International Journal of Scientific Engineering and Research (IJSER)
ISSN (Online): 2347-3878
SJIF (2025): 8.036

Investigating Dimensionality Reduction Strategies for Future Optical Neural Computing Systems

Nishtha Gandhi

B. Tech, IIT Delhi
Email: *es1240841[at]iitd.ac.in*

Abstract: ***Artificial intelligence workloads continue to grow in computational complexity, motivating research into alternative computing paradigms capable of delivering improved scalability and energy efficiency. Optical neural computing has emerged as a promising candidate due to the inherent parallelism of light and the possibility of performing large-scale linear operations at extremely high speed. However, practical realization of optical neural systems remains challenging because high-dimensional data representations often require a correspondingly large number of physical optical modes. This study investigates whether intelligent feature-compression techniques can reduce dimensionality while preserving classification performance. Using handwritten digit recognition as a benchmark task, four dimensionality reduction approaches are evaluated: direct resolution reduction, learned bottleneck compression, Principal Component Analysis (PCA), and autoencoder-based encoding. The work establishes quantitative dimensionality targets for future optical computing systems and provides the theoretical foundation for future Jacobi Time-Wave Packet implementations.***

Keywords: Optical Neural Computing, Machine Learning, Neural Networks, Feature Compression, Dimensionality Reduction, Principal Component Analysis (PCA), Autoencoders, Representation Learning, MNIST, Handwritten Digit Recognition, Optical Machine Learning, Jacobi Time-Wave Packets.

1. Introduction

The rapid advancement of machine learning has transformed numerous scientific and industrial domains. Modern neural networks routinely process high-dimensional data and require substantial computational resources. Consequently, the search for alternative hardware platforms has become increasingly important. Optical computing offers an attractive solution due to its capability for parallel information processing, low propagation losses, and potentially superior energy efficiency.

Despite these advantages, optical systems face practical limitations associated with representing high-dimensional data. Each dimension may correspond to a separate optical mode, channel, or degree of freedom. Therefore, understanding the minimum dimensionality required for accurate machine-learning performance is a critical research problem.

This work investigates dimensionality reduction from both a machine-learning and optical-computing perspective. Instead of focusing solely on maximizing classification accuracy, the study aims to identify compact representations that preserve useful information while minimizing dimensionality. Such representations may ultimately enable practical optical neural computing architectures.

2. Literature Review

Dimensionality reduction has been studied extensively in machine learning. Classical approaches such as Principal Component Analysis seek directions of maximum variance within a dataset. More recent approaches rely on neural networks, including autoencoders and learned bottleneck representations.

PCA remains one of the most widely used dimensionality reduction techniques due to its simplicity, interpretability, and strong empirical performance. Autoencoders extend this concept by learning nonlinear latent representations. These techniques have demonstrated success across image recognition, signal processing, and data compression applications.

In parallel, optical neural computing has gained attention as a potential accelerator for AI workloads. Recent work has explored diffractive optical neural networks, photonic matrix multipliers, and multimode optical systems. However, many studies assume high-dimensional optical representations without explicitly investigating how many dimensions are truly required. This gap motivates the present work.

3. Theoretical Background

Artificial neural networks consist of interconnected layers of neurons. Each neuron computes a weighted sum of inputs followed by a nonlinear activation function. The output of a neuron can be expressed as:

$$y = f(\Sigma w_i x_i + b)$$

where w_i represents trainable weights, x_i denotes input features, b is a bias term, and f(.) is the activation function.

Dimensionality reduction aims to transform high-dimensional data into a lower-dimensional representation while preserving informative characteristics. In PCA, data is projected onto orthogonal directions corresponding to maximum variance. The explained variance ratio is given by:

$$\text{Variance Ratio} = \lambda_i / \Sigma \lambda_j$$

where λ_i represents the eigenvalue associated with the i-th principal component.

