

Research Article

Next-Generation K-Means Clustering: Mojo-Driven Performance for Big Data

Touhidul Alam Seyam¹ , Md. Sazzad Hossain² , Rajib Ghose³ ,
Mekhriddin Nurmatov⁴ , Nazarov Fayzullo⁴ , Zarin Hadika⁵ ,
Abhijit Pathak^{5,*} 

¹Department of Computer Science and Engineering, Begum Gul Chemonara Trust University Bangladesh, Chattogram, Bangladesh

²Faculty of Intelligent Systems and Computer Technologies, Samarkand State University, Samarkand, Uzbekistan

³Department of Computer Science and Engineering, Military Institute of Science and Technology (MIST), Dhaka, Bangladesh

⁴Faculty of Artificial Intelligence and Information Systems, Samarkand State University, Samarkand, Uzbekistan

⁵Department of Computer Science and Engineering, Sonargaon University, Dhaka, Bangladesh

Abstract

K-means clustering, a fundamental unsupervised machine learning technique, is widely used in anomaly detection, image recognition, and customer segmentation. Traditional Python implementations, especially those using NumPy, face performance challenges with large, high-dimensional datasets due to Python's interpreted nature and dynamic typing. This paper introduces an innovative approach using the Mojo programming language, designed for AI development, to significantly improve the performance of the k-means clustering. Mojo combines Python's usability with the performance of system programming languages by offering features like vectorization, parallelization, and strong typing. We compare a NumPy-based Python implementation with an optimized Mojo implementation, detailing the translation process and optimization techniques, including Mojo's support for Single Instruction, Multiple Data (SIMD) operations, explicit memory management, and efficient data structures. These features significantly accelerate distance calculations crucial to the k-means algorithm. Benchmarks on synthetic datasets with varying sample sizes, feature counts, and cluster numbers demonstrate that the Mojo implementation consistently outperforms both the standard Python implementation and the highly optimized sci-kit-learn k-means, achieving speedups of 6x to 250x. These results highlight Mojo's potential as a powerful tool for high-performance data analysis, particularly for computationally demanding algorithms like k-means clustering, and contribute to high-performance computing in machine learning. This research sets the stage for further exploration of Mojo's applicability to other algorithms and hardware-specific optimizations for modern computing architectures.

Keywords

K-means Clustering, Mojo Programming Language, Scikit-learn, Machine Learning, NumPy

*Corresponding author: a.pathak_cse@su.edu.bd (Abhijit Pathak)

Received: 6 March 2025; **Accepted:** 18 March 2025; **Published:** 31 March 2025



Copyright: © The Author(s), 2025. Published by Science Publishing Group. This is an **Open Access** article, distributed under the terms of the Creative Commons Attribution 4.0 License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited.

1. Introduction

Unsupervised machine learning relies heavily on clustering, a fundamental technique that helps reveal underlying structures, relationships, and patterns in data without pre-assigned labels. K-means clustering is a well-known and widely used algorithm in this field due to its ease of use and effectiveness, with numerous applications across several domains, including:

Anomaly detection: Identifying data points that deviate substantially from defined clusters, such as fraudulent transactions, network intrusions, or manufacturing flaws [3].

- (1) Image Recognition
- (2) Customer Segmentation

Although Python has become popular for implementing k-means and other algorithms due to its rich ecosystem of machine learning libraries like NumPy [4] and scikit-learn [5], its interpreted nature frequently causes performance issues when working with large and complex contemporary datasets. These issues arise primarily from:

- (1) *Interpreted Execution:* Python code is executed line by line, introducing significant overhead compared to compiled languages that translate the entire codebase into machine instructions beforehand.
- (2) *Dynamic Typing:* Variables in Python do not require explicit type declarations, leading to runtime type checking that hinders performance, particularly in computationally intensive loops.
- (3) *Memory Management:* Python's dynamic memory allocation and garbage collection, while convenient, can introduce overhead and impact performance predictability, especially with large datasets.

These drawbacks are particularly noticeable in k-means clustering, as the algorithm's iterative nature and reliance on calculating distances between centroids and data points demand high processing efficiency. This work proposes a novel solution: a k-means algorithm implementation using the Mojo programming language [6]. Mojo, specifically designed for AI development, seamlessly combines Python's usability with the performance of systems programming languages like C++. Mojo achieves this by leveraging:

- (1) *Vectorization:* Performing operations on multiple data points simultaneously using Single Instruction, Multiple Data (SIMD) instructions, available on modern CPUs.
- (2) *Parallelization:* Distributing computations across multiple processor cores to accelerate execution, especially beneficial for large datasets and high cluster counts.
- (3) *Strong Typing:* Enforcing explicit type declarations allows for more efficient code generation and compiler optimization.
- (4) *Explicit Memory Management:* Providing greater control over memory allocation and access patterns, reducing overhead, and improving cache locality.

