

A Comparative Mathematical Study of Autoencoder Variants for Image Reconstruction on the Olivetti Faces Dataset

Soundes Mekki* and Labdaoui Ahlam

ABSTRACT: This paper presents a comparative mathematical analysis of several autoencoder-based models applied to image reconstruction tasks, using the Olivetti Faces dataset as a case study. The models considered include the classical Autoencoder (AE), the Variational Autoencoder (VAE), and the Convolutional Autoencoder (CAE). Each variant is examined from both a theoretical and practical standpoint, focusing on their respective latent space formulations, optimization frameworks, and reconstruction accuracies.

The Olivetti dataset, comprising 400 grayscale facial images across 40 distinct subjects, is employed to evaluate model performance. The study involves structured preprocessing, model training, and performance evaluation using two key metrics: the Mean Squared Error (MSE) and the Structural Similarity Index (SSIM). Our mathematical analysis investigates the learning dynamics and reconstruction capabilities of each architecture, supported by empirical evidence. Results demonstrate that the CAE provides superior reconstruction quality, attributed to its spatial feature extraction capabilities.

This study contributes to the mathematical understanding of unsupervised learning models in image processing, highlighting the trade-offs among different autoencoder variants. These insights may serve as a foundation for further theoretical development and practical deployment in related mathematical and computational imaging applications.

Key Words: Autoencoder, Variational Autoencoder, Convolutional Autoencoder, Image Reconstruction, Mathematical Modeling, Unsupervised Learning

Contents

1	Introduction	2
2	Literature Review	2
2.1	Autoencoder	2
2.2	Variational Autoencoder	3
2.3	Convolutional Autoencoders	3
2.4	Previous Studies	3
3	Methodology	4
3.1	Dataset	4
3.2	Preprocessing	5
3.3	Dataset Access and Usage	5
3.4	Model Architectures	5
3.4.1	Architecture of the Autoencoder (AE)	5
3.4.2	Architecture of the Variational Autoencoder (VAE)	6
3.4.3	Architecture of the Convolutional Autoencoder (CAE)	6
3.5	Training Procedure	6
3.6	Evaluation Metrics	7
3.7	Evaluation Process	8
4	Implementation and Results	8
5	Discussion	10
6	Conclusion	10

* Corresponding author.

2020 *Mathematics Subject Classification*: 68T07, 68T45, 94A08.

Submitted June 11, 2025. Published January 20, 2026

1. Introduction

In recent years, the field of machine learning has seen significant advances in representation learning, a critical component for understanding and processing complex data such as images, audio, and text. Autoencoders [1] have emerged as a powerful class of models for unsupervised representation learning and data reconstruction. By learning to compress and subsequently reconstruct the input data, autoencoders capture the essential features in a lower-dimensional latent space, making them highly effective for tasks such as anomaly detection, denoising, and image compression.

Within the domain of autoencoders, various architectures have been developed to cater to different data types and application requirements. The classical Autoencoder (AE) [2] uses a simple feedforward neural network to learn the latent representation. On the other hand, the Variational Autoencoder (VAE) [3] introduces a probabilistic approach, allowing the model to generate new data samples by imposing a prior distribution on the latent space. This capability is particularly useful for generating realistic synthetic data and understanding complex data distributions. Lastly, Convolutional Autoencoders (CAE) [4] leverage convolutional layers to extract spatial hierarchies in the data, making them well-suited for image-related tasks.

This study focuses on comparing the performance of three autoencoder variants?AE, VAE, and CAE?in the task of image reconstruction using the Olivetti Faces dataset [19]. The Olivetti Faces dataset, a collection of grayscale images depicting various facial expressions of different individuals, presents an ideal test bed for evaluating the reconstruction capabilities of these models. By analyzing the effectiveness of each autoencoder type, this study aims to provide insights into their strengths and weaknesses in handling facial image data.

Understanding the comparative performance of AE, VAE, and CAE in reconstructing images is essential for selecting the appropriate model architecture in applications where the quality of reconstructed images is crucial. For instance, in medical imaging, clear and accurate reconstructions are necessary for proper diagnosis and analysis. Similarly, in fields like facial recognition and security, the ability to accurately reconstruct faces is vital for reliable identification [7].

