning models are increasingly being deployed in making high-stake decision including recruitment (Freire and de Castro, 2021), criminal justice (Dass et al., 2023) and credit scoring (Gicić et al., 2023). There are several advantages of deep learning-based systems in which, unlike humans, machines are able to process vast amounts of data and applications quickly and consistently. However, deep learning-based systems have the risk of being prone to biases

present in the data used for training which can lead to unfairness and injustice. Numerous efforts have been made to mitigate this issue which can be categorized into modeling bias detection and modeling bias mitigation. Detection of modeling bias refers to the process of identifying and quantifying biases that may present in predictive models. This approach involves the use of statistical analysis, fairness metrics, counterfactual testing and human review to detect bias in the models. For instance, visualization-based methods such as attribution maps are used to indicate which regions are significant to the predictions (Schaaf et al., 2021). This in turn can be used to detect and quantify bias using metrics such as Relevance Mass Accuracy, Relevance Rank Accuracy and Area over the perturbation curve (AOPC). In Giloni et al. (2022), two modules are presented for estimating bias in predictive models. The first module utilizes an unsupervised deep neural network with a custom loss function to generate hidden representation of the input data called bias vectors, revealing the underlying bias of each feature. The second module combines these bias vectors into a single vector representing the bias estimation of each feature, achieved by aggregating them using the absolute averaging operation.

Bias mitigation refers to the process of reducing the presence of bias in predictive models, which can be done in three stages. The first stage combats bias by modifying the training data, either relabeling the labels or perturbing the feature values (Iosifidis et al., 2019; Kehrenberg et al., 2020). The second stage addresses bias during the training of the model by applying regularization terms to the loss function to penalize discrimination. In Jain et al. (2023), a loss function based on bias parity score (BPS) is introduced to measure the degree of similarity of a statistical measure such as accuracy across different subgroups. The BPS term is added to the loss function as a regularizer to the original prediction task. The last stage mitigates bias after the predictive models have been successfully trained. This stage applies post-processing approaches such as reinforcement learning to obtain a fairer model (Yang et al., 2023). For instance, the detection of minority classes is rewarded to prevent bias towards the majority class. This allows the model to generalize well across different patient demographics.

## 5.2. Technical issues

Even though deep learning architectures have achieved state-of-the-art across various computer vision tasks, they often come with large model parameters (Raiaan et al., 2024). The architecture and complexity of a deep learning network determine the number of model parameters. The deeper the network, the larger the number of model parameters. However, deep learning models with large parameters often suffer limitations when deploying on end devices. For instance, a deep learning model developed for security monitoring by analyzing video data using 3D-CNNs might suffer deployment issues when deploying such models on low-resourced systems like smartphones or small-scale IoT devices. Model training and inference for deep learning models with large parameters demands substantial processing power. As the number of parameters increases, so does the computational complexity, resulting in longer model training duration and more hardware needs. Furthermore, large parameter sizes translate to increased memory requirements, limiting their deployment on end devices. This is because these end devices often have battery, processor, or memory capacity limitations. To address these challenges, it is important to develop sophisticated but lightweight architectures that can achieve state-of-the-art with few model parameters. Such lightweight models will be characterized by their ability to deliver competitive performance while mitigating computational complexity and memory requirements, making them well-suited for deployment on resource-constrained devices. An approach would be to develop novel lightweight plug-and-play modules that can be plugged to a few layered deep learning architectures to improve feature learning without incurring additional model complexity. Other approaches could involve

leveraging model compression techniques to reduce the size and computational complexity of deep learning models. Researchers can focus on improving pruning methods (Li et al., 2019), which can identify and eliminate redundant parameters or connections, thereby reducing the model's footprint without compromising performance. Furthermore, quantization techniques (Yang et al., 2019) can be further explored to reduce the precision of weights and activations, therefore, enabling efficient representation with lower memory requirements. Also, knowledge distillation techniques (Stanton et al., 2021) can be further investigated to facilitate the transfer of knowledge from a complex teacher model to a simpler student model, therefore, enabling compact yet effective representations.

Deep learning training involves complex processes that require efficient optimization to ensure fast convergence and resource management. As models grow in size and complexity, challenges such as slow convergence, vanishing gradients, and computational limitations become more pronounced. Researchers have explored various methods to overcome these issues, proposing innovative approaches like predicting parameter change and incorporating them into training, thus reducing training time and improving the model performance (Ying et al., 2024). Researchers also leverage prior experience to optimize parameter adjustments to reduce forward and back propagation steps, thus reducing computational costs and making training more efficient (Wang et al., 2024). Others have proposed alternative training algorithms such as alternating minimization (Yan et al., 2025) and random search based on the annealing method (Krasnoprosin and Matskevich, 2024). Overall, the researchers aim to accelerate learning and improve efficiency either by decomposing the optimization problem into sequential sub-problems or by refining the search space based on parameter value ranges.

Adversarial attacks and defense mechanisms in deep learning represent a critical area of research and development, particularly as deep learning models become increasingly integrated into various applications. Adversarial attacks involve the deliberate manipulation of input data to mislead or deceive deep learning models, leading to incorrect predictions or behavior (Akhtar and Mian, 2018). Szegedy et al. (2013) was the first to identify this intriguing shortcoming of deep neural networks in image classification. They showed that even with their great accuracy, deep learning models are surprisingly vulnerable to adversarial attacks that take the form of tiny image changes that are (almost) invisible to human vision systems. A neural network classifier may radically alter its prediction about an image as a result of such an attack. Furthermore, such a model can indicate high confidence in wrong predictions, which can be catastrophic for deep learning models deployed in medical or security fields, among many others. In generative models, several studies have investigated how adversarial attacks affect autoencoders and GANs, as seen in Tabacof et al. (2016) where a method to manipulate input images in a way that deceives variational autoencoders into reconstructing an entirely different image was introduced. In recent times, the focus of adversarial attack research has been on images, but studies have shown that adversarial attacks are not limited to image data; they can also affect other types of data such as text, signals, audio, and video (Zhang et al., 2020; Jiang et al., 2019; Esmailpour et al., 2019).

## 5.3. Domain-specific issues

Building and employing deep learning models face several challenges. The training of deep learning requires a large number of instances (examples) to achieve high accuracy and generalization (Munappy et al., 2022). Furthermore, the complexity of deep neural networks may lead to overfitting, where the model performs well on training data but fails to generalize on new, unseen data. This phenomenon frequently arises when the models are trained on insufficient data, highlighting the importance of diverse and extensive datasets. However, the data collection is time-consuming and often require domain experts, specialized training and standardization (Luca et al.,

2022). Moreover, this process is prone to error and has the risk of introducing biases into the dataset which can significantly impact the performance of the trained model. One of the approaches to address this issue is transfer learning. Transfer learning involves the use of a deep learning model (known as pre-trained model) that is trained on a large dataset for solving a specific task (with a small dataset) (Zhuang et al., 2020). The pre-trained model serves as a basis for the model training by fine-tuning the weights of the pre-trained model and adapting it to the new prediction task. This approach helps to mitigate the lack of training data in the target domain. Furthermore, transfer learning reduces computational resources required to train the model and helps faster convergence.

Another approach that can be employed to address the lack of data is data augmentation. Data augmentation is a convenient method that increases the number of instances by performing transformation functions on the existing instances without changing the labels (Mumuni and Mumuni, 2022). In the domain of computer vision, image transformation such as rotation, translation and cropping. However, it is important to consider the output of the transformation because the resultant may not represent the actual data. For example, flipping or adding noise to a signal may introduce distortion or changing the characteristics (trend, seasonality and cyclic variations) of the signal. Thus, careful consideration must be given to ensure that the generated instances still accurately represent the underlying patterns present in the data. Data augmentation can also be realized by generating synthetic data to supplement the training set. Synthetic data is artificially created data that resembles real data but is generated using statistical methods or deep generative models (Hu et al., 2023; Murtaza et al., 2023). The generated data can complement the less-diverse, limited datasets, providing a broader range of examples for the model to learn from. However, generating synthetic data that accurately reflects the characteristics of the real-world data is challenging. Careful consideration must be given to the choice of models and parameters used to ensure the synthetic data is realistic and representative of the real-world data.

Data annotation is expensive and time-consuming. The problem is exacerbated when the data is of low quality corrupted with noise which may lead to bad, resulting in unreliable training data. One of the approaches addresses this issue is active learning. Active learning is an approach where the learning algorithm selectively queries the most informative data points for labeling. The aim is to improve predictive model efficiency with fewer training data, thus reducing the overall cost of data annotation. The approach involves training a model using an initial training set, and then a subset of unlabeled data is selected for labeling by external annotators. The newly labeled instances are appended to the training set for retraining or fine-tuning. This process is repeated until the model performance reaches a desired threshold or the labeling budget has exhausted. The core principle behind active learning is the query strategy or, essentially how to select the most informative data for labeling that will be beneficial to the model training. In general, the query may select the most ambiguous instances based on the model predictions, instances that will have a significant impact on the model performance, or instances that cover the distribution of the entire feature space (Li et al., 2024). The dynamic and sequential nature of real-world applications presents different challenges, which led to the development of online active learning. This approach was introduced to address specific issues, including data streams, concept drift and environmental changes (Cacciarelli and Kulahci, 2024). Recently, researchers have investigated the use of LLMs for data annotation. The study presents methods for generating annotations, assessing annotations and utilizing annotations (Tan et al., 2024).

Self-supervised learning is another approach that was introduced to mitigate the issue of limited labeled datasets. In self-supervised learning, the model learns from unlabeled data by generating its own labels from the input data (Gui et al., 2024). This approach essentially creates a pretext task which the model solves without requiring manual

annotations or labeled datasets. By solving the pretext task, the model leverages the underlying structure in the data and learn useful representations that can be used for specific tasks. For example, in computer vision, a model might predict the speed (Altawabree and Mohd Noor, 2025) or repeated scenes (Altawabree and Mohd Noor, 2024) in a video. In natural language processing, it could predict missing words in sentences to learn the dependencies between words (Lee et al., 2021). The key advantage of self-supervised learning is that it allows the model to take advantage of large amounts of unlabeled data, which is often much more abundant and cheaper to collect than labeled data. The learned representations can then be used for downstream tasks like classification, regression, or segmentation.

## 6. Summary and future directions

We have discussed the state-of-the-art applications and challenges of deep learning in computer vision, natural language processing and time series analysis. In this section, we summarize the advancements made which can be grouped into model architecture, contextual enhancement, and loss function and optimization. Finally, we discuss possible future works in these domains.

**Model Architecture:** The recent advancements in computer vision focus on vision transformers (ViT) for scalable representation learning (Dosovitskiy et al., 2021). In image classification and image segmentation, ViTs are utilized to capture global context and relationships between different parts of an image, while in object detection, ViTs are used to streamline the detection process by eliminating the traditional components like anchor boxes and non-maximum suppression (Carion et al., 2020). However, ViTs lack the ability to exploit local spatial features and struggle with hyperparameter sensitivity and performance on smaller datasets. To address this issue, researchers have developed hybrid deep learning architectures such as conformer (Peng et al., 2023) and MaxViT (Tu et al., 2022) that combine both transformers and convolutional neural networks to capture both local and global features. Similar trend is observed in modern natural language processing whereby the basis of the deep learning models is transformer architectures such as BERT (Devlin et al., 2018) and its variants to achieve remarkable accuracy in tasks such as text classification, machine translation and text generation.

While recent advancements in computer vision and natural language processing have predominantly relied on transformer architectures, the field of time series analysis focuses on a different trajectory, where transformers research is lagging due to fundamental challenges in data structure, sequence length and dataset size and availability. Notably, areas such as HAR and financial prediction have seen significant progress through the development of hybrid architectures that do not depend on transformers. Instead, these advancements leverage recurrent neural networks and convolutional neural networks, often incorporating attention mechanisms. The hybrid architectures allow the models to capture both local features and temporal dependencies in the data, which is crucial for time series analysis (Khan and Ahmad, 2021; Ige and Mohd Noor, 2023). The hybrid attention mechanisms improve the feature extraction process by dynamically weighting important channel, spatial and temporal features, enhancing model performance (Gao et al., 2021; Agac and Incel, 2024; Tang et al., 2022). GAN models are used to address data scarcity in HAR, where the generator consists of convolutional layers (Lupi3n et al., 2024; Kia et al., 2024), with LSTM (Chan and Noor, 2021) while the discriminator consists of convolutional layers.

Similar trends can be observed in ECG and EEG classification, where convolutional and recurrent neural networks are commonly used to build hybrid models (Alamatsaz et al., 2024; Hermawan et al., 2024), with attention mechanisms (Sun et al., 2024; Chen et al., 2022). Generative models such as variational autoencoder (Dai et al., 2019) and generative adversarial network are used to generate data to address challenges posed by limited datasets (Yang et al., 2024; Msigwa et al.,

2024; Song et al., 2024), poor quality (Corley and Huang, 2025), noisy signals (Cai et al., 2025), and overfitting (Zhu et al., 2019). The generative approach employs GAN only or with encoder–decoder (Yang et al., 2024), where the generator consists of convolutional layers with attention mechanisms or bidirectional LSTM, while the discriminator consists of convolutional layers with attention mechanisms.

**Contextual Enhancement:** Recent advancements in text classification focus on capturing contextual information utilizing approaches such as joint embeddings of labels (Wang et al., 2018) and words as well as aspect-aware methods that enhance feature extraction (Zhu et al., 2023). Conversely, in machine translation, although transformers have proven effective, they often struggle to capture nuanced language features and contexts. To address this issue, various strategies have been introduced such as concatenating contextual sentences and context-sensitive training (Wu et al., 2022; Kim et al., 2023). In text generation, additional contexts are provided to improve both the quality and diversity of the generated sentences such as incorporating the category context (Li et al., 2023) and style context (Kwon et al., 2024) into text generation process.

In image generation, the researchers focus on improving the quality and relevance of generated images by enhancing the alignment between input text and images, while also addressing issues like fuzzy shapes and diversity through innovative deep learning architectures such as incorporating attention mechanisms and fusion modules into the model. Furthermore, recent models like DF-GAN (Tao et al., 2022) and DMF-GAN (Yang et al., 2024) showcase a shift towards single-stage generators with regularization strategies that maintain details while enhancing diversity. In text generation, BERT and its variants are the basis of the deep learning models. The research focuses on hierarchical architectures, capturing long text dependencies (Ma et al., 2021) and leveraging domain knowledge through label sequence (Xie et al., 2022) and conversation sequence (Li et al., 2019) to bridge semantic gaps. Overall, the advancements show a clear trend towards improving contextual awareness and feature representation.

LLMs are increasingly being leveraged in robotics for improving human–robot interaction. The main challenge of integrating LLMs is bridging the gap between robots and the physical world to provide real-world contexts, a process known as LLM grounding. The concept of LLM grounding refers to how machines make sense of abstract language such as words and ideas by connecting them to objects they can actually experience or sense in the real world. To this end, sensors such as cameras and microphones, along with object detection, pose estimation and speech recognition methods, are used for semantic and perceptual grounding. The robot data are streamed at varying sampling rates, so they need to be aligned and carefully selected to identify key events (Wang et al., 2024). The detected objects and their relationships are then modeled using 3D scene graph for scene understanding (Chen et al., 2024). Central to the grounding system is a module that collates and constructs the text prompts which are sent to the LLM (Allgeuer et al., 2024). The text prompts are typically formulated by integrating a set of predefined robot skills that can be executed (Asuzu et al., 2025).

**Loss Function and Optimization:** Researchers are increasingly focused on refining loss functions to improve classification accuracy. For instance, in image classification, some studies introduce additive terms to the cross-entropy loss to reward well-classified instances (Zhao et al., 2022), while others propose asymmetric polynomial loss functions that prioritize positive instances to tackle class imbalance (Huang et al., 2023) and adaptive loss function that dynamically adjust the weights assigned to class-level components based on model performance (Maldonado et al., 2023). Another approach to address class imbalanced is data augmentation through generating virtual samples (Zhu et al., 2024), oversampling with target-aware autoencoders for estimating target values for new features (Belhaouiari et al., 2024). In time series analysis, similar approach has been proposed such as penalizing misclassification of minority classes (Wang et al., 2024), maximizing the minimum recall of the classes (Ircio et al., 2023). Improved training

strategies have been proposed to address class imbalanced data such as iteratively selecting the most informative instances (Moles et al., 2024) and contrastive learning to keep the instances of each fine-grained clusters away from the minority class (Zhu et al., 2024).

**Future Directions:** Looking ahead, the future of deep learning presents numerous opportunities for growth and innovation. Future models could explore non-sequential hybrid architectures of transformer and convolutional neural networks to leverage the strengths of both approaches, enhancing performance in image classification, object detection and image segmentation. Furthermore, researchers could investigate simpler or lightweight architectures and new training schemes to address the long training time associated with transformers. Future research in image generation may explore advanced frameworks or architectures that further integrate semantic understanding, perhaps by employing hierarchical attention modules and/or fusion modules to capture both local and global features more effectively. Additionally, incorporating unsupervised learning or self-supervised learning approaches could reduce reliance on labeled datasets, allowing models to learn from a diverse range of inputs.