This paper demonstrates how Mojo's language features

translate into concrete performance gains for k-means clustering. We provide a detailed comparison with a NumPy-based Python implementation, highlighting key optimization techniques employed during the code translation process [1]. Through comprehensive benchmarks on synthetic datasets, we quantify the performance improvements, showcasing Mojo's potential as a powerful tool for developing high-performance data analysis solutions. The findings have broader implications for accelerating other computationally intensive machine learning algorithms and pave the way for efficient and scalable data analysis on modern hardware architectures.

2. Literature Review

The BigVNSClust algorithm enhances K-means clustering for big data by integrating Variable Neighborhood Search (VNS), significantly improving accuracy and efficiency. This novel heuristic fuses data streaming with global optimization, optimizing partial objective function landscapes to refine clustering results. It employs local search and Variable Neighborhood Descent (VND) procedures, achieving a threefold accuracy improvement over traditional K-means. Through 7,366 experiments, cyclic neighborhood changes demonstrated slight accuracy enhancements. By combining global optimization with big data clustering techniques, BigVNSClust represents a significant advancement in scalable clustering algorithms [1].

The study proposes a simplified Map-Reduce architecture for implementing the K-means algorithm on FPGA, achieving a 28.74 Gbps throughput and a 3.93x speedup over existing FPGA-based implementations. By leveraging algorithmic segmentation, data path elaboration, and high-level synthesis techniques, the approach enhances computational efficiency and meets the performance demands of big data applications. This optimization significantly improves the scalability and execution speed of K-means clustering in FPGA environments [2].

Li et al., introduces a coarse-grained Map-Reduce architecture to optimize the K-means algorithm on FPGA, achieving 28.74 Gbps throughput and a 3.93x speedup over existing implementations. By leveraging algorithmic segmentation, data path elaboration, and automatic control optimization, the approach enhances computational efficiency and scalability, effectively addressing the performance demands of big data applications [3].

The study presents a privacy-preserving K-means algorithm designed for horizontally partitioned data, utilizing a cryptography-free multiparty additive scheme to enhance efficiency and security. The proposed sk-means algorithm maintains constant running time and stable computational overhead compared to traditional K-means while ensuring privacy against passive adversaries in distributed environ-

ments. This approach effectively balances data security and computational efficiency without relying on cryptographic methods [4].

The study introduces KMSR, a novel K-means formulation that enhances clustering stability through joint spectral embedding and spectral rotation, effectively improving discretization. Additionally, the generalized model (KMSR-G) integrates advanced data similarity measures, further refining clustering performance. Experimental results demonstrate that both models outperform existing K-means and spectral clustering methods on benchmark datasets, highlighting their effectiveness in improving clustering accuracy and stability [5].

The study focuses on enhancing K-means clustering by optimizing the entanglement of learned latent representations using soft nearest neighbor loss. By introducing an annealing temperature factor, the approach improves the clustering quality of autoencoder-learned representations. Experimental results demonstrate superior performance, achieving 96.2% accuracy on the MNIST dataset and outperforming baseline models across multiple datasets, highlighting its effectiveness in clustering optimization [6].

The study introduces PNN-smoothing, a meta-method for K-means initialization, which enhances clustering efficiency and cost-effectiveness. This approach involves splitting the dataset into random subsets, clustering them individually, and merging results using the pairwise-nearest-neighbor (PNN) method. Experimental results show that PNN-smoothing consistently improves clustering costs, outperforming K-means++ in both speed and effectiveness, making it a superior seeding strategy for K-means clustering [7].

The study introduces a neural-processor-based K-means clustering technique utilizing mobile machine learning to enhance clustering efficiency for big data. By leveraging parallel clustering and distributed processing, the approach achieves up to two times faster performance compared to traditional laptop/desktop processors. Experimental results demonstrate that neural engine processors significantly improve K-means clustering speed, making them a promising solution for real-time, high-performance clustering on mobile devices [8].

The study presents an enhanced K-means clustering algorithm integrating the Equilibrium Optimization Algorithm (EOA) to dynamically adjust the number of clusters and select optimal attributes. The proposed EOAK-means algorithm outperforms traditional K-means and CVK-means methods, demonstrating improved intra-cluster distances (ICD) and Rand index (RI) scores across five benchmark datasets. Experimental results confirm its superior clustering effectiveness compared to conventional approaches [9].

