mathematics

MDPI

Article

Density-Based Outlier Detection of Atypical Timeout Request Contexts in EuroLeague Basketball

José Miguel Contreras García [1,*], Elena Molina Portillo [1] and Juan Manuel Fernández Luna [2]

[1] Department of Didactics of Mathematics, Faculty of Education Sciences, University of Granada, 18071 Granada, Spain; elemo@ugr.es

[2] Department of Computer Science and Artificial Intelligence, School of Computer and Telecommunication Engineering, University of Granada, 18071 Granada, Spain; jmfluna@ugr.es

* Correspondence: jmcontreras@ugr.es

Abstract

Timeouts in elite basketball are scarce, time-constrained interventions requested under evolving pressure and contextual imbalance. Although previous clustering approaches have identified recurrent timeout profiles, less attention has been paid to structurally atypical timeout request contexts that do not conform to dominant patterns. This study proposes a density-based outlier detection framework for EuroLeague basketball. Using play-by-play-derived data from 3743 games across 15 seasons, 25,031 timeout events were represented through differential and rate-adjusted indicators of scoring, shooting, free throws, rebounding, assists, turnovers, fouls, blocks, and score imbalance. Principal component analysis was used to construct a compact multivariate feature space. DBSCAN and OPTICS served as diagnostic density-based comparators, whereas HDBSCAN was used to derive observation-level outlier scores. Atypical contexts were defined using the highest outlier score events, with the top 10% group as the primary threshold. Flat density-based partitions did not support a compact timeout taxonomy, whereas HDBSCAN provided a stable and interpretable outlier score ranking. The most atypical timeouts were associated with larger score imbalances and rate-adjusted performance differentials, but an ablation analysis excluding score margin and score rate variables showed that the outlier ranking structure was not solely determined by direct scoreboard information.

Keywords: density-based clustering; outlier scoring; anomaly detection; HDBSCAN; DBSCAN; OPTICS; principal component analysis; timeout; EuroLeague; sports analytics

MSC: 62H30; 62H25; 62-07

Check for updates

Academic Editor: Juan Ramón Torregrosa Sánchez

Received: 14 June 2026

Revised: 10 July 2026

Accepted: 24 July 2026

Published: 27 July 2026

1. Introduction

In many sequential decision environments, interventions are not made at random but under evolving states of pressure, uncertainty, and incomplete information. These interventions are often scarce, time-constrained, and context-sensitive, making their analysis relevant not only from a domain-specific perspective but also from the standpoint of mathematical modeling, data mining, and decision support systems. In such settings, the analytical problem is not merely to determine whether an intervention is followed by a favorable outcome, but to characterize the types of states in which interventions are requested and to identify whether some of these states correspond to recurrent patterns, rare configurations, or atypical decision contexts. This issue is especially relevant in high-frequency event streams, where each decision can be represented as a point in a multidimensional feature space constructed from contemporaneous state descriptors.

Basketball timeouts provide a particularly suitable empirical context for this problem. In elite basketball, timeouts are limited and strategically valuable interruptions that allow for coaches to reorganize offensive and defensive structures, manage fatigue, disrupt opponent runs, communicate tactical adjustments, or stabilize play under pressure [1–4]. Previous research has shown that timeout usage and post-timeout performance are associated with score margin, game period, momentum-related dynamics, and situational performance indicators [3–6]. However, the interpretation of timeouts remains methodologically challenging because they are typically requested precisely when the team is experiencing some form of negative state. Consequently, descriptive post-timeout improvements may reflect regression to the mean or contextual selection rather than a causal effect of the timeout itself [5,6]. For this reason, before estimating timeout effectiveness, it is necessary to build robust descriptive frameworks that characterize the game contexts in which timeouts are requested.

This challenge is particularly relevant in EuroLeague basketball. Under FIBA/EuroLeague rules, timeout availability is limited and constrained by game phase, which makes timeout requests a resource allocation decision intertwined with time remaining, score state, foul situation, offensive efficiency, and possession control [7]. Recent studies have examined timeout effectiveness, after-timeout offensive patterns, and critical moment actions in European basketball competitions [8,9]. In addition, previous work has shown that hierarchical clustering can recover meaningful structural patterns in EuroLeague timeout requests, suggesting that unsupervised learning can be used to describe recurring timeout contexts at scale [10]. Nevertheless, most existing work has focused either on post-timeout outcomes, specific situational correlates, or recurrent cluster structures. Less attention has been paid to the unsupervised identification of structurally atypical timeout-request contexts, especially when such contexts do not form compact or balanced groups.

Unsupervised learning offers a principled framework for this purpose. Clustering methods compress complex, high-dimensional representations into interpretable structures and have been widely used in data mining, pattern recognition, and exploratory multivariate analysis [11,12]. In sports analytics, and particularly in basketball, clustering has been applied to performance profiles, player roles, lineup structures, playing styles, and game-related patterns, usually with the aim of identifying recurrent and internally coherent groups [13–15]. In the specific context of timeout analysis, hierarchical and partitional approaches are useful when the goal is to derive broad structures or compact taxonomies of recurrent timeout profiles. However, such approaches are less directly aligned with the detection of unusual observations because they tend to organize the full sample into global structures or force each observation into a cluster.

Density-based methods provide a complementary perspective. Unlike centroid-based or hierarchical procedures, density-based algorithms define clusters as regions of high point concentration separated by lower-density areas and can explicitly identify observations located in sparse regions of the feature space [16–20]. This property is central for the present study. In the analysis of timeout requests, low-density observations should not be treated only as algorithmic noise or failure cases; they may represent structurally unusual game states, special tactical circumstances, or atypical timeout contexts that are absorbed into broader groups by methods that force every observation into a cluster. DBSCAN identifies dense regions using the parameters ε and MinPts and is especially useful when clusters may have irregular shapes and when sparse observations are substantively relevant [16,20]. OPTICS extends this logic by producing an ordering that captures density structure across a range of neighborhood thresholds, reducing the dependence on a single global ε value [17].

https://doi.org/10.3390/math14152696

HDBSCAN further generalizes density-based clustering by constructing a hierarchy of density-connected components and providing outlier scores that quantify the degree to which observations depart from dense regions of the data [18,19].

This methodological distinction is important for timeout analysis. If the objective is to obtain a compact taxonomy of recurrent timeout profiles, partition-based methods such as k-means, k-medoids, or fuzzy c-means are appropriate. If the objective is to explore nested structure, hierarchical methods are informative. However, if the objective is to identify unusual timeout requests and observations that do not conform to dominant profiles, density-based outlier detection is better aligned with the research problem. In this sense, the present study shifts the analytical focus from recurrent timeout profiles to low-density timeout events. Rather than asking only which common types of timeout situations exist, it asks which timeout requests appear structurally unusual in the multivariate game state space and whether these unusual cases correspond to interpretable basketball situations.

Accordingly, this study applies density-based methods to a large-scale EuroLeague play-by-play dataset comprising timeout events from 3743 games over 15 seasons. Each timeout is represented through play-by-play-derived performance indicators describing the game state at the moment of the request, including score-related, shooting, free-throw, rebounding, assist, turnover, foul, and block differentials. To avoid the dominance of raw cumulative game volume, the feature space is constructed using differential and rate-adjusted indicators, while temporal and contextual variables are reserved for contextual characterization and interpretive assessment. Principal component analysis is used as a dimensionality reduction step to mitigate collinearity and construct a compact feature space for density-based analysis [21]. DBSCAN and OPTICS are used as diagnostic comparators, whereas HDBSCAN is used primarily to derive a density-based outlier score ranking. Atypical timeout contexts are operationalized through the highest outlier score observations, with top 5%, top 10%, and top 15% thresholds used to assess robustness. Internal validation, silhouette-based summaries, external contextual validation, and stability analyses are used to evaluate the consistency and interpretability of the resulting atypical profiles [22].

The study has three main contributions. First, it extends previous timeout context clustering research by introducing a density-based outlier detection framework designed to identify structurally atypical timeout situations rather than only recurrent profiles. Second, it develops a reproducible analytical workflow combining engineered game state representations, dimensionality reduction, density-based modeling, internal validation, outlier score thresholding, and stability assessment. Third, it characterizes the resulting low-density timeout events using contextual indicators not directly entered as binary predictors in the density-based feature space, including game phase and score margin conditions. Importantly, the study is descriptive rather than causal: it does not estimate whether timeouts improve subsequent performance, but it characterizes the multivariate contexts in which unusual timeout decisions occur and shows how density-based outlier scoring can complement hierarchical and partitional timeout taxonomies.

2. Materials and Methods

2.1. Study Design

This study was designed as a descriptive, unsupervised data-mining analysis aimed at identifying atypical timeout request contexts in EuroLeague basketball. The purpose was not to estimate the causal effect of timeouts on subsequent performance, nor to derive an optimal timeout policy. Instead, the study focused on characterizing the multivariate game state conditions under which timeout requests appear structurally unusual relative to the dominant density structure of the data.

The analytical strategy was based on density-based outlier detection. In contrast to hierarchical or partitional clustering approaches, which are primarily oriented toward recovering recurrent groups or compact taxonomies [11,12], the present approach uses density information to identify timeout events located in sparse regions of the multivariate feature space [18,19]. Accordingly, flat cluster labels were treated as diagnostic outputs, whereas HDBSCAN outlier scores were used as the primary basis for identifying atypical timeout contexts.

The workflow comprised six stages: (i) construction of timeout-level game state representations from play-by-play data; (ii) engineering of differential and rate-adjusted performance indicators; (iii) dimensionality reduction using principal component analysis (PCA); (iv) density-based modeling using DBSCAN, OPTICS, and HDBSCAN; (v) operational definition of atypical timeout contexts through HDBSCAN outlier score thresholds; and (vi) internal, external, and temporal validation of the resulting atypical profiles.

