



Reinventing DISC personality assessment: Machine learning approaches for deeper insights and greater efficiency

Fatima Kalabi¹; Mohammad Hossein Amirhosseini^{2*}

¹Queen's Hospital, Havering and Redbridge University Hospitals NHS Trust, London, United Kingdom.

²Department of Computer Science and Digital Technologies, School of Architecture, Computing and Engineering, University of East London, London, United Kingdom.

***Corresponding Author: Mohammad Hossein Amirhosseini**

Department of Computer Science and Digital Technologies, School of Architecture, Computing and Engineering, University of East London, London, United Kingdom.

Email: m.h.amirhosseini@uel.ac.uk

Received: Jan 20, 2026

Accepted: Feb 11, 2026

Published Online: Feb 18, 2026

Website: www.joaiar.org

License: © Amirhosseini MH (2026). This Article is distributed under the terms of Creative Commons Attribution 4.0 International License

Volume 3 [2026] Issue 1

Abstract

The DISC personality framework, while widely adopted in applied settings, relies on a fixed rule-based classification method that may oversimplify individual behavioural profiles. This study explores whether machine learning can offer a more flexible, efficient, and accurate approach to DISC classification. Using a dataset of over 1,000 participants, we evaluated multiple supervised models—including Logistic Regression, XGBoost, SVM, MLP, Random Forest, and K-Nearest Neighbours—alongside unsupervised clustering techniques. Logistic Regression emerged as the top-performing model, achieving 93.53% accuracy and demonstrating superior cross-validation stability.

Recursive Feature Elimination identified a reduced set of ten key questionnaire items, maintaining over 91% accuracy and enabling the development of a concise assessment tool. Such a shortened questionnaire offers substantial practical benefits for real-world applications, particularly in fast-paced organisational contexts like recruitment, leadership coaching, and team composition, where rapid yet reliable personality insights are invaluable.

Clustering analysis further revealed alignment with traditional DISC categories, while uncovering potential hybrid profiles. A comparative clustering analysis between the full 40-item and reduced 10-item questionnaires confirmed that the same behavioural trait structures could be recovered using fewer items. Despite minor differences in cluster alignment, DISC trait patterns remained consistent across both models. These findings confirm that machine learning can replicate and enhance conventional DISC assessments, not only in terms of classification accuracy but also by preserving the conceptual integrity of the DISC framework.

The study validates that the reduced DISC assessment captures the latent personality structure of the original model, offering a scalable and empirically grounded solution for modern psychological evaluation. The complete modelling pipeline, including feature selection and clustering insights, contributes to the growing field of data-driven psychometrics.

Keywords: DISC personality assessment; Machine learning; Feature selection; Clustering analysis; Psychometrics; Computational psychometrics; Short-form assessment; Data-driven modelling; Personality classification; Questionnaire optimisation.

Introduction

Personality assessment frameworks play a pivotal role across diverse domains, including leadership development, talent acquisition, team optimisation, and career counselling. Among these, the DISC model—categorising individuals into four behavioural archetypes: Dominance (D), Influence (I), Steadiness (S), and Conscientiousness (C)—has garnered widespread use in applied organisational contexts. The model's appeal lies in its operational simplicity and interpretability, making it highly attractive for workplace environments where fast, actionable insights are essential. Unlike academically dominant frameworks such as the Five-Factor Model (FFM) or the Myers-Briggs Type Indicator (MBTI), DISC is especially valued in business settings for its immediate applicability to communication strategies, management styles, and team dynamics.

Despite its widespread practical adoption, the DISC model's traditional scoring system reveals important methodological constraints. Typically, participant responses across the four behavioural categories are aggregated, and the highest-scoring dimension is used to assign a categorical label. While this deterministic process streamlines classification, it inherently assumes linearity, equal weighting of questionnaire items, and strict categorical boundaries. These assumptions risk oversimplifying the complexity of human personality, particularly for individuals who display balanced profiles or hybrid behavioural tendencies. Such profiles, characterised by closely tied scores across multiple dimensions, are often inadequately represented within the rigid rule-based structure, leading to potential misclassification and loss of behavioural nuance.

Machine Learning (ML), with its ability to model complex, non-linear relationships, offers a promising avenue to enhance the utility and flexibility of rule-based personality classification systems like DISC. Rather than seeking to replace the traditional model, this study aims to augment and refine it — improving classification accuracy, handling ambiguous cases more effectively, and providing deeper insights into underlying behavioural structures. Moreover, ML techniques can assist in optimising questionnaire design, reducing the number of required items while retaining predictive power, thus enhancing user experience and application efficiency. Equally important, ML offers tools to validate whether shortened assessments can preserve the underlying trait structure of full-length instruments, an essential consideration for psychometric integrity.

Accordingly, this study pursues four core objectives:

1. To assess whether machine learning models can replicate and incrementally improve upon traditional DISC classifications, especially in ambiguous or borderline cases.
2. To investigate the feasibility of reducing the questionnaire length through feature selection, aiming to maintain high predictive accuracy with fewer items.
3. To utilise unsupervised clustering to explore potential latent structures or hybrid profiles within DISC categories.
4. To compare the personality groupings produced from the full and reduced questionnaires, evaluating whether the reduced-item model retains the latent behavioural architecture of the original DISC framework.

Through a rigorous experimental pipeline integrating both supervised and unsupervised learning, this research seeks to complement and refine the established DISC framework. By demonstrating that a shortened questionnaire can both predict DISC types and replicate the personality structure found in the full model, this study bridges the gap between classification accuracy and conceptual validity. In doing so, it advances the evolving intersection of data science and psychometrics, offering both methodological innovations and practical enhancements for organisational behaviour assessment.

Literature review

Personality assessment continues to play a foundational role in psychological research and applied behavioural sciences, informing diverse fields such as organisational behaviour, leadership development, and psychometric evaluation [1]. Among the major personality models, the Five-Factor Model (FFM) has achieved substantial academic traction for its trait-based dimensional approach [2,3]. Similarly, the Myers-Briggs Type Indicator (MBTI) enjoys widespread use, particularly in applied contexts, despite criticisms of its psychometric robustness [4,5].

By contrast, the DISC personality model, conceptualised by Marston in the early twentieth century, has been primarily adopted in non-clinical, workplace, and organisational settings [6,7]. It segments personalities into four behavioural dimensions: Dominance, Influence, Steadiness, and Conscientiousness, offering a pragmatic framework with straightforward interpretability [7,8]. While DISC assessments are popular in commercial environments, they remain underrepresented in academic research compared to the FFM or MBTI frameworks [9,10]. This academic underrepresentation can be attributed to perceptions of DISC as a simplified, business-oriented tool, which, while accessible, has traditionally lacked the empirical depth sought in academic psychometrics. This gap presents an opportunity for systematic computational exploration.

Although the DISC model enjoys widespread practical use, it has also faced academic critique for its limited construct validity, typological rigidity, and lack of empirical grounding when compared to trait-based frameworks like the FFM [8,11]. This study does not seek to validate or defend these theoretical assumptions, but rather to explore how machine learning can enhance the interpretability, adaptability, and practical utility of the DISC model.

Traditional DISC classification systems rely on a rule-based methodology, where questionnaire responses are grouped, summed, and assigned to the highest scoring category [7,12]. This deterministic structure, while user-friendly, presumes linearity, equal weighting of items, and mutually exclusive categories, potentially obscuring hybrid profiles or nuanced behavioural subtleties [1,11]. Consequently, methodological rigidity may limit the model's explanatory power, particularly for individuals whose profiles straddle multiple DISC categories. Indeed, prior work in personality psychology has consistently demonstrated that most individuals exhibit a blend of traits rather than fitting exclusively into discrete categories, with overlapping or hybrid configurations being common and psychologically meaningful [13]. This underscores the importance of adopting analytical techniques capable of capturing behavioural continua rather than relying solely on categorical distinctions.

The emergence of Machine Learning (ML) offers unprecedented opportunities to revisit such frameworks. Recent scholarship has demonstrated the efficacy of ML algorithms in personality prediction, particularly with the MBTI model [14-16]. Techniques such as Support Vector Machines (SVM), decision trees, random forests, and ensemble models like XGBoost have been effectively employed to predict personality traits from diverse data sources, including text and behavioural data [16-19].