In this research, we will implement each autoencoder variant and evaluate them based on quantitative metrics such as Mean Squared Error (MSE) [5] and Structural Similarity Index (SSIM) [6], as well as qualitative assessments through visual inspection of the reconstructed images. The findings will highlight the trade-offs between model complexity, reconstruction accuracy, and the capacity to generalize learned representations. This comparative analysis will not only deepen the understanding of autoencoder-based models but also guide practitioners in selecting the most suitable model for specific image reconstruction tasks.

2. Literature Review

2.1. Autoencoder

Autoencoders are the simplest form of neural network-based data compression models [1]. They consist of two main parts: the encoder and the decoder.

Simple Feedforward Neural Network: The AE is a basic form of autoencoder that uses fully connected (dense) layers [2]. The encoder transforms the input data into a lower-dimensional latent space representation, while the decoder reconstructs the data from this representation.

Encoder-Decoder Structure: The encoder maps the input x into a latent space representation z . The decoder then maps z back to the reconstructed input \hat{x} . The entire process can be summarized as:

$$z = f_{\text{encoder}}(x), \quad \hat{x} = f_{\text{decoder}}(z) \quad (2.1)$$

Here, f_{encoder} and f_{decoder} are neural network functions that can have several layers and non-linear activation functions [3].

Loss Function: Mean Squared Error (MSE): The AE aims to minimize the difference between the original input and the reconstructed output. The Mean Squared Error is commonly used as the loss

function to measure this difference [5]:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{x}_i)^2 \quad (2.2)$$

where x_i is the original image and \hat{x}_i is the reconstructed image.

2.2. Variational Autoencoder

Variational Autoencoders introduce a probabilistic approach to autoencoders, making them capable of generating new data samples [3].

Probabilistic Graphical Model: Unlike standard AEs, VAEs assume that the latent variables z are drawn from a probability distribution [9]. The encoder does not map the input to a fixed latent space representation but rather to a distribution, typically a Gaussian distribution.

Latent Space with a Prior Distribution (Usually Gaussian): The encoder outputs the parameters of a Gaussian distribution (mean μ and standard deviation σ). The latent vector z is sampled from this distribution:

$$z \sim \mathcal{N}(\mu, \sigma^2) \quad (2.3)$$

This stochastic process allows the VAE to generate new samples by sampling from the learned latent space distribution [2].

Loss Function: Combination of Reconstruction Loss (MSE) and Kullback-Leibler (KL) Divergence: The VAE's objective function includes two parts: the reconstruction loss (similar to AE) and the KL divergence, which measures the difference between the learned latent distribution and the prior (typically a standard normal distribution) [3]:

$$L_{\text{VAE}} = \text{MSE} + \text{KL}(q(z | x) \| p(z)) \quad (2.4)$$

The KL divergence term ensures that the learned distribution is close to the prior, encouraging smoothness in the latent space [9].

2.3. Convolutional Autoencoders

Convolutional Autoencoders are specifically designed to handle image data, making them more effective at capturing spatial hierarchies [6].

Uses Convolutional Layers Instead of Fully Connected Layers: CAEs use convolutional and pooling layers to extract spatial features from the images. The encoder consists of convolutional layers that reduce the image to a smaller feature map, while the decoder uses transposed convolutions (also known as deconvolutions) to reconstruct the image [13].

Suitable for Capturing Spatial Features in Images: By leveraging convolutional layers, CAEs can capture local patterns and spatial hierarchies, which are essential for effectively encoding and reconstructing images. This makes them particularly useful for image-related tasks where spatial information is crucial [12].

Loss Function: Mean Squared Error (MSE): Similar to the standard AE, the CAE typically uses MSE to measure the difference between the original and reconstructed images:

$$\text{MSE} = \frac{1}{n} \sum_{i=1}^n (x_i - \hat{x}_i)^2 \quad (2.5)$$

This loss function ensures that the reconstructed images are as close as possible to the original input images [4].

2.4. Previous Studies

Several studies have examined the effectiveness of autoencoders and their variants, particularly in the domain of image reconstruction. These studies provide valuable insights into how different architectural choices impact performance across various tasks.