2.2. Dataset and Unit of Analysis

The empirical dataset consisted of EuroLeague play-by-play records from 3743 official games across 15 seasons. The unit of analysis was the timeout event. Each timeout was represented by the observed game state at the moment the timeout was requested, using only information available up to that point in the game. This event-level representation was designed to capture contemporaneous timeout request contexts rather than post-timeout outcomes.

The original data source included play-by-play-derived indicators related to scoring, shooting, free throws, rebounding, assists, turnovers, fouls, blocks, and time-related information. After preprocessing and exclusion of incomplete or non-analytic records, the final analytic sample comprised 25,031 timeout events. Overtime timeouts were retained because they may represent distinctive decision contexts.

For each timeout event, the requesting side was identified from the event-level timeout indicators. The requesting team was successfully detected for all 25,031 timeout events: 11,961 timeouts were requested by the local team and 13,070 by the visiting team, with no undetected cases.

2.3. Feature Engineering

A central methodological decision was to avoid representing timeout contexts only through raw cumulative indicators, since cumulative game variables tend to be strongly associated with elapsed time. If raw accumulated points, rebounds, assists, and shot attempts are used directly, density-based algorithms may primarily detect game progression rather than genuinely atypical performance contexts. Therefore, the clustering feature space was constructed using differential and rate-adjusted variables.

First, for each timeout event, the team requesting the timeout was identified using the event-level timeout indicators. Signed performance differentials were then expressed from the behavioral perspective of the requesting team rather than from the fixed local-versus-visiting perspective. Specifically, for each local-minus-visiting differential, the corresponding requesting-team-minus-opponent differential was computed by retaining the sign when the timeout was requested by the local team and reversing the sign when it was requested by the visiting team. Therefore, positive signed values indicate that the requesting team was ahead of the opponent on the corresponding indicator, whereas negative values indicate that the requesting team was behind.

This transformation was applied to scoring, two-point field goals, three-point field goals, free throws, offensive and defensive rebounds, assists, turnovers, fouls, blocks, and valuation indicators. Rate-adjusted versions of selected differentials were then created by

https://doi.org/10.3390/math14152696

dividing cumulative requesting-team-minus-opponent differences by elapsed game time. Score imbalance indicators were also computed, including signed score differential from the requesting team perspective, absolute score differential, requesting team score rate differential, and absolute score rate differential. Absolute variables were retained as direction-free measures of imbalance, whereas signed variables were interpreted behaviorally from the perspective of the team that requested the timeout.

Formally, let ${D_i}^{LocVis}$ denote the original local-minus-visiting differential for timeout event i. The behaviorally oriented differential was computed as

$$D_i^{ReqOpp} = s_i D_i^{LocVis},$$

where $s_i = 1$ when the timeout was requested by the local team and $s_i = -1$ when the timeout was requested by the visiting team. Thus, D_i^{ReqOpp} always represents the requesting team minus the opponent. Positive values indicate an advantage for the team requesting the timeout, whereas negative values indicate a disadvantage for that team. Absolute score difference and absolute rate difference variables were retained as unsigned measures of imbalance magnitude.

Temporal and contextual variables such as game minute, late-game status, final-two-minute status, overtime, close-score situations, and large-margin situations were not used to construct the density-based feature space as binary contextual predictors. Instead, these variables were reserved for contextual characterization and interpretive assessment.

2.4. Standardization and Dimensionality Reduction

All engineered numerical features were z-standardized before dimensionality reduction. Principal component analysis was then applied to the standardized feature matrix. PCA was used for two reasons: first, to mitigate collinearity among related basketball indicators; and second, to construct a compact Euclidean feature space suitable for density-based analysis [21].

The number of retained components was selected to preserve a substantial proportion of the original variance while avoiding unnecessary dimensionality. In the final specification, seven principal components were retained, explaining more than 80% of the total variance. Density-based models were fitted on the retained PCA scores rather than on the original standardized variables. The original engineered variables were preserved for post hoc interpretation of atypical timeout contexts.

2.5. Density-Based Algorithms

Three density-based methods were considered: DBSCAN, OPTICS, and HDBSCAN.

DBSCAN identifies dense regions using two parameters: the neighborhood radius ε and the minimum number of points required to form a dense region, MinPts [16]. Observations that do not satisfy the density connectivity criteria are labeled as noise. DBSCAN was used as a baseline density-based method and as a diagnostic tool for examining sensitivity to ε and MinPts. This use is consistent with later methodological work emphasizing that DBSCAN remains useful when its parameterization and noise structure are interpreted appropriately [20].

OPTICS extends the DBSCAN logic by producing an ordering of observations according to reachability distance, allowing for density structure to be inspected across a range of neighborhood thresholds rather than relying on a single global ε value [17]. OPTICS was used as a diagnostic method to assess whether the timeout context space exhibited heterogeneous density structure.

HDBSCAN generalizes density-based clustering by constructing a hierarchy of density-connected components and selecting clusters according to stability [18,19]. In the present

https://doi.org/10.3390/math14152696

study, HDBSCAN was not selected because its flat partition provided a compact taxonomy of timeout contexts. Rather, it was selected because it provides observation-level outlier scores that quantify the degree to which each timeout event departs from dense regions of the feature space [19]. This distinction is central to the analysis: flat density-based partitions were treated as diagnostic summaries of the density structure, whereas substantive interpretation was based on the continuous outlier score ranking.

These methods were selected because timeout request data are expected to contain heterogeneous density regions, sparse observations, and context-dependent imbalance rather than a small number of compact, equally dense clusters. In basketball, timeout events arise under diverse combinations of score margin, elapsed time, scoring rates, rebounds, assists, turnovers, fouls, and shooting indicators. This makes density-based methods appropriate for assessing whether some timeout events lie in low-density regions of the multivariate game state space. DBSCAN and OPTICS were therefore used as diagnostic tools for examining density heterogeneity and noise structure, whereas HDBSCAN was retained as the primary outlier scoring method because it can represent hierarchical density structure and provides observation-level outlier scores.

2.6. Parameter Search and Internal Validation

For DBSCAN, a grid search was conducted across ε and MinPts values. For OPTICS, the parameter search included ε, MinPts, and ξ. For HDBSCAN, MinPts was varied across a predefined range. Each candidate solution was evaluated using internal validation summaries, including the number of non-noise clusters, number and proportion of noise observations, silhouette values computed on non-noise observations, and density-aware validation measures where applicable.

Silhouette values were interpreted cautiously because they are not specifically designed for density-based outlier detection and may be sensitive to the treatment of noise observations [22]. Therefore, internal validation was not used to select a conventional "best flat clustering" solution. Instead, the validation stage was used to assess whether density-based modeling produced a meaningful and stable outlier score structure. DBSCAN and OPTICS were retained as diagnostic comparators, whereas HDBSCAN was selected as the primary method because it provides an interpretable outlier score ranking [18,19].

Because the substantive objective was to identify structurally atypical timeout events rather than to derive strategically meaningful flat clusters, internal cluster validation indices such as silhouette values were used only as diagnostic summaries. Domain-based assessment was instead conducted by examining whether high outlier score events corresponded to interpretable basketball contexts, including game phase, score imbalance, close-score and large-margin situations, and performance rate differentials.

The final HDBSCAN configuration was selected for outlier score analysis rather than for flat cluster interpretation. The selected MinPts value was retained as a pragmatic low-neighborhood configuration that preserved local density sensitivity and yielded stable outlier score rankings in subsequent robustness analyses. The flat HDBSCAN labels were not interpreted as a taxonomy, because density-based partitions either became highly fragmented and noise-dominated in DBSCAN and OPTICS or collapsed into a dominant component under the selected HDBSCAN configuration. Thus, the role of the density-based diagnostics was to justify an outlier scoring interpretation rather than to identify a small number of substantive clusters.

2.7. Operational Definition of Atypical Timeout Contexts

Atypical timeout contexts were operationalized using the HDBSCAN outlier score distribution. Timeout events with the highest outlier scores were interpreted as structurally atypical because they were located in lower-density regions of the multivariate game state space.

Three percentile-based thresholds were used: top 5%, top 10%, and top 15% highest outlier scores. The top 10% threshold was retained as the primary operational definition because it provides a pragmatic balance between selectivity and sample size for descriptive characterization. This threshold should not be interpreted as a natural, empirically calibrated boundary separating atypical from non-atypical timeouts. Instead, it is an operating convention used to summarize the upper tail of the HDBSCAN outlier score distribution. Robustness was assessed by comparing the top 10% results with the more selective top 5% and broader top 15% thresholds, and by conducting the score margin and score rate ablation analysis.

2.8. Characterization of Atypical Timeout Contexts

To characterize atypical timeout contexts, events in the top outlier score groups were compared with all other timeout events. Descriptive statistics were computed for engineered performance indicators and external contextual variables. For each variable, group means, medians, standardized mean differences, and non-parametric tests were calculated. Cohen's d was used to quantify the magnitude of differences between atypical and non-atypical timeout contexts [23]. Where multiple variables were compared, adjusted p-values were computed using a false-discovery-rate procedure [24].

The primary interpretation focused on the top 10% outlier score group. Variables with the largest absolute standardized differences were used to describe the multivariate profile of atypical timeout requests. This interpretation emphasized effect size patterns rather than statistical significance alone, given the large sample size.

