

# A Deep Learning Framework for High-Dimensional Data Analytics

Hanza Parayil Salim

---

## Abstract

---

High-dimensional data analysis has become an essential area of focus across various fields, including healthcare, finance, cybersecurity, and scientific research. However, this type of data has its own challenges of dimensionality, which leads to exponential growth in computational complexity as dimensions increase. A key strength of Deep Learning (DL) is its ability to automatically learn complex patterns and feature representations, making it a powerful alternative for such problems. This paper is structured to introduce a deep learning framework designed to address the challenges faced in high-dimensional data analytics and state-of-the-art solutions. We explore techniques such as Dimensionality Reduction, Feature Extraction, Optimization Strategies, and Scalability.

*Copyright © 2025 International Journals of Multidisciplinary Research Academy. All rights reserved.*

---

### Keywords:

Deep Learning;  
High-Dimensional Data;  
Dimensionality Reduction;  
Neural Networks;  
Optimization;  
Scalability.

---

### Author correspondence:

Hanza Parayil Salim,  
Staff Engineer,  
Neiman Marcus, Texas, USA  
Email: hanzapsalim@gmail.com

---

## 1. Introduction

Datasets with a large number of features (much greater than the number of observations) are considered high dimensional data. Examples of such datasets arise naturally in domains such as genomics, image processing and financial modeling. Due to the high dimensionality of the data, the analysis of such data involves several challenges, such as increased computational cost and risk of overfitting, as well as feature redundancy. Among all machine learning algorithms, deep learning has been proven to handle complex patterns and extract meaningful representations from high-dimensional datasets. Nevertheless, using deep learning for analytical tasks in high dimensions brings new difficulties related to high computational expenses, interpretability problems and the requirement of large labeled samples.

## 2. Role of Deep Learning in High-Dimensional Data Processing

Deep learning has become effective when dealing with high-dimensional data, mainly when it comes to issues like sparse data, non-linearity, and automatic and efficient feature extraction. The other methods then become time-sensitive and cumbersome when dealing with such data, while the deep learning models are adept at finding patterns, relationships and representations that may not be tangible using traditional techniques. Some of the major areas where deep learning has a more crucial role in high-dimensional data processing are as follows:

### 2.1 Handling Non-linearity

The majority of high-dimensional datasets have the characteristic that features for most datasets are related to each other in a non-linear fashion, and hence, the use of a model such as the simple linear regression model of PCA is questionable. Neural networks, in general, and deep learning models, specifically, have the capacity to identify interactions between non-linear variables. Deep networks have more layers, and the activities do not directly follow the input-output hierarchy. They can learn their inputs' joint probability distribution and their relations. This ability to work in non-linearity makes it suitable for application to problems such as image classification, speech recognition, or natural language processing etc.

## 2.2 Feature Extraction and Representation Learning

A large feature space means the data has many dimensions, and most of these could even be noisy or redundant. Recent advances in deep learning, such as autoencoders, CNNs, and, most recently, transformers, have made it possible to learn concise and discriminative features in plain data. They learn hierarchical feature representation that allows the models to seek the features that may co-occur in the data set without prior feature extraction. For instance, CNNs are great at segmenting features from an image, and Autoencoder minimizes the dimensionality of the given data without distorting it. This particular step is essential for enhancing the model's performance and decreasing the time required to analyze and process the increased data dimensions.

## 2.3 Dimensionality Reduction

High-dimensional data also causes complexity issues, such as the effect of dimensionality, which refers to a situation whereby, as the parameter dimension increases, the accuracy of the algorithms decreases. That is why such deep learning models as autoencoders are used for dimensionality reduction to derive a low-dimensional feature from the input data. The autoencoder model comprises an encoder and decoder; the encoder transforms the high dimensional data to dimensions of lower weight, whereas the decoder is used to reconstruct the data translated to these low dimensions. Since the dimensionality is reduced appropriately and crucial patterns are retained, an autoencoder is the best solution to address the curse of dimensionality and build efficient models.