The study explores parallelizing the K-means algorithm using OpenMP and OpenACC to improve big data clustering performance by reducing computation time while maintaining accuracy. The research compares OpenMP flat synchronous and GPU-based OpenACC parallelization, addressing chal-

lenges like thread safety and race conditions. Results indicate that OpenACC outperforms OpenMP in computation time savings, demonstrating significant speedup and efficiency improvements in large-scale clustering tasks [10].

The study introduces a stochastic optimization approach to K-means clustering, redefining its optimization process by focusing on individual sample-level adjustments. A new target function is proposed to minimize pairwise distances within clusters, leading to a faster convergence and improved local minimum quality. Experimental results demonstrate significant performance improvements over traditional K-means variants, including hierarchical and sequential K-means, across various datasets [11].

The study presents an improved K-means algorithm designed for enhanced performance and efficiency, particularly for large datasets. The proposed approach incorporates a data formatting step to optimize cluster selection based on frame size and distance between means, leading to faster convergence and improved clustering accuracy. Experimental results demonstrate that the method outperforms both standard and modified K-means algorithms, making it a more effective solution for large-scale clustering tasks [12].

The study introduces the Boundary Weighted K-means (BWKM) algorithm, a recursive and parallel approximation to K-means clustering, designed to enhance scalability and efficiency for large datasets. By utilizing small, weighted sets of representative points, BWKM reduces the number of distance computations while maintaining high solution quality, particularly in challenging clustering regions. Experimental results demonstrate that BWKM outperforms state-of-the-art methods, achieving an optimal balance between computational efficiency and clustering accuracy [13].

3. Methodology

3.1. Understanding K-Means Clustering

K-means clustering is an iterative algorithm that partitions a dataset into k distinct, non-overlapping clusters. The core principle is to group data points based on their proximity to centroids, which represent the center of each cluster. The algorithm aims to minimize the total inertia, defined as the sum of squared distances between each data point and its assigned centroid. This section details the key steps and concepts underlying the k-means algorithm [14].

3.1.1. Algorithm

Given a dataset with M data points and N features, and a desired number of clusters k :

- (1) *Initialization*: Select k initial centroids. A common method is the k-means++ algorithm [4], which strategically selects well-separated initial centroids, promoting faster and better convergence.
- (2) *Lloyd's Iteration*: Repeat until convergence:

- a. *Cluster Assignment*: For each data point, calculate its distance to all k centroids. Assign the point to the cluster with the closest centroid.
- b. *Centroid Update*: For each cluster, recalculate the centroid by computing the mean of all data points assigned to that cluster.

3.1.2. K-Means++ Initialization

Randomly selecting initial centroids can lead to suboptimal convergence. K-means++ addresses this with a probabilistic approach:

- (1) Choose one data point uniformly at random as the first centroid.
- (2) For each remaining centroid:
 - a. Calculate the squared Euclidean distance between each data point and the closest existing centroid.
 - b. Choose a new data point as a centroid with probability proportional to its squared distance. This favors points farther from existing centroids [16].

3.1.3. Convergence Criteria

The iterative process stops when a predefined convergence criterion is met:

- (1) *Maximum Iterations*: Limits the maximum number of iterations to prevent infinite loops.
- (2) *Inertia Change Threshold*: Stops when the change in inertia between consecutive iterations falls below a specified threshold, indicating minimal improvement.

3.1.4. Inertia

Inertia measures cluster compactness and serves as a performance metric. It is calculated as:

$$\text{Inertia} = \sum (\text{distance}(\text{data point}, \text{centroid})^2)$$

where the summation iterates over all data points, and distance typically refers to the Euclidean distance. Lower inertia generally indicates tighter, more well-separated clusters.

K-means clustering iteratively refines cluster assignments

and centroids to minimize inertia. Understanding these components—initialization, iteration, convergence criteria, and inertia—is crucial for effective application and optimization.

3.2. Implementation in Python and Mojo

This section details the practical implementation of the k-means algorithm in both Python and Mojo. We compare the code side-by-side, highlighting key differences and demonstrating how Mojo's features enable significant performance gains.

Both implementations center around a K-means class (Python) and a K-means struct (Mojo). These structures encapsulate the algorithm's hyperparameters and provide a fit method for clustering [15].

- (1) *Type System*: Mojo uses a strong, static type system, unlike Python's dynamic typing. This enables the Mojo compiler to optimize code more effectively and perform compile-time type checks, improving safety and performance.
- (2) *Memory Management*: Mojo offers a more explicit memory management model, granting finer control over data structures and memory allocation. This fine-grained control reduces memory overhead and improves cache locality.
- (3) *Vectorization & Parallelization*: Mojo supports low-level optimizations like vectorization and parallelization for concurrent data processing. Python often relies on external libraries like NumPy, which can introduce overhead.