2.9. Contextual Characterization and Validation

Contextual characterization and validation were conducted using game state indicators that were not directly entered as binary contextual predictors in the density-based modeling stage. These indicators included game minute, early-game phase, mid-game phase, late-game status, final-two-minute status, overtime, close-score situations, and large-margin situations. Because close-score and large-margin indicators are derived from score differential information, they were interpreted as contextual summaries of competitive imbalance rather than as fully independent validation variables. Accordingly, this stage was used to assess whether the density-based outlier score ranking corresponded to interpretable basketball contexts, not to provide an external validation based on independent labels.

Close-score situations were defined as timeout events occurring when the absolute score differential was five points or fewer. Large-margin situations were defined as events with an absolute score differential of 15 points or more. Early-game situations were defined as events occurring between minutes 0 and 15, mid-game situations as events between minutes 15 and 30, late-game situations as events occurring from minute 35 onward, and final-two-minute situations as events occurring from minute 38 onward. Overtime events were identified as events occurring after minute 40.

Two complementary contextual assessment strategies were used. First, enrichment analyses compared the prevalence of atypical timeout contexts inside and outside each external context. Risk ratios, odds ratios, and Fisher or chi-square tests were computed to assess whether atypical timeout events were overrepresented or underrepresented in specific game situations [25]. Second, logistic regression models were fitted to examine whether the top outlier score indicators were associated with interpretable contextual conditions [26]. To avoid collinearity among time-related variables, separate models were estimated: one model included continuous minute and absolute score differential, whereas another model included categorical contextual indicators such as early phase, late-game status, final-two-minute status, overtime, close-score status, and large-margin status. Model discrimination was summarized using the area under the receiver operating characteristic curve [27].

2.10. Ablation and External Enrichment Analyses

To address the possibility that the association between outlier scores and score margin contexts could be mechanically driven by the inclusion of score-related variables in the feature space, an ablation analysis was conducted. In this ablation model, all score-margin and score rate variables were removed before standardization and PCA. Specifically, requesting team score differential, absolute score differential, requesting team score rate differential, and absolute score rate differential were excluded. The complete PCA–HDBSCAN pipeline was then refitted using the remaining performance-based variables.

The resulting ablated HDBSCAN outlier score ranking was compared with the full requesting-team-minus-opponent HDBSCAN ranking. Spearman rank correlation was used to compare the continuous outlier score rankings, and Jaccard similarity was used to compare the overlap of the top 5%, top 10%, and top 15% outlier score sets. This analysis evaluated whether the density-based atypicality structure persisted after removing direct score margin and score rate information.

In addition, enrichment analyses were extended to variables external to the engineered feature space. These variables included season, requesting team, opponent team, and requesting side. Where available, coach-related variables were also inspected as external grouping factors. These variables were not used to construct the PCA–HDBSCAN feature space. Their association with top 10% outlier status was summarized using prevalence differences and Cramer's V for categorical association. Lineup-level information and possession sequence descriptors were not available in the play-by-play-derived dataset and were therefore not included.

2.11. Stability Analysis

The stability of the density-based outlier ranking was assessed using two complementary procedures.

First, a random subsampling analysis was performed. In each iteration, a fixed proportion of the data was sampled without replacement, the PCA and HDBSCAN pipeline was refitted, and the resulting outlier score ranking was compared with the full-sample ranking on the overlapping observations. Stability was quantified using Spearman rank correlation for outlier scores and Jaccard similarity for the top 10% outlier set [28,29].

Second, a leave-one-season-out validation analysis was conducted. For each season, the model was refitted after excluding that season, and the resulting outlier score ranking was compared with the full-sample solution for the remaining observations. Again, Spearman rank correlation and Jaccard similarity for the top 10% outlier set were used as stability indicators. This analysis was included to assess whether atypical timeout contexts were stable across seasons rather than being driven by a single competitive year.

2.12. Comparison with a Recurrent Timeout Taxonomy

To examine whether density-based outlier detection captured information distinct from recurrent timeout profiles, the HDBSCAN-derived results were compared with the three-cluster k-means taxonomy previously established and validated by Contreras García et al. [30]. In that study, Euclidean k-means with k = 3 was selected through a consensus model selection procedure as a compact and interpretable representation of recurrent EuroLeague timeout contexts.

In the present study, k = 3 was retained as an externally established benchmark. However, to avoid confounding the comparison with differences in dimensionality, the k-means model was refitted using the same seven retained principal components used in the HDBSCAN analysis. Thus, the comparison held the feature space constant and assessed whether the density-based output reproduced the recurrent three-profile partitional structure.

Agreement between the flat HDBSCAN labels and the k-means labels was assessed using the Adjusted Rand Index [31]. In addition, the prevalence of events belonging to the top 10% of the HDBSCAN outlier score distribution was calculated within each k-means cluster. This analysis examined whether atypical timeout contexts were concentrated within a single recurrent profile or distributed across multiple recurrent timeout situations.

2.13. Software and Reproducibility

All analyses were conducted in R version 4.4.2 [32]. Data preprocessing, feature engineering, descriptive summaries, and visualization were performed using the tidyverse ecosystem [33]. PCA was conducted using the FactoMineR package [34]. Density-based modeling was performed using the dbscan package, including DBSCAN, OPTICS, and HDBSCAN implementations [35]. Cluster validation and silhouette calculations were performed using the cluster package [36]. Agreement between partitions was assessed using the adjustedRandIndex function from the mclust package [37]. Logistic regression models, enrichment analyses, ablation analyses, and stability analyses were implemented through reproducible R scripts. The main random seed was set to 5665; random subsampling iterations used 5665 plus the iteration number.

To improve reproducibility, detailed preprocessing and computational information is provided in the Supplementary Materials. Table S1 reports the raw data source, included seasons, timeout identification rule, requesting-side identification, exclusion log, final sample size, feature orientation rule, missing value handling, outlier handling, PCA retention rule, HDBSCAN parameter grid, selected HDBSCAN parameter, atypicality thresholds, ablation design, k-means benchmark, random seeds, R version, and package-version file. Table S2 provides the variable dictionary for the engineered requesting-team-minus-opponent feature space. Table S3 reports the parameter grids used for DBSCAN, OPTICS, and HDBSCAN. Table S4 lists the package versions used in the revised analysis.

3. Results

3.1. Engineered Feature Space and Principal Component Structure

All results reported below are based on the revised requesting-team-minus-opponent feature space. The requesting side was identified for all 25,031 timeout events, and all signed performance differentials were oriented from the perspective of the team that requested the timeout before standardization, PCA, and density-based modeling.

The final density-based analysis was conducted on an engineered feature space designed to reduce the dominance of raw cumulative game volume. Instead of using only accumulated box score indicators, the analytical representation incorporated requesting-team-minus-opponent differentials, rate-adjusted differentials, signed and absolute score differential indicators, and score rate imbalance measures. Temporal variables such as game minute, late-game status, final-two-minute status, overtime, close-score status, and large-margin status were excluded from the density-based feature space and retained for contextual characterization and interpretive assessment.

Principal component analysis provided a compact representation of this engineered feature space. As shown in Table 1 and Figure 1, seven principal components were retained, explaining 81.00% of the total variance. The variance was more evenly distributed than in a raw cumulative representation: PC1 explained 18.77%, PC2 17.55%, PC3 14.40%, PC4

11.53%, PC5 7.11%, PC6 6.55%, and PC7 5.10% of the variance. This supported the use of the PCA space for density-based analysis, since no single temporal or cumulative dimension dominated the representation.

Table 1. Variance explained by the retained principal components.

Component	Eigenvalue	Variance Explained (%)	Cumulative Variance (%)
PC1	6.0069	18.77	18.77
PC2	5.6146	17.55	36.32
PC3	4.6071	14.40	50.71
PC4	3.6894	11.53	62.24
PC5	2.2749	7.11	69.35
PC6	2.0949	6.55	75.90
PC7	1.6319	5.10	81.00

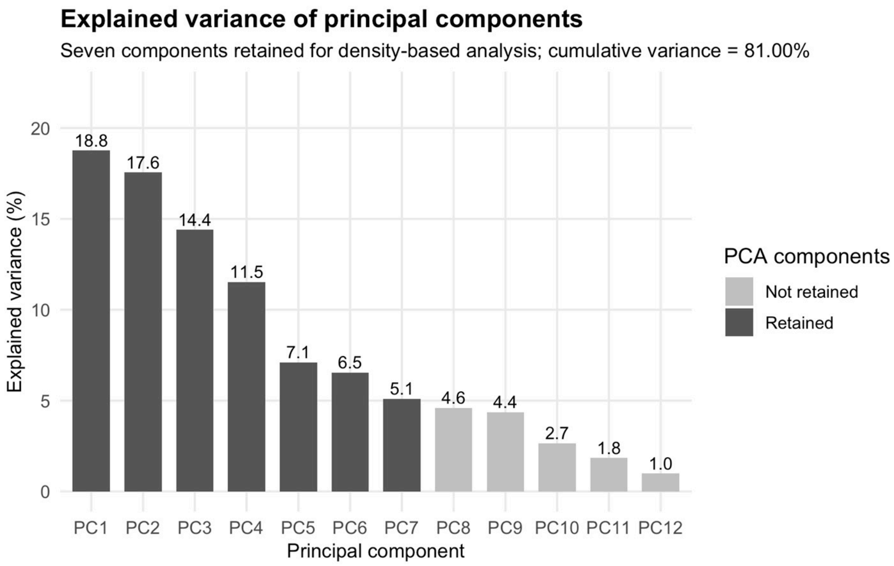

Figure 1. Explained variance of the engineered timeout context feature space. Note. Dark bars indicate the seven principal components retained for density-based analysis; light bars indicate additional non-retained components shown for reference. The retained seven-component solution explained 81.00% of the total variance and avoided dominance by a single cumulative temporal dimension.