Autoencoders learn compressed representations through an

International Journal of Scientific Engineering and Research (IJSER)
ISSN (Online): 2347-3878
SJIF (2025): 8.036

encoder-decoder architecture. The encoder maps the input into a latent space, while the decoder reconstructs the original data. Reconstruction quality is commonly measured using Mean Squared Error (MSE):

$$MSE = (1/N) \Sigma(x_i - \hat{x}_i)^2$$

These techniques are highly relevant to optical neural computing because they provide mechanisms for reducing the number of effective dimensions that must be physically represented within an optical system.

4. Methodology

The MNIST handwritten digit dataset was selected as the primary benchmark due to its widespread use in machine-learning research. Images were resized from 28×28 pixels to 14×14 pixels, producing 196-dimensional input vectors. A fully connected neural network architecture was trained using cross-entropy loss and gradient-based optimization.

Four experimental phases were conducted. Phase 1 evaluated direct resolution reduction. Phase 2 investigated learned bottleneck compression through constrained hidden layers. Phase 3 applied Principal Component Analysis at multiple dimensionality levels. Phase 4 utilized an autoencoder with an 8-dimensional latent space.

Performance was measured using validation accuracy. Additional analysis focused on identifying dimensionality thresholds beyond which significant information loss occurs. The results provide quantitative evidence regarding the minimum dimensionality required for effective handwritten digit classification and establish practical targets for future low-dimensional optical implementations.

Figure Placeholders

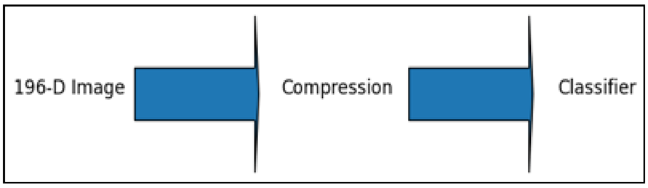

Figure 1: Conceptual pipeline investigated in this study

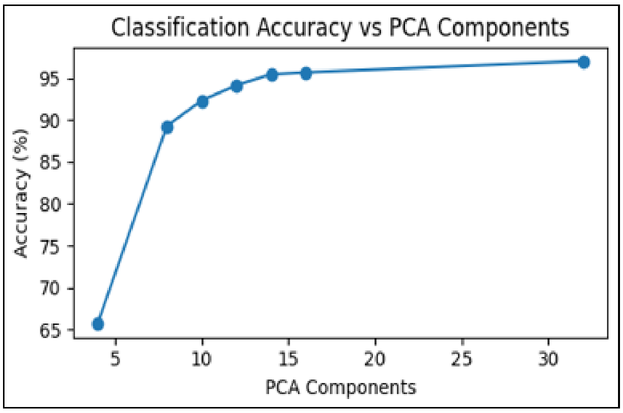

Figure 2: PCA dimensionality-reduction performance obtained during experimentation.

5. Results

The experimental evaluation revealed significant differences among the investigated dimensionality reduction techniques. Direct resolution reduction demonstrated that aggressive pixel removal rapidly degrades classification accuracy. While reducing dimensionality from 196 to 64 features produced only a modest decline in performance, further reduction to 16 and 9 dimensions resulted in substantial information loss.

Learned bottleneck compression consistently outperformed direct resolution reduction at equivalent dimensionalities. This observation indicates that neural networks can learn compact internal representations that preserve discriminative information more effectively than simple downsampling.

Principal Component Analysis produced the strongest overall performance. Remarkably, PCA maintained more than 95% accuracy using only 16 components, suggesting that the majority of useful information within handwritten digit images is concentrated within a relatively small subspace. Autoencoder-based compression also demonstrated strong performance but remained slightly inferior to PCA under the conditions investigated in this study.

6. Discussion

One of the most significant findings of this work is the existence of substantial redundancy within handwritten digit images. Although the original input representation contains 196 features, the results demonstrate that accurate classification can be achieved using a fraction of this dimensionality. This observation has important implications for machine learning and hardware implementation alike.