In the area of natural language processing, future research in text classification and machine translation could explore the integration of external knowledge bases and domain-specific embeddings may further improve context understanding and label alignment. Additionally, refining context incorporation methods such as dynamic context updating during translation may enhance the model performance. In text generation, advancing towards more interactive and adaptive systems that can maintain context over extended dialogs and narratives is crucial. Furthermore, frameworks that can leverage multi-modalities (text, audio, and visual) for richer contextual understanding may facilitate the development of more advanced applications across these domains. Future research can also focus on exploring adversarial attacks in these domains and developing tailored defense mechanisms. Furthermore, researchers can further investigate the practical implications of adversarial attacks in real-world scenarios, such as in autonomous vehicles and medical imaging. Understanding the potential impact of adversarial attacks in these applications can inform the development of more robust and secure systems.

Real-time processing capabilities are paramount, future work can develop efficient models or algorithms for real-time processing of HAR, speech recognition, ECG and EEG signals. This could involve optimizing existing architectures and leveraging hardware acceleration techniques to enable real-time inference on resource-constrained devices such as wearable sensors and implantable devices. Furthermore, future models could explore lightweight transformers to enhance feature representation for time series analysis. Due to the sensitive nature of financial research, future work can focus on enhancing the interpretability of deep learning models in financial predictions. Researchers should explore techniques to explain the predictions of models, to improve trust and understanding of model decisions, which is essential for adoption in HAR, speech recognition and finance.

Future research in data generation for time series analysis and pervasive computing should focus on developing more advanced techniques to address data scarcity and improve model generalization. In HAR, frameworks that can generate diverse signals for subjects of different ages, genders, as well as physical abilities, may facilitate the development of models that can generalize across different populations and real-world scenarios (Jimale and Mohd Noor, 2023). Moreover, the generated data should capture sensor noise, variations in user behavior, and diverse environmental conditions. In finance, synthetic market data should model complex dependencies, volatility patterns and anomalies to improve trend and risk predictions. For EEG and ECG classification, researchers could explore deep generative models that can create physiologically meaningful synthetic signals while preserving individual variability. Future models could explore methods for adapting or personalizing models to account for inter-subject variability and improve performance on individual subjects. Furthermore, EEG electrodes cover

only a fraction of the brain's surface, resulting in limited coverage of neural activity. Deep learning models could investigate strategies to infer activity from unobserved brain regions or integrate information from multiple modalities to provide more comprehensive coverage. These areas can still be further explored.

Future research in robot object identification and navigation will focus on enhancing real-time perception, generalization across diverse environments, and increasing robustness to occlusions and dynamic obstacles. Furthermore, researchers can explore multimodal learning, integrating vision with LiDAR and tactile sensing to improve scene understanding. Advances in self-supervised learning and few-shot learning could enable robots to identify and navigate complex environments with minimal labeled data. LLMs will play a crucial role in human-robot interactions, facilitating natural language understanding, intent recognition and context-aware decision-making (Zhang et al., 2023). Researchers should optimize LLMs for real-time processing on robots and improve their grounding in physical environments. Additionally, future work can explore a deep learning-based navigation with LLMs to develop more autonomous, assistive and socially intelligent robots.

The future of AI lies beyond deep learning, moving towards a conscious intelligent system that mimics human cognition (Butlin et al., 2023). Achieving this requires a significant advancement in hardware and software architectures, enabling systems to process information with greater efficiency, adaptability and reasoning capabilities. Beyond mere pattern recognition, future of AI will need to integrate symbolic reasoning, causal inference and common sense to truly emulate human thought (Marra et al., 2024; Colelough and Regli, 2025; Bhuyan et al., 2024). Future research will likely focus on hybrid AI models that combine deep learning with symbolic AI, enabling machines to understand abstract concepts, reason about cause and effect and apply knowledge across domains. Quantum computing, with its unique property of entanglement could exponentially accelerate deep learning training and inference (Wu et al., 2025; Klusch et al., 2024). Future models could explore quantum neural networks which could process high dimensional data more effectively and discover patterns that are not possible by classical methods. Quantum-based optimizers could also solve problems that are infeasible for classical graphical processing units or tensor processing units. A more radical approach is neuromorphic computing that uses specialized hardware to mimic the structure and function of the human brain, enabling low-power inference, parallel processing and real-time learning and adaptation (Shrestha et al., 2022; Kudithipudi et al., 2025). This results in autonomous AI systems that can solve complex tasks requiring quick response time and adapt to changing situations. Finally, multi-modal intelligence enables AI to process and integrate multiple types of sensory input, including vision, speech and touch, providing a more holistic view of the world, similar to human perception (Fei et al., 2022).

## 7. Conclusion

Deep learning has become the prominent data-driven approach in various state-of-the-art applications. Its importance lies in its ability to revolutionize many aspects of research and industries and tackle complex problems which were once impossible to overcome. Numerous surveys have been published on deep learning, reviewing the concepts, model architectures and applications. However, the studies do not discuss the emerging trends in the state-of-the-art applications of deep learning and emphasize the important traits and elements in the models. This paper presents a structured and comprehensive survey of deep learning, focusing on the latest trends and advancements in state-of-the-art applications such as computer vision, natural language processing, time series analysis and pervasive computing, and robotics. It explores key elements and traits in modern deep learning models, highlighting their significance in addressing complex challenges across diverse domains. The discussion also covers future research directions

in advancing these fields. Furthermore, this paper presents a comprehensive review of the deep learning fundamentals, which is essential for understanding the core principles behind modern deep learning models. The survey finishes by discussing the critical challenges and future directions in deep learning.

## CRedit authorship contribution statement

**Mohd Halim Mohd Noor:** Writing – review & editing, Writing – original draft. **Ayokunle Olalekan Ige:** Writing – original draft.

## Ethics approval

Not applicable.

## Funding

This work was supported by the Ministry of Higher Education Malaysia (FRGS/1/2023/ICT02/USM/02/2).

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgment

This work has been supported in part by the Ministry of Higher Education Malaysia for Fundamental Research Grant Scheme with Project Code: FRGS/1/2023/ICT02/USM/02/2.

## Data availability

No data was used for the research described in the article.

## References

- Abdulwahhab, A.H., Abdulaal, A.H., Thary Al-Ghrai, A.H., Mohammed, A.A., Valizadeh, M., 2024. Detection of epileptic seizure using EEG signals analysis based on deep learning techniques. *Chaos Solitons Fractals* 181, 114700. <http://dx.doi.org/10.1016/j.chaos.2024.114700>, URL <https://www.sciencedirect.com/science/article/pii/S0960077924002522>.
- Abedin, M.Z., Moon, M.H., Hassan, M.K., Hajek, P., 2021. Deep learning-based exchange rate prediction during the COVID-19 pandemic. *Ann. Oper. Res.* 1–52.
- Acharya, U.R., Oh, S.L., Hagiwara, Y., Tan, J.H., Adam, M., Gertych, A., San Tan, R., 2017. A deep convolutional neural network model to classify heartbeats. *Comput. Biol. Med.* 89, 389–396.
- Agac, S., Incel, O.D., 2024. Resource-efficient, sensor-based human activity recognition with lightweight deep models boosted with attention. *Comput. Electr. Eng.* 117, 109274. <http://dx.doi.org/10.1016/j.compeleceng.2024.109274>, URL <https://www.sciencedirect.com/science/article/pii/S0045790624002027>.
- Aghaomidi, P., Wang, G., 2024. ECG-SleepNet: Deep learning-based comprehensive sleep stage classification using ECG signals. <http://dx.doi.org/10.48550/arXiv.2412.01929>, arXiv:2412.01929 [cs]. URL <http://arxiv.org/abs/2412.01929>.
- Ahnouch, M., Elaachak, L., Ghadi, A., 2023. Model risk in financial derivatives and the transformative impact of deep learning: A systematic review. In: *The Proceedings of the International Conference on Smart City Applications*. Springer, pp. 155–165.
- Akhtar, N., Mian, A., 2018. Threat of adversarial attacks on deep learning in computer vision: A survey. *IEEE Access* 6, 14410–14430. <http://dx.doi.org/10.1109/ACCESS.2018.2807385>.
- Alamatsaz, N., Tabatabaei, L., Yazdchi, M., Payan, H., Alamatsaz, N., Nasimi, F., 2024. A lightweight hybrid CNN-LSTM explainable model for ECG-based arrhythmia detection. *Biomed. Signal Process. Control.* 90, 105884. <http://dx.doi.org/10.1016/j.bspc.2023.105884>, URL <https://www.sciencedirect.com/science/article/pii/S1746809423013174>.
- Allgeuer, P., Ali, H., Wermter, S., 2024. When robots get chatty: Grounding multimodal human-robot conversation and collaboration. pp. 306–321. [http://dx.doi.org/10.1007/978-3-031-72341-4\\_21](http://dx.doi.org/10.1007/978-3-031-72341-4_21), Vol. 15019. arXiv:2407.00518 [cs]. URL <http://arxiv.org/abs/2407.00518>.

- Alnuaim, A.A., Zakariah, M., Shashidhar, C., Hatamleh, W.A., Tarazi, H., Shukla, P.K., Ratna, R., 2022. Speaker gender recognition based on deep neural networks and ResNet50. *Wirel. Commun. Mob. Comput.* 2022, 1–13.
- Alom, M.Z., Taha, T.M., Yakopcic, C., Westberg, S., Sidike, P., Nasrin, M.S., Hasan, M., Van Eszen, B.C., Awwal, A.A.S., Asari, V.K., 2019. A state-of-the-art survey on deep learning theory and architectures. *Electronics* 8 (3), 292. <http://dx.doi.org/10.3390/electronics8030292>, Number: 3 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/2079-9292/8/3/292>.
- Alotaibi, A., Alatawi, H., Binnouh, A., Duwayriat, L., Alhmiedat, T., Alia, O.M., 2024. Deep learning-based vision systems for robot semantic navigation: An experimental study. *Technologies* 12 (9), 157. <http://dx.doi.org/10.3390/technologies12090157>, Number: 9 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/2227-7080/12/9/157>.
2025. AlphaFold. URL <https://deepmind.google/technologies/alphafold/>.
- Alrowili, S., Vijay-Shanker, K., 2021. Biom-transformers: building large biomedical language models with BERT, ALBERT and ELECTRA. In: *Proceedings of the 20th Workshop on Biomedical Language Processing*. pp. 221–227.
- Alrowili, S., Vijay-Shanker, K., 2022. Exploring biomedical question answering with BioM-transformers at BioASQ10B challenge: Findings and techniques. In: *CLEF (Working Notes)*. pp. 222–234.
- Altawabee, H., Mohd Noor, M.H., 2024. Repeat and learn: Self-supervised visual representations learning by Repeated Scene Localization. *Pattern Recognit.* 156, 110804. <http://dx.doi.org/10.1016/j.patcog.2024.110804>, URL <https://www.sciencedirect.com/science/article/pii/S0031320324005557>.
- Altawabee, H., Mohd Noor, M.H., 2025. STCLR: Sparse temporal contrastive learning for video representation. *Neurocomputing* 630, 129694. <http://dx.doi.org/10.1016/j.neucom.2025.129694>, URL <https://www.sciencedirect.com/science/article/pii/S0925231225003662>.
- Altaheri, H., Muhammad, G., Alsulaiman, M., Amin, S.U., Altuwajiri, G.A., Abdul, W., Bencherif, M.A., Faisal, M., 2023. Deep learning techniques for classification of electroencephalogram (EEG) motor imagery (MI) signals: A review. *Neural Comput. Appl.* 35 (20), 14681–14722.
- Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaria, J., Fadhel, M.A., Al-Amidie, M., Farhan, L., 2021. Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *J. Big Data* 8 (1), 53. <http://dx.doi.org/10.1186/s40537-021-00444-8>.
- Ang, K.K., Guan, C., 2017. EEG-based strategies to detect motor imagery for control and rehabilitation. *IEEE Trans. Neural Syst. Rehabil. Eng.* 25 (4), 392–401. <http://dx.doi.org/10.1109/TNSRE.2016.2646763>.
- Araabi, A., Monz, C., 2020. Optimizing transformer for low-resource neural machine translation. In: Scott, D., Bel, N., Zong, C. (Eds.), *Proceedings of the 28th International Conference on Computational Linguistics*. International Committee on Computational Linguistics, pp. 3429–3435. <http://dx.doi.org/10.18653/v1/2020.coling-main.304>, URL <https://aclanthology.org/2020.coling-main.304>.
- Arevian, G., 2007. Recurrent neural networks for robust real-world text classification. In: *IEEE/WIC/ACM International Conference on Web Intelligence*. WI'07, pp. 326–329. <http://dx.doi.org/10.1109/WI.2007.126>, URL <https://ieeexplore.ieee.org/abstract/document/4427112>.
- Asuzu, K., Singh, H., Idrissi, M., 2025. Human–robot interaction through joint robot planning with large language models. *Intell. Serv. Robot.* <http://dx.doi.org/10.1007/s11370-024-00570-1>.
- Ba, J.L., Kiros, J.R., Hinton, G.E., 2016. Layer normalization. URL <https://arxiv.org/abs/1607.06450v1>.
- Badrinarayanan, V., Kendall, A., Cipolla, R., 2017. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 39 (12), 2481–2495, Publisher: IEEE.
- Baevski, A., Zhou, H., Mohamed, A., Auli, M., 2020. Wav2vec 2.0: A framework for self-supervised learning of speech representations. <http://dx.doi.org/10.48550/arXiv.2006.11477>, arXiv:2006.11477 [cs]. URL <http://arxiv.org/abs/2006.11477>.
- Bahdanau, D., Cho, K., Bengio, Y., 2014. Neural machine translation by jointly learning to align and translate. URL <https://arxiv.org/abs/1409.0473v7>.
- Baloglu, U.B., Talo, M., Yildirim, O., San Tan, R., Acharya, U.R., 2019. Classification of myocardial infarction with multi-lead ECG signals and deep CNN. *Pattern Recognit. Lett.* 122, 23–30.
- Banerjee, R., Ghose, A., Muthana Mandana, K., 2020. A hybrid CNN-LSTM architecture for detection of coronary artery disease from ECG. In: *2020 International Joint Conference on Neural Networks*. IJCNN, pp. 1–8. <http://dx.doi.org/10.1109/IJCNN48605.2020.9207044>.
- Banjarey, K., Sahu, S.P., Dewangan, D.K., 2022. Human activity recognition using 1D convolutional neural network. In: *Sentimental Analysis and Deep Learning: Proceedings of ICSADL 2021*. Springer, pp. 691–702.
- Baraka, A.M.A., Mohd Noor, M.H., 2023. Similarity segmentation approach for sensor-based activity recognition. *IEEE Sens. J.* 23 (17), 19704–19716. <http://dx.doi.org/10.1109/JSEN.2023.3295778>, Conference Name: IEEE Sensors Journal. URL <https://ieeexplore.ieee.org/document/10188611>.
- Baraka, A., Mohd Noor, M.H., 2024. Deep similarity segmentation model for sensor-based activity recognition. *Multimedia Tools Appl.* <http://dx.doi.org/10.1007/s11042-024-18933-2>.
- Belhaouari, S.B., Islam, A., Kassoul, K., Al-Fuqaha, A., Bouzardoum, A., 2024. Over-sampling techniques for imbalanced data in regression. *Expert Syst. Appl.* 252, 124118.
- Berti, L., Kasneci, G., 2025. TLOB: A novel transformer model with dual attention for stock price trend prediction with limit order book data. <http://dx.doi.org/10.48550/arXiv.2502.15757>, arXiv:2502.15757 [q-fin]. URL <http://arxiv.org/abs/2502.15757>.
- Bhuyan, B.P., Ramdane-Cherif, A., Tomar, R., Singh, T.P., 2024. Neuro-symbolic artificial intelligence: a survey. *Neural Comput. Appl.* 36 (21), 12809–12844. <http://dx.doi.org/10.1007/s00521-024-09960-z>.
- Bochkovskiy, A., Wang, C.-Y., Liao, H.-Y.M., 2020. YOLOv4: Optimal speed and accuracy of object detection. URL <https://arxiv.org/abs/2004.10934v1>.
- Bodla, N., Singh, B., Chellappa, R., Davis, L.S., 2017. Soft-NMS—improving object detection with one line of code. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5561–5569.
- Boonyakitanont, P., Lek-Uthai, A., Chomtho, K., Songsiri, J., 2020. A review of feature extraction and performance evaluation in epileptic seizure detection using EEG. *Biomed. Signal Process. Control.* 57, 101702.
- Brunner, C., Naeem, M., Leeb, R., Graimann, B., Pfurtscheller, G., 2007. Spatial filtering and selection of optimized components in four class motor imagery EEG data using independent components analysis. *Pattern Recognit. Lett.* 28 (8), 957–964.
- Butlin, P., Long, R., Elmoznino, E., Bengio, Y., Birch, J., Constant, A., Deane, G., Fleming, S.M., Frith, C., Ji, X., Kanai, R., Klein, C., Lindsay, G., Michel, M., Mudrik, L., Peters, M.A.K., Schwitzgebel, E., Simon, J., VanRullen, R., 2023. Consciousness in artificial intelligence: Insights from the science of consciousness. <http://dx.doi.org/10.48550/arXiv.2308.08708>, arXiv:2308.08708 [cs]. URL <http://arxiv.org/abs/2308.08708>.
- Cacciarelli, D., Kulahci, M., 2024. Active learning for data streams: a survey. *Mach. Learn.* 113 (1), 185–239. <http://dx.doi.org/10.1007/s10994-023-06454-2>.
- Cai, S., Feng, X., Deng, Z., Ming, Z., Shan, Z., 2018. Financial news quantization and stock market forecast research based on CNN and LSTM. In: *Smart Computing and Communication: Third International Conference, SmartCom 2018, Tokyo, Japan, December 10–12, 2018, Proceedings 3*. Springer, pp. 366–375.
- Cai, Y., Meng, Z., Huang, D., 2025. DHCT-GAN: Improving EEG signal quality with a dual-branch hybrid CNN-transformer network. *Sensors* 25 (1), 231. <http://dx.doi.org/10.3390/s25010231>, Number: 1 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/25/1/231>.
- Cai, Y., Wang, X., Yu, Z., Li, F., Xu, P., Li, Y., Li, L., 2019. Dualattn-GAN: Text to image synthesis with dual attentional generative adversarial network. *IEEE Access* 7, 183706–183716. <http://dx.doi.org/10.1109/ACCESS.2019.2958864>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/8930532?denied=>.
- Calvo, R.A., Ceccatto, H.A., 2000. Intelligent document classification. *Intell. Data Anal.* 4 (5), 411–420. <http://dx.doi.org/10.3233/IDA-2000-4503>, Publisher: IOS Press. URL <https://content.iospress.com/articles/intelligent-data-analysis/ida00028>.
- Cao, Z., Gong, C., Meng, J., Liu, L., Rao, Y., Hou, W., 2024. Orchard vision navigation line extraction based on YOLOv8-trunk detection. *IEEE Access* 12, 104126–104137. <http://dx.doi.org/10.1109/ACCESS.2024.3422422>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/10583867>.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S., 2020. End-to-end object detection with transformers. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (Eds.), *Computer Vision – ECCV 2020*. In: *Lecture Notes in Computer Science*, Springer International Publishing, pp. 213–229. [http://dx.doi.org/10.1007/978-3-030-58452-8\\_13](http://dx.doi.org/10.1007/978-3-030-58452-8_13).
- Chan, W., Jaitly, N., Le, Q.V., Vinyals, O., 2015. Listen, attend and spell. URL <https://arxiv.org/abs/1508.01211>.
- Chan, M.H., Noor, M.H.M., 2021. A unified generative model using generative adversarial network for activity recognition. *J. Ambient. Intell. Humaniz. Comput.* 12 (7), 8119–8128. <http://dx.doi.org/10.1007/s12652-020-02548-0>.
- Chaurasia, A., Culurciello, E., 2017. Linknet: Exploiting encoder representations for efficient semantic segmentation. In: *2017 IEEE Visual Communications and Image Processing*. VCIP, IEEE, pp. 1–4.
- Chen, P., Boukouvalas, Z., Corizzo, R., 2024. A deep fusion model for stock market prediction with news headlines and time series data. *Neural Comput. Appl.* 36 (34), 21229–21271. <http://dx.doi.org/10.1007/s00521-024-13030-1>.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., Abbeel, P., 2016. InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. <http://dx.doi.org/10.48550/arXiv.1606.03657>, arXiv:1606.03657 [cs, stat] version: 1. URL <http://arxiv.org/abs/1606.03657>.
- Chen, L.-C., Hermans, A., Papandreou, G., Schroff, F., Wang, P., Adam, H., 2018. MaskLab: Instance segmentation by refining object detection with semantic and direction features. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4013–4022.
- Chen, C.-Y., Lin, Y.-T., Lee, S.-J., Tsai, W.-C., Huang, T.-C., Liu, Y.-H., Cheng, M.-C., Dai, C.-Y., 2022. Automated ECG classification based on 1D deep learning network. *Methods* 202, 127–135. <http://dx.doi.org/10.1016/j.jymeth.2021.04.021>, URL <https://www.sciencedirect.com/science/article/pii/S1046202321001134>.
- Chen, L.-C., Papandreou, G., Schroff, F., Adam, H., 2017. Rethinking atrous convolution for semantic image segmentation. <http://dx.doi.org/10.48550/arXiv.1706.05587>, arXiv:1706.05587 [cs] version: 3. URL <http://arxiv.org/abs/1706.05587>.
- Chen, L., Pelger, M., Zhu, J., 2024. Deep learning in asset pricing. *Manag. Sci.* 70 (2), 714–750.