3.2. Diagnostic Comparison of Density-Based Methods

DBSCAN, OPTICS, and HDBSCAN were evaluated to assess the density structure of the revised requesting-team-minus-opponent timeout-context space. The three methods provided complementary diagnostic information but did not support the interpretation of density-based flat labels as a compact timeout taxonomy.

DBSCAN and OPTICS yielded highly fragmented, noise-dominated partitions. In the selected diagnostic summaries, DBSCAN produced 128 non-noise clusters and labeled 97.36% of observations as noise, whereas OPTICS produced 256 non-noise clusters and labeled 93.47% of observations as noise. These results indicate that the timeout context space contains substantial density heterogeneity, and that flat density-based extraction is highly sensitive to sparse regions.

The selected HDBSCAN configuration showed a different limitation. Rather than producing hundreds of interpretable clusters, it generated a flat partition dominated by one overwhelmingly large component, with two non-noise clusters and 129 noise observations. The largest cluster accounted for almost all non-noise observations. Therefore, the HDBSCAN flat partition was also unsuitable as a substantive taxonomy of timeout contexts. However, HDBSCAN remained appropriate for the purpose of the study because it provided observation-level outlier scores, which were used to rank timeout events according to their degree of departure from dense regions of the multivariate feature space.

The variation in the HDBSCAN diagnostic metrics across MinPts values is shown in Figure 2. These diagnostic results are summarized in Table 2. Overall, the density-based models did not recover a compact alternative taxonomy of timeout requests. Instead, they supported the interpretation of density-based modeling as an outlier scoring procedure.

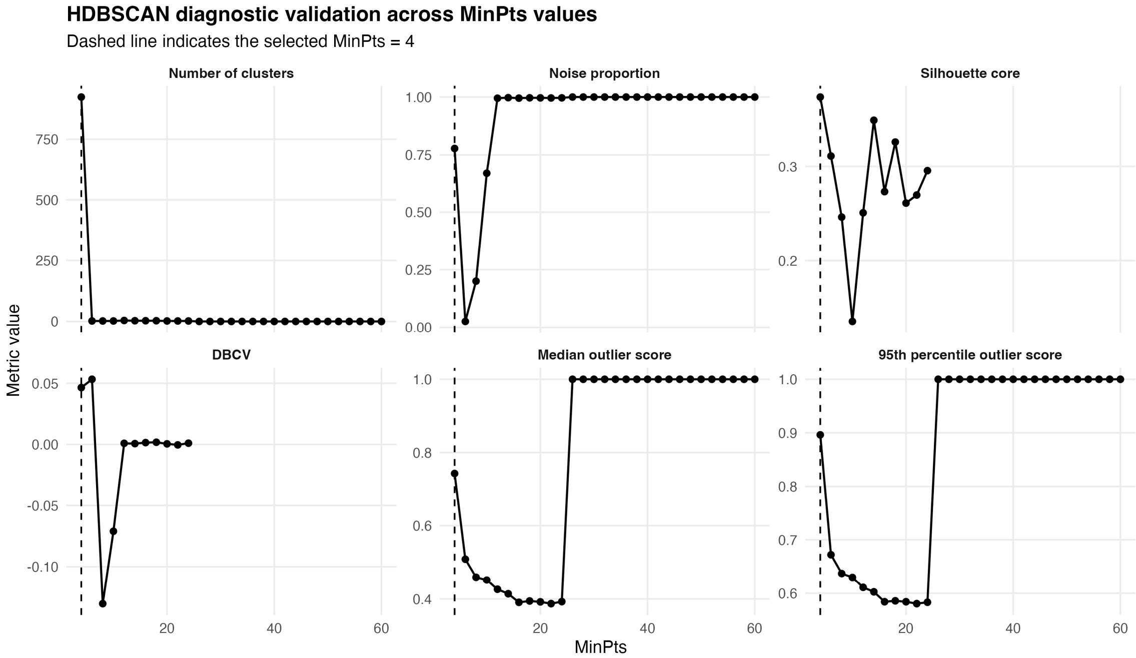

Figure 2. HDBSCAN diagnostic validation across MinPts values. Note. The figure summarizes how HDBSCAN diagnostic metrics changed across MinPts values. The selected configuration did not yield a compact taxonomic solution; instead, the flat partition was dominated by one large component, supporting the decision to interpret HDBSCAN through outlier scores rather than through flat labels.

Table 2. Diagnostic comparison of density-based methods in the revised requesting-team-minus-opponent feature space.

Method	Parameters	Non-Noise Clusters	Noise observations	Noise (%)	Interpretation
DBSCAN	Grid-selected diagnostic configuration	128	24,370	97.36	Highly fragmented and noise-dominated
OPTICS	Grid-selected diagnostic configuration	256	23,398	93.47	Heterogeneous density structure
HDBSCAN	MinPts = 4	2	129	0.52	Flat partition dominated by one large component; used for outlier scoring

Note. The density-based flat partitions were treated as diagnostic outputs rather than as substantive timeout taxonomies. DBSCAN and OPTICS yielded fragmented, noise-dominated solutions, whereas HDBSCAN produced a non-taxonomic flat partition but provided observation-level outlier scores for the main analysis.

3.3. *Outlier-Score Thresholds and Distribution of Atypical Timeout Contexts*

Atypical timeout contexts were defined using the upper tail of the HDBSCAN outlier score distribution. Three percentile thresholds were considered: top 5%, top 10%, and top 15% highest outlier scores. The top 10% threshold was retained as the primary operational definition because it provided an adequate balance between selectivity and sample size.

As shown in Table 3, the top 5% group comprised 1252 timeout events, the top 10% group comprised 2504 events, and the top 15% group comprised 3755 events. The corresponding outlier score thresholds were 0.7105, 0.6828, and 0.6613, respectively.

Table 3. Operational thresholds for atypical timeout contexts.

Atypical Context Definition	Outlier Score Threshold	*n*	Percentage of Sample
Top 5% outlier score	0.7105	1252	5.00
Top 10% outlier score	0.6828	2504	10.00
Top 15% outlier score	0.6613	3755	15.00

The spatial distribution of the top 10% outlier score events in the retained PCA space is shown in Figure 3. These events did not form a compact alternative cluster. Instead, they were located in lower-density regions of the multivariate game state space, supporting their interpretation as structurally atypical timeout contexts.

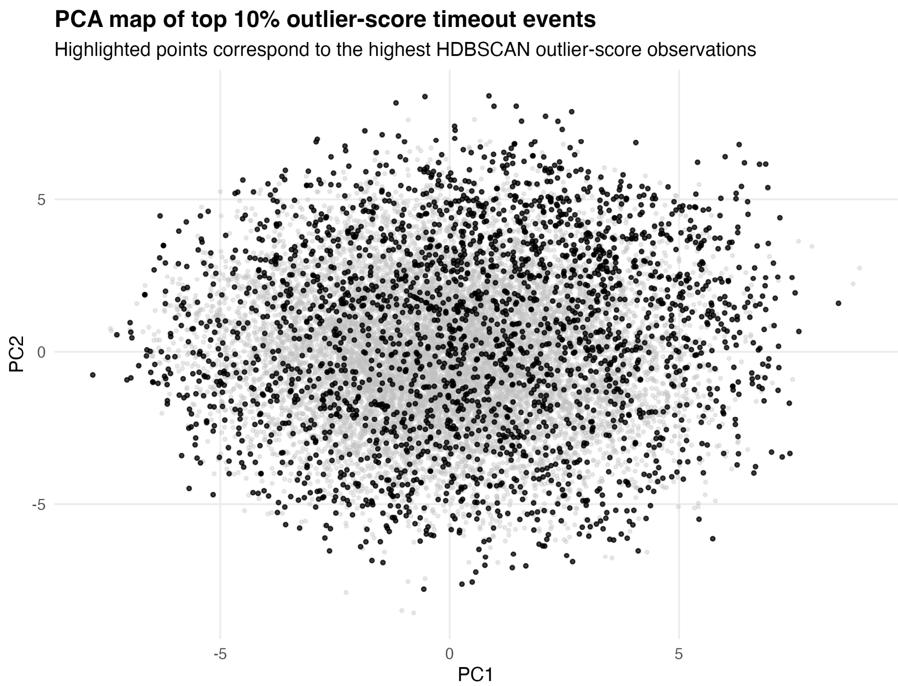

Figure 3. PCA map of top 10% outlier score timeout events. Note. The figure represents timeout events in the retained PCA space. Events in the top 10% of the HDBSCAN outlier score distribution are highlighted. These observations occupy sparse regions of the multivariate game state space rather than forming a compact cluster. Gray dots represent the remaining timeout events, whereas black dots represent the events in the top 10% of the HDBSCAN outlier-score distribution.

3.4. Multivariate Profile of the Top 10% Atypical Timeout Contexts

The comparison between top 10% outlier score events and all other timeout events showed that atypical contexts were primarily structured by score imbalance intensity. The strongest standardized differences were observed for absolute requesting team score rate differential and absolute score differential. These were followed by signed requesting team score rate differential, requesting team score differential, requesting team assist differentials, defensive rebounding differentials, and shooting rate differentials.

The largest effect sizes are reported in Table 4. The two strongest discriminating variables were absolute requesting team score rate differential (Cohen's d = 0.804) and absolute score differential (Cohen's d = 0.712), indicating that atypical timeout contexts were characterized by more pronounced competitive imbalance. Secondary contributors included requesting team score rate differential (d = 0.343), requesting team score differential (d = 0.323), assist rate differential (d = 0.259), defensive rebound rate differential (d = 0.242), requesting team assist differential (d = 0.231), requesting team defensive rebound differential (d = 0.210), three-point made rate differential (d = 0.175), and two-point made rate differential (d = 0.170).