The superiority of PCA over autoencoders is particularly noteworthy. Several factors may contribute to this outcome. First, MNIST is a relatively simple dataset whose dominant structure can often be captured using linear projections. Second, the autoencoder architecture employed in this work utilized a compact latent space and limited training depth. More sophisticated architectures may yield improved performance but would introduce additional complexity.

Another important observation concerns dimensionality thresholds. Performance remains relatively stable between 16 and 12 dimensions but declines more rapidly below 10 dimensions. This suggests the existence of a critical operating region where information preservation remains effective. Identifying such thresholds is essential for future optical implementations because each retained dimension may correspond to a physical optical mode.

The comparison between raw reduction, bottleneck compression, PCA, and autoencoders highlights a broader principle: intelligent feature extraction is substantially more effective than naive dimensionality reduction. Simply discarding pixels removes information indiscriminately, whereas learned and statistical representations preferentially preserve informative features.

7. Implications for Optical Neural Computing

The results obtained in this work provide quantitative guidance for future optical neural computing systems. Practical optical implementations often require physical modes, waveguides, or channels to represent information. Consequently, reducing dimensionality directly translates into reduced system complexity.

A conventional implementation based on 196 input dimensions may require hundreds of optical modes. In contrast, the present study demonstrates that classification accuracy can remain above 94% using approximately 10–16 dimensions. Such reductions may significantly improve the feasibility of future optical architectures.

The findings also motivate exploration of Jacobi Time-Wave Packet representations. Unlike PCA, which is derived statistically from a dataset, Jacobi-based modes may provide physically realizable basis functions suitable for optical implementation. Establishing whether Jacobi modes can achieve performance comparable to PCA remains an important research direction.

8. Future Work and Conclusion

Future research will focus on integrating Jacobi Time-Wave Packet representations directly into the classification pipeline. Images will be projected onto low-order Jacobi modes and represented using a small number of coefficients. Classification performance will then be compared against PCA-based representations.

Additional investigations may include multimode optical simulations, optical signal propagation models, and experimental validation using photonic hardware. The dimensionality targets identified in this study provide a clear benchmark for such future efforts.

In conclusion, this work demonstrates that substantial dimensionality reduction is possible without severe degradation in classification performance. PCA achieved the strongest overall results, while learned bottleneck representations and autoencoders also preserved useful information. The identification of a practical 10–16 dimensional operating range establishes an important foundation for future low-dimensional optical neural computing architectures.

Figures

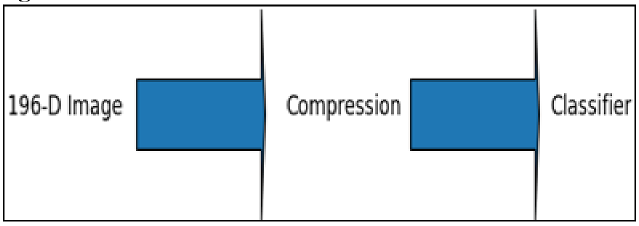

Figure 3: Comparative performance of compression techniques at low dimensionality.

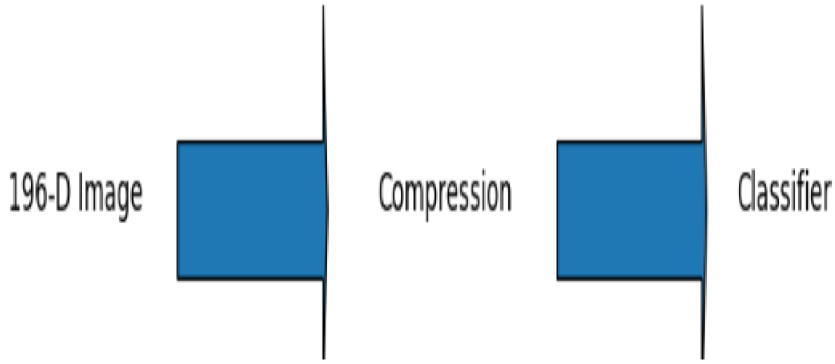

Figure 4: Information loss associated with direct resolution reduction.