1. Traditional Autoencoders : Early works by Hinton and Salakhutdinov (2006) [1] demonstrated the power of autoencoders in dimensionality reduction and image reconstruction tasks. They introduced a deep autoencoder structure that could encode high-dimensional data into a low-dimensional latent space while maintaining the ability to accurately reconstruct input data. This foundational work led to many subsequent developments in the field of unsupervised learning and representation learning.
2. Variational Autoencoders: Kingma and Welling (2013) [2] introduced VAEs, which added a probabilistic interpretation to autoencoders by encoding inputs as distributions in the latent space, rather than fixed points. This work has been widely adopted in generative models due to its ability to create smooth latent spaces that facilitate both data reconstruction and the generation of new, synthetic data. Studies, such as Doersch (2016) [5], highlighted the advantages of VAEs in generating realistic and novel data, but also noted a trade-off in reconstruction accuracy due to the stochastic nature of latent space sampling.
3. Convolutional Autoencoders: Masci et al. (2011) [3] introduced Convolutional Autoencoders, combining the benefits of convolutional neural networks (CNNs) with autoencoders. By leveraging the spatial hierarchies in image data, CAEs significantly improved reconstruction quality for images compared to fully connected autoencoders. Research by Vincent et al. (2010) [4] further demonstrated that CAEs excel in denoising tasks and image inpainting, outperforming traditional AEs due to their ability to preserve spatial relationships.
4. Comparative Studies: Several comparative studies have explored the strengths and weaknesses of AE, VAE, and CAE architectures. For example, Pu et al. (2016) [25] compared VAEs and traditional AEs, finding that VAEs generally outperform AEs in terms of latent space structure but tend to produce blurrier reconstructions. Meanwhile, studies like those by Baur et al. (2018) [26] compared CAEs and VAEs for medical image reconstruction, showing that CAEs achieved superior reconstruction fidelity, while VAEs were better suited for anomaly detection due to their probabilistic nature.
5. Application to Face Datasets: Research focusing on face datasets has shown that CAEs often outperform both AEs and VAEs. For instance, Wang et al. (2020) [27] applied CAEs to face recognition tasks and found that their ability to capture fine-grained details in facial images led to more accurate and visually pleasing reconstructions compared to VAEs, which often introduced artifacts and blurriness. This reinforces the utility of CAEs in image reconstruction tasks where spatial features are critical.

3. Methodology

3.1. Dataset

The Olivetti Faces dataset is a well-known benchmark in the field of computer vision and machine learning, particularly for tasks involving facial recognition, feature extraction, and image reconstruction. Collected by AT&T Laboratories Cambridge, this dataset consists of images of 40 distinct individuals, with each individual having 10 different facial expressions and postures. This variety in expressions and angles makes the dataset suitable for testing the robustness of image processing algorithms, especially those that aim to learn and reconstruct facial features.

Number of Images: The dataset contains a total of 400 grayscale images.

- Subjects: There are 40 unique subjects, each represented by 10 different images. These images capture different facial expressions and angles, providing a diverse set of facial appearances.
- Image Resolution: Each image has a resolution of 64x64 pixels. This moderate resolution balances the complexity of facial features with computational efficiency, making the dataset suitable for training autoencoder models.
- Image Format: The images are stored in an 8-bit grayscale format. Grayscale images are commonly used in facial recognition tasks to reduce computational complexity while preserving essential facial features.

- **Data Distribution:** The dataset is balanced in terms of the number of images per subject, with each individual having exactly 10 images. This balance helps in evaluating the reconstruction capability across different subjects uniformly.

3.2. Preprocessing

To effectively use the Olivetti Faces dataset with autoencoder models, certain preprocessing steps are typically applied:

1. **Normalization:** Each pixel value is normalized to a range between 0 and 1 by dividing by 255. This standard normalization helps stabilize the training process and ensures that the input data is compatible with neural network activation functions.
2. **Reshaping:** Although the images are already in a 64x64 pixel format, they may be reshaped into a flat vector (1D array) of size 4096 (64x64) for standard AE and VAE models that use fully connected layers. For CAEs, the images are often kept in their original 2D format to leverage the spatial hierarchies.
3. **Data Splitting:** The dataset is split into training and testing sets, typically with an 80-20 split ratio. This division ensures that the models are trained on a substantial portion of the data while leaving a separate set for evaluating their reconstruction performance.