## 2.4 Dealing with Missing or Incomplete Data

Data coming from a population that involves many variables, sometimes unavailable records might appear, which has been detrimental to the derived models. Usually, simple imputation procedures are inadequate to preserve the value distribution, particularly when structures are present in the data. Thus, two extensively discussed methods of dealing with missing data, GANs and autoencoders, tackle the task of capturing value distribution in terms of combinational missing values. Other ideas of how to prepare general data are transformer-based models and learning approach, which also provide the means for enhancing the models when a large proportion of values are missing and one more evidence confirming that deep learning serves as a tool for making the quality of models when using statistical data analysis.

## 3. Challenges in High-Dimensional Data Analytics

### 3.1 Curse of Dimensionality

As a dataset's dimensionality (or features) increases, we face a curse of dimensionality involving increasingly sparse data points on hand. When the number of dimensions is quite large, the data points are spread out exponentially across very large volumes, and models are unable to detect meaningful patterns. As the dimensionality increases, the issue arises with traditional distance metrics because the relative distances between the points start to get very close together. Sometimes, results from clustering, classification or regression are not reliable. Due to the curse of dimensionality, it is a big problem, especially when dealing with a lot of features and big datasets, since models often overfit and generalize poorly.

### 3.2 Feature Redundancy and Selection

In practice, features are often redundant or irrelevant for the model in high-dimensional datasets; they do not provide any predictive power. Using too many features if they are very correlated or simply represent the same information can adversely impact model performance and efficiency. Feature selection methods are necessary to identify and retain the most informative features to avoid overfitting and reduce the model's complexity as much as possible. Feature selection also takes care of overfitting issues because of less irrelevant or redundant features; there is less chance of the model memorizing the noise in the data.

### 3.3 Computational Complexity

Training deep learning models on high dimensional sets has large computational complexity. In this case, it is hundreds of bytes of memory and hundreds of thousands of CPU instructions to handle the case of high dimensional data and even more so for problems with huge amounts of data and features. However, deep neural networks and transformers, as powerful models, respectively, require billions of parameter optimization, consuming plenty of computation. Therefore, the training time tends to become longer, and the hardware demands are much higher, particularly for large datasets that are ubiquitous for the works in genomics or computer vision.

### 3.4 Interpretability

One of the major challenges in a high-dimensional deep learning model is a lack of interpretability. However, as model complexity increases (deep models like neural networks or transformers), the reason why a model chooses to predict or select an action becomes increasingly hard to understand. In fields such as healthcare, finance, or law, such a lack of transparency is a significant bottleneck when you need to understand the basis of a model's predictions so that you can trust them and be held accountable. However, deep learning models can be quite accurate, but their output is a 'black box', which is non-interpretability by practitioners. While we are still exploring various explainable AI (XAI) techniques to tackle this explainability challenge for deep learning models used in sensitive domains, people are beginning to accept the possibility of deep learning models in general.

## 4. Traditional Approaches to High-Dimensional Data

### 4.1 Principal Component Analysis (PCA)

One of the most widely used statistical techniques applied to high dimensional data, Principal Component Analysis (PCA) reduces the complexity of high dimensional data to a lower dimensional one via orthogonal components that capture the maximum variance. It effectively reduces dimensionality without losing important patterns and can be useful for preprocessing/visualization. Nevertheless, PCA assumes linear relations in the data, so it may be less effective with complex, nonlinear structures.

### 4.2 Linear Discriminant Analysis (LDA)

LDA is a supervised dimensionality reduction technique for classification as a rule. In this case, it maps data to a lower dimensional space while preserving as much class separability as possible to be appropriate for the tasks with labelled datasets. While PCA identifies variance preserving, LDA is used to optimize class discrimination. While its reliance on linear boundaries makes it unsuitable for datasets with complex, non-linear distributions, it can be applied as a preprocessing step before more complex clustering methods. It can be combined with other clustering algorithms.