In particular, XGBoost has been identified for its robustness in handling high-dimensional data and managing feature interactions efficiently [20,21]. Studies like those of Vaddem and Agarwal [21] have confirmed XGBoost's utility in classifying social media data, demonstrating its adaptability across languages and contexts. Meanwhile, deep learning architectures, have revolutionised NLP tasks, with applications in personality recognition from text proving highly effective [22,23].

Feature selection techniques play an equally critical role. Recursive Feature Elimination (RFE), an iterative method to identify the most predictive features, has shown promise in reducing model complexity while retaining performance [24,25]. In personality research, such methods help identify the most salient behavioural indicators, paving the way for shorter, efficient assessment instruments [24].

Complementary to supervised learning, unsupervised techniques such as K-Means clustering provide valuable insights into latent personality structures [26,27]. Methods like the Elbow criterion and Silhouette analysis assist in determining optimal cluster numbers and evaluating cluster cohesion [28,29]. Previous work has demonstrated that clustering can surface natural groupings that traditional typologies might overlook [27,30].

The quality of data preprocessing, especially for imbalanced datasets, has been extensively documented as a determinant of model success [31,32]. The Synthetic Minority Over-Sampling Technique (SMOTE) is widely recognised for addressing class imbalance in behavioural datasets [31]. Similarly, scaling techniques such as standardisation ensure equitable contribution of features during training [33].

Despite the advancement of ML across psychological frameworks, the integration of these methods with DISC remains nascent. Previous research, such as the work by O'Neill et al. [10], has explored DISC profiles in the context of virtual teams but has not applied machine learning to optimise classification or reduce questionnaire length. Furthermore, no study to date has systematically investigated whether DISC's heuristic approach can be improved through data-driven classification models, or whether unsupervised learning can reveal hybrid or emergent DISC profiles.

This study addresses these research gaps by implementing a comprehensive ML pipeline that combines supervised classification, feature selection, and unsupervised clustering. Our objectives are threefold: first, to evaluate whether modern ML models can enhance traditional rule-based DISC classification; second, to investigate the feasibility of a shorter DISC questionnaire by leveraging RFE; and third, to explore the underlying structure of DISC data through unsupervised clustering. In doing so, this research contributes novel insights to computational psychometrics and offers practical implications for efficient, scalable, and nuanced personality assessment tools.

Methodology

This study employed an end-to-end data-driven methodology to critically examine and enhance DISC personality classification using machine learning techniques. The pipeline integrated both supervised and unsupervised learning, feature selection for dimensionality reduction, robust evaluation protocols, and state-of-the-art tooling for full reproducibility and reporting. The overall aim was to explore not only the replication of DISC outcomes but also to assess whether computational methods could improve classification robustness, flexibility, and practical efficiency.

Research design

The research was structured around four principal objectives:

1. To evaluate the capacity of supervised machine learning models to replicate or outperform the traditional rule-based DISC assignment method.
2. To investigate the feasibility of constructing shorter, yet predictive, DISC assessments by applying feature selection and testing model performance at varying questionnaire lengths.
3. To explore latent patterns and potential hybrid personality structures within the dataset through unsupervised learning techniques.
4. To compare clustering results derived from the full DISC questionnaire and a reduced item set, assessing the structural equivalence of behavioural groupings and validating the integrity of the shortened assessment.
5. To establish a fully automated, reproducible, and professional-grade machine learning pipeline, capable of generating robust and generalisable findings suitable for academic and practical deployment.

The approach was modular and iterative. First, baseline models were developed and optimised using the full feature set to establish a performance benchmark. Subsequently, feature selection was applied, and the best-performing model was retrained using reduced feature sets to assess the impact of dimensionality reduction. In parallel, unsupervised learning techniques were employed to explore natural groupings within the data that could extend or refine DISC categorisations. A dedicated comparison between cluster structures derived from the full 40-item questionnaire and the reduced 10-item set was conducted to evaluate whether personality groupings remained stable across dimensionality reductions. (Figure 1) illustrates the full machine learning pipeline implemented in this study, including data preprocessing, feature selection, supervised classification, and unsupervised clustering for exploratory analysis.

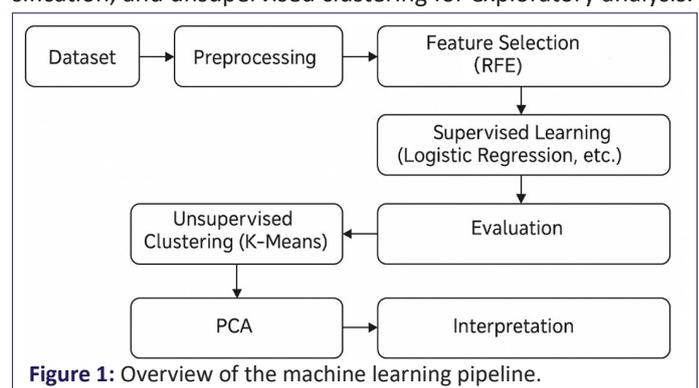


Figure 1: Overview of the machine learning pipeline.

Dataset and preparation

The dataset comprised responses from over 1,000 participants, collected via an online questionnaire disseminated through professional networks and public platforms to ensure diverse participant representation. Participants completed the questionnaire via Microsoft Forms, a secure web-based platform that ensured anonymity and standardised response collection. Participants were voluntarily recruited, provided informed consent, and were not compensated for participation. Although no formal demographic quotas were applied during recruitment, the dataset includes self-reported age and gender information and reflects a balanced distribution across professional backgrounds and life stages, contributing to the generalisability of the findings.

The DISC-based psychometric instrument consisted of 40 self-report items, systematically distributed across four behavioural dimensions: Assertiveness/Influence (AS1–AS10), Steadiness/Calm (SC1–SC10), Conscientiousness/Discipline (AD1–AD10), and Dominance/Decisiveness (DO1–DO10). Participants rated their level of agreement with each behavioural statement using a Likert-type scale, designed to capture nuanced self-perceptions relevant to personality categorisation. The full set of 40 questionnaire items is provided in Appendix A to support transparency and facilitate replication.

As the original dataset included only responses to the questionnaire items, age and gender.

As there was no target variable (personality type label) in the original dataset, for the purposes of supervised learning tasks, personality labels were derived using the conventional rule-based DISC classification method, as established in organisational psychology practice [6,34]. Specifically, total scores across the items corresponding to each behavioural category were aggregated, and the category with the highest score determined the assigned DISC personality type. This rule-generated label served as the dependent variable in all supervised learning experiments conducted in this study.

The age and gender variables were not included in the predictive modelling process, as the study focused on evaluating the internal behavioural dimensions captured by the DISC questionnaire itself. Incorporating demographic variables could have introduced confounding factors or demographic biases, which fell outside the scope of this research. Descriptive statistics of the collected demographics are reported in (Table 1) in the Results section, confirming a reasonable diversity in age and gender distribution across the sample.

Prior to model development, the dataset underwent rigorous preprocessing to ensure analytical integrity. A comprehensive data audit confirmed that the dataset contained no duplicate entries or missing values. All numerical features were standardised using z-score normalisation to maintain scale invariance across variables — an essential step for optimising the performance of algorithms sensitive to feature magnitude, such as Support Vector Machines (SVM) and K-Nearest Neighbours (KNN). Furthermore, the categorical target variable representing DISC types was systematically label-encoded to enable compatibility with a wide range of machine learning classifiers.

For exploratory data analysis and the facilitation of unsupervised learning tasks, Principal Component Analysis (PCA) was applied. This dimensionality reduction technique projected the original 40-dimensional feature space into a two-dimensional

representation, providing valuable visual insights into the distribution of personality types and supporting cluster analysis aimed at uncovering latent behavioural structures.

From an ethical standpoint, this study poses no concerns. The dataset was fully anonymised, containing no personally identifiable information. All responses were de-identified at the point of data collection, ensuring that individual participants cannot be identified or traced.