Table 4. Main engineered variables distinguishing top 10% outlier score timeout events from other timeout events.

Variable	Cohen's d
Absolute requesting team score rate differential	0.804
Absolute score differential	0.712
Requesting team score rate differential	0.343
Score differential	0.323
Assist rate differential	0.259
Defensive rebound rate differential	0.242
Requesting team assist differential	0.231
Requesting team defensive rebound differential	0.210
Three-point made rate differential	0.175
Two-point made rate differential	0.170

These results indicate that atypical timeout contexts were not driven by a single isolated KPI. Rather, they reflected a multivariate configuration dominated by score imbalance and complemented by playmaking, defensive rebounding, and shooting efficiency differentials.

3.5. External Contextual Characterization

Contextual characterization was conducted using variables not entered as binary contextual predictors in the density-based feature space. As reported in Table 5, the top 10% outlier score events occurred earlier in the game and under larger score imbalances than other timeout events.

Table 5. External contextual profile of top 10% outlier score timeout events.

External Context	Other Timeout Events	Top 10% Outlier Events
Mean minute	25.92	22.29
Median minute	26	20
Mean absolute score differential	7.05	11.19
Median absolute score differential	6	10
Close-score situations (%)	46.29	24.48
Large-margin situations (%)	9.60	27.92
Early phase, 0–15 min (%)	19.75	31.55
Late game, 35+ min (%)	32.73	21.49
Final two minutes (%)	22.57	12.06
Overtime (%)	1.03	0.36

Top 10% outlier events occurred at a mean game minute of 22.29, compared with 25.92 for the remaining timeout events. The median minute was 20 for top 10% outlier events and 26 for other timeouts. The mean absolute score differential was also substantially larger in the top 10% group: 11.19 points versus 7.05 points. Similarly, the median absolute score differential was 10 points in the top 10% group and 6 points in the remaining sample.

The contextual distribution also differed clearly. Close-score situations were less frequent among top 10% outlier events (24.48%) than among other timeout events (46.29%). Conversely, large-margin situations were more frequent among top 10% outlier events

(27.92%) than among other timeouts (9.60%). Early-phase events were also more common in the top 10% group (31.55% vs. 19.75%), whereas late-game, final-two-minute, and overtime situations were less common.

3.6. Enrichment Analysis of External Game Contexts

The enrichment analysis confirmed the contextual pattern observed in the descriptive comparison. As shown in Table 6 and Figure 4, large-margin situations were strongly overrepresented among top 10% outlier score events, with a risk ratio of 3.00. Early-phase situations were also overrepresented, with a risk ratio of 1.74. In contrast, close-score situations, late-game situations, final-two-minute situations, and overtime were underrepresented among top 10% outlier score events.

Table 6. Enrichment of top 10% outlier score events across external contexts.

External Context	Risk Ratio	Interpretation
Large-margin situation	3.00	Overrepresented
Early phase, 0–15 min	1.74	Overrepresented
Mid phase, 15–30 min	0.98	Approximately neutral
Late game, 35+ min	0.59	Underrepresented
Final two minutes	0.50	Underrepresented
Close-score situation	0.41	Underrepresented
Overtime	0.37	Underrepresented

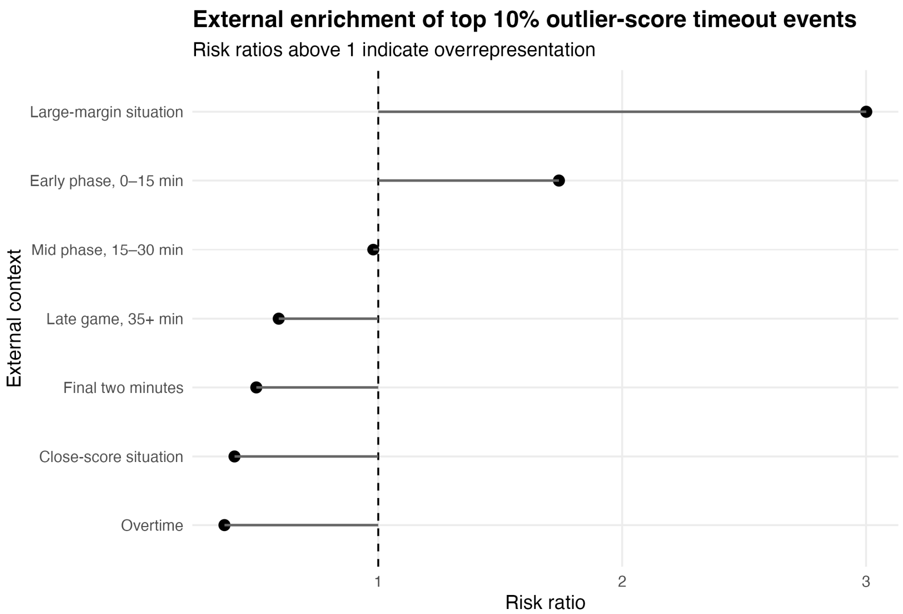

Figure 4. External enrichment risk ratios for top 10% outlier score timeout events. Note. Risk ratios above 1 indicate that top 10% outlier score timeout events were overrepresented in that external context, whereas values below 1 indicate underrepresentation. The strongest positive enrichment was observed for large-margin and early-phase situations. Black dots represent the estimated risk ratios, horizontal solid lines connect each estimate to the reference value, and the vertical dashed line indicates a risk ratio of 1.

Thus, atypical timeout contexts were externally characterized by early or non-late-game occurrence and larger score margins, rather than by the close-score and final possession contexts usually associated with high-leverage timeout decisions.

https://doi.org/10.3390/math14152696

3.7. Ablation Analysis Excluding Score Margin and Score Rate Variables

To assess whether the outlier score ranking was driven mechanically by the inclusion of score margin and score rate variables, the PCA–HDBSCAN pipeline was refitted after excluding requesting team score differential, absolute score differential, requesting team score rate differential, and absolute score rate differential. The ablated model retained the remaining requesting-team-minus-opponent performance indicators.

The ablated outlier score ranking remained strongly aligned with the full model. As shown in Table 7, the Spearman correlation between the full and ablated outlier score rankings was 0.881. The overlap of the highest outlier observations was also substantial: Jaccard similarity was 0.647 for the top 5% set, 0.686 for the top 10% set, and 0.712 for the top 15% set. These results indicate that the density-based atypicality structure was not solely determined by direct score margin or score rate variables. Rather, a substantial part of the outlier ranking persisted when these variables were removed from the feature space.

Table 7. Comparison of the full and ablated HDBSCAN outlier-score rankings.

Comparison	Value
Variables excluded	Requesting team score differential; absolute score differential; requesting team score rate differential; absolute score rate differential
Spearman correlation, full vs. ablated outlier score ranking	0.881
Jaccard similarity, top 5% outlier set	0.647
Jaccard similarity, top 10% outlier set	0.686
Jaccard similarity, top 15% outlier set	0.712

3.8. Enrichment Across External Grouping Variables

Additional enrichment analyses were conducted using grouping variables that were external to the PCA–HDBSCAN feature space, Table 8. These included season, requesting team, opponent team, and requesting side. The associations between these variables and top 10% outlier score status were small. Cramer's V was 0.051 for season, 0.080 for requesting team, 0.080 for opponent team, and 0.016 for requesting side. Thus, top 10% outlier score events were not concentrated in a single season, team, opponent, or home/away requesting side.

Table 8. Association between top 10% outlier score status and external grouping variables.

External Grouping Variable	Role in the Analysis	Cramer's V
Season	External grouping variable	0.051
Requesting team	External grouping variable	0.080
Opponent team	External grouping variable	0.080
Requesting side	External grouping variable	0.016

These results complement the score contextual characterization by showing that the atypicality ranking was not simply attributable to one specific team, opponent, competitive season, or requesting-side category. Lineup level information and possession sequence descriptors were not available in the source data and could not be evaluated.

3.9. Logistic Validation Models

Two logistic validation models were estimated to avoid collinearity among time-related predictors, Table 9.

Table 9. Logistic validation models for top 10% outlier score timeout events.

Model	Predictor	Odds Ratio	Interpretation
Model A	Minute	0.962	Later timeouts less likely to be atypical
Model A	Absolute score differential	1.115	Larger score imbalance increases atypicality
Model A	AUC	0.698	Moderate discrimination
Model B	Early phase, 0–15 min	2.44	Higher odds of atypicality
Model B	Mid phase, 15–30 min	1.22	Slightly higher odds
Model B	Late game, 35+ min	1.01	Approximately neutral
Model B	Final two minutes	0.84	Lower odds
Model B	Overtime	0.86	Lower odds
Model B	Close-score situation	0.51	Lower odds
Model B	Large-margin situation	3.12	Higher odds
Model B	AUC	0.685	Moderate discrimination

In Model A, top 10% outlier status was predicted from continuous game minute and absolute score differential. Game minute had an odds ratio below 1 (OR = 0.962), indicating that later timeouts were less likely to belong to the top 10% outlier score group. Absolute score differential had an odds ratio above 1 (OR = 1.115), indicating that each additional point of score imbalance increased the likelihood of a timeout being classified as atypical. The model achieved an AUC of 0.698.

In Model B, categorical contextual predictors were used. Early-phase status increased the odds of belonging to the top 10% outlier group (OR = 2.44), and large-margin status showed the strongest positive association (OR = 3.12). Close-score status reduced the odds of being classified as atypical (OR = 0.51). The remaining temporal indicators showed weaker associations. Model B achieved an AUC of 0.685.