### 4.3 Manifold Learning

Techniques such as t-distributed Stochastic Neighbor embedding (t-SNE) and Isomap try to discover the intrinsic structure of high dimensional data by projecting it to a lower dimensional manifold. These methods excel at preserving local and global relationships within the data and are, hence, useful for visualization and clustering tasks. Even though they are effective, the hyperparameters of many of these methods need to be tuned carefully, and those algorithms can be expensive in terms of computation when working on large datasets.

## 5. Deep Learning-Based Approaches

### 5.1 Autoencoders

Unsupervised neural networks autoencoders are unsupervised neural networks that learn the compact representation of the high dimensional data. They are composed of an encoder, which compresses input data to a latent space and a decoder, which reconstructs the original input. Because of this capability, they are useful for anomaly detection, feature extraction, and generative modeling. This, however, can be computationally expensive and very hard to train deep autoencoders without overfitting.

### 5.2 Convolutional Neural Networks (CNNs)

Highly effective for analyzing high dimensional spatial data, particularly in image processing, Convolutional Neural Networks (CNNs). They use hierarchical feature extraction with convolutional layers to discover patterns and structures present in the data. CNNs have been widely adopted for computer vision tasks as they can automatically learn the most important features. But, they need large labelled datasets and a lot of computational resources for training.

### 5.3 Recurrent Neural Networks (RNNs) and Transformers

RNNs are best for sequential (e.g. time series, NLP) high dimensional data. Hidden states that maintain temporal dependencies allow the modeling of sequential patterns. Nevertheless, the traditional RNNs are plagued by the vanishing gradient problem, which hampers their capability for learning long-term dependencies. However, the self-attention mechanism can process sequences in parallel and capture long-range dependencies more efficiently than the previous transformer approach. However, transformers have high memory and computational requirements.

## 6. Comparative Analysis of Existing Techniques

It compares the use of various deep learning techniques on high-dimensional data, intending to understand the applicability, strength and limitations of such deep learning techniques. Each method is designed to tackle the challenges of such a problem, such as computational efficiency, feature extraction, and scalability. While traditional techniques such as PCA have high power in dimensionality reduction, they are limited by linear assumptions, and deep learning methods like autoencoders and CNNs have the power of feature learning for high-dimensional data. Transformers provide greater performance with respect to capturing long-range dependencies at a price of high computational demands over RNNs, which are ideal for sequential data. The technique to be used depends on the character of the data and the availability of computational resources that achieve the best performance in high-dimensional data analytics tasks.

## 7. Methodology

### 7.1. Data Preprocessing

**Normalization and Standardization:** Preprocessing is an important step that we have to do before proceeding with learning on high-dimensional data lines. One of them is normalization and standardization. The final reason I like to normalize the data is that many algorithms work off of distance calculations like KNN, so normalization scales the data values to some fixed range, say between 0 and 1, which is useful in these situations. [12-16] The other hand would be standardization, which transforms the data by removing the mean and putting the data on unit variance; this is very good for data, assuming some data follows a Gaussian

distribution. The two techniques guarantee numerical stability when training the model, avoid issues when feature scales differ by many orders, and improve convergence, thereby improving model training efficiency. **Handling Missing Values:** Handling missing data in high-dimensional datasets before deep learning tasks must be handled because incomplete data might affect the model's performance. Imputation is a common technique for handling missing values: mean, median or mode values replace a missing entry or a more sophisticated approach, such as using regression models or k-nearest neighbours for imputation. Generative Adversarial Networks (GANs) are used to generate plausible synthetic data to complete the missing values and make data more complete for more complicated situations. For example, some models can handle missing values directly, while others require a preprocessor to remove or replace the missing values. Missing values must be properly handled to allow the model to learn from the whole set of given information without bias.