Supervised learning framework

A robust supervised learning pipeline was established to evaluate the potential of machine learning algorithms to replicate or enhance the traditional DISC classification method. Six distinct classifiers were selected to ensure methodological diversity and comprehensive exploration of modelling paradigms. These models included Support Vector Machine (SVM), Extreme Gradient Boosting (XGBoost), Logistic Regression, Random Forest, K-Nearest Neighbours (KNN), and Multi-Layer Perceptron (MLP) neural network.

The dataset was partitioned into training and testing subsets using an 80/20 stratified split to maintain class balance. Hyperparameter optimisation was conducted through Randomised Search Cross-Validation, employing a five-fold strategy to robustly identify optimal configurations for each model. This process systematically sampled from defined parameter distributions, enhancing model generalisability and minimising overfitting risks.

Throughout this framework, attention was paid not only to classification accuracy but also to maintaining interpretability and replicability. Probabilistic models were configured to output confidence scores, providing an additional layer of interpretive depth for subsequent evaluation.

Feature selection and dimensionality reduction

To explore the feasibility of developing a shorter yet reliable DISC assessment tool, feature selection was conducted using Recursive Feature Elimination (RFE). RFE is an iterative algorithm that ranks features by recursively removing the least significant attributes and re-evaluating model performance at each iteration. While RFE provides an ordered ranking of feature importance, it does not inherently determine the optimal number of features to retain. Therefore, determining the ideal feature subset size required empirical testing across multiple configurations.

Accordingly, the feature selection process was operationalised by retraining the best-performing model from the baseline full-feature evaluation — Logistic Regression — using progressively reduced feature sets. Specifically, subsets containing the top 5, 10, 15, 20, 30, and all 40 features were tested. This systematic approach enabled a thorough assessment of the trade-off between model complexity and predictive accuracy. By comparing model performance at each feature count, it became possible to identify the point at which questionnaire length could be reduced without compromising predictive efficacy.

The decision to focus feature reduction experiments on best performing model was grounded in its strong baseline performance and interpretability, making it well-suited for evaluating the trade-offs between model simplicity and predictive power. This design ensured that any reduction in questionnaire length would be both methodologically sound and practically meaningful for future applications of the DISC model in organisation-

al and behavioural contexts.

The final set of 10 RFE-selected items was further validated through unsupervised learning techniques, confirming that this reduced questionnaire preserved the core personality structure identified in the full 40-item version.

Unsupervised learning and pattern discovery

Beyond supervised classification, unsupervised learning methods were applied to explore latent structures within the DISC dataset. K-Means clustering was selected for its scalability, ease of interpretability, and alignment with the categorical nature of the DISC model. The clustering algorithm was applied to the full standardised dataset, with the number of clusters empirically determined using the Elbow Method, which evaluates within-cluster inertia to identify the point of diminishing returns.

To further strengthen the cluster validation, Silhouette analysis was also employed. While the Elbow Method focuses on minimising intra-cluster variance, the Silhouette score evaluates both intra-cluster cohesion and inter-cluster separation. This complementary method provided additional insight, particularly valuable given the gradual inflection observed in the Elbow plot — a characteristic often expected in psychometric data due to the fluidity of human behavioural traits. The use of both techniques ensured a more robust, multi-perspective determination of the optimal number of clusters.

To interpret the composition of each cluster, a quantitative profiling step was incorporated. Specifically, post-clustering, the average response scores for each of the four DISC behavioural dimensions were calculated across clusters. These averages provided empirical evidence for mapping clusters to corresponding DISC categories based on dominant behavioural tendencies. For example, clusters exhibiting higher average scores in Assertiveness/Influence items were associated with the Influence (I) dimension, while clusters with elevated Conscientiousness/Discipline scores corresponded to the Conscientiousness (C) type. This approach ensured that cluster labelling was not subjective but grounded in quantitative behavioural data.

A second round of clustering was conducted using only the 10 features selected through RFE. All unsupervised analyses (Elbow, Silhouette, PCA) were replicated to assess whether the reduced-item model recovered similar personality groupings. Trait profile comparisons between clusters from the full and reduced models were performed, alongside Euclidean distance matching and heatmap visualisations to evaluate structural similarity. This enabled validation of the reduced model's capacity to replicate the underlying behavioural structure of the DISC framework, even in the absence of the full feature set.

Additionally, cross-tabulation between cluster assignments and original rule-based DISC labels was conducted to evaluate alignment and divergence between data-driven clusters and traditional categories. Principal Component Analysis (PCA) was further employed to reduce the feature space to two dimensions, providing visual clarity on cluster separation and structure. The resulting plots offered an intuitive visual interpretation of the cluster groupings and corroborated findings from the quantitative cluster profiling.

This multi-faceted approach allowed for both numerical and visual validation of cluster identities, enhancing the interpretability and robustness of the unsupervised learning component.

Model testing and evaluation

Model evaluation constituted a critical component of this research, ensuring that the predictive capabilities of the machine learning models were rigorously assessed. Each trained model was evaluated on the hold-out test set, and multiple performance metrics were computed to provide a multidimensional assessment of classification quality.

Accuracy was selected as the primary evaluation metric, reflecting the proportion of correct predictions. Complementary metrics included precision, recall, and F1 score, offering insight into the models' abilities to balance sensitivity and specificity across multiple classes. Cohen's Kappa was additionally computed to measure agreement between the model predictions and actual labels, adjusted for chance agreement.

To further understand the decision-making processes of the models, confusion matrices were generated for each classifier. These matrices provided granular insights into class-level performance, highlighting patterns of correct classifications and misclassifications across the four DISC categories.

Cross-validation scores were calculated during model training, capturing the mean and standard deviation of accuracy across five folds to assess the models' stability and generalisability. Furthermore, probabilistic outputs from classifiers capable of probability estimation were recorded to derive average confidence scores, adding depth to the interpretative analysis of model predictions.

In addition to classification evaluation, unsupervised clustering structures derived from the full and reduced feature sets were compared through DISC trait profile matching and distance-based cluster alignment. This structural validation step ensured that reductions in questionnaire length did not distort the underlying behavioural patterns originally captured in the full DISC assessment.

This comprehensive evaluation strategy ensured that model comparisons were grounded in robust empirical evidence, supporting the reliability of the findings.

Tools and libraries

All analyses were conducted within a Python 3.9 programming environment, using the Jupyter Notebook interface to facilitate modular code development, documentation, and reproducibility. The selection of libraries was guided by considerations of methodological robustness, compatibility, and widespread acceptance in the scientific community.

Data preprocessing and manipulation were managed using pandas and NumPy, which provided efficient handling of structured data and numerical arrays. Machine learning model development, hyperparameter optimisation, and evaluation were performed using scikit-learn, while XGBoost was employed for its advanced gradient boosting capabilities.

Visualisations of results, including feature selection trends, confusion matrices, and clustering outcomes, were created using matplotlib and seaborn, enabling high-quality graphical representations of complex outputs. Trait heatmaps were used to enhance the interpretability of cluster structures across different feature spaces.

Pairwise distance metrics from scikit-learn's metrics module were applied to compare cluster profiles and compute best-

matching pairs between full and reduced feature models. This additional analysis layer provided a quantitative foundation for evaluating the structural equivalence of the DISC-based clusters.

The combination of these tools ensured that the research adhered to principles of transparency, reproducibility, and scientific rigor, thereby reinforcing the validity of the methodology.

Results

This section presents the empirical findings of the study, structured around model performance evaluation, confusion matrix analyses, feature selection and questionnaire reduction, clustering exploration, and the relationship between feature selection and clustering outcomes. All results derive from the end-to-end machine learning pipeline detailed in the methodology.

Sample demographics

To contextualise the dataset and support the generalisability of findings, basic demographic characteristics of the participants were analysed. The sample comprised a well-distributed range of ages and genders. As summarised in (Table 1), participants represented a diverse cross-section of life stages.

The gender distribution was relatively balanced, with 50.8% of participants identifying as male and 49.2% of participants as female. The age distribution spanned from early adulthood to older age groups, with the largest proportion of participants (42.4%) falling within the 26-35 age range, followed by 28.4% in the 36-45 bracket. Notably, younger participants (18-25) ac-

counted for 18.6%, and participants aged 46 and above represented 10.6% of the sample.