- Chen, M., Xu, Z., Weinberger, K.Q., Sha, F., 2012. Marginalized denoising autoencoders for domain adaptation. In: *Proceedings of the 29th International Conference on International Conference on Machine Learning. ICML '12*, Omnipress, pp. 1627–1634.
- Chen, D., Yongchareon, S., Lai, E.M.-K., Yu, J., Sheng, Q.Z., Li, Y., 2022. Transformer with bidirectional GRU for nonintrusive, sensor-based activity recognition in a multiresident environment. *IEEE Internet Things J.* 9 (23), 23716–23727. <http://dx.doi.org/10.1109/IJOT.2022.3190307>, Conference Name: IEEE Internet of Things Journal. URL <https://ieeexplore.ieee.org/document/9826881>.
- Chen, Y., Zhang, G., Zhang, Y., Xu, H., Zhi, P., Li, Q., Huang, S., 2024. SYNERGAI: Perception alignment for human-robot collaboration. <http://dx.doi.org/10.48550/arXiv.2409.15684>, arXiv:2409.15684 [cs]. URL <http://arxiv.org/abs/2409.15684>.
- Chiu, C.-C., Sainath, T.N., Wu, Y., Prabhavalkar, R., Nguyen, P., Chen, Z., Kannan, A., Weiss, R.J., Rao, K., Gonina, E., Jaitly, N., Li, B., Chorowski, J., Bacchiani, M., 2018. State-of-the-art speech recognition with sequence-to-sequence models. In: 2018 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, pp. 4774–4778. <http://dx.doi.org/10.1109/ICASSP.2018.8462105>.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., Bengio, Y., 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In: Moschitti, A., Pang, B., Daelemans, W. (Eds.), *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing. EMNLP, Association for Computational Linguistics*, pp. 1724–1734. <http://dx.doi.org/10.3115/v1/D14-1179>, URL <https://aclanthology.org/D14-1179>.
- Chougrat, H., Zouaki, H., Alheyane, O., 2018. Deep Convolutional Neural Networks for breast cancer screening. *Comput. Methods Programs Biomed.* 157, 19–30. <http://dx.doi.org/10.1016/j.cmpb.2018.01.011>, URL <https://www.sciencedirect.com/science/article/pii/S0169260717301451>.
- CIFAR-10 and CIFAR-100 datasets. URL <https://www.cs.toronto.edu/~kriz/cifar.html>.
- Cinar, N., Ozcan, A., Kaya, M., 2022. A hybrid DenseNet121-UNet model for brain tumor segmentation from MR Images. *Biomed. Signal Process. Control.* 76, 103647. <http://dx.doi.org/10.1016/j.bspc.2022.103647>, URL <https://www.sciencedirect.com/science/article/pii/S1746809422001690>.
- Clevert, D.-A., Unterthiner, T., Hochreiter, S., 2016. Fast and accurate deep network learning by exponential linear units (ELUs). <http://dx.doi.org/10.48550/arXiv.1511.07289>, arXiv:1511.07289 [cs]. URL <https://arxiv.org/abs/1511.07289>.
- Colelough, B.C., Regli, W., 2025. Neuro-symbolic AI in 2024: A systematic review. <http://dx.doi.org/10.48550/arXiv.2501.05435>, arXiv:2501.05435 [cs] version: 1. URL <http://arxiv.org/abs/2501.05435>.
- Corley, I., Huang, Y., 2025. Deep EEG super-resolution: Upsampling EEG spatial resolution with generative adversarial networks. <http://dx.doi.org/10.48550/arXiv.2502.08803>, arXiv:2502.08803 [cs]. URL <http://arxiv.org/abs/2502.08803>.
- Crippa, P., Curzi, A., Falaschetti, L., Turchetti, C., et al., 2015. Multi-class ECG beat classification based on a Gaussian mixture model of Karhunen-Loève transform. *Int. J. Simul. Syst. Sci. Technol.* 16 (1), 2–1.
- Cybenko, G., 1989. Approximation by superpositions of a sigmoidal function. *Math. Control Signals Systems* 2 (4), 303–314. <http://dx.doi.org/10.1007/BF02551274>.
- Dai, X., Chen, Y., Yang, J., Zhang, P., Yuan, L., Zhang, L., 2021. Dynamic DETR: End-to-end object detection with dynamic attention. pp. 2988–2997, URL [https://openaccess.thecvf.com/content/ICCV2021/html/Dai\\_Dynamic\\_DETR\\_End-to-End\\_Object\\_Detection\\_With\\_Dynamic\\_Attention\\_ICCV\\_2021\\_paper.html](https://openaccess.thecvf.com/content/ICCV2021/html/Dai_Dynamic_DETR_End-to-End_Object_Detection_With_Dynamic_Attention_ICCV_2021_paper.html).
- Dai, Y., Kim, D., Lee, K., 2024. An advanced approach to object detection and tracking in robotics and autonomous vehicles using YOLOv8 and LiDAR data fusion. *Electronics* 13 (12), 2250. <http://dx.doi.org/10.3390/electronics13122250>, Number: 12 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/2079-9292/13/12/2250>.
- Dai, M., Zheng, D., Na, R., Wang, S., Zhang, S., 2019. EEG classification of motor imagery using a novel deep learning framework. *Sensors* 19 (3), 551.
- Dang, L.M., Min, K., Wang, H., Piran, M.J., Lee, C.H., Moon, H., 2020. Sensor-based and vision-based human activity recognition: A comprehensive survey. *Pattern Recognit.* 108, 107561.
- Dass, R.K., Petersen, N., Omori, M., Lave, T.R., Visser, U., 2023. Detecting racial inequalities in criminal justice: towards an equitable deep learning approach for generating and interpreting racial categories using mugshots. *AI Soc.* 38 (2), 897–918, Publisher: Springer.
- Deep, S., Zheng, X., 2019. Hybrid model featuring CNN and LSTM architecture for human activity recognition on smartphone sensor data. In: 2019 20th International Conference on Parallel and Distributed Computing, Applications and Technologies. PDCAT, pp. 259–264. <http://dx.doi.org/10.1109/PDCAT46702.2019.00055>, ISSN: 2640-6721. URL <https://ieeexplore.ieee.org/document/9029136>.
- Delorme, A., Sejnowski, T., Makeig, S., 2007. Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *Neuroimage* 34 (4), 1443–1449.
- Desai, M.P., Caffarena, G., Jevtic, R., Márquez, D.G., Otero, A., 2021. A low-latency, low-power FPGA implementation of ECG signal characterization using hermite polynomials. *Electronics* 10 (19), 2324.
- Devlin, J., Chang, M.-W., Lee, K., Toutanova, K., 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. URL <https://arxiv.org/abs/1810.04805>.
- Dixit, P., Silakari, S., 2021. Deep learning algorithms for cybersecurity applications: A technological and status review. *Comput. Sci. Rev.* 39, 100317. <http://dx.doi.org/10.1016/j.cosrev.2020.100317>, URL <https://www.sciencedirect.com/science/article/pii/S1574013720304172>.
- Dong, S., Wang, P., Abbas, K., 2021. A survey on deep learning and its applications. *Comput. Sci. Rev.* 40, 100379. <http://dx.doi.org/10.1016/j.cosrev.2021.100379>, URL <https://www.sciencedirect.com/science/article/pii/S1574013721000198>.
- Dong, L., Wang, C., Yang, G., Huang, Z., Zhang, Z., Li, C., 2023. An improved ResNet-1d with channel attention for tool wear monitor in smart manufacturing. *Sensors* 23 (3), 1240. <http://dx.doi.org/10.3390/s23031240>, Number: 3 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/23/3/1240>.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N., 2021. An image is worth 16x16 words: Transformers for image recognition at scale. <http://dx.doi.org/10.48550/arXiv.2010.11929>, arXiv:2010.11929 [cs]. URL <http://arxiv.org/abs/2010.11929>.
- Dua, N., Singh, S.N., Semwal, V.B., Challa, S.K., 2023. Inception inspired CNN-GRU hybrid network for human activity recognition. *Multimedia Tools Appl.* 82 (4), 5369–5403. <http://dx.doi.org/10.1007/s11042-021-11885-x>.
- Duchi, J., Hazan, E., Singer, Y., 2011. Adaptive subgradient methods for online learning and stochastic optimization. *J. Mach. Learn. Res.* 12 (61), 2121–2159, URL <http://jmlr.org/papers/v12/duchi11a.html>.
- Dugas, C., Bengio, Y., Bélisle, F., Nadeau, C., Garcia, R., 2000. Incorporating second-order functional knowledge for better option pricing. In: *Advances in Neural Information Processing Systems*. Vol. 13, MIT Press, URL [https://papers.nips.cc/paper\\_files/paper/2000/hash/44968aace94f667e4095002d140b5896-Abstract.html](https://papers.nips.cc/paper_files/paper/2000/hash/44968aace94f667e4095002d140b5896-Abstract.html).
- Eapen, J., Bein, D., Verma, A., 2019. Novel deep learning model with CNN and bi-directional LSTM for improved stock market index prediction. In: 2019 IEEE 9th Annual Computing and Communication Workshop and Conference. CCWC, IEEE, pp. 0264–0270.
- Erdas, C.B., Guney, S., 2021. Human activity recognition by using different deep learning approaches for wearable sensors. *Neural Process. Lett.* 53, 1795–1809.
- Esmailpour, M., Cardinal, P., Koerich, A.L., 2019. A robust approach for securing audio classification against adversarial attacks. *IEEE Trans. Inf. Forensics Secur.* 15, 2147–2159.
- Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., Cui, C., Corrado, G., Thrun, S., Dean, J., 2019. A guide to deep learning in healthcare. *Nature Med.* 25 (1), 24–29.
- Faheem, M.A., Wassif, K.T., Bayomi, H., Abdou, S.M., 2024. Improving neural machine translation for low resource languages through non-parallel corpora: a case study of Egyptian dialect to modern standard Arabic translation. *Sci. Rep.* 14 (1), 2265. <http://dx.doi.org/10.1038/s41598-023-51090-4>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-023-51090-4>.
- Fei, N., Lu, Z., Gao, Y., Yang, G., Huo, Y., Wen, J., Lu, H., Song, R., Gao, X., Xiang, T., Sun, H., Wen, J.-R., 2022. Towards artificial general intelligence via a multimodal foundation model. *Nat. Commun.* 13 (1), 3094. <http://dx.doi.org/10.1038/s41467-022-30761-2>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41467-022-30761-2>.
- Fernandes, L., Fernandes, J.N., Calado, M., Pinto, J.R., Cerqueira, R., Cardoso, J.S., 2023. Intrinsic explainability for end-to-end object detection. *IEEE Access* Publisher: IEEE.
- Freire, M.N., de Castro, L.N., 2021. E-Recruitment recommender systems: a systematic review. *Knowl. Inf. Syst.* 63, 1–20, Publisher: Springer.
- Gao, Z., Dang, W., Wang, X., Hong, X., Hou, L., Ma, K., Perc, M., 2021. Complex networks and deep learning for EEG signal analysis. *Cogn. Neurodynamics* 15 (3), 369–388.
- Gao, Z., Xie, J., Wang, Q., Li, P., 2018. Global second-order pooling convolutional networks. <http://dx.doi.org/10.48550/arXiv.1811.12006>, arXiv:1811.12006 [cs] version: 2. URL <http://arxiv.org/abs/1811.12006>.
- Gao, W., Zhang, L., Teng, Q., He, J., Wu, H., 2021. DanHAR: Dual attention network for multimodal human activity recognition using wearable sensors. *Appl. Soft Comput.* 111, 107728. <http://dx.doi.org/10.1016/j.asoc.2021.107728>, URL <https://www.sciencedirect.com/science/article/pii/S1568494621006499>.
- Ge, L., Sadhu, A., 2025. Deep learning-enhanced smart ground robotic system for automated structural damage inspection and mapping. *Autom. Constr.* 170, 105951. <http://dx.doi.org/10.1016/j.autcon.2024.105951>, URL <https://www.sciencedirect.com/science/article/pii/S0926580524006873>.
- Gezmu, A.M., Nürnberger, A., 2022. Transformers for low-resource neural machine translation. In: *ICAART* (1), pp. 459–466.
- Gicić, A., Donko, D., Subasi, A., 2023. Intelligent credit scoring using deep learning methods. *Concurr. Comput.: Pr. Exp.* 35 (9), e7637.
- Giloni, A., Grolman, E., Hagemann, T., Fromm, R., Fischer, S., Elovici, Y., Shabtai, A., 2022. BENN: Bias estimation using a deep neural network. *IEEE Trans. Neural Netw. Learn. Syst.* Publisher: IEEE.
- Girshick, R., 2015. Fast R-CNN. In: 2015 IEEE International Conference on Computer Vision. ICCV, pp. 1440–1448. <http://dx.doi.org/10.1109/ICCV.2015.169>, ISSN: 2380-7504. URL <https://ieeexplore.ieee.org/document/7410526>.