As summarized in Table 9, the logistic models support the contextual characterization results: density-based atypicality was positively associated with early-game occurrence and larger score imbalances and negatively associated with close-score situations.

3.10. Stability of the Outlier Score Ranking

The stability analysis showed that the HDBSCAN outlier score ranking was robust. In the random subsampling analysis, the mean Spearman correlation between refitted and full-sample outlier score rankings was 0.844, and the Jaccard similarity for the top 10% outlier set was 0.709. This indicates that, under random subsampling, both the ranking of atypicality and the selected top 10% set were reasonably stable.

The leave-one-season-out validation produced stronger stability results. The mean Spearman correlation was 0.957, and the Jaccard similarity for the top 10% outlier set was 0.898. This indicates that the density-based outlier ranking was highly stable across seasons and was not driven by a single season.

The stability results are summarized in Table 10, and the leave-one-season-out stability pattern is displayed in Figure 5.

Table 10. Stability of the HDBSCAN outlier score ranking.

Stability Procedure	Spearman Correlation	Jaccard Similarity, Top 10%
Random subsampling	0.844	0.709
Leave-one-season-out	0.957	0.898

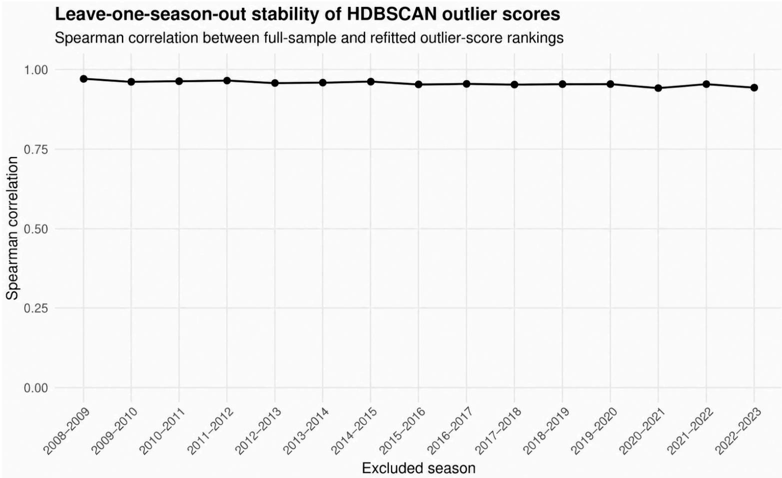

Figure 5. Leave-one-season-out stability of the HDBSCAN outlier score ranking. Note. The figure reports the Spearman correlation between the full-sample HDBSCAN outlier score ranking and the ranking obtained after excluding each season. Higher values indicate greater temporal stability of the density-based outlier structure.

3.11. Comparison with a Recurrent k-Means Timeout Taxonomy

The HDBSCAN-derived results were compared with the previously established three-cluster k-means taxonomy [30]. To avoid confounding the comparison with differences in dimensionality, k-means with (k = 3) was refitted using the same seven retained principal components used in the HDBSCAN analysis.

Agreement between the flat HDBSCAN partition and the recurrent k-means taxonomy was near zero. As shown in Table 11, the Adjusted Rand Index was 0.0010 when all observations, including HDBSCAN noise points, were considered and 0.0000 when the analysis was restricted to non-noise HDBSCAN observations. This near-zero agreement indicates that the density-based output did not reproduce the recurrent partitional taxonomy.

Table 11. Comparison between density-based atypicality and k-means (k = 3) taxonomy using the same seven-component PCA space.

Comparison	Value
ARI, HDBSCAN flat labels vs. k-means k = 3	0.0010
ARI, HDBSCAN core labels vs. k-means k = 3	0.0000
Top 10% outlier prevalence in k-means Cluster 1	8.00%
Top 10% outlier prevalence in k-means Cluster 2	8.21%
Top 10% outlier prevalence in k-means Cluster 3	15.26%

Note. The number of k-means clusters was retained from the previously validated recurrent timeout taxonomy, but the k-means model was refitted using the same seven retained principal components used for HDBSCAN in the present analysis.

The prevalence of top 10% outlier score events differed across the three recurrent profiles. Cluster 3 showed the highest prevalence of top 10% outlier score events (15.26%), compared with 8.00% in Cluster 1 and 8.21% in Cluster 2. Nevertheless, atypical events occurred in all three clusters, indicating that density-based atypicality was not restricted to a single recurrent timeout profile. These results support the interpretation that HDBSCAN outlier scoring captured a complementary analytical dimension: the extent to which individual timeout events deviated from dense regions of the multivariate game state space.

4. Discussion

This study applied density-based outlier detection to identify atypical timeout request contexts in EuroLeague basketball. The main finding is that density-based methods did not yield a compact alternative taxonomy of timeout situations. Instead, the HDBSCAN outlier score ranking provided a stable and interpretable measure of structural atypicality in the multivariate timeout context space. Here, structural atypicality refers to departure from dense regions of the multivariate timeout context space and should not be interpreted as evidence of tactical importance, timeout optimality, or timeout effectiveness. This distinction is central to the contribution of the study. Whereas hierarchical and partitional clustering approaches are useful for recovering recurrent timeout profiles, the present approach captures a different analytical layer: timeout events that occupy low-density regions of the ordinary timeout request space.

4.1. From Recurrent Timeout Profiles to Atypical Timeout Contexts

Previous timeout context analyses have primarily focused on recurrent patterns, broad situational structures, or post-timeout effects. In this sense, clustering has usually been interpreted as a way to derive compact groups of observations. The present results suggest that this logic is not sufficient when the research objective is to identify atypical timeout contexts. The diagnostic comparison of DBSCAN, OPTICS, and HDBSCAN showed that the timeout context space was heterogeneous, sparse, and highly fragmented when represented through differential and rate-adjusted performance indicators. Flat density-based partitions produced large numbers of clusters and substantial proportions of low-density observations. Therefore, interpreting these outputs as a conventional cluster taxonomy would be misleading.

This is not a methodological weakness of the study but a substantive result. It indicates that atypical timeout requests are not necessarily organized into a small number of clearly separated groups. Rather, they appear as observations located in sparse regions of the multivariate game state space. This justifies the use of HDBSCAN primarily as an outlier scoring procedure. In this framework, atypicality is not defined by membership in a particular cluster, but by the degree to which a timeout event departs from high-density regions of ordinary timeout contexts.

This approach complements previous hierarchical and partitional timeout clustering studies. Hierarchical methods are well suited to describing nested structure and global segmentation, while partitional methods such as k-means or k-medoids are useful for deriving compact, coach-oriented profiles. By contrast, the present density-based approach is designed to detect timeout events that do not conform to those dominant profiles. The near-zero agreement between the flat HDBSCAN labels and the k-means $k = 3$ taxonomy reinforces this interpretation: density-based outlier scoring does not reproduce recurrent timeout profiles but identifies a different dimension of structure.

4.2. Atypical Timeouts as Early or Intermediate Competitive Imbalance Events

A key empirical finding is that the most atypical timeout contexts were not concentrated in close-score, final-two-minute, or overtime situations. Instead, the top 10% outlier score events were more frequent in earlier or intermediate phases of the game and in larger-margin situations. This result is important because it challenges a common intuitive expectation that the most unusual or strategically meaningful timeouts should occur primarily in late-game, high-leverage scenarios.

The contextual characterization results showed that top 10% outlier score events were overrepresented in large-margin situations and early-game contexts, and underrepresented in close-score, final-two-minute, late-game, and overtime contexts. Logistic validation models supported the same pattern: larger absolute score differentials increased the odds of belonging to the top 10% outlier group, while later game minute reduced this likelihood. Similarly, large-margin status and early-phase status were positively associated with atypicality, whereas close-score status was negatively associated with atypicality.

This suggests that density-based atypicality is not equivalent to conventional late-game importance. Rather, it captures structural unusualness in the multivariate game state space. A timeout in the final two minutes of a close game may be tactically important, but it may also be relatively common within elite basketball decision patterns. Conversely, a timeout requested earlier in the game under a pronounced score imbalance or unusual rate-adjusted performance differential may be less frequent and more structurally atypical. In practical terms, these events may reflect early attempts to stop a rapidly deteriorating game state, disrupt an opponent's run before it becomes irreversible, or respond to an unexpected imbalance in offensive efficiency, defensive rebounding, or playmaking.

4.3. Multivariate Nature of Timeout Atypicality

The top 10% outlier score group was primarily characterized by score imbalance intensity. The strongest differences between atypical and non-atypical contexts were observed for absolute requesting team score rate differential and absolute score differential. This indicates that atypical timeouts were associated not only with larger margins, but also with the speed or intensity with which those imbalances emerged relative to game time.

Secondary contributors included signed requesting team score rate differential, requesting team score differential, assist rate differential, defensive rebound rate differential, requesting team assist differential, requesting team defensive rebound differential, and shooting efficiency-related rate differentials. This pattern suggests that atypical timeout requests are not explained by a single KPI. Instead, they arise from multivariate configurations in which scoreboard pressure combines with performance rate imbalance, playmaking differences, rebounding control, and shooting efficiency.

This finding is consistent with the idea that coaching interventions are triggered by composite pressure signatures rather than isolated box score events. A missed shot, a turnover, or a rebound deficit alone may not define an atypical context. However, when these indicators co-occur with a rapidly widening score differential or an unusual rate-adjusted imbalance, the resulting game state may become structurally rare. The density-based framework is particularly useful for detecting these configurations because it evaluates each timeout event relative to the local density of the multivariate feature space, rather than forcing it into a predefined recurrent profile.