These demographic insights confirm the diversity of the participant pool and enhance the representativeness of the study's findings. However, it is important to note that demographic variables were collected solely for descriptive purposes and were not incorporated into the predictive modelling pipeline, as the focus of this research was on the behavioural dimensions captured by the DISC assessment itself.

Table 1: Participant demographic distribution.

Variable	Category	Percentage (%)
Gender	Male	50.8%
	Female	49.2%
Age group	18-25	18.6%
	26-35	42.4%
	36-45	28.4%
	46 and above	10.6%

Overall model performance

Six supervised learning models—Logistic Regression, XGBoost, Support Vector Machine (SVM), Multi-Layer Perceptron (MLP), Random Forest, and k-Nearest Neighbours (KNN)—were evaluated for their ability to predict DISC personality types based on responses to the 40-item questionnaire. The performance metrics for all models are summarised in (Table 2).

Table 2: Supervised learning model performance summary.

Model	Accuracy	F1 score	Precision	Recall	Cohen's kappa	Avg. confidence	CV mean accuracy	CV std dev
Logistic regression	0.9353	0.9356	0.9395	0.9353	0.9105	0.9244	0.9428	0.0213
XGBoost	0.9353	0.9353	0.9364	0.9353	0.9104	0.9307	0.9353	0.0174
SVM	0.9254	0.9251	0.9272	0.9254	0.8968	0.8712	0.9142	0.0271
MLP	0.9104	0.9102	0.9129	0.9104	0.8761	0.9580	0.9291	0.0160
Random forest	0.8159	0.8027	0.8243	0.8159	0.7394	0.6450	0.8296	0.0214
K-Nearest neighbours	0.6915	0.6554	0.6758	0.6915	0.5569	0.6071	0.7313	0.0377

According to (Table 2), both Logistic Regression and XGBoost models achieved the highest overall classification accuracy of 93.53%, confirming excellent predictive performance. However, Logistic Regression demonstrated a slight but meaningful edge over XGBoost across several secondary metrics. Specifically, it achieved higher precision (93.95% versus 93.64%), a superior cross-validation mean accuracy (94.28% versus 93.53%), and a marginally better Cohen's Kappa (0.9105 versus 0.9104), indicating stronger generalisability and alignment with actual labels.

SVM followed closely with an accuracy of 92.54%, maintaining balance across all key metrics. MLP achieved slightly lower accuracy (91.04%) but reported the highest average confidence (95.80%), indicating the model's strong internal certainty despite its lower external validation performance.

Conversely, Random Forest and KNN lagged significantly. Random Forest achieved an accuracy of 81.59%, and KNN delivered the lowest performance at 69.15%, coupled with the lowest average confidence (60.71%) and the highest cross-validation variance (3.77%), reflecting instability.

Given its superior balance of accuracy, cross-validation stability, and interpretability, Logistic Regression was selected for

subsequent feature reduction experiments, serving as the optimal candidate for developing a shorter, efficient DISC assessment tool.

Confusion matrices

To evaluate model performance at the class level, confusion matrices were generated for all models and are presented in (Figure 2).

Logistic Regression and XGBoost produced the most compact and accurate confusion matrices, demonstrating high true positive rates across all four DISC categories and minimal misclassifications. SVM also exhibited commendable class-level performance but showed slightly more confusion between adjacent classes, particularly within the Steadiness (SC) dimension.

MLP, despite high average confidence, exhibited increased class overlaps, notably between Steadiness (SC) and Assertiveness (AS). Random Forest and KNN presented the least desirable confusion matrices, with Random Forest showing moderate misclassification rates and KNN exhibiting substantial confusion, particularly within the Steadiness class, confirming its underperformance.

Table 4: Top 10 features for DISC classification (after RFE).

Feature code	Description
AS3	Assertiveness Item 3
DO6	Decisiveness Item 6
AD4	Discipline Item 4
SC7	Steadiness Item 7
DO2	Decisiveness Item 2
AS9	Assertiveness Item 9
SC1	Steadiness Item 1
AD8	Discipline Item 8
DO8	Decisiveness Item 8
AS1	Assertiveness Item 1

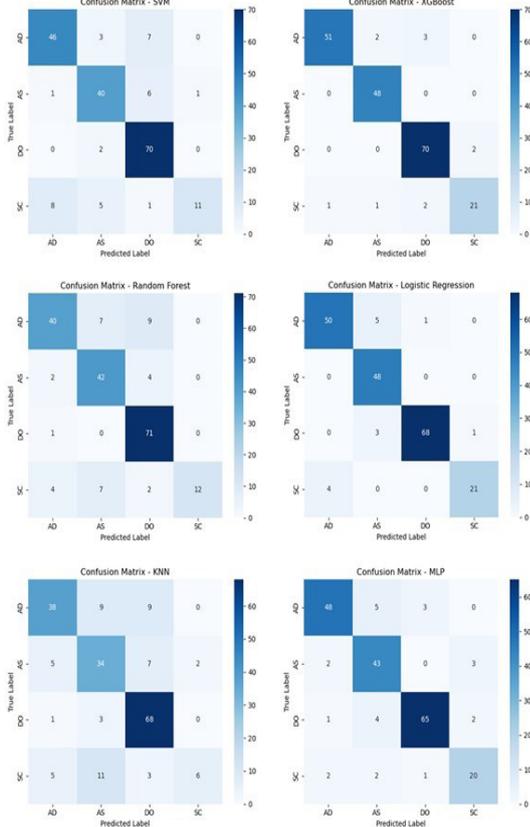


Figure 2: Confusion matrices for the developed models.

Feature selection and questionnaire reduction

To explore the feasibility of a reduced-item DISC assessment, Recursive Feature Elimination (RFE) was performed using the best-performing model, Logistic Regression. Model accuracy was tested across progressively reduced feature sets (5, 10, 15, 20, 30, and all 40 features). The results are summarised in (Table 3).

Table 3: Accuracy of Logistic Regression model at different feature counts.

Number of features	Accuracy
5	0.8806
10	0.9149
15	0.9164
20	0.9189
30	0.9313
40 (Full Set)	0.9353

According to (Table 3), the reduced feature set of 10 items achieved an accuracy of 91.49%, closely approximating the full model's performance. Notably, the incremental gain between 10 and 20 features was minimal (91.49% to 91.89%), highlighting the strong predictive contribution of the top-ranked features.

The top 10 features identified by RFE are reported in (Table 4). These items, drawn from across the behavioural dimensions, collectively formed the foundation of the proposed shortened DISC questionnaire.

Clustering analysis

K-Means clustering was employed to investigate potential natural groupings within the dataset, moving beyond the pre-defined rule-based DISC labels to uncover data-driven behavioural structures. The Elbow Method and Silhouette analysis were used to determine the optimal number of clusters, as visualised in (Figure 3).

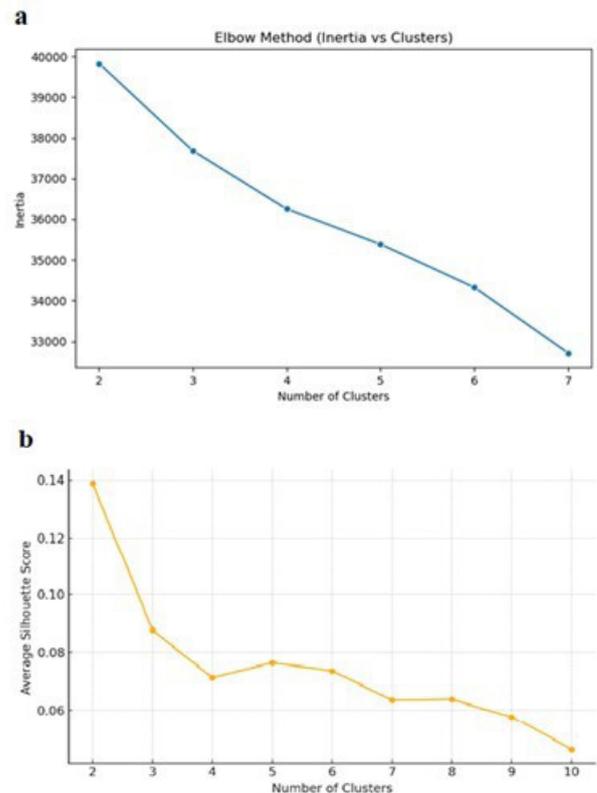


Figure 3: Determining optimal number of clusters. (a) Elbow Method. (b) Silhouette analysis.