- Girshick, R., Donahue, J., Darrell, T., Malik, J., 2014. Rich feature hierarchies for accurate object detection and semantic segmentation. pp. 580–587, URL [https://openaccess.thecvf.com/content\\_cvpr\\_2014/html/Girshick\\_Rich\\_Feature\\_Hierarchies\\_2014\\_CVPR\\_paper.html](https://openaccess.thecvf.com/content_cvpr_2014/html/Girshick_Rich_Feature_Hierarchies_2014_CVPR_paper.html).
- Glorot, X., Bordes, A., Bengio, Y., 2011. Deep sparse rectifier neural networks. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics. JMLR Workshop and Conference Proceedings, pp. 315–323, ISSN: 1938-7228. URL <https://proceedings.mlr.press/v15/glorot11a.html>.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y., 2014. Generative adversarial nets. *Adv. Neural Inf. Process. Syst.* 27.
- Graves, A., 2012. Sequence transduction with recurrent neural networks. URL <https://arxiv.org/abs/1211.3711>.
- Gudelek, M.U., Boluk, S.A., Ozbayoglu, A.M., 2017. A deep learning based stock trading model with 2-D CNN trend detection. In: 2017 IEEE Symposium Series on Computational Intelligence. SSCI, IEEE, pp. 1–8.
- Gui, J., Chen, T., Zhang, J., Cao, Q., Sun, Z., Luo, H., Tao, D., 2024. A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Trans. Pattern Anal. Mach. Intell.* 46 (12), 9052–9071. <http://dx.doi.org/10.1109/TPAMI.2024.3415112>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/10559458>.
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., Han, W., Wang, S., Zhang, Z., Wu, Y., Pang, R., 2020. Conformer: Convolution-augmented transformer for speech recognition. <http://dx.doi.org/10.48550/arXiv.2005.08100>, arXiv:2005.08100 [eess]. URL <http://arxiv.org/abs/2005.08100>.
- Guo, J., AI Notes: Initializing neural networks. URL <https://www.deeplearning.ai/ai-notes/initialization/>.
- Guo, M.-H., Liu, Z.-N., Mu, T.-J., Hu, S.-M., 2023. Beyond self-attention: External attention using two linear layers for visual tasks. *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (5), 5436–5447. <http://dx.doi.org/10.1109/TPAMI.2022.3211006>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/9912362>.
- Gupta, S., 2021. Deep learning based human activity recognition (HAR) using wearable sensor data. *Int. J. Inf. Manag. Data Insights* 1 (2), 100046.
- Hajipour, O., Nickabadi, A., Homayounpour, M.M., 2025. GPTGAN: Utilizing the GPT language model and GAN to enhance adversarial text generation. *Neurocomputing* 617, 128865. <http://dx.doi.org/10.1016/j.neucom.2024.128865>, URL <https://www.sciencedirect.com/science/article/pii/S0925231224016369>.
- Han, C., Zhang, L., Tang, Y., Huang, W., Min, F., He, J., 2022. Human activity recognition using wearable sensors by heterogeneous convolutional neural networks. *Expert Syst. Appl.* 198, 116764.
- Hao, S., Zhang, P., Liu, S., Wang, Y., 2023. Sentiment recognition and analysis method of official document text based on BERT-SVM model. *Neural Comput. Appl.* 35 (35), 24621–24632. <http://dx.doi.org/10.1007/s00521-023-08226-4>.
- Hasani, R., Lechner, M., Amini, A., Rus, D., Grosu, R., 2021. Liquid time-constant networks. *Proc. AAAI Conf. Artif. Intell.* 35 (9), 7657–7666. <http://dx.doi.org/10.1609/aaai.v35i9.16936>, Number: 9. URL <https://ojs.aaai.org/index.php/AAAI/article/view/16936>.
- Hassouneh, A., Mutawa, A.M., Murugappan, M., 2020. Development of a real-time emotion recognition system using facial expressions and EEG based on machine learning and deep neural network methods. *Inform. Med. Unlocked* 20, 100372. <http://dx.doi.org/10.1016/j.imu.2020.100372>, URL <https://www.sciencedirect.com/science/article/pii/S235291482030201X>.
- Hatamizadeh, A., Yin, H., Heinrich, G., Kautz, J., Molchanov, P., 2023. Global context vision transformers. In: *International Conference on Machine Learning*. PMLR, pp. 12633–12646.
- He, K., Gkioxari, G., Dollár, P., Girshick, R., 2017. Mask R-CNN. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 2961–2969.
- He, P., Liu, X., Gao, J., Chen, W., 2020. DeBERTa: Decoding-enhanced bert with disentangled attention. URL <https://arxiv.org/abs/2006.03654>.
- He, Y., Nazir, S., Nie, B., Khan, S., Zhang, J., 2020. Developing an efficient deep learning-based trusted model for pervasive computing using an LSTM-based classification model. *Complexity* 2020 (1), 4579495. <http://dx.doi.org/10.1155/2020/4579495>, eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1155/2020/4579495>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1155/2020/4579495>.
- He, K., Zhang, X., Ren, S., Sun, J., 2015. Deep residual learning for image recognition. <http://dx.doi.org/10.48550/arXiv.1512.03385>, arXiv:1512.03385 [cs]. URL <http://arxiv.org/abs/1512.03385>.
- Hema, C., Garcia Marquez, F.P., 2023. Emotional speech Recognition using CNN and Deep learning techniques. *Appl. Acoust.* 211, 109492. <http://dx.doi.org/10.1016/j.apacoust.2023.109492>, URL <https://www.sciencedirect.com/science/article/pii/S0003682X23002906>.
- Hermawan, A.T., Zaeni, I.A.E., Wibawa, A.P., Gunawan, G., Hendrawan, W.H., Kristian, Y., 2024. A multi representation deep learning approach for epileptic seizure detection. *J. Robot. Control. (JRC)* 5 (1), 187–204. <http://dx.doi.org/10.18196/jrc.v5i1.20870>, Number: 1. URL <https://journal.ums.ac.id/index.php/jrc/article/view/20870>.
- Hernández-Blanco, A., Herrera-Flores, B., Tomás, D., Navarro-Colorado, B., et al., 2019. A systematic review of deep learning approaches to educational data mining. *Complexity* 2019.
- Hinton, G.E., 2009. Deep belief networks. *Scholarpedia* 4 (5), 5947. <http://dx.doi.org/10.4249/scholarpedia.5947>, URL [http://www.scholarpedia.org/article/Deep\\_belief\\_networks](http://www.scholarpedia.org/article/Deep_belief_networks).
- Hinton, G.E., 2012. A practical guide to training restricted Boltzmann machines. In: Montavon, G., Orr, G.B., Müller, K.-R. (Eds.), *Neural Networks: Tricks of the Trade: Second Edition*. In: *Lecture Notes in Computer Science*, Springer, pp. 599–619. [http://dx.doi.org/10.1007/978-3-642-35289-8\\_32](http://dx.doi.org/10.1007/978-3-642-35289-8_32).
- Hinton, G.E., Osindero, S., Teh, Y.-W., 2006. A fast learning algorithm for deep belief nets. *Neural Comput.* 18 (7), 1527–1554. <http://dx.doi.org/10.1162/neco.2006.18.7.1527>.
- Hochreiter, S., Bengio, Y., Frasconi, P., Schmidhuber, J., et al., 2001. Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. In: *A Field Guide to Dynamical Recurrent Networks. A field guide to dynamical recurrent neural networks*. IEEE Press Inc.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural Comput.* 9 (8), 1735–1780. <http://dx.doi.org/10.1162/neco.1997.9.8.1735>.
- Hornik, K., Stinchcombe, M., White, H., 1989. Multilayer feedforward networks are universal approximators. *Neural Netw.* 2 (5), 359–366. [http://dx.doi.org/10.1016/0893-6080\(89\)90020-8](http://dx.doi.org/10.1016/0893-6080(89)90020-8), URL <https://www.sciencedirect.com/science/article/pii/0893608089900208>.
- Hou, Y., Wu, Z., Cai, X., Zhu, T., 2024. The application of improved densenet algorithm in accurate image recognition. *Sci. Rep.* 14 (1), 8645. <http://dx.doi.org/10.1038/s41598-024-58421-z>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-024-58421-z>.
- Hu, J., Shen, L., Albanie, S., Sun, G., Wu, E., 2019. Squeeze-and-excitation networks. <http://dx.doi.org/10.48550/arXiv.1709.01507>, arXiv:1709.01507 [cs] version: 4. URL <http://arxiv.org/abs/1709.01507>.
- Hu, C., Sun, Z., Li, C., Zhang, Y., Xing, C., 2023. Survey of time series data generation in IoT. *Sensors* 23 (15), 6976. <http://dx.doi.org/10.3390/s23156976>, Number: 15 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/23/15/6976>.
- Huang, J., Feng, Y., 2020. Optimization of recurrent neural networks on natural language processing. In: *Proceedings of the 2019 8th International Conference on Computing and Pattern Recognition. ICCPR '19, Association for Computing Machinery*, pp. 39–45. <http://dx.doi.org/10.1145/3373509.3373573>, URL <https://dl.acm.org/doi/10.1145/3373509.3373573>.
- Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q., 2017. Densely connected convolutional networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4700–4708.
- Huang, Y., Liu, W., Yin, Z., Hu, S., Wang, M., Cai, W., 2024. ECG classification based on guided attention mechanism. *Comput. Methods Programs Biomed.* 257, 108454. <http://dx.doi.org/10.1016/j.cmpb.2024.108454>, URL <https://www.sciencedirect.com/science/article/pii/S0169260724004474>.
- Huang, L., Lu, K., Song, G., Wang, L., Liu, S., Liu, Y., Li, H., 2023. Teach-DETR: Better training DETR with teachers. *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (12), 15759–15771. <http://dx.doi.org/10.1109/TPAMI.2023.3319387>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/10264211>.
- Huang, Y., Qi, J., Wang, X., Lin, Z., 2023. Asymmetric polynomial loss for multi-label classification. In: *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing. ICASSP, IEEE*, pp. 1–5.
- Huang, L., Yu, W., Ma, W., Zhong, W., Feng, Z., Wang, H., Chen, Q., Peng, W., Feng, X., Qin, B., Liu, T., 2025. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *ACM Trans. Inf. Syst.* 43 (2), 42:1–42:55. <http://dx.doi.org/10.1145/3703155>.
- Husham Al-Badri, A., Azman Ismail, N., Al-Dulaimi, K., Ahmed Salman, G., Sah Hj Salam, M., 2023. Adaptive Non-Maximum Suppression for improving performance of *Rumex* detection. *Expert Syst. Appl.* 219, 119634. <http://dx.doi.org/10.1016/j.eswa.2023.119634>, URL <https://www.sciencedirect.com/science/article/pii/S0957417423001355>.
- Hwang, J., Park, S., Chi, J., 2023. Improving multi-class motor imagery EEG classification using overlapping sliding window and deep learning model. *Electronics* 12 (5), 1186. <http://dx.doi.org/10.3390/electronics12051186>, Number: 5 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/2079-9292/12/5/1186>.
- Ige, A.O., Mohd Noor, M.H., 2022. A survey on unsupervised learning for wearable sensor-based activity recognition. *Appl. Soft Comput.* 127, 109363. <http://dx.doi.org/10.1016/j.asoc.2022.109363>, URL <https://www.sciencedirect.com/science/article/pii/S1568494622005191>.
- Ige, A.O., Mohd Noor, M.H., 2023. A deep local-temporal architecture with attention for lightweight human activity recognition. *Appl. Soft Comput.* 149, 110954. <http://dx.doi.org/10.1016/j.asoc.2023.110954>, URL <https://www.sciencedirect.com/science/article/pii/S1568494623009729>.
- Ige, A.O., Noor, M.H.M., 2023. WSense: A robust feature learning module for lightweight human activity recognition. URL <https://arxiv.org/abs/2303.17845>.
- ImageNet. URL <https://www.image-net.org/>.
- Imran, H.A., Riaz, Q., Hussain, M., Tahir, H., Arshad, R., 2024. Smart-wearable sensors and CNN-BiGRU model: A powerful combination for human activity recognition. *IEEE Sens. J.* 24 (2), 1963–1974. <http://dx.doi.org/10.1109/JSEN.2023.3338264>, Conference Name: IEEE Sensors Journal. URL <https://ieeexplore.ieee.org/document/10348511>.