3.2.1. Code Comparison

This section presents a side-by-side code comparison, focusing on two critical parts of the k-means implementation:

- (1) *Distance Calculation* (distance_norm)
 - (2) *K-Means++ Initialization* (kmeans_plus_plus)
- Python (NumPy)

```
def _kmeans_plus_plus(self, data):
    """Initializes centroids using the
    k-means++ algorithm."""
    centroids = [data[np.random.randint(data.shape[0])]]
    for _ in range(1, self.k):
        distances = np.min(np.linalg.norm(
            data[:, np.newaxis] - centroids, axis=2), axis=1)
        probs = distances / np.sum(distances)
        centroids.append(data[np.random.choice(
            data.shape[0], p=probs)])
    return np.array(centroids)
```

Figure 1. K-Means++ Centroid Initialization in Python.

```

fn _kmeans_plus_plus(self, data: Matrix[ dtype ]) ->
List[ Array[ dtype ] ]:
    """Initializes centroids using the k-means++
    algorithm ."""
    var centroids = List[ Array[ dtype ] ]([ data . data
    [ random_si64 ( 0, data . rows ) ] ])
    var distances = Matrix[ dtype ]( data . rows , 1)
    for i in range( 1, self . k ):
        self . distance_norm ( data , len ( centroids ) - 1,
        & distances )
        # Calculates distances to the last added centroid
        # Rest of the logic for probabilistic centroid
        selection
    return centroids

```

Figure 2. K-Means++ Centroid Initialization Using Typed Data Structures.

3.2.2. Differences and Similarities Analysis

Similarities

- (1) *Initialization*: Both start by randomly selecting the first centroid.
- (2) *Loop for Subsequent Centroids*: Both iterate $k-1$ times to select the remaining centroids, using distances to determine selection probabilities.
- (3) *Probabilistic Selection*: Both calculate probabilities based on squared distances from the nearest centroid.
- (4) *Difference*:
 - 1). *Syntax and Data Structures*:
 - a. *Python (NumPy)*: Uses NumPy arrays and leverages Python's random. choice and broadcasting.
 - b. *Mojo*: Employs specific data structures (Matrix, Array), includes a type hint system (dtype), and implements a custom random number generator (random si64).
 - 2). *Distance Calculation*:
 - a. *Python (NumPy)*: Calculates distances directly using np.linalg.norm and NumPy broadcasting.
 - b. *Mojo*: Uses a custom function (distance norm), potentially with more detailed manual calculation or optimization specific to Mojo's data structures.
 - 3). *Intermediate Data Storage*:
 - a. *Python (NumPy)*: Uses a temporary variable distance.
 - b. *Mojo*: Initializes a Matrix to store distances, possibly optimizing memory layout and access.
 - 4). *Code Structure*:
 - a. *Python (NumPy)*: Implements the logic inline within the loop.
 - b. *Mojo*: Likely uses more helper functions and explicit memory management.

Overall, both implementations aim for efficient centroid initialization using k-means++. However, they utilize different paradigms and optimizations suited to their respective programming environments. The Python implementation is straightforward, leveraging high-level abstractions. The Mojo

version is more detailed and potentially optimized for higher performance [17].

3.3. Benchmarking and Performance Evaluation

To quantify Mojo's performance benefits, we conducted benchmarks comparing it against a NumPy-based Python implementation. We evaluated the impact of three key parameters: number of clusters, dataset size, and data dimensionality.

Benchmark Setup We used synthetically generated datasets with varying numbers of samples, features, and clusters using scikit-learn's make_blobs function, ensuring a controlled environment. Benchmarks were performed on:

- (1) *Processor*: Apple M2 Air
- (2) *Memory*: 16 GB

Benchmarking Parameters To isolate each parameter's impact, we varied one parameter at a time while keeping others constant:

- (1) Number of Clusters (k): 5, 10, 15,... 180 (incrementing by 5)
- (2) Number of Samples (M): 2000, 4000, 6000,... 22000 (incrementing by 2000)
- (3) Number of Features (N): 200, 400, 600,... 3800 (incrementing by 200)

Metrics

We recorded the execution time of the *fit* method for both Mojo and Python implementations (in milliseconds). To demonstrate performance gains, we calculated the speedup using the following formula:

$$\text{Speedup} = \frac{\text{Execution Time (Python)}}{\text{Execution Time (Mojo)}} \quad (1)$$

4. Results

The following figures depict the benchmark results.

4.1. Result Plots




Figure 3. Execution time Mojo vs Python + NumPy K-Means (Samples 2000 Features 200).




Figure 4. Speedup Mojo vs Python + NumPy K-Means (Samples 2000 Features 200).