The Elbow Method curve (Figure 3a) exhibited a gradual reduction in within-cluster inertia, with a noticeable inflection observed between two and four clusters. Beyond the four-cluster mark, inertia reduction plateaued, indicating diminishing returns from additional clusters. While the elbow was not sharply defined — a phenomenon commonly observed in personality research due to the fluid and continuous nature of human behavioural traits — the selection of four clusters remains both conceptually and empirically justified. The absence of a sharp elbow reflects the inherent complexities of psychometric data, where individuals often exhibit overlapping characteristics rather than falling neatly into discrete categories. Nonetheless, the four-cluster solution aligns well with the theoretical structure of the DISC model, reinforcing its validity within this context.

To complement the Elbow Method, Silhouette analysis was also conducted to strengthen the determination of the optimal cluster count. The Silhouette plot, presented in (Figure 3b), further supported the selection of four clusters by showing a local

maximum around this point. The average silhouette score peaks at $k=2$, but such a low number of clusters may oversimplify the underlying data structure. While the silhouette score decreases slightly for $k>2$, it remains reasonably stable up to $k=4$, after which it declines more noticeably. Although the overall Silhouette scores were modest — consistent with expectations for behavioural data — the analysis provided valuable confirmation that the four-cluster solution balances intra-cluster cohesion with inter-cluster separation. The convergence of these two methods reinforces the robustness of the chosen cluster configuration.

To visualise the clustering structure, PCA was applied to project the high-dimensional data into two dimensions. The resulting plot is presented in (Figure 4).

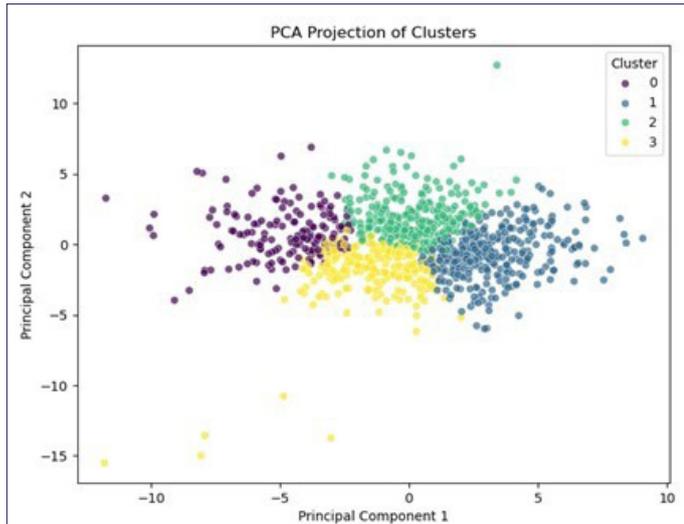


Figure 4: PCA projection of K-Means clusters showing four distinct groupings aligned with DISC categories.

The PCA visualisation revealed distinct and well-defined cluster separations. Cluster profiling based on average scores across the four DISC dimensions provided empirical justification for this mapping, as shown in (Table 5).

Table 5: Cluster profile averages by DISC dimension.

Cluster	Dominance (DO)	Influence (AS)	Steadiness (SC)	Conscientiousness (AD)
0	3.9	2.1	2.5	2.7
1	2.8	3.8	3.1	2.6
2	2.4	2.3	3.9	2.7
3	2.3	2.5	2.6	4.1

As illustrated in (Table 5), each cluster exhibits a clear dominance of one behavioural category: Cluster 0 demonstrated the highest mean scores in Dominance (DO), Cluster 1 in Influence (AS), Cluster 2 in Steadiness (SC), and Cluster 3 in Conscientiousness (AD). This quantitative profiling provides robust justification for mapping clusters to the DISC dimensions, moving beyond visual interpretation alone.

Further validation was achieved through cross-tabulation of cluster assignments and rule-based DISC labels, demonstrating strong alignment between unsupervised clusters and established categories. These findings provide empirical support for the underlying structure of the DISC model while revealing the potential of unsupervised learning to capture nuanced personality patterns.

Comparative validation of full vs reduced clustering

To assess whether the shortened questionnaire preserved the behavioural structure of the original DISC model, a comparative clustering analysis was conducted. K-Means clustering was re-applied using only the top 10 features selected through Recursive Feature Elimination (RFE). The goal was to determine whether the same latent behavioural groupings could emerge from a reduced input space.

Despite dimensionality reduction, the reduced-item clustering yielded a similar four-cluster solution, supported by both Elbow and Silhouette analyses (Figure 5).

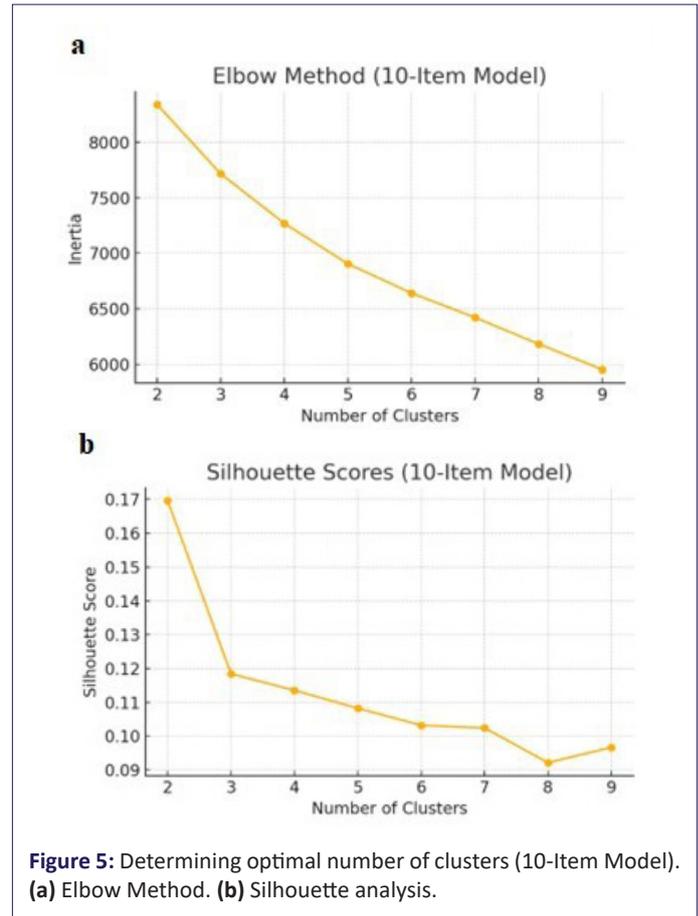


Figure 5: Determining optimal number of clusters (10-Item Model). (a) Elbow Method. (b) Silhouette analysis.

In (Figure 5a), the Elbow Method plots inertia (within-cluster sum of squares) against the number of clusters. A clear inflection point is not sharply defined, but a noticeable reduction in the rate of inertia decline begins around $k=4$, suggesting that increasing the number of clusters beyond this point yields only marginal improvements in compactness. This flattening of the curve indicates diminishing returns in terms of variance explained by additional clusters.

In (Figure 5b), the average Silhouette Score—an indicator of how well-separated and cohesive the resulting clusters are—peaks at $k=2$, with a gradual decline thereafter. Although the silhouette score is highest at $k=2$, this number of clusters may underrepresent meaningful groupings in more complex datasets. Notably, the silhouette score stabilizes around $k=4$, which corresponds with the elbow region observed in the inertia plot.

Considering both the inertia reduction and the clustering cohesion, $k=4$ emerges as a balanced choice. It offers a compromise between model compactness and interpretability, aligning with both analytical methods while avoiding oversimplification. PCA visualisation (Figures 6) clearly shows distinct personality groupings emerging from the reduced 10-

item feature space, reinforcing the structural validity of the optimized DISC questionnaire.