- Ioffe, S., Szegedy, C., 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. In: Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37. ICML '15, JMLR.org, pp. 448–456.
- Iosifidis, V., Tran, T.N.H., Ntoutsi, E., 2019. Fairness-enhancing interventions in stream classification. In: Database and Expert Systems Applications: 30th International Conference, DEXA 2019, Linz, Austria, August 26–29, 2019, Proceedings, Part I 30. Springer, pp. 261–276.
- Ircio, J., Lojo, A., Mori, U., Malinowski, S., Lozano, J.A., 2023. Minimum recall-based loss function for imbalanced time series classification. *IEEE Trans. Knowl. Data Eng.* 35 (10), 10024–10034.
- Jafarifarmand, A., Badamchizadeh, M.A., 2019. EEG artifacts handling in a real practical brain-computer interface controlled vehicle. *IEEE Trans. Neural Syst. Rehabil. Eng.* 27 (6), 1200–1208.
- Jain, B., Huber, M., Elmasri, R., 2023. Increasing fairness in predictions using bias parity score based loss function regularization. *Int. FLAIRS Conf. Proc.* 36, <http://dx.doi.org/10.32473/flairs.36.133311>, URL <https://journals.flvc.org/FLAIRS/article/view/133311>.
- Jaitly, N., Le, Q.V., Vinyals, O., Sutskever, I., Sussillo, D., Bengio, S., 2016. An online sequence-to-sequence model using partial conditioning. *Adv. Neural Inf. Process. Syst.* 29.
- Jia, D., Yuan, Y., He, H., Wu, X., Yu, H., Lin, W., Sun, L., Zhang, C., Hu, H., 2023. DETRs with hybrid matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19702–19712.
- Jiang, L., Ma, X., Chen, S., Bailey, J., Jiang, Y.-G., 2019. Black-box adversarial attacks on video recognition models. In: Proceedings of the 27th ACM International Conference on Multimedia. pp. 864–872.
- Jiang, B., Zeng, W., Yang, C., Wang, R., Zhang, B., 2024. DE-GAN: Text-to-image synthesis with dual and efficient fusion model. *Multimedia Tools Appl.* 83 (8), 23839–23852. <http://dx.doi.org/10.1007/s11042-023-16377-8>.
- Jiang, H., Zhang, X., Xiang, S., 2024. Non-maximum suppression guided label assignment for object detection in crowd scenes. *IEEE Trans. Multimed.* 26, 2207–2218. <http://dx.doi.org/10.1109/TMM.2023.3293333>, Conference Name: IEEE Transactions on Multimedia. URL <https://ieeexplore.ieee.org/document/10175568>.
- Jiao, Y., Tu, M., Berisha, V., Liss, J.M., 2016. Accent identification by combining deep neural networks and recurrent neural networks trained on long and short term features. In: Interspeech. pp. 2388–2392.
- Jimale, A.O., Mohd Noor, M.H., 2022. Fully connected generative adversarial network for human activity recognition. *IEEE Access* 10, 100257–100266. <http://dx.doi.org/10.1109/ACCESS.2022.3206952>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/9893100>.
- Jimale, A.O., Mohd Noor, M.H., 2023. Subject variability in sensor-based activity recognition. *J. Ambient. Intell. Humaniz. Comput.* 14 (4), 3261–3274. <http://dx.doi.org/10.1007/s12652-021-03465-6>.
- Karras, T., Aila, T., Laine, S., Lehtinen, J., 2018. Progressive growing of GANs for improved quality, stability, and variation. <http://dx.doi.org/10.48550/arXiv.1710.10196>, arXiv:1710.10196 [cs, stat] version: 3. URL <http://arxiv.org/abs/1710.10196>.
- Karras, T., Laine, S., Aila, T., 2019. A style-based generator architecture for generative adversarial networks. <http://dx.doi.org/10.48550/arXiv.1812.04948>, arXiv:1812.04948 [cs, stat] version: 3. URL <http://arxiv.org/abs/1812.04948>.
- Kehrenberg, T., Chen, Z., Quadrianto, N., 2020. Tuning fairness by balancing target labels. *Front. Artif. Intell.* 3, 33, Publisher: Frontiers Media SA.
- Khalil, R.A., Jones, E., Babar, M.I., Jan, T., Zafar, M.H., Alhussain, T., 2019. Speech emotion recognition using deep learning techniques: A review. *IEEE Access* 7, 117327–117345. <http://dx.doi.org/10.1109/ACCESS.2019.2936124>.
- Khan, Z.N., Ahmad, J., 2021. Attention induced multi-head convolutional neural network for human activity recognition. *Appl. Soft Comput.* 110, 107671. <http://dx.doi.org/10.1016/j.asoc.2021.107671>, URL <https://www.sciencedirect.com/science/article/pii/S1568494621005925>.
- Kia, Z., Yadollahzaeh-Tabari, M., Motameni, H., 2024. Human activity recognition by body-worn sensor data using bi-directional generative adversarial networks and frequency analysis techniques. *J. Supercomput.* 81 (1), 342. <http://dx.doi.org/10.1007/s11227-024-06743-0>.
- Kim, D., Baek, Y., Yang, S., Choo, J., 2023. Towards formality-aware neural machine translation by leveraging context information. In: Findings of the Association for Computational Linguistics: EMNLP 2023. pp. 7384–7392.
- Kingma, D.P., Ba, J., 2017. Adam: A method for stochastic optimization. <http://dx.doi.org/10.48550/arXiv.1412.6980>, arXiv:1412.6980 [cs]. URL <http://arxiv.org/abs/1412.6980>.
- Kingma, D.P., Welling, M., 2022. Auto-encoding variational Bayes. <http://dx.doi.org/10.48550/arXiv.1312.6114>, arXiv:1312.6114 [stat]. URL <http://arxiv.org/abs/1312.6114>.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.-Y., Dollár, P., Girshick, R., 2023. Segment anything. <http://dx.doi.org/10.48550/arXiv.2304.02643>, arXiv:2304.02643 [cs]. URL <http://arxiv.org/abs/2304.02643>.
- Klusch, M., Lässig, J., Müssig, D., Macaluso, A., Wilhelm, F.K., 2024. Quantum artificial intelligence: A brief survey. *KI - Künstliche Intell.* 38 (4), 257–276. <http://dx.doi.org/10.1007/s13218-024-00871-8>.
- Krasnoprosin, V.V., Matskevich, V.V., 2024. Random search in neural networks training. *Pattern Recognit. Image Anal.* 34 (2), 309–316. <http://dx.doi.org/10.1134/S105466182470010X>.
- Krizhevsky, A., Sutskever, I., Hinton, G.E., 2012. ImageNet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems. Vol. 25, Curran Associates, Inc., URL [https://proceedings.neurips.cc/paper\\_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html).
- Kudithipudi, D., Schuman, C., Vineyard, C.M., Pandit, T., Merkel, C., Kubendran, R., Aimone, J.B., Orchard, G., Mayr, C., Benosman, R., Hays, J., Young, C., Bartolozzi, C., Majumdar, A., Cardwell, S.G., Payvand, M., Buckley, S., Kulkarni, S., Gonzalez, H.A., Cauwenberghs, G., Thakur, C.S., Subramoney, A., Furber, S., 2025. Neuromorphic computing at scale. *Nature* 637 (8047), 801–812. <http://dx.doi.org/10.1038/s41586-024-08253-8>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41586-024-08253-8>.
- Kumar, S., Sharma, A., Mamun, K., Tsunoda, T., 2016. A deep learning approach for motor imagery EEG signal classification. In: 2016 3rd Asia-Pacific World Congress on Computer Science and Engineering. APWC on CSE, IEEE, pp. 34–39.
- Kusuma, S., Jothi, K.R., 2022. ECG signals-based automated diagnosis of congestive heart failure using Deep CNN and LSTM architecture. *Biocybern. Biomed. Eng.* 42 (1), 247–257. <http://dx.doi.org/10.1016/j.bbe.2022.02.003>, URL <https://www.sciencedirect.com/science/article/pii/S0208521622000067>.
- Kwon, N., Yoo, Y., Lee, B., 2024. Class conditioned text generation with style attention mechanism for embracing diversity. *Appl. Soft Comput.* 163, 111893. <http://dx.doi.org/10.1016/j.asoc.2024.111893>, URL <https://www.sciencedirect.com/science/article/pii/S1568494624006677>.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., Soricut, R., 2019. ALBERT: A lite bert for self-supervised learning of language representations. URL <https://arxiv.org/abs/1909.11942>.
- Lattanzi, E., Calisti, L., Contoli, C., 2025. Are transformers a useful tool for tiny devices in human activity recognition? In: Proceedings of the 2024 8th International Conference on Advances in Artificial Intelligence. ICAAI '24, Association for Computing Machinery, New York, NY, USA, pp. 339–344. <http://dx.doi.org/10.1145/3704137.3704171>, URL <https://dl.acm.org/doi/10.1145/3704137.3704171>.
- Le, H.T.N., Ngo, H.Q.T., 2025. Application of the vision-based deep learning technique for waste classification using the robotic manipulation system. *Int. J. Cogn. Comput. Eng.* 6, 391–400. <http://dx.doi.org/10.1016/j.ijcce.2025.02.005>, URL <https://www.sciencedirect.com/science/article/pii/S2666307425000154>.
- LeCun, Y., Bengio, Y., Hinton, G., 2015. Deep learning. *Nature* 521 (7553), 436–444. <http://dx.doi.org/10.1038/nature14539>, Number: 7553 Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/nature14539>.
- Lecun, Y., Bottou, L., Bengio, Y., Haffner, P., 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86 (11), 2278–2324. <http://dx.doi.org/10.1109/5.726791>, Conference Name: Proceedings of the IEEE. URL <https://ieeexplore.ieee.org/document/726791>.
- LeCun, Y.A., Bottou, L., Orr, G.B., Müller, K.-R., 2012. Efficient BackProp. In: Montavon, G., Orr, G.B., Müller, K.-R. (Eds.), *Neural Networks: Tricks of the Trade: Second Edition*. Springer Berlin Heidelberg, pp. 9–48. [http://dx.doi.org/10.1007/978-3-642-35289-8\\_3](http://dx.doi.org/10.1007/978-3-642-35289-8_3).
- Lee, J.D., Lei, Q., Saunshi, N., Zhuo, J., 2021. Predicting what you already know helps: provable self-supervised learning. In: Proceedings of the 35th International Conference on Neural Information Processing Systems. NIPS '21, Curran Associates Inc., Red Hook, NY, USA, pp. 309–323.
- Lei, Y., Peng, Q., Shen, Y., 2020. Deep learning for algorithmic trading: enhancing MACD strategy. In: Proceedings of the 2020 6th International Conference on Computing and Artificial Intelligence. pp. 51–57.
- Li, Z., Cai, J., He, S., Zhao, H., 2018. Seq2seq dependency parsing. In: Proceedings of the 27th International Conference on Computational Linguistics. pp. 3203–3214.
- Li, F., He, F., Wang, F., Zhang, D., Xia, Y., Li, X., 2020. A novel simplified convolutional neural network classification algorithm of motor imagery EEG signals based on deep learning. *Appl. Sci.* 10 (5), 1605.
- Li, X., Mao, K., Lin, F., Feng, Z., 2023. Feature-aware conditional GAN for category text generation. *Neurocomputing* 547, 126352. <http://dx.doi.org/10.1016/j.neucom.2023.126352>, URL <https://www.sciencedirect.com/science/article/pii/S09252312230004757>.
- Li, J., Monroe, W., Ritter, A., Jurafsky, D., Galley, M., Gao, J., 2016. Deep reinforcement learning for dialogue generation. In: Su, J., Duh, K., Carreras, X. (Eds.), *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pp. 1192–1202. <http://dx.doi.org/10.18653/v1/D16-1127>, URL <https://aclanthology.org/D16-1127/>.
- Li, Z., Niu, C., Meng, F., Feng, Y., Li, Q., Zhou, J., 2019. Incremental transformer with deliberation decoder for document grounded conversations. In: Korhonen, A., Traum, D., Màrquez, L. (Eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, pp. 12–21. <http://dx.doi.org/10.18653/v1/P19-1002>, URL <https://aclanthology.org/P19-1002/>.
- Li, P., Pei, Y., Li, J., 2023. A comprehensive survey on design and application of autoencoder in deep learning. *Appl. Soft Comput.* 138, 110176. <http://dx.doi.org/10.1016/j.asoc.2023.110176>, URL <https://www.sciencedirect.com/science/article/pii/S1568494623001941>.

- Li, W., Qi, F., Tang, M., Yu, Z., 2020. Bidirectional LSTM with self-attention mechanism and multi-channel features for sentiment classification. *Neurocomputing* 387, 63–77. <http://dx.doi.org/10.1016/j.neucom.2020.01.006>, URL <https://www.sciencedirect.com/science/article/pii/S0925231220300254>.
- Li, S., Schulwol, Z.B., Mikkiläinen, R., 2025. Transformer based time-series forecasting for stock. <http://dx.doi.org/10.48550/arXiv.2502.09625>, arXiv:2502.09625 [q-fin]. URL <http://arxiv.org/abs/2502.09625>.
- Li, X., Shen, Q., 2024. A hybrid framework based on knowledge distillation for explainable disease diagnosis. *Expert Syst. Appl.* 238, 121844, Publisher: Elsevier.
- Li, D., Wang, Z., Chen, Y., Jiang, R., Ding, W., Okumura, M., 2024. A survey on deep active learning: Recent advances and new frontiers. <http://dx.doi.org/10.48550/arXiv.2405.00334>, arXiv:2405.00334 [cs]. URL <http://arxiv.org/abs/2405.00334>.
- Li, B., Weng, Y., Xia, F., Deng, H., 2024. Towards better Chinese-centric neural machine translation for low-resource languages. *Comput. Speech Lang.* 84, 101566, Publisher: Elsevier.
- Li, X., Zhang, T., Zhao, X., Yi, Z., 2020. Guided autoencoder for dimensionality reduction of pedestrian features. *Appl. Intell.* 50 (12), 4557–4567. <http://dx.doi.org/10.1007/s10489-020-01813-1>.
- Li, L., Zhu, J., Sun, M.-T., 2019. Deep learning based method for pruning deep neural networks. In: 2019 IEEE International Conference on Multimedia and Expo Workshops. ICMEW, pp. 312–317. <http://dx.doi.org/10.1109/ICMEW.2019.00-68>.
- Li, M., Zhu, W., Zhang, M., Sun, Y., Wang, Z., 2017. The novel recognition method with Optimal Wavelet Packet and LSTM based Recurrent Neural Network. In: 2017 IEEE International Conference on Mechatronics and Automation. ICMA, pp. 584–589. <http://dx.doi.org/10.1109/ICMA.2017.8015882>.
- Lin, M., Chen, Q., Yan, S., 2014. Network in network. <http://dx.doi.org/10.48550/arXiv.1312.4400>, arXiv:1312.4400 [cs] version: 3. URL <http://arxiv.org/abs/1312.4400>.
- Lin, T.-Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S., 2017. Feature pyramid networks for object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 2117–2125.
- Lin, Z., Feng, M., Santos, C.N.d., Yu, M., Xiang, B., Zhou, B., Bengio, Y., 2017. A structured self-attentive sentence embedding. URL <https://arxiv.org/abs/1703.03130>.
- Lin, T.-Y., Goyal, P., Girshick, R., He, K., Dollár, P., 2020. Focal loss for dense object detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 42 (2), 318–327. <http://dx.doi.org/10.1109/TPAMI.2018.2858826>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/8417976>.
- Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.-Y., Berg, A.C., 2016. SSD: Single shot MultiBox detector. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (Eds.), *Computer Vision – ECCV 2016*. In: Lecture Notes in Computer Science, Springer International Publishing, pp. 21–37. [http://dx.doi.org/10.1007/978-3-319-46448-0\\_2](http://dx.doi.org/10.1007/978-3-319-46448-0_2).
- Liu, Y., Chen, Y., Lasang, P., Sun, Q., 2022. Covariance attention for semantic segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.* 44 (4), 1805–1818. <http://dx.doi.org/10.1109/TPAMI.2020.3026069>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/9206128>.
- Liu, W., Chen, J., Lyu, Z., Feng, R., Hu, T., Deng, L., 2025. Automatic tile position and orientation detection combining deep-learning and rule-based computer vision algorithms. *Autom. Constr.* 171, 106001. <http://dx.doi.org/10.1016/j.autcon.2025.106001>, URL <https://www.sciencedirect.com/science/article/pii/S092658052500041X>.
- Liu, Q., Dong, Y., Li, X., 2023. Multi-stage context refinement network for semantic segmentation. *Neurocomputing* 535, 53–63. <http://dx.doi.org/10.1016/j.neucom.2023.03.006>, URL <https://www.sciencedirect.com/science/article/pii/S0925231223002254>.
- Liu, G., Guo, J., 2019. Bidirectional LSTM with attention mechanism and convolutional layer for text classification. *Neurocomputing* 337, 325–338. <http://dx.doi.org/10.1016/j.neucom.2019.01.078>, URL <https://www.sciencedirect.com/science/article/pii/S0925231219301067>.
- Liu, S., Huang, D., Wang, Y., 2019. Adaptive NMS: Refining pedestrian detection in a crowd. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6459–6468.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. RoBERTa: A robustly optimized bert pretraining approach. URL <https://arxiv.org/abs/1907.11692>.
- Liu, S., Qi, L., Qin, H., Shi, J., Jia, J., 2018. Path aggregation network for instance segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8759–8768.
- Liu, X., Wang, H., Li, Z., Qin, L., 2021. Deep learning in ECG diagnosis: A review. *Knowl.-Based Syst.* 227, 107187. <http://dx.doi.org/10.1016/j.knsys.2021.107187>, URL <https://www.sciencedirect.com/science/article/pii/S0950705121004494>.
- Liu, Z., Wang, L., Wu, W., Qian, C., Lu, T., 2021. TAM: Temporal adaptive module for video recognition. <http://dx.doi.org/10.48550/arXiv.2005.06803>, arXiv:2005.06803 [cs]. URL <http://arxiv.org/abs/2005.06803>.
- Liu, Z., Xu, Z., Jin, J., Shen, Z., Darrell, T., 2023. Dropout reduces underfitting. In: Proceedings of the 40th International Conference on Machine Learning. ICML '23, Vol. 202, JMLR.org, pp. 22233–22248.
- Liu, H., Zeng, X., Shen, Y., Xu, J., Khan, Z., 2025. A single-stage navigation path extraction network for agricultural robots in orchards. *Comput. Electron. Agric.* 229, 109687. <http://dx.doi.org/10.1016/j.compag.2024.109687>, URL <https://www.sciencedirect.com/science/article/pii/S0168169924010780>.
- Llamedo, M., Martínez, J.P., 2011. Heartbeat classification using feature selection driven by database generalization criteria. *IEEE Trans. Biomed. Eng.* 58 (3), 616–625. <http://dx.doi.org/10.1109/TBME.2010.2068048>.
- Loh, C., Han, S., Sudalairaj, S., Dangovski, R., Xu, K., Wenzel, F., Soljagic, M., Srivastava, A., 2023. Multi-symmetry ensembles: Improving diversity and generalization via opposing symmetries. In: Proceedings of the 40th International Conference on Machine Learning. PMLR, pp. 22614–22630, ISSN: 2640-3498. URL <https://proceedings.mlr.press/v202/loh23a.html>.
- Long, J., Shelhamer, E., Darrell, T., 2015. Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3431–3440.
- Lu, X., Li, S., Fujimoto, M., 2020. Automatic speech recognition. *Speech-to-Speech Transl.* 21–38.
- Luca, A.R., Ursuleanu, T.F., Gheorghe, L., Grigoriu, R., Iancu, S., Hlusuac, M., Grigoriu, A., 2022. Impact of quality, type and volume of data used by deep learning models in the analysis of medical images. *Inform. Med. Unlocked* 29, 100911. <http://dx.doi.org/10.1016/j.imu.2022.100911>, URL <https://www.sciencedirect.com/science/article/pii/S2352914822000612>.
- Luong, T., Pham, H., Manning, C.D., 2015. Effective approaches to attention-based neural machine translation. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing. Association for Computational Linguistics, Lisbon, Portugal, pp. 1412–1421. <http://dx.doi.org/10.18653/v1/D15-1166>, URL <https://aclanthology.org/D15-1166/>.
- Lupión, M., Cruciani, F., Cleland, I., Nugent, C., Ortigosa, P.M., 2024. Data augmentation for human activity recognition with generative adversarial networks. *IEEE J. Biomed. Heal. Inform.* 28 (4), 2350–2361. <http://dx.doi.org/10.1109/JBHI.2024.3364910>, Conference Name: IEEE Journal of Biomedical and Health Informatics. URL <https://ieeexplore.ieee.org/document/10432776>.
- Lupo, L., Dinarelli, M., Besacier, L., 2023. Encoding sentence position in context-aware neural machine translation with concatenation. In: Proceedings of the Fourth Workshop on Insights from Negative Results in NLP. Association for Computational Linguistics, Dubrovnik, Croatia, pp. 33–44. <http://dx.doi.org/10.18653/v1/2023.insights-1.4>, URL <https://aclanthology.org/2023.insights-1.4/>.
- Luwe, Y.J., Lee, C.P., Lim, K.M., 2022. Wearable sensor-based human activity recognition with hybrid deep learning model. *Informatics* 9 (3), 56. <http://dx.doi.org/10.3390/informatics9030056>, Number: 3 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/2227-9709/9/3/56>.
- Ma, T., Pan, Q., Rong, H., Qian, Y., Tian, Y., Al-Nabhan, N., 2021. T-BERTSum: Topic-aware text summarization based on bert. *IEEE Trans. Comput. Soc. Syst.* 9 (3), 879–890, Publisher: IEEE.
- Ma, Y., Wang, R., Zong, M., Ji, W., Wang, Y., Ye, B., 2024. Convolutional transformer network for fine-grained action recognition. *Neurocomputing* 569, 127027. <http://dx.doi.org/10.1016/j.neucom.2023.127027>, URL <https://www.sciencedirect.com/science/article/pii/S0925231223011505>.
- Maas, A.L., Hannun, A.Y., Ng, A.Y., et al., 2013. Rectifier nonlinearities improve neural network acoustic models. In: *Proc. Icm. Vol. 30*, Atlanta, GA, p. 3, Issue: 1.
- Maldonado, S., Vairetti, C., Jara, K., Carrasco, M., López, J., 2023. OWAdept: An adaptive loss function for deep learning using OWA operators. *Knowl.-Based Syst.* 280, 111022.
- Marra, G., Dumančić, S., Manhaeve, R., De Raedt, L., 2024. From statistical relational to neurosymbolic artificial intelligence: A survey. *Artificial Intelligence* 328, 104062. <http://dx.doi.org/10.1016/j.artint.2023.104062>, URL <https://www.sciencedirect.com/science/article/pii/S0004370223002084>.
- Mathews, S.M., Kamthammettu, C., Barner, K.E., 2018. A novel application of deep learning for single-lead ECG classification. *Comput. Biol. Med.* 99, 53–62. <http://dx.doi.org/10.1016/j.combiomed.2018.05.013>, URL <https://www.sciencedirect.com/science/article/pii/S0010482518301264>.
- Matsugu, M., Mori, K., Mitari, Y., Kaneda, Y., 2003. Subject independent facial expression recognition with robust face detection using a convolutional neural network. *Neural Netw.* 16 (5–6), 555–559. [http://dx.doi.org/10.1016/S0893-6080\(03\)00115-1](http://dx.doi.org/10.1016/S0893-6080(03)00115-1).
- Meetei, L.S., Singh, A., Singh, T.D., Bandyopadhyay, S., 2023. Do cues in a video help in handling rare words in a machine translation system under a low-resource setting? *Nat. Lang. Process. J.* 3, 100016, Publisher: Elsevier.
- Misir, O., 2024. Drivable path detection for a mobile robot with differential drive using a deep learning based segmentation method for indoor navigation. *PeerJ Comput. Sci.* 10, e2514. <http://dx.doi.org/10.7717/peerj-cs.2514>, Publisher: PeerJ Inc. URL <https://peerj.com/articles/cs-2514>.
- Misra, D., 2020. Mish: A self regularized non-monotonic activation function. <http://dx.doi.org/10.48550/arXiv.1908.08681>, arXiv:1908.08681 [cs, stat]. URL <http://arxiv.org/abs/1908.08681>.
- Misra, D., Nalamada, T., Arasanipalai, A.U., Hou, Q., 2020. Rotate to attend: Convolutional triplet attention module. <http://dx.doi.org/10.48550/arXiv.2010.03045>, arXiv:2010.03045 [cs]. URL <http://arxiv.org/abs/2010.03045>.
- Modir, A., Shamekhi, S., Ghaderyan, P., 2023. A systematic review and methodological analysis of EEG-based biomarkers of Alzheimer's disease. *Measurement* 113274.