4.4. Methodological Implications for Sports Analytics

The present findings have methodological implications for sports analytics and decision support systems. First, they show that density-based methods should not always be evaluated according to whether they produce compact and balanced clusters. In many applied settings, especially those involving decision events, the most relevant output may be the identification of sparse or atypical observations. Treating noise-labeled or low-density observations as failures may lead analysts to discard precisely the cases that are most informative.

Second, the results show the importance of feature engineering before applying density-based methods. When raw cumulative indicators are used, clustering may detect elapsed game time or accumulated volume rather than meaningful contextual structure. By using differential and rate-adjusted indicators and reserving temporal variables for contextual characterization, the present study reduced the risk of confounding density structure with simple game progression. The more balanced PCA variance structure supports this representation choice.

Third, the study highlights the value of separating model diagnosis from substantive interpretation. DBSCAN and OPTICS were useful for diagnosing heterogeneous density structure, but their flat partitions were not suitable as final timeout taxonomies. HDBSCAN did not provide a compact taxonomy either; under the selected configuration, the flat partition was dominated by one large component. However, it provided observation-level outlier scores that could be evaluated, validated, and interpreted. This distinction allowed the analysis to avoid overinterpreting hundreds of small clusters while still extracting meaningful density-based information.

Fourth, the stability analyses support the robustness of the outlier score approach. The HDBSCAN outlier ranking remained stable under random subsampling and was especially stable under leave-one-season-out validation. This suggests that the identified atypical timeout contexts were not artifacts of a single season or a particular data split. For applied sports analytics, this is important because a decision support measure must be robust across competitive periods if it is to be useful for longitudinal monitoring or post-game review.

4.5. Practical Implications for Coaching and Performance Analysis

From a practical perspective, the results suggest that atypical timeout detection may be useful for identifying moments of early competitive disruption. Traditional timeout analysis often emphasizes late-game management, close-score situations, and final possession tactics. These contexts are undoubtedly important, but they are not necessarily the most structurally atypical. The present analysis indicates that unusual timeout contexts may occur earlier, when a team faces a larger-than-expected imbalance in score and rate-adjusted performance indicators.

For coaching staff, this type of analysis could support post-game review by identifying timeout events that deserve closer tactical inspection. For example, an early timeout associated with a rapidly increasing score deficit, reduced assist production, defensive-rebounding imbalance, and unfavorable shooting rate differentials may signal a breakdown in offensive organization, transition defense, or possession control. Such events may not be captured by standard late-game or close-score filters, but they may be highly relevant for understanding how a game became destabilized.

During live performance analysis, the same logic could be used as a contextual alert rather than as an automatic decision rule. For instance, an analyst could monitor whether the current game state resembles historically high outlier score timeout contexts, such as an early or intermediate phase combined with a rapidly widening score imbalance, adverse rate-adjusted scoring differential, reduced assist production, defensive rebounding disadvantage, or unfavorable shooting rate pattern. If such a configuration appears, the information could be communicated to the coaching staff as evidence that the game has entered an unusual state relative to historical timeout request patterns. The coach would still interpret this information together with tactical, psychological, lineup, fatigue, and opponent-specific considerations before deciding whether to call a timeout.

In a decision support environment, density-based outlier scores could be used as a complementary diagnostic indicator. They should not be interpreted as automatic recommendations to call a timeout. Rather, they could flag unusual game states for analyst review, video tagging, or coach-facing dashboards. This is especially useful because the outlier score framework does not assume that all relevant situations must belong to one of a small number of recurrent profiles. Instead, it highlights events that are structurally unusual relative to historical game state patterns.

4.6. Relationship with Timeout Effectiveness and Causal Interpretation

The study is descriptive and should not be interpreted as evidence about timeout effectiveness. The fact that a timeout event is structurally atypical does not imply that the timeout was effective, ineffective, optimal, or suboptimal. Timeout requests are endogenous decisions: coaches call them in response to evolving game conditions, often precisely when performance is deteriorating. As previous causal and quasi-causal studies have emphasized, post-timeout improvements may reflect regression to the mean, contextual selection, or the natural interruption of extreme sequences rather than a direct causal effect of the timeout itself [5,6].

The contribution of the present study is therefore prior to causal estimation. It provides a framework for identifying the contexts in which unusual timeout decisions occur. Such a framework can inform future causal or predictive studies by defining meaningful strata of timeout contexts. For example, post-timeout outcomes could be compared separately for recurrent profiles and high-outlier-score events, or causal designs could examine whether the effect of timeouts differs between ordinary and atypical game states. However, those analyses require additional assumptions and are beyond the scope of the present study.

4.7. Limitations

Several limitations should be acknowledged. First, the analysis is restricted to EuroLeague basketball. The tactical environment, timeout rules, game rhythm, and competitive structure differ from other basketball contexts such as the NBA or youth competitions. Therefore, the identified atypical profiles should not be generalized without replication.

Second, the feature space was based on play-by-play-derived box score indicators. Although these variables are interpretable and consistently available across seasons, they do not capture all relevant aspects of timeout decisions. Lineups, defensive schemes, player fatigue, injuries, travel, psychological states, scouting plans, coach preferences, and video-based tactical information were not included. These unobserved factors may explain some timeout decisions that appear atypical in the available data.

Third, the outlier score thresholds were percentile-based. The top 10% threshold was selected as a pragmatic operating definition rather than as an empirically calibrated cutoff. The top 5% and top 15% thresholds were retained as sensitivity definitions, and the ablation analysis excluding score margin and score rate variables provided additional evidence that the upper-tail outlier structure was not solely driven by direct scoreboard information. Nevertheless, future research should explore alternative thresholding strategies based on expert annotation, outcome validation, or stability-optimized cutoff selection.

Fourth, although the outlier score ranking was stable, the flat HDBSCAN partition was not substantively interpretable as a taxonomy, because it was dominated by one large component under the selected configuration. This reinforces the methodological choice to focus on outlier scores, but it also means that the study does not provide a small set of density-based timeout clusters. Researchers seeking compact typologies should use hierarchical or partitional methods instead, while treating density-based methods as tools for identifying sparse or atypical observations.

Fifth, contextual characterization was based on indicators such as game phase and score margin. Because close-score and large-margin indicators are derived from score differential information, they should be interpreted as contextual summaries rather than independent external validation labels. These variables were intentionally excluded from the density-based feature space, which strengthens their use for validation. However, they are still descriptive indicators and do not constitute independent expert labels of atypical coaching decisions. Future work should incorporate coach annotations, video analysis, or expert judgement to validate whether high-outlier-score events correspond to tactically meaningful special situations.

4.8. Future Research

Future research should extend the present framework in several directions. First, the method should be replicated in other competitions and rule environments to examine whether unusual timeout contexts differ across basketball cultures and tactical systems. Comparing EuroLeague, NBA, national leagues, and youth competitions could reveal whether early large-margin atypicality is specific to elite European basketball or a broader property of timeout decision-making.

Second, future studies should connect density-based atypicality with post-timeout outcomes. This should be done cautiously using designs that address selection bias and regression to the mean. For instance, outcome analyses could compare ordinary timeout contexts and high-outlier-score contexts using matched samples, event study designs, or causal inference frameworks.

Third, the outlier score framework could be combined with video-based and tracking data. Player spacing, shot quality, defensive matchups, lineup structure, pace, and possession-level sequences may provide richer explanations for why certain timeout contexts are structurally atypical. Integrating these data sources would allow a more tactical interpretation of the outlier events.

Fourth, expert validation should be incorporated. Coaches, analysts, or performance staff could review high-outlier-score timeout events and classify whether they represent tactical breakdowns, psychological resets, injury-related interruptions, foul-management decisions, or planned strategic adjustments. This would provide an external semantic layer that cannot be obtained from play-by-play data alone.

5. Conclusions

This study proposed a density-based outlier detection framework for identifying atypical timeout request contexts in EuroLeague basketball. The results show that density-based methods should not be interpreted here as producing a compact alternative taxonomy of timeout profiles. DBSCAN and OPTICS confirmed the heterogeneous density structure of the timeout context space, whereas HDBSCAN provided a stable and interpretable ranking of structural atypicality through outlier scores.

The most atypical timeout requests were not primarily concentrated in late-game, close-score, or overtime situations. Instead, they appeared more frequently in earlier or intermediate phases of the game and were characterized by larger score imbalances and unusual multivariate performance rate differentials. This suggests that density-based atypicality captures a different analytical dimension from conventional tactical importance or late-game leverage.

The main contribution of this study is therefore not another classification of timeout profiles, but a reproducible procedure for identifying timeout events that deviate from ordinary game state structures. This approach complements previous hierarchical and partitional clustering analyses by adding an outlier-oriented perspective. In applied settings, HDBSCAN outlier scores may help analysts and coaches flag unusual intervention contexts for post-game review, video analysis, and future outcome-based research. However, these scores should be interpreted as descriptive indicators of structural atypicality, not as causal evidence of timeout effectiveness or as automatic recommendations for coaching decisions.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/math14152696/s1, Table S1: Reproducibility information and preprocessing decisions; Table S2: Variable dictionary for the engineered requesting-team-minus-opponent feature space; Table S3: Parameter grids used for density-based diagnostics; Table S4: R package versions used in the revised analysis.