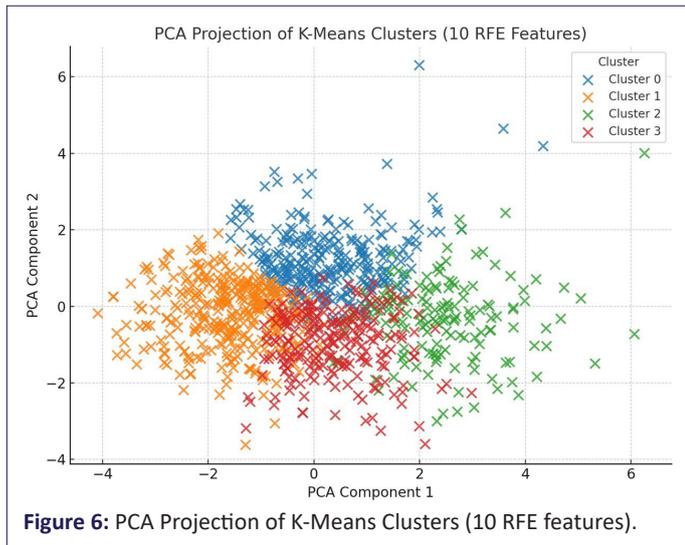


Figure 6: PCA Projection of K-Means Clusters (10 RFE features).

To evaluate the alignment between clusters from both models, trait profiles were generated by calculating the mean DISC dimension scores (Dominance, Influence, Steadiness, Conscientiousness) for each cluster. (Table 6) presents these profiles side-by-side for the 40-item and 10-item clusters.

Table 6: Cluster profile averages by DISC trait – full vs reduced models.

	AS	SC	AD	DO
F0	30.869	29.630	30.205	27.373
F1	32.660	29.856	29.847	37.839
F2	31.308	29.770	32.687	35.296
F3	27.234	28.617	30.873	23.895
R0	30.243	29.300	29.763	25.826
R1	32.710	29.997	30.964	36.047
R2	27.914	29.223	31.703	25.684
R3	29.538	29.137	31.672	33.297

As shown in (Table 6), cluster-level patterns were consistent across models. Clusters in the reduced feature set mirrored those in the full feature set in terms of dominant traits. For instance, Cluster F1 (full model) and R1 (reduced model) both reflected high Dominance traits, while F3 and R2 were similarly high in Conscientiousness and low in Dominance. A heatmap of these trait scores (Figure 7) visually confirms the close alignment.

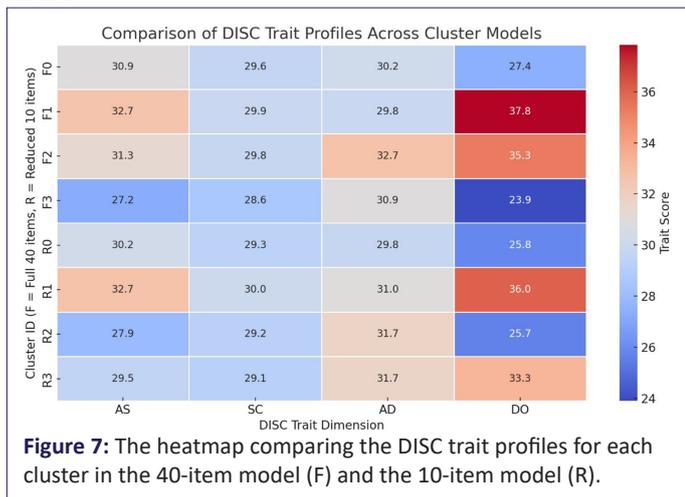


Figure 7: The heatmap comparing the DISC trait profiles for each cluster in the 40-item model (F) and the 10-item model (R).

To further quantify the structural similarity, Euclidean distances between trait profiles were calculated. Each full-model cluster was matched to its closest reduced-model counterpart based on these distances (Table 7).

Table 7: Cluster matching table (Full vs Reduced).

	R0	R1	R2	R3	Best match (Reduced)	Distance
F0	1.757	8.907	3.741	6.265	R0	1.757
F1	12.265	2.116	13.195	5.849	R1	2.116
F2	9.978	2.355	10.255	2.925	R1	2.355
F3	3.805	13.399	2.171	9.726	R2	2.171

The resulting matches—such as F0 ≈ R0 and F1 ≈ R1—validate that the reduced clustering preserves the psychological structure of the full model, not just its classifications. In addition, (Table 8) summarises the dominant behavioural patterns represented in each cluster match captured in both models.

Table 8: Conceptual alignment between full and reduced clustering models based on DISC trait profiles.

Full Cluster	Best match	Interpretation
F0	R0	Balanced or mild traits
F1	R1	High Dominance
F2	R1 (close tie)	Dominance–Conscientious mix
F3	R2	High Conscientiousness, low Dominance

While the reduced and full-cluster models exhibit strong alignment in overall trait profiles, not all clusters mapped one-to-one. Specifically, both F1 and F2 aligned most closely with R1 in the reduced model. This is likely due to overlapping Dominance traits present in both clusters, which were emphasized more heavily in the reduced feature space. Similarly, although F3 and R2 showed strong alignment on Conscientiousness, R2 also captured traits present in F0. These overlaps reflect the fluidity and interdependence of personality dimensions in behavioural data, where strict categorical boundaries are uncommon.

Importantly, despite these variations in cluster-to-cluster mapping, the underlying DISC trait profiles remained consistent across both models, reinforcing the structural validity of the 10-item version. This additional layer of unsupervised validation complements the high classification accuracy observed in earlier sections and provides strong evidence for the conceptual robustness of the 10-item DISC assessment.

Alignment between clustering and feature selection

To further substantiate the robustness of the reduced questionnaire, an empirical alignment was conducted between the behavioural dimensions of the most influential features

Table 9: Alignment between top 10 features and cluster-dominant behavioural dimensions.

Feature name	Behavioural dimension	Corresponding cluster
DO7	Dominance	Cluster 0
DO5	Dominance	Cluster 0
AD9	Conscientiousness	Cluster 3
AD7	Conscientiousness	Cluster 3
AS4	Influence	Cluster 1
AS6	Influence	Cluster 1
SC8	Steadiness	Cluster 2
SC3	Steadiness	Cluster 2
AD2	Conscientiousness	Cluster 3
DO3	Dominance	Cluster 0

and the dominant traits of each cluster. (Table 9) presents this alignment.

This alignment reveals a compelling correspondence between the clusters and the most influential features selected by RFE. Features reflecting Dominance predominantly align with Cluster 0, Influence features with Cluster 1, Steadiness features with Cluster 2, and Conscientiousness features with Cluster 3.

The convergence of these independent analytical approaches provides robust empirical support for the validity of the reduced DISC assessment. The clustering results not only validate the feature selection outcomes but also reinforce the potential of a shorter, optimised questionnaire to preserve the core behavioural structures of the full DISC model.

Discussion

The findings of this study offer a data-driven re-evaluation of the traditional DISC personality classification system, yielding several critical insights. Rather than replacing the established rule-based method, this research demonstrates how modern machine learning techniques can complement, enhance, and extend its capabilities.

The supervised learning experiments confirmed that machine learning models can reliably replicate the conventional rule-based DISC classifications. Both Logistic Regression and XGBoost emerged as top performers, each attaining an accuracy of 93.53%. Notably, Logistic Regression marginally outperformed XGBoost in precision, cross-validation mean accuracy, and Cohen's Kappa, indicating not only strong predictive capability but also greater stability and generalisability. The robust performance of Logistic Regression—a comparatively simple, linear model—suggests that the DISC classification task, as operationalised in this study, possesses a high degree of linear separability when paired with appropriate preprocessing and feature selection strategies. This finding is especially valuable for applied contexts, where model interpretability and deployment simplicity are essential.

The feature selection analysis demonstrated that substantial questionnaire length reduction is achievable without materially compromising predictive accuracy. Specifically, reducing the questionnaire to 10 carefully selected items maintained a high accuracy of 91.49%, closely approximating the performance of the full 40-item set. This finding underscores the feasibility and practicality of a more efficient DISC assessment tool. Such a reduction is particularly valuable in time-constrained environments such as organisational settings, corporate training programmes, or rapid talent screening contexts, where minimising respondent burden enhances both engagement and data quality.