Figure 5. Execution time Mojo vs Python + NumPy K-Means (Samples 4000 Clusters 5).




Figure 6. Speedup Mojo vs Python + NumPy K-Means (Samples 4000 Clusters 5).




Figure 7. Execution time Mojo vs Python + NumPy K-Means (Cluster 80 Features 200).




Figure 8. Speedup Mojo vs Python + NumPy K-Means (Cluster 80 Features 200).

4.2. Benchmark Results

The benchmark results, visualized in Figures 3-8, demonstrate Mojo's significant performance advantage over the NumPy-based Python implementation across a range of scenarios.

- (1) **Figures 3 & 4** (Samples: 2000, Features: 200):
 - a. **Figure 3** shows that Mojo's execution time increases much more slowly than Python's as the number of clusters grows.
 - b. **Figure 4** quantifies this, showing speedups ranging from approximately 11x to over 250x as the number of clusters increases from 5 to 180. This indicates Mojo's superior efficiency, especially with a larger number of clusters.
- (2) **Figures 5 & 6** (Samples: 4000, Clusters: 5):
 - a. **Figure 5** illustrates that increasing the number of features has a less pronounced impact on Mojo's execution time compared to Python.
 - b. **Figure 6** shows speedups consistently around 6x to 13x, demonstrating Mojo's ability to handle higher dimensionality relatively well, although the speedup is less dramatic than in the previous case.
- (3) **Figures 7 & 8** (Clusters: 80, Features: 200):
 - a. **Figure 7** highlights Mojo's efficiency with a large number of clusters and a moderate number of features. Mojo's execution time remains significantly lower than Python's as the number of samples increases.
 - b. **Figure 8** shows substantial speedups, ranging from roughly 130x to nearly 200x, emphasizing Mojo's advantage in scenarios with many clusters.

4.3. Analysis

The benchmark results consistently demonstrate the superior performance of the Mojo K-means implementation.

- (1) *Impact of Number of Clusters:* As the number of clusters increases, Mojo's speedup becomes more pronounced, especially with larger datasets. This highlights Mojo's efficient vectorization and parallelization, effectively handling increased distance calculations.
- (2) *Impact of Dataset Size:* Mojo's performance advantage is increasingly evident with larger datasets, indicating efficient memory management and scalability. This is crucial for handling the growing size of real-world datasets.

- (3) *Impact of Data Dimensionality:* While Mojo maintains a significant advantage across varying feature counts, the speedup decreases slightly as dimensionality increases. This suggests that data movement and memory access overhead, which typically grow with higher dimensionality, may impact Mojo's performance, though it still outperforms the Python implementation. This is an area for potential future optimization.

A performance comparison between Mojo, Python with NumPy, and scikit-learn for K-means clustering reveals significant differences. Mojo consistently outperforms Python with NumPy in execution time and speedup across all data configurations. This is attributed to Mojo's compiled nature, efficient memory management, vectorization, and parallelization capabilities.

While scikit-learn, a highly optimized library, offers better performance than basic NumPy implementations, Mojo still maintains a considerable performance edge, particularly as data scale and complexity increase. This highlights Mojo's design advantages for high-performance computing and AI workloads.

4.4. Cluster Visualization

To illustrate the correctness of the K-means implementations, we visualize the clusters generated from a sample dataset with 2000 samples, 10 features, and 5 clusters. We applied Principal Component Analysis (PCA) to reduce the dimensionality to two for visualization. The plot (**Figure 9**) shows the data points colored according to their assigned cluster, along with the centroids identified by both the Mojo and Python implementations.

The close alignment of the centroids and the clear separation of clusters visually confirm that Mojo's implementation produces accurate clustering results, mirroring the results of the Python implementation.

5. Discussion and Future Work

This research presents a compelling case for Mojo as a high-performance language for implementing data-intensive algorithms like K-means clustering. The benchmarks demonstrate a significant performance advantage over traditional Python implementations, with Mojo achieving speedups ranging from 6x to 250x. This substantial improvement is largely attributed to Mojo's core language features and design choices.




Figure 9. Scatter Plot of Data Points with Centroid.

5.1. A. Mojo's Performance Advantage

- (1) *Efficient Vectorization*: Mojo leverages SIMD instructions through vectorization, accelerating distance calculations—the computational bottleneck of K-means. Processing multiple data points concurrently within a single CPU instruction reduces overhead.
- (2) *Effective Parallelization*: Mojo's parallelization features, while not extensively explored in this implementation, offer avenues for further performance gains, especially with larger datasets and more clusters. Distributing computations across multiple cores can significantly reduce execution time.
- (3) *Optimized Memory Management*: Mojo's explicit memory management allows finer-grained control over data structures and memory allocation, leading to reduced overhead compared to Python's garbage collection. This enables optimization strategies like data locality, improving cache utilization, and reducing memory access times.
- (4) *Strong Typing Benefits*: Mojo's strong typing system plays a crucial role in enabling performance optimizations. Knowing data types at compile time allows the

compiler to generate more efficient machine code, eliminating runtime type checks and enabling more aggressive optimizations.