- Mohammadzadeh, M., Ghadami, A., Taheri, A., Behzadipour, S., 2025. cGAN-based high dimensional IMU sensor data generation for enhanced human activity recognition in therapeutic activities. *Biomed. Signal Process. Control.* 103, 107476. <http://dx.doi.org/10.1016/j.bspc.2024.107476>, URL <https://www.sciencedirect.com/science/article/pii/S1746809424015349>.
- Mohd Noor, M.H., 2021. Feature learning using convolutional denoising autoencoder for activity recognition. *Neural Comput. Appl.* 33 (17), 10909–10922. <http://dx.doi.org/10.1007/s00521-020-05638-4>.
- Mohd Noor, M.H., Tan, S.Y., Ab Wahab, M.N., 2022. Deep temporal Conv-LSTM for activity recognition. *Neural Process. Lett.* 54 (5), 4027–4049. <http://dx.doi.org/10.1007/s11063-022-10799-5>.
- Moles, L., Andres, A., Echegaray, G., Boto, F., 2024. Exploring data augmentation and active learning benefits in imbalanced datasets. *Mathematics* 12 (12), 1898.
- Mozaffari, L., Zhang, J., 2024. Predictive modeling of stock prices using transformer model. In: *Proceedings of the 2024 9th International Conference on Machine Learning Technologies. ICLMT '24*, Association for Computing Machinery, New York, NY, USA, pp. 41–48. <http://dx.doi.org/10.1145/3674029.3674037>, URL <https://dl.acm.org/doi/10.1145/3674029.3674037>.
- Msiqwa, C., Bernard, D., Yun, J., 2024. IoT-driven wearable devices enhancing healthcare: ECG classification with cluster-based GAN and meta-features. *Internet Things* 28, 101405. <http://dx.doi.org/10.1016/j.iot.2024.101405>, URL <https://www.sciencedirect.com/science/article/pii/S2542660524003469>.
- Mukhamadiyev, A., Khujayarov, I., Djuraev, O., Cho, J., 2022. Automatic speech recognition method based on deep learning approaches for Uzbek language. *Sensors* 22 (10), 3683. <http://dx.doi.org/10.3390/s22103683>.
- Mumuni, A., Mumuni, F., 2022. Data augmentation: A comprehensive survey of modern approaches. *Array* 16, 100258. <http://dx.doi.org/10.1016/j.array.2022.100258>, URL <https://www.sciencedirect.com/science/article/pii/S2590005622000911>.
- Munappy, A.R., Bosch, J., Olsson, H.H., Arpteg, A., Brinne, B., 2022. Data management for production quality deep learning models: Challenges and solutions. *J. Syst. Softw.* 191, 111359. <http://dx.doi.org/10.1016/j.jss.2022.111359>, URL <https://www.sciencedirect.com/science/article/pii/S0164121222000905>.
- Murfi, H., Syamsyuriani, Gowandi, T., Ardanawari, G., Nurrohman, S., 2024. BERT-based combination of convolutional and recurrent neural network for Indonesian sentiment analysis. *Appl. Soft Comput.* 151, 111112. <http://dx.doi.org/10.1016/j.asoc.2023.111112>, URL <https://www.sciencedirect.com/science/article/pii/S1568494623011304>.
- Murtaza, H., Ahmed, M., Khan, N.F., Murtaza, G., Zafar, S., Bano, A., 2023. Synthetic data generation: State of the art in health care domain. *Comput. Sci. Rev.* 48, 100546. <http://dx.doi.org/10.1016/j.cosrev.2023.100546>, URL <https://www.sciencedirect.com/science/article/pii/S1574013723000138>.
- Nafea, O., Abdul, W., Muhammad, G., Alsulaiman, M., 2021. Sensor-based human activity recognition with spatio-temporal deep learning. *Sensors* 21 (6), 2141. <http://dx.doi.org/10.3390/s21062141>, Number: 6 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/21/6/2141>.
- Nakamura, K., Hong, B.-W., 2019. Adaptive weight decay for deep neural networks. *IEEE Access* 7, 118857–118865. <http://dx.doi.org/10.1109/ACCESS.2019.2937139>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/8811458>.
- Nanni, L., Loreggia, A., Barcellona, L., Ghidoni, S., 2023. Building ensemble of deep networks: Convolutional networks and transformers. *IEEE Access* 11, 124962–124974. <http://dx.doi.org/10.1109/ACCESS.2023.3330442>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/10309107>.
- Narotamo, H., Dias, M., Santos, R., Carreiro, A.V., Gamboa, H., Silveira, M., 2024. Deep learning for ECG classification: A comparative study of 1D and 2D representations and multimodal fusion approaches. *Biomed. Signal Process. Control.* 93, 106141. <http://dx.doi.org/10.1016/j.bspc.2024.106141>, URL <https://www.sciencedirect.com/science/article/pii/S174680942400199X>.
- Nassif, A.B., Shahin, I., Attili, I., Azzeh, M., Shaalan, K., 2019. Speech recognition using deep neural networks: A systematic review. *IEEE Access* 7, 19143–19165. <http://dx.doi.org/10.1109/ACCESS.2019.2896880>.
- Ng, A., et al., 2011. Sparse autoencoder. *CS294A Lect. Notes* 72 (2011), 1–19.
- Nie, Y.-p., Han, Y., Huang, J.-m., Jiao, B., Li, A.-p., 2017. Attention-based encoder-decoder model for answer selection in question answering. *Front. Inf. Technol. Electron. Eng.* 18 (4), 535–544, Publisher: Springer.
- Nikou, M., Mansourfar, G., Bagherzadeh, J., 2019. Stock price prediction using DEEP learning algorithm and its comparison with machine learning algorithms. *Intell. Syst. Account. Financ. Manag.* 26 (4), 164–174.
- Noh, H., Hong, S., Han, B., 2015. Learning deconvolution network for semantic segmentation. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 1520–1528.
- Odena, A., Olah, C., Shlens, J., 2017. Conditional image synthesis with auxiliary classifier gans. In: *International Conference on Machine Learning*. PMLR, pp. 2642–2651.
- Okta, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., Glocker, B., Rueckert, D., 2018. Attention U-net: Learning where to look for the pancreas. <http://dx.doi.org/10.48550/arXiv.1804.03999>, arXiv:1804.03999 [cs]. URL <http://arxiv.org/abs/1804.03999>.
- Otter, D.W., Medina, J.R., Kalita, J.K., 2021. A survey of the usages of deep learning for natural language processing. *IEEE Trans. Neural Netw. Learn. Syst.* 32 (2), 604–624. <http://dx.doi.org/10.1109/TNNLS.2020.2979670>.
- Ozbayoglu, A.M., Gudelek, M.U., Sezer, O.B., 2020. Deep learning for financial applications: A survey. *Appl. Soft Comput.* 93, 106384.
- Padmanabhan, J., Premkumar, M.J.J., 2015. Machine learning in automatic speech recognition: A survey. *IETE Tech. Rev.* 32 (4), 240–251. <http://dx.doi.org/10.1080/02564602.2015.1010611>.
- Pan, J., Fang, W., Zhang, Z., Chen, B., Zhang, Z., Wang, S., 2024. Multimodal emotion recognition based on facial expressions, speech, and EEG. *IEEE Open J. Eng. Med. Biol.* 5, 396–403. <http://dx.doi.org/10.1109/OJEMB.2023.3240280>, Conference Name: IEEE Open Journal of Engineering in Medicine and Biology. URL <https://ieeexplore.ieee.org/document/10026861>.
- Pandey, P., Seeja, K.R., 2022. Subject independent emotion recognition from EEG using VMD and deep learning. *J. King Saud Univ. - Comput. Inf. Sci.* 34 (5), 1730–1738. <http://dx.doi.org/10.1016/j.jksuci.2019.11.003>, URL <https://www.sciencedirect.com/science/article/pii/S1319157819309991>.
- Park, S., Kwak, N., 2017. Analysis on the dropout effect in convolutional neural networks. In: *Lai, S.-H., Lepetit, V., Nishino, K., Sato, Y. (Eds.), Computer Vision – ACCV 2016*. In: *Lecture Notes in Computer Science*, Springer International Publishing, pp. 189–204. [http://dx.doi.org/10.1007/978-3-319-54184-6\\_12](http://dx.doi.org/10.1007/978-3-319-54184-6_12).
- Park, W., Park, I., Kim, S., Ryu, J., 2023. Robust asymmetric loss for multi-label long-tailed learning. pp. 2711–2720, URL [https://openaccess.thecvf.com/content/ICCV2023W/CVAMD/html/Park\\_Robust\\_Asymmetric\\_Loss\\_for\\_Multi-Label\\_Long-Tailed\\_Learning\\_ICCVW\\_2023\\_paper.html](https://openaccess.thecvf.com/content/ICCV2023W/CVAMD/html/Park_Robust_Asymmetric_Loss_for_Multi-Label_Long-Tailed_Learning_ICCVW_2023_paper.html).
- Peng, Z., Guo, Z., Huang, W., Wang, Y., Xie, L., Jiao, J., Tian, Q., Ye, Q., 2023. Conformer: Local features coupling global representations for recognition and detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 45 (8), 9454–9468. <http://dx.doi.org/10.1109/TPAMI.2023.3243048>, Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. URL <https://ieeexplore.ieee.org/document/10040235>.
- Pierson, H.A., Gashler, M.S., 2017. Deep learning in robotics: a review of recent research. *Adv. Robot.* 31 (16), 821–835. <http://dx.doi.org/10.1080/01691864.2017.1365009>, Publisher: Taylor & Francis eprint: <https://doi.org/10.1080/01691864.2017.1365009>.
2025. Poloclub/cnn-explainer. original-date: 2019-11-03T23:15:24Z. URL <https://github.com/poloclub/cnn-explainer>.
2025. Poloclub/transformer-explainer. original-date: 2024-05-16T21:42:10Z. URL <https://github.com/poloclub/transformer-explainer>.
- Pouyanfar, S., Sadiq, S., Yan, Y., Tian, H., Tao, Y., Reyes, M.P., Shyu, M.-L., Chen, S.-C., Iyengar, S.S., 2018. A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput. Surv.* 51 (5), 92:1–92:36. <http://dx.doi.org/10.1145/3234150>.
- Prabhakararao, E., Dandapat, S., 2020. Attentive RNN-based network to fuse 12-lead ECG and clinical features for improved myocardial infarction diagnosis. *IEEE Signal Process. Lett.* 27, 2029–2033.
- Prabhavalkar, R., Rao, K., Sainath, T.N., Li, B., Johnson, L., Jaitly, N., 2017. A comparison of sequence-to-sequence models for speech recognition. In: *Interspeech*. pp. 939–943.
- Prechelt, L., 2012. Early stopping — But when? In: *Montavon, G., Orr, G.B., Müller, K.-R. (Eds.), Neural Networks: Tricks of the Trade: Second Edition*. In: *Lecture Notes in Computer Science*, Springer, pp. 53–67. [http://dx.doi.org/10.1007/978-3-642-35289-8\\_5](http://dx.doi.org/10.1007/978-3-642-35289-8_5).
- Qian, N., 1999. On the momentum term in gradient descent learning algorithms. *Neural Netw.* 12 (1), 145–151. [http://dx.doi.org/10.1016/S0893-6080\(98\)00116-6](http://dx.doi.org/10.1016/S0893-6080(98)00116-6), URL <https://www.sciencedirect.com/science/article/pii/S0893608098001166>.
- Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I., 2022. Robust speech recognition via large-scale weak supervision. <http://dx.doi.org/10.48550/arXiv.2212.04356>, arXiv:2212.04356 [eess]. URL <http://arxiv.org/abs/2212.04356>.
- Radford, A., Metz, L., Chintala, S., 2015. Unsupervised representation learning with deep convolutional generative adversarial networks. URL <https://arxiv.org/abs/1511.06434>.
- Raffel, C., Luong, M.-T., Liu, P.J., Weiss, R.J., Eck, D., 2017. Online and linear-time attention by enforcing monotonic alignments. In: *International Conference on Machine Learning*. PMLR, pp. 2837–2846.
- Ragab, M.G., Abdulkadir, S.J., Aziz, N., 2020. Random search one dimensional CNN for human activity recognition. In: *2020 International Conference on Computational Intelligence. ICCI, IEEE*, pp. 86–91.
- Rai, H.M., Chatterjee, K., 2022. Hybrid CNN-LSTM deep learning model and ensemble technique for automatic detection of myocardial infarction using big ECG data. *Appl. Intell.* 52 (5), 5366–5384.
- Raiaan, M.A.K., Sakib, S., Fahad, N.M., Mamun, A.A., Rahman, M.A., Shatabda, S., Mukta, M.S.H., 2024. A systematic review of hyperparameter optimization techniques in Convolutional Neural Networks. *Decis. Anal. J.* 11, 100470. <http://dx.doi.org/10.1016/j.dajour.2024.100470>, URL <https://www.sciencedirect.com/science/article/pii/S2772662224000742>.
- Redmon, J., Divvala, S., Girshick, R., Farhadi, A., 2016. You only look once: Unified, real-time object detection. pp. 779–788, URL [https://www.cv-foundation.org/openaccess/content\\_cvpr\\_2016/html/Redmon\\_You\\_Only\\_Look\\_CVPR\\_2016\\_paper.html](https://www.cv-foundation.org/openaccess/content_cvpr_2016/html/Redmon_You_Only_Look_CVPR_2016_paper.html).
- Redmon, J., Farhadi, A., 2018. YOLOv3: An incremental improvement. URL <https://arxiv.org/abs/1804.02767v1>.