The clustering analysis provided complementary insights, revealing that the DISC framework broadly aligns with natural groupings in the data. Although the Elbow method yielded a gradual rather than sharply defined inflection point—a common occurrence in personality and psychometric data, where human behaviour typically spans fluid continua rather than rigid divisions—the four-cluster solution remains conceptually robust and empirically justified. Similarly, while the Silhouette scores were moderate rather than high, this is entirely consistent with expectations for behavioural data, where overlapping traits and spectrum-based characteristics naturally lead to modest cohesion scores. The convergence of both methods reinforces the appropriateness of the selected cluster count and enhances the

credibility of the findings. This outcome reflects the nuanced reality of behavioural data, where clear-cut separations are rare, yet meaningful groupings can still be identified with appropriate analytical techniques.

A key innovation of this study lies in the comparative clustering analysis between the full 40-item and reduced 10-item feature sets. Despite the dimensionality reduction, the reduced model produced a similar four-cluster solution, confirmed through Elbow and Silhouette diagnostics, PCA visualisation, and trait profile analysis. Notably, DISC trait distributions within each cluster remained consistent across models, validating the structural integrity of the shortened instrument. While not all clusters aligned one-to-one—e.g., F2 and F1 both mapping to R1—these overlaps were explainable through shared Dominance characteristics. Such nuances reflect the continuous and overlapping nature of personality traits, reinforcing the validity of using reduced data representations in psychometric clustering.

Importantly, this finding also addresses longstanding concerns from personality psychology regarding the artificiality of rigid typological systems. DISC users and researchers have long observed that individuals often exhibit traits from more than one category, especially when score distributions are near-equal across dimensions. The proximity observed between clusters—particularly between Assertiveness and Steadiness—provides empirical evidence of such behavioural overlap. Rather than enforcing a strict one-label assignment, the machine learning models employed in this study offer probabilistic outputs that can reflect hybrid personality profiles. This flexibility introduces a more nuanced and inclusive interpretation of personality data, extending the DISC framework without discarding its practical strengths.

PCA visualisation and cluster profiling by average DISC dimension scores further validated this alignment. Each identified cluster exhibited clear dominance of a single DISC dimension: Cluster 0 (Dominance), Cluster 1 (Influence), Cluster 2 (Steadiness), and Cluster 3 (Conscientiousness). These same patterns emerged in the reduced-feature clustering model, where R0 through R3 showed parallel trait dominance. This empirical mapping provides quantifiable justification for the cluster-to-DISC alignment and reinforces the theoretical validity of the DISC model. Importantly, this step adds analytical depth to the study by confirming that the behavioural structures captured by unsupervised learning naturally reflect the established framework of DISC, even without pre-imposed labels.

It is worth noting that while PCA offered helpful visual confirmation of cluster distinctions, its linear nature may limit its ability to capture deeper, non-linear behavioural dynamics. Advanced techniques such as t-SNE or UMAP could provide more nuanced representations and should be considered in future visualisation efforts.

An especially noteworthy insight emerged from the relationship between feature selection outcomes and clustering profiles. The top-ranked features identified by RFE corresponded closely with the dominant behavioural dimensions of the clusters identified via unsupervised learning. For example, features representing Dominance primarily aligned with Cluster 0, while those associated with Conscientiousness corresponded with Cluster 3. This cross-validation between independent analytical processes provides compelling evidence that the reduced item set not only preserves overall predictive performance but also

captures the underlying behavioural structure of the data. Such alignment underscores the integrity of the proposed shortened questionnaire and reinforces its practical utility.

Collectively, these findings advance both theoretical understanding and practical application of DISC assessment. While acknowledging the practical utility of the DISC model, we also recognise its theoretical limitations, particularly its simplified typological structure and contested validity in academic psychology. This research does not aim to resolve these longstanding critiques, but instead demonstrates how computational enhancements can improve the interpretive depth of DISC in applied contexts.

From a theoretical standpoint, the results affirm the foundational validity of the four-category DISC model while illuminating areas of nuance, such as potential trait proximities and the existence of behavioural continua. From an applied perspective, the development of an empirically validated, reduced-item DISC assessment offers significant promise for enhancing efficiency and scalability in real-world settings.

Furthermore, the integration of supervised and unsupervised techniques exemplifies the value of hybrid methodological approaches in psychometric research. While supervised learning ensured fidelity to established classifications, unsupervised methods offered exploratory insights, confirming and extending the traditional model's assumptions. The alignment between these methods substantiates the robustness of the study's conclusions and illustrates how modern computational tools can enrich classical psychological frameworks.

Overall, this study does not propose to discard the established DISC method but rather to augment its effectiveness and flexibility. By integrating data-driven approaches with traditional psychometric frameworks, the research advances the development of more efficient, adaptive, and empirically validated personality assessment tools, with meaningful implications for both academic inquiry and organisational practice. The reduced-item DISC assessment offers valuable time efficiencies, particularly for use in fast-paced organisational contexts where rapid decision-making is critical. While the demographic diversity of the current sample supports generalisability, future research could explore broader cultural or professional samples to further enhance external validity. Future studies may also refine these insights by leveraging larger, more diverse samples, incorporating longitudinal data, or applying these methodologies to other psychometric instruments, thereby advancing the evolving field of computational psychometrics. While this study employed rigorous internal validation, including stratified data splitting and cross-validation, the generalisability of findings could be further enhanced by evaluating model performance on fully independent datasets. Future research could extend this work by testing the reduced DISC assessment and clustering structure across different populations or longitudinal samples to assess stability and external applicability.

The alignment between these methods substantiates the robustness of the study's conclusions and illustrates how modern computational tools can enrich classical psychological frameworks. Rather than aiming to supersede established methods, this research complements the rule-based approach, enhancing its diagnostic precision while preserving its foundational simplicity. Importantly, this balance ensures that improvements in analytical accuracy do not come at the cost of the intuitive interpretability that remains essential for practical applications of

the DISC model.

Conclusion

This study undertook a comprehensive examination of the DISC personality framework through the application of advanced machine learning methodologies. By integrating both supervised and unsupervised learning techniques, the research critically assessed the robustness of the traditional heuristic DISC scoring system, explored the feasibility of questionnaire optimisation, and provided empirical insights into the underlying behavioural structures captured within the data.

The supervised learning evaluation demonstrated that machine learning models, particularly Logistic Regression and XGBoost, can successfully replicate and potentially enhance the predictive performance of the conventional rule-based DISC classification method. Among the models tested, Logistic Regression emerged as the optimal choice, not only matching XGBoost in overall accuracy but also surpassing it in precision, cross-validation stability, and interpretability. This outcome underscores the suitability of Logistic Regression for real-world deployment, where both accuracy and explainability are vital.

The feature selection experiments yielded particularly impactful findings. A reduced questionnaire consisting of just ten strategically selected items preserved over 91% classification accuracy, closely mirroring the full 40-item assessment. This result confirms the feasibility of developing a more efficient, time-conscious DISC assessment tool, enhancing practicality without sacrificing diagnostic validity. Such a streamlined instrument holds clear potential for application in organisational settings, leadership development, and other time-sensitive environments where respondent burden is a critical consideration.

Complementing these findings, the unsupervised clustering analysis revealed a natural four-cluster structure within the data, closely aligning with the established four dimensions of the DISC model. The clear behavioural profiles exhibited by each cluster, as validated through empirical profiling of average item scores, reinforce the foundational validity of the DISC framework while also highlighting nuanced behavioural proximities—particularly between Assertiveness and Steadiness dimensions. To further validate the reduced-item model, a comparative clustering analysis was conducted, confirming that the same DISC trait structures were recoverable from both the full 40-item and reduced 10-item assessments. Despite minor mismatches in cluster-to-cluster mapping, the dominant behavioural dimensions within each grouping remained consistent, underscoring the structural integrity of the reduced form.

This structural equivalence was quantified through DISC trait profile comparisons and distance-based cluster alignment, providing empirical support for the conceptual fidelity of the shortened questionnaire. Such validation demonstrates not only that the reduced form replicates DISC labels accurately but also that it preserves the latent personality architecture inherent in the original assessment.