5.2. B. Comparative Analysis with Existing Solutions

While Mojo demonstrates superior performance compared to a basic NumPy-based Python implementation, highly optimized libraries like scikit-learn leverage sophisticated algorithms and data structures. However, even compared to scikit-learn, Mojo maintains a considerable edge, particularly as the data scale grows. This suggests Mojo's performance advantages stem not only from low-level optimizations but also from its design as a language tailored for high-performance computing and AI workloads.

5.3. C. Trade-offs and Considerations

Performance often involves trade-offs. While Mojo delivers impressive speedups, there's a learning curve associated with mastering its systems programming features. Developers accustomed to Python's ease of use might find Mojo's explicit

memory management and type annotations initially less intuitive. However, the performance gains can justify this learning investment, especially for performance-critical applications.

5.4. D. Implications for Data Analysis and Machine Learning

Mojo's performance improvements have broader implications for data analysis and machine learning. As datasets grow, the need for high-performance computing solutions becomes increasingly paramount. Mojo's ability to bridge the gap between Python's expressiveness and the performance of systems programming languages positions it as a valuable tool.

5.5. E. Future Research Directions

This research serves as a starting point for further exploration of Mojo's capabilities. Future research directions include:

- (1) *Investigating Mojo's performance on diverse hardware platforms:* Exploring Mojo's performance scaling on different CPU architectures, GPUs, and potentially specialized AI accelerators.
- (2) *Applying Mojo to other machine learning algorithms:* Evaluating Mojo's performance on other computationally intensive algorithms (e.g., support vector machines, deep learning models, graph algorithms) to demonstrate its generalizability.
- (3) *Developing a comprehensive benchmarking suite:* Creating a standardized benchmark suite specifically for evaluating Mojo's performance across a diverse range of machine learning tasks and datasets.

6. Conclusion

This work presented a performance-oriented implementation of the k-means clustering algorithm in Mojo, demonstrating significant speedups over traditional Python implementations. By leveraging Mojo's unique combination of Python-like syntax and systems programming features (vectorization, parallelization, explicit memory management), we achieved substantial reductions in execution time, particularly for larger datasets and a higher number of clusters. The benchmarks highlight Mojo's ability to bridge the gap between Python's ease of use and the performance demands of data-intensive workloads. Speedups achieved by the Mojo implementation, ranging up to 250x compared to the baseline Python implementation, underscore the potential of this emerging language for developing high-performance data analysis solutions. While this work focused on k-means clustering, the core principles and optimization techniques can be readily applied to other machine learning algorithms. As Mojo matures, we anticipate even greater performance gains through ongoing compiler optimizations and expanded hardware support. The development of a comprehensive ecosystem of libraries and tools for Mojo will further solidify

its position as a compelling alternative for AI practitioners seeking to unlock the full potential of their hardware and accelerate their data analysis pipelines. Mojo represents an exciting step forward in programming languages for AI, empowering developers to write high-level, expressive code without sacrificing performance. The ability to seamlessly integrate with existing Python codebases lowers the barrier to entry, enabling incremental adoption and facilitating a smooth transition to a high-performance environment.

Abbreviations

KMC	K-Means Clustering
ML	Machine Learning
HPC	High-Performance Computing
SIMD	Single Instruction, Multiple Data
VNS	Variable Neighborhood Search
PCA	Principal Component Analysis
FPGA	Field-Programmable Gate Array
AI	Artificial Intelligence
CPU	Central Processing Unit
GPU	Graphics Processing Unit
UV	Ultraviolet

Author Contributions

Touhidul Alam Seyam: Conceptualization, Resources, Project Administration, Writing - Review & Editing

Md. Sazzad Hossain: Data Curation, Methodology, Software, Visualization

Rajib Ghose: Formal Analysis, Investigation, Validation

Mekhriddin Nurmatov: Software, Writing - Original Draft, Data Curation

Nazarov Fayzullo: Methodology, Investigation, Supervision

Zarin Hadika: Validation, Writing - Original Draft, Visualization

Abhijit Pathak: Supervision, Funding Acquisition, Writing - Review & Editing

Funding

This work is not supported by any external funding.

Data Availability Statement

The data supporting the outcome of this research work has been reported in this manuscript.