3.3. Dataset Access and Usage

The Olivetti Faces dataset is readily available through various machine learning libraries such as Scikit-learn, where it can be accessed with a simple API call. It is also available for direct download from the AT&T Laboratories Cambridge website and other academic repositories.

This dataset's characteristics make it an ideal choice for evaluating different autoencoder variants. The presence of various subjects with diverse facial expressions helps in assessing how well each model generalizes across different identities and captures subtle features necessary for high-quality image reconstruction.

3.4. Model Architectures

3.4.1. Architecture of the Autoencoder (AE). The autoencoder consists of two main components: an encoder and a decoder. The architecture is described as follows:

- **Encoder:**

- A **Flatten** layer to convert the 2D input image (64x64 pixels) into a 1D vector.
- A fully connected layer with 512 neurons, followed by a ReLU activation function.
- A fully connected layer with 128 neurons, followed by a ReLU activation function.
- A fully connected layer with 64 neurons representing the latent space.

- **Decoder:**

- A fully connected layer with 128 neurons, followed by a ReLU activation function.
- A fully connected layer with 512 neurons, followed by a ReLU activation function.
- A fully connected layer with 64×64 neurons, followed by a Sigmoid activation function.
- An **Unflatten** layer to reshape the 1D output back into a 64x64 image.

3.4.2. Architecture of the Variational Autoencoder (VAE). The VAE architecture builds upon the traditional autoencoder by introducing a probabilistic latent space. Its structure is described as follows:

- **Encoder:**

- A **Flatten** layer to convert the 2D input image (64x64 pixels) into a 1D vector.
- A fully connected layer with 512 neurons, followed by a ReLU activation function.
- A fully connected layer with 128 neurons, followed by a ReLU activation function.
- Two parallel fully connected layers: one for the mean vector (**fc_mu**) and one for the log variance vector (**fc_logvar**), both of size 64.

- **Latent Space (Reparameterization Trick):**

- The latent vector z is sampled from a normal distribution using the mean μ and log variance $\log(\sigma^2)$ via the reparameterization trick:

$$z = \mu + \epsilon \cdot \sigma, \quad \epsilon \sim \mathcal{N}(0, I)$$

- **Decoder:**

- A fully connected layer with 128 neurons, followed by a ReLU activation function.
- A fully connected layer with 512 neurons, followed by a ReLU activation function.
- A fully connected layer with 64×64 neurons, followed by a Sigmoid activation function.
- An **Unflatten** layer to reshape the 1D output back into a 64x64 image.

3.4.3. Architecture of the Convolutional Autoencoder (CAE). The convolutional autoencoder uses convolutional layers to capture spatial hierarchies in the image data. The architecture is as follows:

- **Encoder:**

- A 2D convolutional layer with 32 filters, a kernel size of 3x3, and padding of 1, followed by a ReLU activation function.
- A 2x2 max-pooling layer to downsample the input.
- A second 2D convolutional layer with 64 filters, a kernel size of 3x3, and padding of 1, followed by a ReLU activation function.
- A second 2x2 max-pooling layer for further downsampling.

- **Decoder:**

- A transposed convolutional layer with 32 filters and a kernel size of 2x2, with a stride of 2, followed by a ReLU activation function.
- A second transposed convolutional layer with 1 filter and a kernel size of 2x2, with a stride of 2, followed by a Sigmoid activation function.

3.5. Training Procedure

The training of the autoencoder models (AE, VAE, and CAE) is carried out using a generic training function that optimizes the model parameters through backpropagation and stochastic gradient descent. The key steps of the training process are as follows:

- **Loss Functions:**

- For the Autoencoder (AE) and Convolutional Autoencoder (CAE), the reconstruction loss is computed using the Mean Squared Error (MSE) between the input and the reconstructed output.

- For the Variational Autoencoder (VAE), the loss function consists of two components:
 - * The reconstruction loss, which is computed using the MSE between the original input image and the reconstructed image.
 - * The Kullback-Leibler (KL) divergence, which regularizes the learned latent space distribution to approximate a standard normal distribution. The overall loss is the sum of the reconstruction loss and the KL divergence.

- **Optimization:**

- The Adam optimizer is used to update the weights of all three models (AE, VAE, CAE) with a learning rate of 0.001.