Collectively, these findings illustrate that integrating machine learning with traditional psychometric frameworks offers significant advantages. The study not only confirms the resilience of the DISC model but also demonstrates how data-driven methodologies can optimise its application, delivering a more agile and empirically grounded assessment approach. By validating both predictive accuracy and structural alignment, this study enhances confidence in the use of reduced-form assessments

for practical deployment without compromising psychological rigor. Beyond DISC specifically, this research contributes to the broader field of computational psychometrics by exemplifying how classical psychological models can be re-examined and enhanced through modern analytical techniques.

Looking ahead, future research could extend this work by testing the reduced-item DISC tool across diverse populations and settings, incorporating longitudinal data to examine stability over time, or applying similar methodologies to other personality models. Such endeavours would further refine the application of machine learning in psychological assessment and continue advancing the development of efficient, scalable, and scientifically robust personality profiling tools.

Declarations

Data availability: All anonymised data used in this study are publicly available via Zenodo at <https://doi.org/10.5281/zenodo.15226178>. This open-access repository ensures full transparency and enables replication or secondary analysis by other researchers.

Competing interests: The authors declare no competing interests.

Funding declaration: The authors received no specific funding for this work.

Correspondence and requests for materials should be addressed to Mohammad Hossein Amirhosseini.

References

- Furnham A. The bright and dark side of personality: The relationship between personality traits and work behaviours. *Pers Individ Dif*. 2015; 78: 1–12.
- John OP, Naumann LP, Soto CJ. Paradigm shift to the integrative Big Five trait taxonomy. In: John OP, Robins RW, Pervin LA, eds. *Handbook of Personality: Theory and Research*. New York: Guilford Press; 2008: 114–158.
- McCrae RR, Costa PT Jr. The Five-Factor Theory of Personality. In: John OP, Robins RW, Pervin LA, eds. *Handbook of Personality: Theory and Research*. New York: Guilford Press; 2008: 159–181.
- Myers and Briggs Foundation. MBTI basics. 1998. Available at: <https://www.myersbriggs.org>. Accessed 8 Apr 2025.
- Pittenger DJ. Cautionary comments regarding the Myers-Briggs Type Indicator. *Consult Psychol J Pract Res*. 2005;57: 210–221.
- Marston W. *Emotions of Normal People*. New York: Harcourt, Brace and Company; 1928.
- Bradberry T. *TalentSmart DISC Profile*. San Diego: TalentSmart; 2005.
- Furnham A. *The Personality at Work: The Role of Individual Differences in the Workplace*. London: Routledge; 1990.
- Rentfrow PJ. The role of music in everyday life. *Soc Pers Psychol Compass*. 2010; 4: 1090–1103.
- O’Neill TA, Hambley LA, Bercovich A. Prediction of personality profiles in virtual teams: Combining behavioural observations with text analysis. *Comput Hum Behav*. 2014; 39: 213–222.
- Gough HG. Assessment of personality. *Annu Rev Psychol*. 1992; 43: 417–444.
- DISC Profile. What is DISC? 2024. Available at: <https://www.discprofile.com>.
- Paulhus DL, John OP. Egoistic and moralistic biases in self-perception. *J Pers*. 1998; 66: 1025–1060.
- Yarkoni T. Personality in 100,000 words. *J Res Pers*. 2010; 44: 363–373.
- Plonsky O, Osher Y. Predicting personality using social networks and other textual data. In: *Proc Int Conf Weblogs Soc Media (ICWSM-14)*. Palo Alto: AAAI Press; 2014: 8–17.
- Amirhosseini MH, Kazemian H. Machine learning approach to personality type prediction based on the Myers–Briggs Type Indicator. *Multimodal Technol Interact*. 2020; 4: 9.
- Roy A, Das A. Personality detection in social media text using ensemble machine learning models. *J Comput Soc Sci*. 2022; 5: 1–20.
- Ontoum Z, Chan KL. Personality prediction using online textual data with ensemble methods. *Appl Intell*. 2022; 52: 2000–2018.
- Sheth P, Pandhare R. Comparative analysis of machine learning models for personality prediction. In: *Proc Int Conf Data Sci Eng. Piscataway: IEEE*; 2021: 1–8.
- Chen T, Guestrin C. XGBoost: A scalable tree boosting system. In: *Proc 22nd ACM SIGKDD Int Conf Knowledge Discovery and Data Mining*. New York: ACM; 2016: 785–794.
- Vaddem SB, Agarwal M. Personality prediction from non-English textual data using advanced machine learning algorithms. *Expert Syst Appl*. 2023; 214: 119022.
- Devlin J, Chang MW, Lee K, Toutanova K. BERT: Pre-training of deep bidirectional transformers for language understanding. In: *Proc NAAACL-HLT 2019*. Stroudsburg: ACL; 2019: 4171–4186.
- Garcia dos Santos M, Paraboni I. Investigating personality recognition from online text with pre-trained language models. *Appl Soft Comput*. 2021; 107: 107402.
- Hasan A, et al. Feature selection in high-dimensional data: A comprehensive review. *Artif Intell Rev*. 2022; 55: 451–506.
- Zhao Z, et al. A survey on feature selection: Performance, stability, and knowledge discovery. *IEEE Trans Pattern Anal Mach Intell*. 2022; 44: 6679–6699.
- MacQueen J. Some methods for classification and analysis of multivariate observations. In: *Proc 5th Berkeley Symp Math Stat Probab. Vol 1*. Berkeley: University of California Press; 1967: 281–297.
- Jain AK. Data clustering: 50 years beyond K-means. *Pattern Recognit Lett*. 2010; 31: 651–666.
- Rousseeuw PJ. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J Comput Appl Math*. 1987; 20: 53–65.
- Ketchen DJ Jr, Shook CL. The application of cluster analysis in strategic management research: An analysis and critique. *Strategic Manage J*. 1996; 17: 441–458.
- Ackerman PL, Beier ME, Boyle MO. Working memory and intelligence: The same or different constructs? *Psychol Bull*. 2021; 147: 8–29.
- Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: Synthetic minority over-sampling technique. *J Artif Intell Res*. 2002; 16: 321–357.
- Fernández A, et al. *Learning from Imbalanced Data Sets*. Cham: Springer; 2018.
- Han J, Kamber M, Pei J. *Data Mining: Concepts and Techniques*. 3rd ed. Burlington: Morgan Kaufmann; 2011.

34. Inscape Publishing. Everything DiSC Manual. Hoboken: John Wiley & Sons; 2001.

Appendix A – Full DISC questionnaire items

Participants rated the following 40 items using a 5-point Likert scale, where 1 = Strongly Disagree and 5 = Strongly Agree. All statements began with “I ...”.

Assertiveness / Influence (AS1–AS10):

- AS1: Express myself easily
- AS2: Try to lead others
- AS3: Automatically take charge
- AS4: Know how to convince others
- AS5: Am the first to act
- AS6: Take control of things
- AS7: Wait for others to lead the way
- AS8: Let others make the decisions
- AS9: Am not highly motivated to succeed
- AS10: Can’t come up with new ideas

Steadiness / Calm (SC1–SC10):

- SC1: Feel comfortable around people
- SC2: Don’t mind being the center of attention
- SC3: Am good at making impromptu speeches
- SC4: Express myself easily
- SC5: Have a natural talent for influencing people
- SC6: Hate being the center of attention
- SC7: Lack the talent for influencing people
- SC8: Often feel uncomfortable around others
- SC9: Don’t like to draw attention to myself
- SC10: Have little to say

Adventurousness / Discipline (AD1–AD10):

- AD1: Prefer variety to routine
- AD2: Like to visit new places
- AD3: Interested in many things
- AD4: Like to begin new things
- AD5: Prefer to stick with things that I know
- AD6: Dislike changes
- AD7: Don’t like the idea of change
- AD8: Am a creature of habit
- AD9: Dislike new foods
- AD10: Am attached to conventional ways

Dominance / Decisiveness (DO1–DO10):

- DO1: Try to surpass others' accomplishments
- DO2: Try to outdo others
- DO3: Am quick to correct others
- DO4: Impose my will on others
- DO5: Demand explanations from others
- DO6: Want to control the conversation
- DO7: Am not afraid of providing criticism
- DO8: Challenge others' points of view
- DO9: Lay down the law to others
- DO10: Put people under pressure