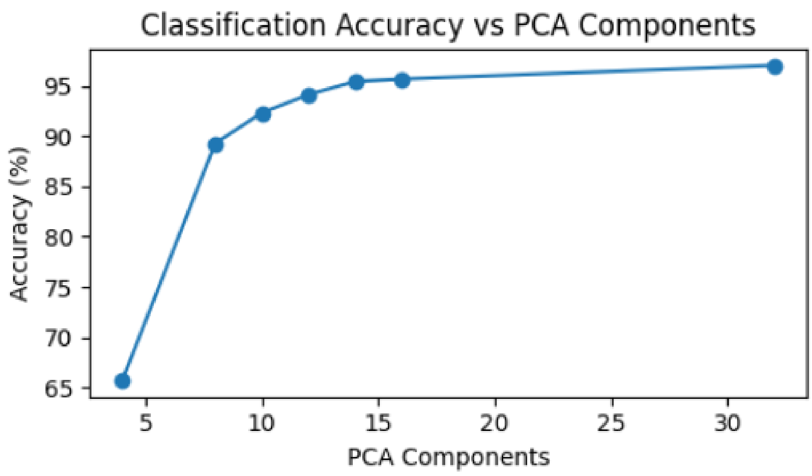

Figure 5: Proposed Jacobi Time-Wave Packet classification pipeline.

9. Limitations and Research Scope

While the present study provides valuable insight into dimensionality reduction for future optical neural computing systems, several limitations should be acknowledged. First, all experiments were conducted within a simulation environment using software-based neural networks rather than physical optical hardware. Consequently, practical implementation constraints such as optical noise, propagation losses, modal cross-talk, and fabrication tolerances were not considered.

Second, the study focused primarily on MNIST handwritten digit classification. Although MNIST remains a widely accepted benchmark, it is a relatively simple dataset. More complex datasets may exhibit different dimensionality requirements and could potentially alter the observed performance hierarchy between PCA, bottleneck compression, and autoencoder-based methods.

Third, the Jacobi Time-Wave Packet representation that motivated this investigation was not directly integrated into the classification pipeline. Instead, the present work establishes quantitative dimensionality targets that can guide future Jacobi-mode implementations. Future studies should investigate whether physically realizable optical basis functions can achieve performance comparable to PCA-derived features.

References

[1] Y. LeCun et al., 'Gradient-Based Learning Applied to Document Recognition,' Proceedings of the IEEE, 1998.
[2] I. T. Jolliffe, Principal Component Analysis, Springer, 2002.
[3] G. Hinton and R. Salakhutdinov, 'Reducing the Dimensionality of Data with Neural Networks,' Science, 2006.
[4] I. Goodfellow, Y. Bengio, and A. Courville, Deep Learning, MIT Press, 2016.
[5] D. Rumelhart et al., 'Learning Representations by Back-Propagating Errors,' Nature, 1986.
[6] X. Lin et al., 'All-Optical Machine Learning Using Diffractive Deep Neural Networks,' Science, 2018.
[7] D. Psaltis and F. Mok, 'Holographic Memories,' Scientific American, 1995.
[8] M. Reck et al., 'Experimental Realization of Any Discrete Unitary Operator,' PRL, 1994.
[9] A. Krizhevsky, 'Learning Multiple Layers of Features from Tiny Images,' 2009.
[10] J. Schmidhuber, 'Deep Learning in Neural Networks: An Overview,' Neural Networks, 2015.

Appendix

The appendix summarizes the principal experimental observations. PCA consistently achieved the highest accuracy at comparable dimensionality levels. A practical operating region of approximately 10–16 dimensions was identified, providing a quantitative target for future low-mode optical implementations. Autoencoder and bottleneck representations preserved substantially more information than naive pixel reduction, confirming the importance of intelligent feature extraction

Paper ID: SE26624170150 DOI: https://dx.doi.org/10.70729/SE26624170150