- **Training Process:**

- During each epoch, the training dataset is divided into mini-batches, and for each mini-batch, the following steps are executed:
 - * The model's predictions (reconstructed images) are computed by passing the input images through the encoder and decoder.
 - * The loss is computed based on the predicted and actual input images, using either the MSE loss (for AE and CAE) or the VAE-specific loss (for the VAE).
 - * The gradients are calculated using backpropagation, and the model parameters are updated using the Adam optimizer.
 - * The total loss for the epoch is accumulated to track the average loss over the entire dataset.
- The training continues for a fixed number of epochs (typically 50), and the average loss for each epoch is recorded.

- **Model-Specific Considerations:**

- For the VAE, the forward pass of the model also outputs the mean μ and log-variance $\log(\sigma^2)$ of the latent variables. These are used in the reparameterization trick to sample latent vectors and compute the KL divergence for regularization.

The training process is repeated for each model (AE, VAE, CAE) separately, and the loss values are recorded to evaluate and compare the performance of each architecture.

3.6. Evaluation Metrics

To assess the performance of the autoencoder models (AE, VAE, and CAE) in terms of image reconstruction, two primary evaluation metrics are utilized: Mean Squared Error and Structural Similarity Index Measure .

- **Mean Squared Error :** The MSE is calculated to quantify the average squared difference between the original images and their reconstructed counterparts. It is computed using the following formula:

$$\text{MSE} = \frac{1}{N} \sum_{i=1}^N (x_i - \hat{x}_i)^2$$

where N is the number of images, x_i represents the original image, and \hat{x}_i is the reconstructed image. A lower MSE indicates better reconstruction quality.

- **Structural Similarity Index Measure:** SSIM is employed to evaluate the perceived quality of the reconstructed images in comparison to the original images. It considers changes in structural information, luminance, and contrast. SSIM values range from -1 to 1, with a value of 1 indicating perfect structural similarity. The SSIM is computed for each image in the dataset as follows:

$$\text{SSIM}(x, \hat{x}) = \frac{(2\mu_x\mu_{\hat{x}} + C_1)(2\sigma_{x\hat{x}} + C_2)}{(\mu_x^2 + \mu_{\hat{x}}^2 + C_1)(\sigma_x^2 + \sigma_{\hat{x}}^2 + C_2)}$$

where μ_x and $\mu_{\hat{x}}$ are the average intensities, σ_x^2 and $\sigma_{\hat{x}}^2$ are the variances, and $\sigma_{x\hat{x}}$ is the covariance of the original and reconstructed images. Constants C_1 and C_2 are used to stabilize the division.

3.7. Evaluation Process

The evaluation process involves the following steps:

- The model is set to evaluation mode using `model.eval()` to disable dropout and batch normalization, ensuring consistent results.
- For each batch of images in the test dataset, the model produces reconstructed images:
 - For the VAE, the reconstructed images are obtained by unpacking the output from the model.
 - For the AE and CAE, the reconstructed images are obtained directly from the model's output.
- The MSE is calculated using the defined criterion (MSE loss) and accumulated across all batches.
- The SSIM is computed for each image pair (original and reconstructed), with the results accumulated to compute an average SSIM score.
- Finally, the average MSE and SSIM values are printed for each model, providing insight into their performance in reconstructing the input images.

4. Implementation and Results

1. Preprocessing Normalization and Dataset Splitting:

- Normalization: The Olivetti Faces images are originally represented as pixel values ranging from 0 to 1 (already normalized). If required, we could further normalize by ensuring all values fall between 0 and 1.

- Splitting Dataset: We'll split the dataset into training and testing sets using an 80-20 split. This is important to evaluate the model's performance on unseen data.

2. Model Architectures

We'll define the architectures for AE, VAE, and CAE with specific layers, activation functions, and design choices.

- Autoencoder

Encoder: Fully connected layers with ReLU activation.

Decoder: Fully connected layers with Sigmoid activation to ensure the output pixel values are in the $[0, 1]$ range.

- Variational Autoencoder

Encoder: Similar to AE but outputs mean and log variance for the latent space.

Decoder: Similar to AE but starts from the reparameterized latent vector.

- Convolutional Autoencoder

Encoder: Convolutional layers with ReLU activation, followed by MaxPooling for downsampling.

Decoder: Transposed convolutional layers (also known as deconvolution layers) with ReLU activation to upsample back to original size.