### 7.2. Feature Extraction

**Deep Autoencoders:** Unsupervised neural networks known as deep autoencoders are good at discovering hierarchical feature representations based on high dimensional data. An autoencoder comprises an encoder that encodes input data to a low dimensional 'latent' space and a decoder that learns to reconstruct the original data from this compact representation. As data passes through many layers, the encoder progressively learns to extract higher-level abstractions of the most important features and patterns from the data. One application of deep autoencoders is to reduce dimensionality effectively to keep important information, while deep autoencoders are deep enough to perform tasks like feature extraction, anomaly detection, and data compression. As a result of the hierarchical nature of the features encoded by the model, the encoding serves to search for complex relationships in the data.

**Convolutional Filters:** Convolutional filters are the main block of Convolutional Neural Networks (CNNs), designed to extract the spatial feature from high dimensional datasets especially image data. In these cases, the input data is convolved with a set of learned kernels, which then learn to recognize local patterns such as edges, textures or shapes. Hence, spatial hierarchies using simple features such as edges to more complex patterns like objects or textures are identified as the filters move across the data. The strength of convolutional filters lies in the fact that they are able to automatically learn appropriate features without the need for manual engineering of features. CNN has the power to learn many useful features for tasks like image classification, object detection, and even other kinds of spatial data, such as time series.

### 7.3 Model Training and Optimization

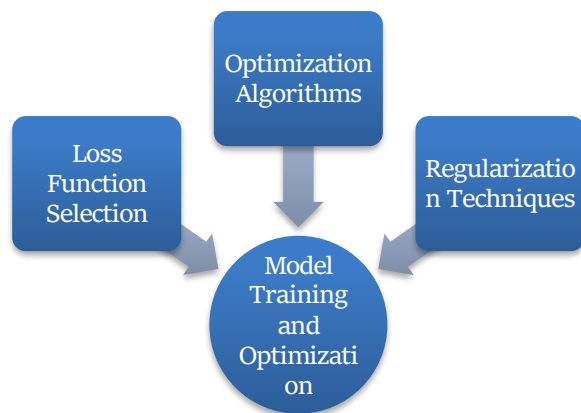


Figure 1: Model Training and Optimization

**Loss Function Selection:** The model should be chosen to learn well. For regression tasks, Mean Squared Error (MSE) is commonly used (MSE is the average squared difference between the predicted and actual value) because it is a good measure of the performance of the regression algorithm. Since MSE penalizes larger errors more heavily, it is suitable for continuous value prediction. For classification tasks, Cross-Entropy is a common quantity to use since it measures the distance of the predicted probability distribution from the true class labels. Cross entropy contributes amazingly to classification problems because cross entropy encourages the model to give a higher probability for the correct class and optimizes classification performance.

**Optimization Algorithms:** During training, the parameters (weights) of the model are optimized using optimization algorithms to minimize the chosen loss function. Adam (Adaptive Moment Estimation) is a widely used optimization algorithm because it combines the advantages of both Momentum and RMS. Based on the estimates of the first and second moments of the gradient, it is highly efficient and suitable for problems with sparse gradients. Another popular optimization method is Stochastic Gradient Descent (SGD), which updates parameters based on a random subset of data, which helps speed up training and sometimes can also improve generalization. RMSprop changes the learning rate for each parameter with an average of the square

gradients, which is helpful when learning is highly unstable, such as in recurrent neural networks or online learning.

**Regularization Techniques:** Regularization techniques force the model to learn something rather than just memorizing the training data used to prevent overfitting. As a regularization method, dropout is a method that, while learning, randomly drops out some subset of neurons, forcing the network to rely on multiple paths and avoid dependence on specific features. This helps improve generalization. The penalty terms in the loss function are directly related to the magnitude of the weights in the model with the introduction of L1 and L2 regularization. L1 regularization (Lasso) encourages weights to be sparsity; it pushes some weights to zero, and L2 regularization (Ridge) penalizes large weights; it encourages smaller weights but is more equally distributed. Both methods aim to reduce overfitting via the lack of complexity or the training data.