Conflicts of Interest

The authors declare no conflicts of interest.

References

- [1] Shah, S. M., Sonawane, H. N., & Patil, D. D. (2015). A brief survey on clustering techniques. *International Journal of Science and Research (IJSR)*, 4(6), 2319-7064.
- [2] Datta, R., Joshi, D., Li, J., & Wang, J. Z. (2008). Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys (CSUR)*, 40(2), 1-60.
- [3] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3), 1-58.
- [4] Harris, C. R., Millman, K. J., Van Der Walt, S. J., et al. (2020). Array programming with NumPy. *Nature*, 585(7825), 357-362.
- [5] Pedregosa, F., Varoquaux, G., Gramfort, A., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825-2830.
- [6] Modular AI. (2023). Mojo Programming Language. <https://docs.modular.com/mojo/>
- [7] Cap ó, M., P érez, A., & Lozano, J. A. (2020). An efficient K-means clustering algorithm for tall data. *Data Mining and Knowledge Discovery*. <https://doi.org/10.1007/S10618-020-00678-9>
- [8] Islam, S., Balasubramaniam, S., Goyal, P., Sultana, A., Bhutani, L., Raje, S., & Goyal, N. (2019). A rapid prototyping approach for high performance density-based clustering. <https://doi.org/10.1109/DSAA.2019.00041>
- [9] Liu, Y., Du, X., & Ma, S. (2021). Innovative study on clustering center and distance measurement of K-means algorithm: Mapreduce efficient parallel algorithm based on user data of JD mall. *Electronic Commerce Research*. <https://doi.org/10.1007/S10660-021-09458-Z>
- [10] Kruli's, M., & Kratochv'íl, M. (2020). Detailed analysis and optimization of CUDA K-means algorithm. <https://doi.org/10.1145/3404397.3404426>
- [11] Wang, X., Chen, R.-C., Yan, F., Zeng, Z., & Hong, C. (2019). Fast adaptive K-means subspace clustering for high-dimensional data. *IEEE Access*. <https://doi.org/10.1109/ACCESS.2019.2907043>
- [12] Clustering big data based on distributed fuzzy K-medoids: An application to geospatial informatics. (2022). *IEEE Access*. <https://doi.org/10.1109/access.2022.3149548>
- [13] Dai, D.-B., Ma, Y., & Zhao, M. (2021). Analysis of big data job requirements based on K-means text clustering in China. *PLOS ONE*. <https://doi.org/10.1371/JOURNAL.PONE.0255419>
- [14] Wen, Z., & Tzerpos, V. (2003). An optimal algorithm for Mo-Jo distance. <https://doi.org/10.1109/WPC.2003.1199206>
- [15] Alam, T. S., Jowthi, C. B. & Pathak, A. Comparing pre-trained models for efficient leaf disease detection: a study on custom CNN. *Journal of Electrical Systems and Inf Technol* 11, 12 (2024). <https://doi.org/10.1186/s43067-024-00137-1>
- [16] Seyam, T. A., Pathak, A. AgriScan: Next.js powered cross-platform solution for automated plant disease diagnosis and crop health management. *Journal of Electrical Systems and Inf Technol* 11, 45 (2024). <https://doi.org/10.1186/s43067-024-00169-7>
- [17] A. Pathak et al., "Application of Machine Learning K-Means Clustering and Linear Regression in Determining the Risk Level of Pulmonary Tuberculosis," 2024 IEEE International Conference on Computing, Applications and Systems (COMPAS), Cox's Bazar, Bangladesh, 2024, pp. 1-6, <https://doi.org/10.1109/COMPAS60761.2024.10796963>

Biography



Touhidul Alam Seyam is a passionate Software Engineer and Research Assistant with expertise in full-stack development, machine learning, and applied research. He works at BGC Trust University Bangladesh and Hello World Communications Limited. His research focuses on AI in healthcare and agriculture, with publications on leaf disease detection, tuberculosis risk, and cardiovascular prediction in SCOPUS and Springer journals. Proficient in Next.js, Django, React, Golang, and PostgreSQL, he builds robust web applications. Seyam holds certifications from IBM, Google, and HackerRank, and actively engages in competitive programming. His work reflects a strong dedication to solving real-world problems through AI and software innovation.



Md. Sazzad Hossain is a distinguished professor, researcher, and academic advisor with vast experience in the ICT sector. He is currently a Visiting Professor at Samarkand State University, Uzbekistan, and Head of International Education Development at Synergy University, Moscow. He earned his Ph.D. and M.S. in Electrical and Computer Engineering from Portland State University, USA, with his doctoral work focusing on Quantum Computing, integrating concepts from physics, math, computer science, and biology. He completed his B.Sc. in Electrical System Network Engineering from Moscow Technical University. He is also an accomplished writer contributing to global academic and scientific communities.