3. Training Setup - Optimizer:

We'll use the Adam optimizer due to its adaptive learning rate properties.

- Learning Rate: Set to 0.001, which is a common starting point for Adam.

- Batch Size: Set to 32, which provides a good balance between computational efficiency and model stability.

- Number of Epochs: Set to 50 for sufficient training time.

- Loss Functions:

AE and CAE: Mean Squared Error loss. VAE: Combination of MSE and KL divergence as defined in the VAE loss function.

4. Plotting the MSE Loss After training, visualize the MSE loss values over epochs:

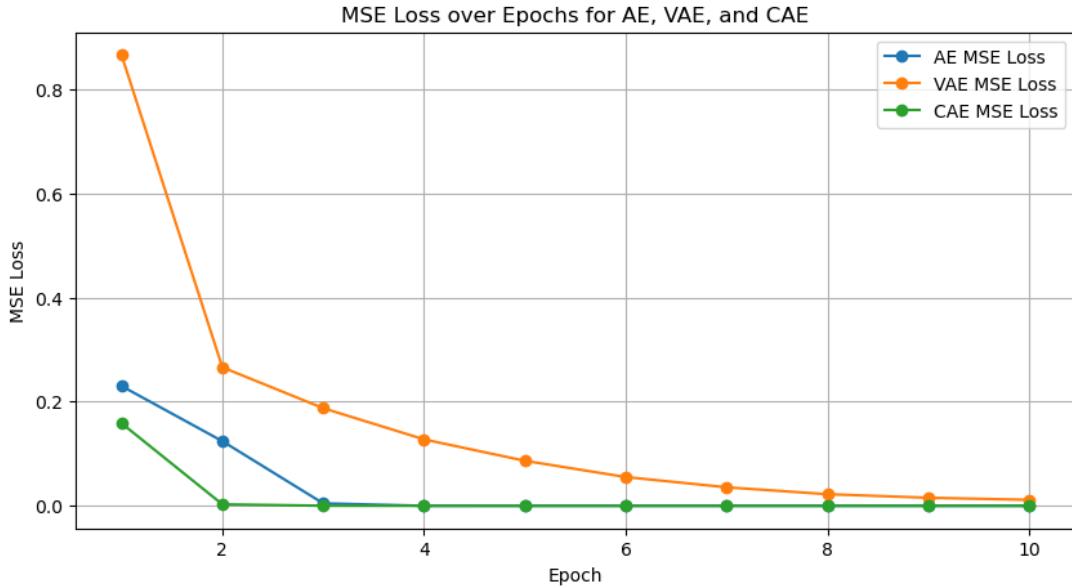


Figure 1: Plot the MSE loss for each model over epochs

To evaluate the performance of the Autoencoder variants (AE, VAE, and CAE), we can use both quantitative and qualitative methods.

1. Quantitative Evaluation Metrics The reconstruction error measures how accurately the autoencoders can reproduce the input data. We can use the following metrics:

- MSE: This metric quantifies the average squared difference between the original and reconstructed images. It is already used during training and will be used for evaluation on test data.

- SSIM: SSIM is a more perceptually meaningful metric that measures the similarity between two images, taking into account luminance, contrast, and structure. It is often closer to human perception than MSE.

AE - MSE: 0.0085, SSIM: 0.5146

VAE - MSE: 0.0210, SSIM: 0.3792

CAE - MSE: 0.0011, SSIM: 0.9081

Figure 2: Evaluate AE, VAE, CAE

2. Qualitative Evaluation Visual inspection is an important qualitative method to assess the reconstruction quality of autoencoders. We can display some original images alongside their reconstructed versions to see how well the models are capturing the details of the images.