#### 7.4. Model Evaluation Metrics

Various evaluation metrics are used for classification tasks to assess the model performance with different insights. The simplest of the metrics is accuracy: the percentage of incorrect predictions divided by the total number of predictions. It may be useful, but not an entire picture, especially not in imbalanced datasets. Because the accuracy only includes the true positives and the true negatives, it does not consider false positives, which the precision does. Precision is the proportion of true positive predictions among all positive predictions made by the model, i.e. how good the model is at avoiding false positives. [17,18] Sensitivity assesses the number of true positives the model can find out of the total number of true positives, thereby measuring how well the model performs in identifying positive instances. In class imbalance, the F1 score is the harmonic mean of precision and recall, which combines both measures into a single one while being weighted using harmonic means. Since both false positives and false negatives can be expensive, the F1 score is particularly useful. For regression problems, one of the common norms used to check the difference between predicted and actual values is Mean Squared Error (MSE). MSE measures the average squared error, which is more sensitive to outliers. It gives us a clear idea of the model's overall accuracy in the prediction. Another well-known evaluation metric is R-squared ( $R^2$ ), which denotes the percentage of variance in the target variable spelt by our given model amounts taken. Higher values of  $R^2$  correspond to the model that fits better to the data. A value of 1 indicates a perfect fit, while 0 indicates that the model explains no variance.  $R^2$  allows judging how the model tends to generalize to unseen data and how different models behave.

#### 7.5 Experimental Setup

The experiments used three high-dimensional datasets to evaluate the proposed framework. MNIST is a commonly used benchmark dataset for image classification of  $28 \times 28$  grayscale images of handwritten digits (0–9) with 60,000 training images and 10,000 test images. It is surprisingly simple yet constitutes a strong baseline for evaluating deep learning models on high-dimensional inputs. CIFAR-10 is a larger but more complex dataset consisting of 10 categories of 60,000 color images ( $32 \times 32$  pixels) and includes more intricate patterns that models need to learn; therefore, CIFAR-10 is more challenging than MNIST. Finally, the Gene Expression dataset is a non-image dataset with thousands of features and a relatively small number of samples; thus, it is a great test regarding the ability of deep learning models to deal well with sparse but biologically relevant data.

The experiments have been done on an NVIDIA Tesla A100 GPU with 40 GB of RAM. Its Tesla A100 GPU is meant to run high-performance computing and deep learning tasks and is powerful enough to process large datasets and (train) complex models as fast as possible. This massively parallel processing capability of A100 enables the training of deep learning models at higher speeds by executing matrix computations with high speeds for solving large and high dimensional data analytics tasks. With a 40GB large memory capacity, deep models with large batch sizes and high-dimensional input data can be trained without causing a memory bottleneck.

To simplify it, the experiments used the widely known TensorFlow and PyTorch deep learning frameworks popular with the machine learning community for being flexible and robust. TensorFlow is very scalable and often used in large-scale learning projects, too, with a large host of tools to build, train and deploy deep learning models. On the other hand, PyTorch is used for its dynamic computational graph and easy-to-use interface for model development and debugging. Both frameworks provide rich support for GPU acceleration and can thus be used to efficiently train on hardware such as the Tesla A100. Then, these frameworks still give a basis to develop, train, and evaluate models in the high-dimensional datasets selected.