- Ren, S., He, K., Girshick, R., Sun, J., 2015. Faster R-CNN: Towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*. Vol. 28, Curran Associates, Inc., URL [https://proceedings.neurips.cc/paper\\_files/paper/2015/hash/14bfa6bb14875e45bba028a21ed38046-Abstract.html](https://proceedings.neurips.cc/paper_files/paper/2015/hash/14bfa6bb14875e45bba028a21ed38046-Abstract.html).
- Rifai, S., Mesnil, G., Vincent, P., Muller, X., Bengio, Y., Dauphin, Y., Glorot, X., 2011. Higher order contractive auto-encoder. In: *Gunopulos, D., Hofmann, T., Malerba, D., Vazirgiannis, M. (Eds.), Machine Learning and Knowledge Discovery in Databases*. In: *Lecture Notes in Computer Science*, Springer, pp. 645–660. [http://dx.doi.org/10.1007/978-3-642-23783-6\\_41](http://dx.doi.org/10.1007/978-3-642-23783-6_41).
- Rippeth, E., Carpuat, M., Duh, K., Post, M., 2023. Improving word sense disambiguation in neural machine translation with salient document context. URL <https://arxiv.org/abs/2311.15507v1>.
- Rodrawangpai, B., Daungjaiboon, W., 2022. Improving text classification with transformers and layer normalization. *Mach. Learn. Appl.* 10, 100403. <http://dx.doi.org/10.1016/j.mlwa.2022.100403>, URL <https://www.sciencedirect.com/science/article/pii/S2666827022000792>.
- Romero, J., Olson, J.P., Aspuru-Guzik, A., 2017. Quantum autoencoders for efficient compression of quantum data. *Quantum Sci. Technol.* 2 (4), 045001. <http://dx.doi.org/10.1088/2058-9565/aa8072>, Publisher: IOP Publishing.
- Sak, H., Shannon, M., Rao, K., Beaufays, F., 2017. Recurrent neural aligner: An encoder-decoder neural network model for sequence to sequence mapping. In: *Interspeech*. Vol. 8, pp. 1298–1302.
- Sánchez-Hevia, H.A., Gil-Pita, R., Utrilla-Manso, M., Rosa-Zurera, M., 2022. Age group classification and gender recognition from speech with temporal convolutional neural networks. *Multimedia Tools Appl.* 81 (3), 3535–3552.
- Sarker, I.H., 2021. Deep learning: A comprehensive overview on techniques, taxonomy, applications and research directions. *SN Comput. Sci.* 2 (6), 420. <http://dx.doi.org/10.1007/s42979-021-00815-1>.
- Schaaf, N., de Mitri, O., Kim, H.B., Windberger, A., Huber, M.F., 2021. Towards measuring bias in image classification. In: *Artificial Neural Networks and Machine Learning-ICANN 2021: 30th International Conference on Artificial Neural Networks*, Bratislava, Slovakia, September 14–17, 2021, *Proceedings, Part III* 30. Springer, pp. 433–445.
- Schirmeister, R.T., Springenberg, J.T., Fiederer, L.D.J., Glasstetter, M., Eggenberger, K., Tangermann, M., Hutter, F., Burgard, W., Ball, T., 2017. Deep learning with convolutional neural networks for EEG decoding and visualization. *Hum. Brain Mapp.* 38 (11), 5391–5420. <http://dx.doi.org/10.1002/hbm.23730>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/hbm.23730>.
- Shakhadri, S.A.G., KR, K., Angadi, K.B., 2025. Samba-ASR: State-of-the-art speech recognition leveraging structured state-space models. <http://dx.doi.org/10.48550/arXiv.2501.02832>, arXiv:2501.02832 [cs]. URL <http://arxiv.org/abs/2501.02832>.
- Shamshirband, S., Fathi, M., Dehngani, A., Chronopoulos, A.T., Alinejad-Rokny, H., 2021. A review on deep learning approaches in healthcare systems: Taxonomies, challenges, and open issues. *J. Biomed. Inform.* 113, 103627. <http://dx.doi.org/10.1016/j.jbi.2020.103627>, URL <https://www.sciencedirect.com/science/article/pii/S1532046420302550>.
- Sharma, M., Tiwari, J., Patel, V., Acharya, U.R., 2021. Automated identification of sleep disorder types using triplet half-band filter and ensemble machine learning techniques with EEG signals. *Electronics* 10 (13), 1531.
- Shen, Z., Li, G., Fang, J., Zhong, H., Wang, J., Sun, Y., Shen, X., 2022. Aberrated multidimensional EEG characteristics in patients with generalized anxiety disorder: A machine-learning based analysis framework. *Sensors* 22 (14), 5420.
- Shen, F., Zhao, X., Kou, G., Alsaadi, F.E., 2021. A new deep learning ensemble credit risk evaluation model with an improved synthetic minority oversampling technique. *Appl. Soft Comput.* 98, 106852.
- Shewalkar, A., Nyavanandi, D., Ludwig, S.A., 2019. Performance evaluation of deep neural networks applied to speech recognition: RNN, LSTM and GRU. *J. Artif. Intell. Soft Comput. Res.* 9 (4), 235–245.
- Shi, X., Chen, Z., Wang, H., Yeung, D.-Y., Wong, W.-k., Woo, W.-c., 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. <http://dx.doi.org/10.48550/arXiv.1506.04214>, arXiv:1506.04214 [cs] version: 2. URL <http://arxiv.org/abs/1506.04214>.
- Shi, S., Fang, Q., Xu, X., Zhao, T., 2024. Similarity distance-based label assignment for tiny object detection. In: *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems. IROS*, pp. 13711–13718. <http://dx.doi.org/10.1109/IROS58592.2024.10801448>, ISSN: 2153-0866. URL <https://ieeexplore.ieee.org/document/10801448>.
- Shi, H., Hayat, M., Cai, J., 2023. Transformer scale gate for semantic segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3051–3060.
- Shi, L.-F., Liu, Z.-Y., Zhou, K.-J., Shi, Y., Jing, X., 2023. Novel deep learning network for gait recognition using multimodal inertial sensors. *Sensors* 23 (2), 849. <http://dx.doi.org/10.3390/s23020849>, Number: 2 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/23/2/849>.
- Shrestha, A., Fang, H., Mei, Z., Rider, D.P., Wu, Q., Qiu, Q., 2022. A survey on neuro-morphic computing: Models and hardware. *IEEE Circuits Syst. Mag.* 22 (2), 6–35. <http://dx.doi.org/10.1109/MCAS.2022.3166331>, Conference Name: IEEE Circuits and Systems Magazine. URL <https://ieeexplore.ieee.org/document/9782767>.
- Shuvo, M.M.H., Ahmed, N., Nouduri, K., Palaniappan, K., 2020. A hybrid approach for human activity recognition with support vector machine and 1D convolutional neural network. In: *2020 IEEE Applied Imagery Pattern Recognition Workshop. AIPR*, IEEE, pp. 1–5.
- Simard, P., Steinkraus, D., Platt, J., 2003. Best practices for convolutional neural networks applied to visual document analysis. In: *Seventh International Conference on Document Analysis and Recognition*, 2003. *Proceedings*. pp. 958–963. <http://dx.doi.org/10.1109/ICDAR.2003.1227801>, URL <https://ieeexplore.ieee.org/document/1227801>.
- Simonyan, K., Zisserman, A., 2015. Very deep convolutional networks for large-scale image recognition. <http://dx.doi.org/10.48550/arXiv.1409.1556>, arXiv:1409.1556 [cs] version: 6. URL <http://arxiv.org/abs/1409.1556>.
- Singh, S., Pandey, S.K., Pawar, U., Janghel, R.R., 2018. Classification of ECG arrhythmia using recurrent neural networks. *Procedia Comput. Sci.* 132, 1290–1297.
- Singh, G., Sharma, S., Kumar, V., Kaur, M., Baz, M., Masud, M., et al., 2021. Spoken language identification using deep learning. *Comput. Intell. Neurosci.* 2021.
- Singh, R., Srivastava, S., 2017. Stock prediction using deep learning. *Multimedia Tools Appl.* 76 (18), 18569–18584. <http://dx.doi.org/10.1007/s11042-016-4159-7>.
- Smilov, D., Carter, S., Tensorflow — Neural Network Playground. URL <http://playground.tensorflow.org>.
- Song, J., Zhai, Q., Wang, C., Liu, J., 2024. EEGAN-Net: enhancing EEG signal classification through data augmentation. *Front. Hum. Neurosci.* 18, <http://dx.doi.org/10.3389/fnhum.2024.1430086>, Publisher: Frontiers. URL <https://www.frontiersin.org/journals/human-neuroscience/articles/10.3389/fnhum.2024.1430086/full>.
- Soori, M., Arezoo, B., Dastres, R., 2023. Artificial intelligence, machine learning and deep learning in advanced robotics, A review. *Cogn. Robot.*
- Sowmya, S., Jose, D., 2022. Contemplate on ECG signals and classification of arrhythmia signals using CNN-LSTM deep learning model. *Meas.: Sens.* 24, 100558.
- Srivastava, R.K., Greff, K., Schmidhuber, J., 2015. Highway networks. <http://dx.doi.org/10.48550/arXiv.1505.00387>, arXiv:1505.00387 [cs]. URL <http://arxiv.org/abs/1505.00387>.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15 (56), 1929–1958, URL <http://jmlr.org/papers/v15/srivastava14a.html>.
- Srivastava, R., Pandey, D., 2022. Speech recognition using HMM and Soft Computing. *Mater. Today: Proc.* 51, 1878–1883. <http://dx.doi.org/10.1016/j.matpr.2021.10.097>.
- Stahlberg, F., 2020. Neural machine translation: A review. *J. Artificial Intelligence Res.* 69, 343–418. <http://dx.doi.org/10.1613/jair.1.12007>, URL <https://www.jair.org/index.php/jair/article/view/12007>.
- Stanton, S., Izmailov, P., Kirichenko, P., Alemi, A.A., Wilson, A.G., 2021. Does knowledge distillation really work? In: *Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (Eds.), Advances in Neural Information Processing Systems*. Vol. 34, Curran Associates, Inc., pp. 6906–6919, URL [https://proceedings.neurips.cc/paper\\_files/paper/2021/file/376c6b9ff3bedbba56751a84fffc10c-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2021/file/376c6b9ff3bedbba56751a84fffc10c-Paper.pdf).
- Stooke, A., Prabhavalkar, R., Sim, K.C., Mengibar, P.M., 2025. Aligner-encoders: Self-attention transformers can be self-transducers. <http://dx.doi.org/10.48550/arXiv.2502.05232>, arXiv:2502.05232 [cs]. URL <http://arxiv.org/abs/2502.05232>.
- Strudel, R., Garcia, R., Laptev, I., Schmid, C., 2021. Segmenter: Transformer for semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 7262–7272.
- Sun, X., Fu, J., Wei, B., Li, Z., Li, Y., Wang, N., 2023. A self-attentional ResNet-LightGBM model for IoT-enabled voice liveness detection. *IEEE Internet Things J.* 10 (9), 8257–8270. <http://dx.doi.org/10.1109/IJOT.2022.3230992>, Conference Name: IEEE Internet of Things Journal. URL <https://ieeexplore.ieee.org/document/9996089>.
- Sun, A., Hong, W., Li, J., Mao, J., 2024. An arrhythmia classification model based on a CNN-LSTM-SE algorithm. *Sensors* 24 (19), 6306. <http://dx.doi.org/10.3390/s24196306>, Number: 19 Publisher: Multidisciplinary Digital Publishing Institute. URL <https://www.mdpi.com/1424-8220/24/19/6306>.
- Sun, S., Wang, S., Wei, Y., 2020. A new ensemble deep learning approach for exchange rates forecasting and trading. *Adv. Eng. Inform.* 46, 101160. <http://dx.doi.org/10.1016/j.aei.2020.101160>, URL <https://www.sciencedirect.com/science/article/pii/S1474034620301312>.
- Sun, Y., Xu, X., Tian, X., Zhou, L., Li, Y., 2024. Efficient human activity recognition: A deep convolutional transformer-based contrastive self-supervised approach using wearable sensors. *Eng. Appl. Artif. Intell.* 135, 108705. <http://dx.doi.org/10.1016/j.engappai.2024.108705>, URL <https://www.sciencedirect.com/science/article/pii/S0952197624008637>.
- Sun, J., Zhou, Y., Zhang, B., 2019. ResFPA-GAN: Text-to-image synthesis with generative adversarial network based on residual block feature pyramid attention. In: *2019 IEEE International Conference on Advanced Robotics and Its Social Impacts. ARSO*, IEEE, pp. 317–322.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks. *Adv. Neural Inf. Process. Syst.* 27.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2014. Going deeper with convolutions. <http://dx.doi.org/10.48550/arXiv.1409.4842>, arXiv:1409.4842 [cs] version: 1. URL <http://arxiv.org/abs/1409.4842>.

- Szegedy, C., Zaremba, W., Sutskever, I., Bruna, J., Erhan, D., Goodfellow, I., Fergus, R., 2013. Intriguing properties of neural networks. URL <https://arxiv.org/abs/1312.6199>.
- Szűcs, G., Huszti, D., 2019. Seq2seq deep learning method for summary generation by LSTM with two-way encoder and beam search decoder. In: 2019 IEEE 17th International Symposium on Intelligent Systems and Informatics. SISY, pp. 221–226. <http://dx.doi.org/10.1109/SISY47553.2019.9111502>.
- Tabacof, P., Tavares, J., Valle, E., 2016. Adversarial images for variational autoencoders. URL <https://arxiv.org/abs/1612.00155>.
- Talaei Khoei, T., Ould Slimane, H., Kaabouch, N., 2023. Deep learning: systematic review, models, challenges, and research directions. *Neural Comput. Appl.* 35 (31), 23103–23124. <http://dx.doi.org/10.1007/s00521-023-08957-4>.
- Tan, M., Le, Q., 2019. EfficientNet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning. PMLR, pp. 6105–6114.
- Tan, Z., Li, D., Wang, S., Beigi, A., Jiang, B., Bhattacharjee, A., Karami, M., Li, J., Cheng, L., Liu, H., 2024. Large language models for data annotation: A survey. <http://dx.doi.org/10.48550/arXiv.2402.13446>, arXiv:2402.13446 [cs] version: 2. URL <http://arxiv.org/abs/2402.13446>.
- Tan, M., Pang, R., Le, Q.V., 2020. EfficientDet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR.
- Tan, M., Pang, R., Le, Q.V., 2020. Efficientdet: Scalable and efficient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10781–10790.
- Tang, Y., Zhang, L., Teng, Q., Min, F., Song, A., 2022. Triple cross-domain attention on human activity recognition using wearable sensors. *IEEE Trans. Emerg. Top. Comput. Intell.* 6 (5), 1167–1176. <http://dx.doi.org/10.1109/TETCI.2021.3136642>, Conference Name: IEEE Transactions on Emerging Topics in Computational Intelligence. URL <https://ieeexplore.ieee.org/document/9669910>.
- Tao, M., Tang, H., Wu, F., Jing, X.-Y., Bao, B.-K., Xu, C., 2022. DF-GAN: A simple and effective baseline for text-to-image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16515–16525.
- Termritthikun, C., Umer, A., Suwanwimolkul, S., Xia, F., Lee, I., 2023. Explainable knowledge distillation for on-device chest x-ray classification. *IEEE/ACM Trans. Comput. Biol. Bioinform.* Publisher: IEEE.
- Tibrewal, N., Leeuwis, N., Alimardani, M., 2022. Classification of motor imagery EEG using deep learning increases performance in inefficient BCI users. *Plos One* 17 (7), e0268880.
- Tibshirani, R., 1996. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 58 (1), 267–288, Publisher: [Royal Statistical Society, Wiley]. URL <https://www.jstor.org/stable/2346178>.
- Tirumala, S.S., Shahamiri, S.R., 2016. A review on deep learning approaches in speaker identification. In: Proceedings of the 8th International Conference on Signal Processing Systems. Association for Computing Machinery, pp. 142–147. <http://dx.doi.org/10.1145/3015166.3015210>.
- Tjoa, E., Khok, H.J., Chouhan, T., Guan, C., 2023. Enhancing the confidence of deep learning classifiers via interpretable saliency maps. *Neurocomputing* 562, 126825, Publisher: Elsevier.
- Tonekaboni, S., Joshi, S., McCradden, M.D., Goldenberg, A., 2019. What clinicians want: contextualizing explainable machine learning for clinical end use. In: Machine Learning for Healthcare Conference. PMLR, pp. 359–380.
- Tsao, C.W., Aday, A.W., Almarzooq, Z.I., Anderson, C.A., Arora, P., Avery, C.L., Baker-Smith, C.M., Beaton, A.Z., Boehme, A.K., Buxton, A.E., et al., 2023. Heart disease and stroke statistics—2023 update: a report from the American Heart Association. *Circulation* 147 (8), e93–e621.
- Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y., 2022. MaxViT: Multi-axis vision transformer. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (Eds.), *Computer Vision – ECCV 2022*. In: Lecture Notes in Computer Science, Springer Nature Switzerland, pp. 459–479. [http://dx.doi.org/10.1007/978-3-031-20053-3\\_27](http://dx.doi.org/10.1007/978-3-031-20053-3_27).
- Van Houdt, G., Mosquera, C., Nápoles, G., 2020. A review on the long short-term memory model. *Artif. Intell. Rev.* 53 (8), 5929–5955. <http://dx.doi.org/10.1007/s10462-020-09838-1>.
- Vaquerizo-Villar, F., Gutiérrez-Tobal, G.C., Calvo, E., Álvarez, D., Kheirandish-Gozal, L., Del Campo, F., Gozal, D., Hornero, R., 2023. An explainable deep-learning model to stage sleep states in children and propose novel EEG-related patterns in sleep apnea. *Comput. Biol. Med.* 165, 107419.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2023. Attention is all you need. <http://dx.doi.org/10.48550/arXiv.1706.03762>, arXiv:1706.03762 [cs]. URL <http://arxiv.org/abs/1706.03762>.
- Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.-A., 2008. Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th International Conference on Machine Learning. ICML '08, Association for Computing Machinery, pp. 1096–1103. <http://dx.doi.org/10.1145/1390156.1390294>.
- Vukicevic, A.M., Petrovic, M., Jurisevic, N., Djapan, M., Knezevic, N., Novakovic, A., Jovanovic, K., 2025. Versatile waste sorting in small batch and flexible manufacturing industries using deep learning techniques. *Sci. Rep.* 15 (1), 3756. <http://dx.doi.org/10.1038/s41598-025-87226-x>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-025-87226-x>.
- Wang, P., hamme, H.V., 2024. Disentangled-transformer: An explainable end-to-end automatic speech recognition model with speech content-context separation. <http://dx.doi.org/10.48550/arXiv.2411.17846>, arXiv:2411.17846 [eess]. URL <http://arxiv.org/abs/2411.17846>.
- Wang, K., Huang, X., Tan, C., Yiu, S.-M., Chen, Z., Lei, X., 2024. Optimizing neural network training: A Markov chain approach for resource conservation. *IEEE Trans. Artif. Intell.* 1–14. <http://dx.doi.org/10.1109/TAI.2024.3413688>, Conference Name: IEEE Transactions on Artificial Intelligence. URL <https://ieeexplore.ieee.org/document/10557251>.
- Wang, R., Li, Z., Cao, J., Chen, T., Wang, L., 2019. Convolutional recurrent neural networks for text classification. In: 2019 International Joint Conference on Neural Networks. IJCNN, pp. 1–6. <http://dx.doi.org/10.1109/IJCNN.2019.8852406>, ISSN: 2161-4407. URL <https://ieeexplore.ieee.org/document/8852406>.
- Wang, G., Li, C., Wang, W., Zhang, Y., Shen, D., Zhang, X., Henao, R., Carin, L., 2018. Joint embedding of words and labels for text classification. In: Gurevych, I., Miyao, Y. (Eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Melbourne, Australia, pp. 2321–2331. <http://dx.doi.org/10.18653/v1/P18-1216>, URL <https://aclanthology.org/P18-1216/>.
- Wang, W., Li, W., Zhang, N., Liu, K., 2020. Portfolio formation with preselection using deep learning from long-term financial data. *Expert Syst. Appl.* 143, 113042.
- Wang, Z., Liang, B., Dhat, V., Brumbaugh, Z., Walker, N., Krishna, R., Cakmak, M., 2024. I can tell what I am doing: Toward real-world natural language grounding of robot experiences. <http://dx.doi.org/10.48550/arXiv.2411.12960>, arXiv:2411.12960 [cs] version: 1. URL <http://arxiv.org/abs/2411.12960>.
- Wang, G., Ma, J., Wang, Y., Tao, T., Ren, G., Zhu, H., 2023. SUDF-RS: A new foreign exchange rate prediction method considering the complementarity of supervised and unsupervised deep representation features. *Expert Syst. Appl.* 214, 119152. <http://dx.doi.org/10.1016/j.eswa.2022.119152>, URL <https://www.sciencedirect.com/science/article/pii/S0957417422021704>.
- Wang, J., Ma, Y., Zhang, L., Gao, R.X., Wu, D., 2018. Deep learning for smart manufacturing: Methods and applications. *J. Manuf. Syst.* 48, 144–156. <http://dx.doi.org/10.1016/j.jmsy.2018.01.003>, URL <https://www.sciencedirect.com/science/article/pii/S0278612518300037>.
- Wang, Q., Ma, Y., Zhao, K., Tian, Y., 2022. A comprehensive survey of loss functions in machine learning. *Ann. Data Sci.* 9 (2), 187–212. <http://dx.doi.org/10.1007/s40745-020-00253-5>.
- Wang, J., Qiao, X., Liu, C., Wang, X., Liu, Y., Yao, L., Zhang, H., 2021. Automated ECG classification using a non-local convolutional block attention module. *Comput. Methods Programs Biomed.* 203, 106006. <http://dx.doi.org/10.1016/j.cmpb.2021.106006>, URL <https://www.sciencedirect.com/science/article/pii/S016926072100081X>.
- Wang, S., Qu, J., Zhang, Y., Zhang, Y., 2023. Multimodal emotion recognition from EEG signals and facial expressions. *IEEE Access* 11, 33061–33068. <http://dx.doi.org/10.1109/ACCESS.2023.3263670>, Conference Name: IEEE Access. URL <https://ieeexplore.ieee.org/document/10089483>.
- Wang, M., Rahardja, S., Fränti, P., Rahardja, S., 2023. Single-lead ECG recordings modeling for end-to-end recognition of atrial fibrillation with dual-path RNN. *Biomed. Signal Process. Control* 79, 104067.
- Wang, Q., Wu, B., Zhu, P., Li, P., Zuo, W., Hu, Q., 2020. ECA-net: Efficient channel attention for deep convolutional neural networks. <http://dx.doi.org/10.48550/arXiv.1910.03151>, arXiv:1910.03151 [cs]. URL <http://arxiv.org/abs/1910.03151>.
- Wang, X., Zhang, Y., Bai, N., Yu, Q., Wang, Q., 2024. Class-imbalanced time series anomaly detection method based on cost-sensitive hybrid network. *Expert Syst. Appl.* 238, 122192.
- Weiser, M., 1999. The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev.* 3 (3), 3–11. <http://dx.doi.org/10.1145/329124.329126>, URL <https://dl.acm.org/doi/10.1145/329124.329126>.
- Weng, L., 2019. From GAN to WGAN. <http://dx.doi.org/10.48550/arXiv.1904.08994>, arXiv:1904.08994 [cs, stat]. URL <http://arxiv.org/abs/1904.08994>.
- Widrow, B., Rumelhart, D.E., Lehr, M.A., 1994. Neural networks: applications in industry, business and science. *Commun. ACM* 37 (3), 93–105. <http://dx.doi.org/10.1145/175247.175257>, URL <https://dl.acm.org/doi/10.1145/175247.175257>.
- Wortsman, M., Ilharco, G., Gadre, S.Y., Roelofs, R., Gontijo-Lopes, R., Morcos, A.S., Namkoong, H., Farhadi, A., Carmon, Y., Kornblith, S., et al., 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In: International Conference on Machine Learning. PMLR, pp. 23965–23998.
- Wu, J., Mu, T., Thiyyagalingam, J., Goulermas, J.Y., 2020. Building interactive sentence-aware representation based on generative language model for community question answering. *Neurocomputing* 389, 93–107, Publisher: Elsevier.
- Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J., 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. <http://dx.doi.org/10.48550/arXiv.1609.08144>, arXiv:1609.08144 [cs]. URL <http://arxiv.org/abs/1609.08144>.
- Wu, X., Xia, Y., Zhu, J., Wu, L., Xie, S., Qin, T., 2022. A study of BERT for context-aware neural machine translation. *Mach. Learn.* 111 (3), 917–935, Publisher: Springer.

- Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L., 2021. CvT: Introducing convolutions to vision transformers. In: 2021 IEEE/CVF International Conference on Computer Vision. ICCV, pp. 22–31. <http://dx.doi.org/10.1109/ICCV48922.2021.00009>, ISSN: 2380-7504. URL <https://ieeexplore.ieee.org/document/9710031?denied=>.
- Wu, H., Zhang, J., Wang, L., Li, D., Kong, D., Han, Y., 2025. A survey on quantum deep learning. *J. Supercomput.* 81 (4), 564. <http://dx.doi.org/10.1007/s11227-025-07083-3>.
- Xia, W., Zhang, R., Zhang, X., Usman, M., 2023. A novel method for diagnosing Alzheimer's disease using deep pyramid CNN based on EEG signals. *Heliyon* 9 (4), e14858. <http://dx.doi.org/10.1016/j.heliyon.2023.e14858>, URL <https://www.sciencedirect.com/science/article/pii/S2405844023020650>.
- Xiao, T., Singh, M., Mintun, E., Darrell, T., Dollár, P., Girshick, R., 2021. Early convolutions help transformers see better. *Adv. Neural Inf. Process. Syst.* 34, 30392–30400.
- Xie, Q., Bishop, J.A., Tiwari, P., Ananiadou, S., 2022. Pre-trained language models with domain knowledge for biomedical extractive summarization. *Knowl.-Based Syst.* 252, 109460, Publisher: Elsevier.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K., 2017. Aggregated residual transformations for deep neural networks. <http://dx.doi.org/10.48550/arXiv.1611.05431>, arXiv: 1611.05431 [cs]. URL <http://arxiv.org/abs/1611.05431>.
- Xiong, W., Xiong, Z., Cui, Y., 2022. An explainable attention network for fine-grained ship classification using remote-sensing images. *IEEE Trans. Geosci. Remote Sens.* 60, 1–14. <http://dx.doi.org/10.1109/TGRS.2022.3162195>, Conference Name: IEEE Transactions on Geoscience and Remote Sensing. URL <https://ieeexplore.ieee.org/abstract/document/9741720>.
- Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X., 2018. AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1316–1324.
- Yan, Y., Liu, F., Zhuang, X., Ju, J., 2023. An R-Transformer-BiLSTM model based on attention for multi-label text classification. *Neural Process. Lett.* 55 (2), 1293–1316. <http://dx.doi.org/10.1007/s11063-022-10938-y>.
- Yan, C., Xu, J., Wang, Q., Peng, Z., 2025. A triple-inertial accelerated alternating optimization method for deep learning training. <http://dx.doi.org/10.48550/arXiv.2503.08489>, arXiv:2503.08489 [cs]. URL <http://arxiv.org/abs/2503.08489>.
- Yang, Y., Lan, T., Wang, Y., Li, F., Liu, L., Huang, X., Gao, F., Jiang, S., Zhang, Z., Chen, X., 2024. Data imbalance in cardiac health diagnostics using CECG-GAN. *Sci. Rep.* 14 (1), 14767. <http://dx.doi.org/10.1038/s41598-024-65619-8>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-024-65619-8>.
- Yang, J., Shen, X., Xing, J., Tian, X., Li, H., Deng, B., Huang, J., Hua, X.-s., 2019. Quantization networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. CVPR*.
- Yang, J., Soltan, A.A., Eyre, D.W., Clifton, D.A., 2023. Algorithmic fairness and bias mitigation for clinical machine learning with deep reinforcement learning. *Nat. Mach. Intell.* 5 (8), 884–894, Publisher: Nature Publishing Group UK London.
- Yang, B., Xiang, X., Kong, W., Zhang, J., Peng, Y., 2024. DMF-GAN: Deep multi-modal fusion generative adversarial networks for text-to-image synthesis. *IEEE Trans. Multim.* 1–13. <http://dx.doi.org/10.1109/TMM.2024.3358086>, Conference Name: IEEE Transactions on Multimedia. URL <https://ieeexplore.ieee.org/document/10413630>.
- Ye, F., Yang, J., 2021. A deep neural network model for speaker identification. *Appl. Sci.* 11 (8), <http://dx.doi.org/10.3390/app11083603>, URL <https://www.mdpi.com/2076-3417/11/8/3603>.
- Yin, J., Jiang, X., Lu, Z., Shang, L., Li, H., Li, X., 2016. Neural generative question answering. In: Iyyer, M., He, H., Boyd-Graber, J., Daumé III, H. (Eds.), *Proceedings of the Workshop on Human-Computer Question Answering*. Association for Computational Linguistics, San Diego, California, pp. 36–42. <http://dx.doi.org/10.18653/v1/W16-0106>, URL <https://aclanthology.org/W16-0106/>.
- Ying, H., Song, M., Tang, Y., Xiao, S., Xiao, Z., 2024. Enhancing deep neural network training efficiency and performance through linear prediction. *Sci. Rep.* 14 (1), 15197. <http://dx.doi.org/10.1038/s41598-024-65691-0>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-024-65691-0>.
- Yu, B., Xu, Z.-b., Li, C.-h., 2008. Latent semantic analysis for text categorization using neural network. *Knowl.-Based Syst.* 21 (8), 900–904. <http://dx.doi.org/10.1016/j.knosys.2008.03.045>, URL <https://www.sciencedirect.com/science/article/pii/S0950705108000993>.
- Zagoruyko, S., Komodakis, N., 2017. Wide residual networks. <http://dx.doi.org/10.48550/arXiv.1605.07146>, arXiv:1605.07146 [cs] version: 4. URL <http://arxiv.org/abs/1605.07146>.
- Zeiler, M.D., Fergus, R., 2013. Visualizing and understanding convolutional networks. <http://dx.doi.org/10.48550/arXiv.1311.2901>, arXiv:1311.2901 [cs] version: 3. URL <http://arxiv.org/abs/1311.2901>.
- Zhang, C., Chen, J., Li, J., Peng, Y., Mao, Z., 2023. Large language models for human-robot interaction: A review. *Biomim. Intell. Robot.* 3 (4), 100131. <http://dx.doi.org/10.1016/j.birob.2023.100131>, URL <https://www.sciencedirect.com/science/article/pii/S2667379723000451>.
- Zhang, X., Han, Y., Xu, W., Wang, Q., 2021. HOBA: A novel feature engineering methodology for credit card fraud detection with a deep learning architecture. *Inform. Sci.* 557, 302–316. <http://dx.doi.org/10.1016/j.ins.2019.05.023>, URL <https://www.sciencedirect.com/science/article/pii/S00202551930427X>.
- Zhang, J., Liu, A., Gao, M., Chen, X., Zhang, X., Chen, X., 2020. ECG-based multi-class arrhythmia detection using spatio-temporal attention-based convolutional recurrent neural network. *Artif. Intell. Med.* 106, 101856. <http://dx.doi.org/10.1016/j.artmed.2020.101856>, URL <https://www.sciencedirect.com/science/article/pii/S0933365719312606>.
- Zhang, W.E., Sheng, Q.Z., Alhazmi, A., Li, C., 2020. Adversarial attacks on deep-learning models in natural language processing: A survey. *ACM Trans. Intell. Syst. Technol. (TIST)* 11 (3), 1–41.
- Zhang, Z., Xie, Y., Yang, L., 2018. Photographic text-to-image synthesis with a hierarchically-nested adversarial network. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 6199–6208.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N., 2017. StackGAN: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 5907–5915.
- Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N., 2018. StackGAN++: Realistic image synthesis with stacked generative adversarial networks. *IEEE Trans. Pattern Anal. Mach. Intell.* 41 (8), 1947–1962, Publisher: IEEE.
- Zhang, C.-J., Yuan-Lu, Tang, F.-Q., Cai, H.-P., Qian, Y.-F., Chao-Wang, 2024. Heart failure classification using deep learning to extract spatiotemporal features from ECG. *BMC Med. Inform. Decis. Mak.* 24 (1), 17. <http://dx.doi.org/10.1186/s12911-024-02415-4>.
- Zhang, W., Zhang, T.-H., Luo, C., Zhou, H., Yang, C., Qian, X., Yin, X.-C., 2025. Breaking through the spike: Spike window decoding for accelerated and precise automatic speech recognition. <http://dx.doi.org/10.48550/arXiv.2501.03257>, arXiv:2501.03257 [eess]. URL <http://arxiv.org/abs/2501.03257>.
- Zhang, H.-B., Zhang, Y.-X., Zhong, B., Lei, Q., Yang, L., Du, J.-X., Chen, D.-S., 2019. A comprehensive survey of vision-based human action recognition methods. *Sensors* 19 (5), 1005.
- Zhang, R., Zong, Q., Dou, L., Zhao, X., 2019. A novel hybrid deep learning scheme for four-class motor imagery classification. *J. Neural Eng.* 16 (6), 066004.
- Zhao, Y., He, L., 2015. Deep learning in the EEG diagnosis of Alzheimer's disease. In: Jawahar, C., Shan, S. (Eds.), *Computer Vision - ACCV 2014 Workshops*. Springer International Publishing, Cham, pp. 340–353. [http://dx.doi.org/10.1007/978-3-319-16628-5\\_25](http://dx.doi.org/10.1007/978-3-319-16628-5_25).
- Zhao, M., Wei, Y., Lu, Y., Wong, K.K.L., 2020. A novel U-Net approach to segment the cardiac chamber in magnetic resonance images with ghost artifacts. *Comput. Methods Programs Biomed.* 196, 105623. <http://dx.doi.org/10.1016/j.cmpb.2020.105623>, URL <https://www.sciencedirect.com/science/article/pii/S0169260720314565>.
- Zhao, G., Yang, W., Ren, X., Li, L., Wu, Y., Sun, X., 2022. Well-classified examples are underestimated in classification with deep neural networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 36, pp. 9180–9189, Issue: 8.
- Zhou, H., Zhang, S., Peng, J., Zhang, S., Li, J., Xiong, H., Zhang, W., 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. <http://dx.doi.org/10.48550/arXiv.2012.07436>, arXiv:2012.07436 [cs]. URL <http://arxiv.org/abs/2012.07436>.
- Zhu, M., Pan, P., Chen, W., Yang, Y., 2019. DM-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5802–5810.
- Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J., 2020. Deformable DETR: Deformable transformers for end-to-end object detection. URL <https://arxiv.org/abs/2010.04159v4>.
- Zhu, Y., Wang, S., Yu, H., Li, W., Tian, J., 2024. SFPL: Sample-specific fine-grained prototype learning for imbalanced medical image classification. *Med. Image Anal.* 97, 103281.
- Zhu, W., Wu, O., Yang, N., 2024. IRDA: Implicit data augmentation for deep imbalanced regression. *Inform. Sci.* 120873.
- Zhu, F., Ye, F., Fu, Y., Liu, Q., Shen, B., 2019. Electrocardiogram generation with a bidirectional LSTM-CNN generative adversarial network. *Sci. Rep.* 9 (1), 6734. <http://dx.doi.org/10.1038/s41598-019-42516-z>, Publisher: Nature Publishing Group. URL <https://www.nature.com/articles/s41598-019-42516-z>.
- Zhu, X., Zhu, Y., Zhang, L., Chen, Y., 2023. A BERT-based multi-semantic learning model with aspect-aware enhancement for aspect polarity classification. *Appl. Intell.* 53 (4), 4609–4623. <http://dx.doi.org/10.1007/s10489-022-03702-1>.
- zhu wenliang wenliang, w., Chen, X., Wang, Y., Wang, L., 2019. Arrhythmia recognition and classification using ECG morphology and segment feature analysis. *IEEE/ACM Trans. Comput. Biol. Bioinform.* 16 (1), 131–138. <http://dx.doi.org/10.1109/TCBB.2018.2846611>.
- Zhuang, F., Qi, Z., Duan, K., Xi, D., Zhu, Y., Zhu, H., Xiong, H., He, Q., 2020. A comprehensive survey on transfer learning. *Proc. IEEE* 109 (1), 43–76, Publisher: IEEE.
- Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 67 (2), 301–320. <http://dx.doi.org/10.1111/j.1467-9868.2005.00503.x>.