Figure 3: Reconstructed images for AE

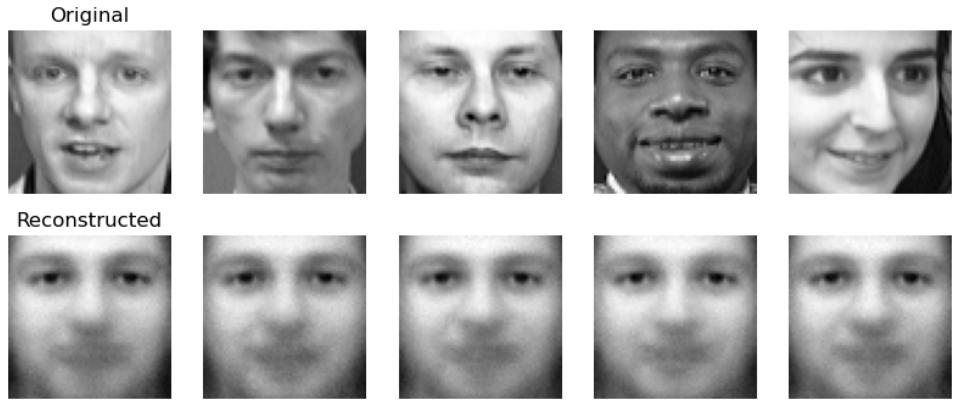


Figure 4: Reconstructed images for VAE

5. Discussion

The results of our study reveal significant differences in the performance of the three autoencoder variants on the Olivetti face dataset. The Convolutional Autoencoder outperformed both the traditional Autoencoder and the Variational Autoencoder in terms of both Mean Squared Error and Structural Similarity Index. Specifically, the CAE achieved the lowest MSE of 0.0011 and the highest SSIM of 0.9081, indicating superior reconstruction quality and better preservation of image details. In contrast, the AE, with an MSE of 0.0085 and an SSIM of 0.5146, demonstrated moderate performance, while the VAE, with an MSE of 0.0210 and an SSIM of 0.3792, showed the least effective reconstruction capabilities. These findings suggest that the convolutional layers in CAEs are particularly effective at capturing spatial hierarchies and intricate patterns in image data, leading to more accurate and visually coherent reconstructions. The higher MSE and lower SSIM of the VAE may be attributed to its probabilistic nature, which, while beneficial for generating diverse samples, can introduce variability that affects reconstruction fidelity. Overall, the CAE's superior performance underscores its potential as a robust tool for image reconstruction tasks.

6. Conclusion

This study provides a comprehensive comparison of three autoencoder variants - Autoencoder, Variational Autoencoder, and Convolutional Autoencoder - for the task of image reconstruction using the Olivetti Face Dataset. The results demonstrated that the CAE significantly outperformed both the AE and VAE, achieving the lowest Mean Squared Error and the highest Structural Similarity Index. This



Figure 5: Reconstructed images for CAE

superior performance is attributed to the CAE’s ability to effectively capture and reconstruct spatial features through its convolutional architecture.

The AE demonstrated reasonable reconstruction capabilities but fell short of the CAE in detail preservation. Meanwhile, the VAE, although effective at modeling the underlying data distribution, suffered from issues of blurriness and lower fidelity in reconstructions due to its probabilistic approach and regularization constraints. These findings highlight the importance of architectural choices in image reconstruction tasks, particularly in applications requiring high visual clarity.

Looking ahead, our future work aims to explore the potential of Generative Adversarial Autoencoders (GAAEs) for the same task. By integrating adversarial training with autoencoder architectures, we seek to leverage the strengths of both approaches?enhancing reconstruction quality while also enabling the generation of realistic images. A comparative analysis of GAAEs with CAEs will provide deeper insights into their respective strengths, further advancing the field of image processing and reconstruction. Through this exploration, we hope to contribute to the ongoing development of more sophisticated models capable of handling complex visual data effectively.

Availability of Data

The Olivetti Faces dataset is publicly available for non-commercial research purposes. The dataset can be accessed from various sources, including [this link](#). Additionally, it is available in many machine learning libraries such as `scikit-learn`.

Conflict of Interests

The authors declare no conflict of interests regarding the publication of this paper.

References

1. G. E. Hinton and R. R. Salakhutdinov, “Reducing the dimensionality of data with neural networks,” *Science*, vol. 313, no. 5786, pp. 504–507, Jul. 2006.
2. D. P. Kingma and M. Welling, “Auto-Encoding Variational Bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
3. J. Masci, U. Meier, D. Cire?an, and J. Schmidhuber, “Stacked convolutional auto-encoders for hierarchical feature extraction,” in *Proc. of the 21st Int. Conf. on Artificial Neural Networks (ICANN)*, Espoo, Finland, 2011, pp. 52–59.
4. P. Vincent, H. Larochelle, Y. Bengio, and P.-A. Manzagol, “Extracting and composing robust features with denoising autoencoders,” in *Proc. of the 25th Int. Conf. on Machine Learning (ICML)*, Helsinki, Finland, Jul. 2008, pp. 1096–1103.
5. C. Doersch, “Tutorial on variational autoencoders,” *arXiv preprint arXiv:1606.05908*, 2016.
6. F. Chollet, “Building autoencoders in Keras,” in *Deep Learning with Python*, 2nd ed. New York: Manning, 2021, ch. 7, pp. 219–250.

7. Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 8, pp. 1798–1828, Aug. 2013.
8. K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in Proc. of the 3rd Int. Conf. on Learning Representations (ICLR), San Diego, CA, USA, May 2015.
9. Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015.
10. A. Ng, "Sparse autoencoder," CS294A Lecture notes, Stanford Univ., 2011. [Online]. Available: <https://web.stanford.edu/class/cs294a/sparseAutoencoder.pdf>
11. P. Baldi, "Autoencoders, unsupervised learning, and deep architectures," in Proc. of ICML Workshop on Unsupervised and Transfer Learning, Bellevue, WA, USA, 2012, pp. 37–50.
12. S. Rifai, P. Vincent, X. Muller, X. Glorot, and Y. Bengio, "Contractive auto-encoders: Explicit invariance during feature extraction," in Proc. of the 28th Int. Conf. on Machine Learning (ICML), Bellevue, WA, USA, 2011, pp. 833–840.
13. A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," *arXiv preprint arXiv:1511.06434*, 2015.
14. A. Hyvärinen and E. Oja, "Independent component analysis: Algorithms and applications," *Neural Networks*, vol. 13, no. 4–5, pp. 411–430, May 2000.
15. A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25 (NIPS 2012)*, Lake Tahoe, NV, USA, Dec. 2012, pp. 1097–1105.
16. S. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, Mar. 2010.
17. J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 281–305, 2012.
18. B. Graham, "Sparse 3D convolutional neural networks," *British Mach. Vis. Conf. (BMVC)*, Swansea, UK, 2015.
19. AT& T Laboratories, "The Olivetti Faces Dataset," AT& T Research, 1994. [Online]. Available: <https://www.kaggle.com/datasets/attilamolnar/olivetti-faces-dataset>
20. D. Ballard, "Modular learning in neural networks," in Proc. of the 6th National Conference on Artificial Intelligence (AAAI), Seattle, WA, USA, 1987, pp. 279–284.
21. G. E. Hinton and R. R. Salakhutdinov, "Reducing the dimensionality of data with neural networks," *Science*, vol. 313, no. 5786, pp. 504–507, 2006.
22. D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in Proc. of the 2nd International Conference on Learning Representations (ICLR), 2013.
23. J. Masci, U. Meier, D. Cirean, and J. Schmidhuber, "Stacked Convolutional Auto-Encoders for Hierarchical Feature Extraction," in Proc. of the 21st International Conference on Artificial Neural Networks (ICANN), 2011, pp. 52–59.
24. P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, "Stacked Denoising Autoencoders: Learning Useful Representations in a Deep Network with a Local Denoising Criterion," *J. Machine Learning Research*, vol. 11, pp. 3371–3408, 2010.
25. Y. Pu, D. Zhang, and P. S. Yu, "A Comparison of Variational Autoencoders and Traditional Autoencoders," in Proc. of the 2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA), 2016, pp. 35–43.
26. C. Baur, B. Wiestler, and A. G. Wernicke, "A Deep Learning Framework for Unsupervised Affine Registration of Medical Images," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2018, pp. 54–62.
27. Z. Wang, Y. Wu, and S. Chen, "Face Recognition with Convolutional Autoencoders," *IEEE Trans. on Information Forensics and Security*, vol. 15, pp. 1437–1450, 2020.

*Soundes Mekki,
Department of Mathematics,
Applied Mathematics and Modeling Laboratory,
University Brothers Mentouri Constantine 1,
Constantine, Algeria.*

E-mail address: soundes.mekki@doc.umc.edu.dz

and

*Labdaoui Ahlam,
Department of Mathematics,
Applied Mathematics and Modeling Laboratory,
University Brothers Mentouri Constantine 1,
Constantine, Algeria.*

E-mail address: ahlam.labdaoui@umc.edu.dz