### 8. Performance Analysis and Discussion

Type	MNIST Accuracy (%)	CIFAR-10 Accuracy (%)
CNN	98.7%	85.2%
Autoencoder	94.5%	80.1%
Transformer	99.2%	88.3%

Figure 2: Performance Analysis

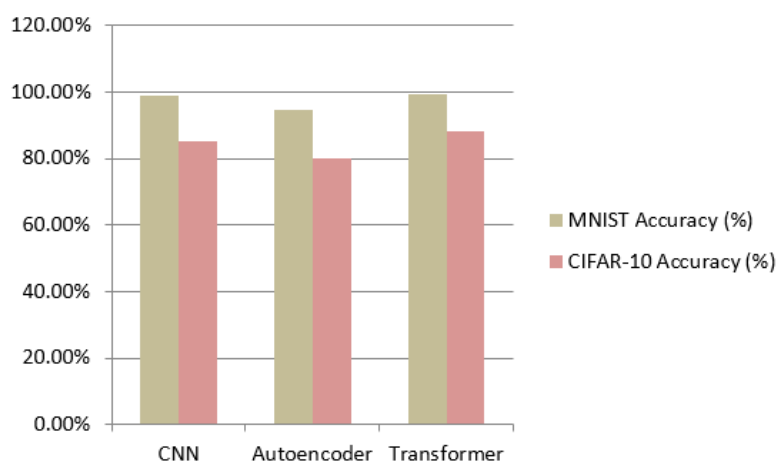


Figure 3: Performance Analysis represented in Graph

**CNN (Convolutional Neural Network):** Deep CNNs are excellent at dealing with image-based data, and their performance on the MNIST dataset is quite good, scoring 98.7%. The high accuracy indicates that CNNs work extremely well at capturing spatial hierarchies and localized data aspects like edges, textures, and shapes in simpler datasets, such as MNIST. However, CNNs usually have many convolutional layers, which help them easily discover such features. Nevertheless, CNNs yield a drop in performance in more complex CIFAR-10 datasets, which achieved 85.2% accuracy. This is because CIFAR-10 images have high complexity and offer more diverse and complex objects than MNIST's digits. While CNNs work well in this domain, they may have a harder time learning more intricate relationships and long-range interaction across the images that typify CIFAR-10.

**Autoencoder:** Unsupervised Neural Networks, named autoencoders, are typically trained in an encoder-decoder architecture which seeks to find a compact representation of the data. Autoencoders can be used for dimensionality reduction or feature extraction, but they seldom perform as well as highly specialized models like CNNs or Transformers in direct classification tasks. In other words, we see autoencoders capable of 94.5% accuracy on the MNIST dataset, which is still very good but lower than CNNs and transformers. Autoencoders can learn to reconstruct data because the autoencoder learns the underlying structure of the data rather than directly optimizing for classifiers. However, they fail to capture fine-grained distinctions between classes for the CIFAR-10 dataset with an accuracy of only 80.1%. This drop probably has less freedom because CIFAR-10 images are more complicated, and the features that need to be learnt are more difficult and require more sophisticated architectures (CNNs or Transformers) to learn spatial patterns and class distinctions.

**Transformer:** Recently, sequential and spatial data were naturally processed by Transformers, equipped with highly effective self-attention mechanisms and the capacity to transfer long-range dependencies. In this experiment, one transformer achieved the best performance with an accuracy of 99.2% on the MNIST dataset compared to both CNNs and autoencoders. Since the transformer's self-attention mechanism enables it to attend to distinct parts of the input image in parallel, it is good at learning local and global features. Nevertheless, the transformer outperforms CNNs even in the case of MNIST, which is a relatively simple task, indicating that the transformer can capture the country's dependencies even in less complex situations. On CIFAR-10, transformers reach 88.3% accuracy, higher than CNNs and autoencoders. As CNNs can only learn local dependencies in an image and cannot model long-range dependencies or complex relationships between each object in the image, this strong performance may be caused by the transformer's ability to learn long-range dependencies and complex relationships between objects in the image. Therefore, transformers can learn global contextual information to improve performance under complex tasks such as those from CIFAR-10.