Author Contributions: Conceptualization, J.M.C.G.; methodology, J.M.C.G. and E.M.P.; software, J.M.C.G.; validation, J.M.C.G., E.M.P. and J.M.F.L.; formal analysis, J.M.C.G.; investigation, J.M.C.G.; resources, J.M.C.G.; data curation, J.M.C.G.; writing—original draft preparation, J.M.C.G.; writing—review and editing, J.M.C.G., E.M.P. and J.M.F.L.; visualization, J.M.C.G. and E.M.P.; supervision, E.M.P. and J.M.F.L.; project administration, J.M.C.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author, subject to the terms and conditions of the original data providers. The data are not publicly available because they were compiled from third-party EuroLeague sources, and the authors do not hold unrestricted redistribution rights for the complete event-level dataset. The analytical code and derived summary outputs are also available from the corresponding author upon reasonable request.

Acknowledgments: The authors also acknowledge the support of the research group SEJ-622 “Data, Education and Society”.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

DBSCAN	Density-Based Spatial Clustering of Applications with Noise
OPTICS	Ordering Points To Identify the Clustering Structure
HDBSCAN	Hierarchical Density-Based Spatial Clustering of Applications with Noise
PCA	Principal Component Analysis
PC	Principal Component
MinPts	Minimum Number of Points
DBCV	Density-Based Clustering Validation
AUC	Area Under the Curve
OR	Odds Ratio
RR	Risk Ratio
ARI	Adjusted Rand Index
GLM	Generalized Linear Model
FDR	False Discovery Rate
KPI	Key Performance Indicator

References

1. Mace, F.C.; Lalli, J.S.; Shea, M.C.; Nevin, J.A. Behavioral momentum in college basketball. *J. Appl. Behav. Anal.* **1992**, *25*, 657–663. [CrossRef] [PubMed]
2. Nikolaidis, Y. Building a basketball game strategy through statistical analysis of data. *Ann. Oper. Res.* **2015**, *227*, 137–159. [CrossRef]
3. Gómez, M.A.; Jiménez, S.; Navarro, R.; Lago-Peñas, C.; Sampaio, J. Effects of coaches' timeouts on basketball teams' offensive and defensive performances according to momentary differences in score and game period. *Eur. J. Sport Sci.* **2011**, *11*, 303–308. [CrossRef]
4. Sampaio, J.; Lago-Peñas, C.; Gómez, M.A. Brief exploration of short and mid-term timeout effects on basketball scoring according to situational variables. *Eur. J. Sport Sci.* **2013**, *13*, 25–30. [CrossRef]
5. Gibbs, C.P.; Elmore, R.; Fosdick, B.K. The causal effect of a timeout at stopping an opposing run in the NBA. *Ann. Appl. Stat.* **2022**, *16*, 1359–1379. [CrossRef]
6. Assis, N.; Assunção, R.; Vaz-de-Melo, P.O.S. Stop the Clock: Are Timeout Effects Real? In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track*; ECML PKDD 2020; Springer: Cham, Switzerland, 2021; pp. 507–523. [CrossRef]

7. FIBA. *Official Basketball Rules 2024*; Fédération Internationale de Basketball: Mies, Switzerland, 2024.
8. Vázquez-Estévez, C.; Prieto-Lage, I.; Reguera-López-de-la-Osa, X.; Silva-Pinto, A.J.; Argibay-González, J.C.; Gutiérrez-Santiago, A. Analysis of Offensive Patterns After Timeouts in Critical Moments in the EuroLeague 2022/23. *Appl. Sci.* **2025**, *15*, 1580. [CrossRef]
9. Vázquez-Estévez, C.; Prieto-Lage, I.; Reguera-López-de-la-Osa, X.; Gutiérrez-Santiago, J.A.; Toledo-González, M.; Gutiérrez-Santiago, A. Offensive Patterns and Performance Analysis in One-Possession Scenarios During the Last Minute and Overtime in the EuroLeague. *Appl. Sci.* **2025**, *15*, 1928. [CrossRef]
10. Contreras, J.M.; Molina Portillo, E.; Fernández Luna, J.M. Evaluation of Hierarchical Clustering Methodologies for Identifying Patterns in Timeout Requests in EuroLeague Basketball. *Mathematics* **2025**, *13*, 2414. [CrossRef]
11. Jain, A.K.; Murty, M.N.; Flynn, P.J. Data clustering: A review. *ACM Comput. Surv.* **1999**, *31*, 264–323. [CrossRef]
12. Xu, R.; Wunsch, D. Survey of clustering algorithms. *IEEE Trans. Neural Netw.* **2005**, *16*, 645–678. [CrossRef] [PubMed]
13. Zhang, S.; Lorenzo, A.; Gómez, M.A.; Mateus, N.; Gonçalves, B.; Sampaio, J. Clustering performances in the NBA according to players' anthropometric attributes and playing experience. *J. Sports Sci.* **2018**, *36*, 2511–2520. [CrossRef] [PubMed]
14. Xu, X.; Zhang, S.; Gómez, M.A.; Lorenzo, A.; Sampaio, J. Clustering Performances in Elite Basketball Matches According to Situational Variables. *Front. Psychol.* **2022**, *13*, 955292. [CrossRef] [PubMed]
15. Chen, R.; Zhang, M.; Xu, X. Modeling the influence of basketball players' offense roles on team performance. *Front. Psychol.* **2023**, *14*, 1256796. [CrossRef] [PubMed]
16. Ester, M.; Kriegel, H.-P.; Sander, J.; Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, Portland, OR, USA, 2–4 August 1996; pp. 226–231.
17. Ankerst, M.; Breunig, M.M.; Kriegel, H.-P.; Sander, J. OPTICS: Ordering points to identify the clustering structure. In Proceedings of the 1999 ACM SIGMOD International Conference on Management of Data, Philadelphia, PA, USA, 1–3 June 1999; pp. 49–60. [CrossRef]
18. Campello, R.J.G.B.; Moulavi, D.; Sander, J. Density-based clustering based on hierarchical density estimates. In *Advances in Knowledge Discovery and Data Mining*; Pei, J., Tseng, V.S., Cao, L., Motoda, H., Xu, G., Eds.; Springer: Berlin/Heidelberg, Germany, 2013; pp. 160–172. [CrossRef]
19. Campello, R.J.G.B.; Moulavi, D.; Zimek, A.; Sander, J. Hierarchical density estimates for data clustering, visualization, and outlier detection. *ACM Trans. Knowl. Discov. Data* **2015**, *10*, 5. [CrossRef]
20. Schubert, E.; Sander, J.; Ester, M.; Kriegel, H.-P.; Xu, X. DBSCAN revisited, revisited: Why and how you should still use DBSCAN. *ACM Trans. Database Syst.* **2017**, *42*, 19. [CrossRef]
21. Jolliffe, I.T.; Cadima, J. Principal component analysis: A review and recent developments. *Philos. Trans. R. Soc. A Math. Phys. Eng. Sci.* **2016**, *374*, 20150202. [CrossRef] [PubMed]
22. Rousseeuw, P.J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* **1987**, *20*, 53–65. [CrossRef]
23. Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed.; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1988.
24. Benjamini, Y.; Hochberg, Y. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B Methodol.* **1995**, *57*, 289–300. [CrossRef]
25. Agresti, A. *An Introduction to Categorical Data Analysis*, 3rd ed.; Wiley: Hoboken, NJ, USA, 2018.
26. Hosmer, D.W.; Lemeshow, S.; Sturdivant, R.X. *Applied Logistic Regression*, 3rd ed.; Wiley: Hoboken, NJ, USA, 2013.
27. Robin, X.; Turck, N.; Hainard, A.; Tiberti, N.; Lisacek, F.; Sanchez, J.C.; Müller, M. pROC: An open-source package for R and S+ to analyze and compare ROC curves. *BMC Bioinform.* **2011**, *12*, 77. [CrossRef] [PubMed]
28. Spearman, C. The proof and measurement of association between two things. *Am. J. Psychol.* **1904**, *15*, 72–101. [CrossRef] [PubMed]
29. Jaccard, P. The distribution of the flora in the Alpine zone. *New Phytol.* **1912**, *11*, 37–50. [CrossRef]
30. Contreras García, J.M.; Molina-Portillo, E.; Fernández-Luna, J.M. Context-aware event clustering with consensus model selection: A large-scale study of EuroLeague timeouts. *Expert Syst. Appl.* **2026**, *331*, 133282. [CrossRef]
31. Hubert, L.; Arabie, P. Comparing partitions. *J. Classif.* **1985**, *2*, 193–218. [CrossRef]
32. R Core Team. *R: A Language and Environment for Statistical Computing*; R Foundation for Statistical Computing: Vienna, Austria, 2025.
33. Wickham, H.; Averick, M.; Bryan, J.; Chang, W.; McGowan, L.D.; François, R.; Grolemund, G.; Hayes, A.; Henry, L.; Hester, J.; et al. Welcome to the tidyverse. *J. Open Source Softw.* **2019**, *4*, 1686. [CrossRef]
34. Lê, S.; Josse, J.; Husson, F. FactoMineR: An R package for multivariate analysis. *J. Stat. Softw.* **2008**, *25*, 1–18. [CrossRef]
35. Hahsler, M.; Piekenbrock, M.; Doran, D. dbscan: Fast density-based clustering with R. *J. Stat. Softw.* **2019**, *91*, 1–30. [CrossRef]

36. Maechler, M.; Rousseeuw, P.; Struyf, A.; Hubert, M.; Hornik, K. *cluster: Cluster Analysis Basics and Extensions*; R Package Version 2.1.8.1; R Foundation for Statistical Computing: Vienna, Austria, 2025.
37. Scrucca, L.; Fop, M.; Murphy, T.B.; Raftery, A.E. mclust 5: Clustering, classification and density estimation using Gaussian finite mixture models. *R J.* **2016**, *8*, 289–317. [CrossRef]