Rajib Ghose is a distinguished researcher in computational intelligence, with expertise in formal verification, cryptography, and cloud security. His research spans optimization algorithms and AI ethics, contributing to advancements in secure computing, distributed systems, and bioinformatics. He has published extensively in reputed scientific journals, with impactful work in knowledge representation and formal analysis. Rajib actively mentors students and researchers in AI-driven cyber security and computational complexity. His interdisciplinary approach bridges theory and application, shaping modern methodologies for secure, intelligent data processing. Through his academic and research contributions, he continues to influence the evolution of trustworthy and efficient computational systems.

tems.



Mekhriddin Nurmamatov, PhD in Technical Sciences, is an Associate Professor specializing in predictive modeling, multi-parameter optimization, and intelligent data analysis using Machine Learning and Deep Learning. He has authored over 50 scientific articles, with more than five indexed in SCOPUS and Web of Science. He has also written five textbooks and manuals, contributing significantly to academic development. His research focuses on forecasting population employment and optimizing complex systems through AI-driven methodologies. Actively engaged in collaboration, he mentors students and partners with researchers to advance computational solutions for real-world challenges in data science and intelligent decision-making systems.



Nazarov Fayzullo Makhmadiyarovich, PhD in Technical Sciences, is an Associate Professor whose research focuses on enhancing payment data reliability using Blockchain and intelligent data analysis through Machine Learning and Deep Learning. He has authored over 100 scientific articles, with more than 20 indexed in SCOPUS and Web of Science. He has also written around 10 textbooks and manuals that support the advancement of technical sciences. His work significantly contributes to secure financial transactions, AI-driven analytics, and the application of emerging technologies. Committed to innovation and education, he continues to drive progress in secure, intelligent, and practical computing solutions.



Zarin Hadika is a dedicated university lecturer in Computer Science and Engineering, with a strong passion for teaching and research. Her primary focus lies in image processing and computer vision, with notable work on cross-modal person re-identification using HOG, addressing complex challenges in pattern recognition. She is currently exploring Cloud Computing and Artificial Intelligence to expand her technical expertise. Zarin is deeply committed to academic growth, inspiring students through innovative teaching and continuous learning. Outside the classroom, she enjoys reading, researching, and engaging in meaningful discussions, aiming to make impactful contributions to the dynamic and evolving world of technology.



Abhijit Pathak is an Assistant Professor in Computer Science and Engineering at Sonargaon University with over 16 years of academic and research experience. His expertise spans IoT, Machine Learning, AI, and Software Development. He has published 30+ papers in top-tier journals and conferences, with a Google Scholar h-index of 9 and over 600 citations. Recognized among the top 7 scientists globally from BGC Trust University Bangladesh by AD Scientific Index 2024, he mentors students and leads AI and IoT projects. A Commonwealth Scholar, he has received multiple awards for academic excellence and innovation in interdisciplinary technological research, academic excellence and leadership in technological innovation.

Research Field

Touhidul Alam Seyam: Artificial intelligence, Machine learning, Cybersecurity, Smart cities, Big data analytics, Fault-tolerant computing, AI-driven surveillance, Human-robot interaction.

Md. Sazzad Hossain: Quantum computing, Data science, Computational intelligence, Internet of Things, Pattern recognition, Deep learning, Computer vision, Reinforcement learning, Edge computing, Secure computing, Privacy-preserving AI, Neural networks.

Rajib Ghose: Formal verification, Computational complexity, Cryptography, Artificial intelligence, Machine learning, AI ethics, Knowledge representation, Bioinformatics, Distributed computing, Algorithmic game theory.

Mekhriddin Nurmamatov: Embedded systems, Cyber-physical systems, Hardware security, FPGA design, Software verification, Quantum cryptography, Real-time systems, High-performance computing, Blockchain technology, Secure AI applications.

Nazarov Fayzullo: Natural language processing, Speech recognition, Computational linguistics, Text mining, AI-powered translation, Sentiment analysis, Explainable AI, Knowledge graphs, Conversational AI, Information retrieval.

Zarin Hadika: Computational neuroscience, AI-assisted healthcare, Biomedical data analysis, Artificial intelligence, Machine learning, Digital signal processing, AI-driven diagnostics, Emotion recognition, Computational psychology, Assistive technologies.

Abhijit Pathak: Artificial intelligence, Machine learning, Parallel computing, Neuromorphic computing, Complex systems modeling, Adaptive algorithms, Evolutionary computation, AI-driven simulations, Computational fluid dynamics, AI for space research.