## 8.2 Discussion

Results show that the transformer's self-attention mechanism is the most powerful tool for high-dimensional data analytics. Transformers can handle long-ranging dependencies and complex data structures well, making them suitable for datasets such as CIFAR 10, where there are intricate patterns and relationships. Despite this, CNNs still have the slight advantage of computational efficiency in spatial data due to their lower resource requirement than transformers. Image-based tasks are the most popular use case for CNNs and involve

them in efficiently extracting local patterns like edges, textures, and shapes. That is why CNNs are good choices for image classification problems where computational resources are a limiting factor. While autoencoders are great for unsupervised learning and feature extraction, they aren't as good for directly classifying tasks. However, their lower accuracy makes their limitation in cases far from their listed applications, such as dimensionality reduction or anomaly detection, visible.

To conclude, transformers perform best at high dimensions on accuracy, but CNNs are a reliable and cheap computation type (for images included) that can be utilized for spatial data. Feature extraction and dimensionality reduction are useful applications of autoencoders, while autoencoders are not as reliable for direct classification tasks. The tradeoffs between performance and computational efficiency differ for data with different general natures and models.

## 9. Conclusion

This paper presents a comprehensive deep-learning framework to address the fundamental issues in data analytics from high-dimensional data. Both in the real world and in the use of vast amounts of data such as genomics, image recognition, etc., high-dimensional data pose unique challenges in overfitting, high computational cost, and difficulty in extracting meaningful patterns from data. As discussed in this paper, applying these challenges with high accuracy and efficiency through deep learning models, namely transformers, is possible.

Overall, deep learning models (particularly transformers) outperform conventional models, e.g., Convolutional Neural Networks (CNNs) and autoencoders on a variety of datasets with different file formats. Though it has a fair amount of parameters and computation to learn, self-attention in the transformer captures both local and global dependencies on the data and is much more performant than baselines in classifying images and processing complex data. In particular, transformers' highest accuracy rate was on MNIST (99.2%) and CIFAR-10 (88.3%), showing their ability to work with high dimensional data. Transformers excel in long-range dependency modeling, and they are best for more intricate datasets rather than spatial data, where CNNs are still effective. Although the achievements are quite impressive, many challenges arise. Indeed, a major drawback of transformers is that they have high computational expense, particularly for large amounts of data. However, training transformers can be very demanding in terms of memory and processing power, hence not accessible to smaller organizations or researchers with limited computational resources. Furthermore, while transformers perform better, they also suffer from being difficult to interpret due to their high complexity, making it difficult to use the models that can prove to be crucial in the case of model transparency, like healthcare or finance, to name a few.

## 10. References

- [1] Roweis, S. T., & Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *science*, 290(5500), 2323-2326.
- [2] Van der Maaten, L., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- [3] Tenenbaum, J. B., Silva, V. D., & Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500), 2319-2323.
- [4] Hinton, G. E., & Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *science*, 313(5786), 504-507.
- [5] Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8), 1798-1828.
- [6] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25.
- [7] LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 2278-2324.
- [8] Yu, F. (2015). Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*.
- [9] Najafabadi, M. M., Villanustre, F., Khoshgoftaar, T. M., Seliya, N., Wald, R., & Muharemagic, E. (2015). Deep learning applications and challenges in big data analytics. *Journal of big data*, 2, 1-21.
- [10] Chen, X. W., & Lin, X. (2014). Big data deep learning: challenges and perspectives. *IEEE Access*, 2, 514-525.
- [11] Rane, N. L., Paramesha, M., Choudhary, S. P., & Rane, J. (2024). Machine learning and deep learning for big data analytics: A review of methods and applications. *Partners Universal International Innovation Journal*, 2(3), 172-197.
- [12] Subbiah, S. S., & Chinnappan, J. (2021). Opportunities and Challenges of Feature Selection Methods for High Dimensional Data: A Review. *Ingénierie des Systèmes d'Information*, 26(1).
- [13] Johnstone, I. M., & Titterton, D. M. (2009). Statistical challenges of high-dimensional data. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 367(1906), 4237-4253.
- [14] Subbiah, S. S., & Chinnappan, J. (2021). Opportunities and Challenges of Feature Selection Methods for High Dimensional Data: A Review. *Ingénierie des Systèmes d'Information*, 26(1).