
Quantifying Era-Relative Dominance in the NBA Using Signed Mahalanobis Distance

Rishabh Ranjan
Stanford University
rranjan2@stanford.edu

Dawen Cheng
Stanford University
dawen@stanford.edu

Abstract

Identifying the greatest NBA player of all time is a question complicated by shifting rules, pace, and playing styles across decades. Standard Mahalanobis distance can identify statistically unusual player-seasons, but it is *directionless*—it rewards unusualness in any direction, including negative ones. We propose **Signed Mahalanobis Distance**, a modification that projects each player’s Mahalanobis-whitened displacement onto the first principal component of era-normalized advanced statistics, which empirically captures an axis of overall playing quality. Applied to 5,674 player-seasons from 1981–2017, our method ranks Michael Jordan’s 1988 season as the most statistically dominant in modern NBA history, with career rankings of Jordan first and LeBron James second. A predictive evaluation shows that team-level Signed Mahalanobis features predict playoff qualification with $AUC = 0.917$, outperforming a traditional VORP baseline ($AUC = 0.881$). We release our full pipeline as a reproducible notebook.

1 Introduction

The question of who is the “greatest of all time” (GOAT) in the NBA has long been debated by fans, journalists, and analysts. A central challenge is that the game has changed substantially over time: rule changes, pace of play, three-point line adoption, and evolving defensive schemes all affect raw box-score statistics, making direct comparisons across eras misleading Kubatko et al. [2007], Hollinger [2004], Oliver [2004]. A player averaging 30 points per game in a high-pace 1980s season may be no more dominant relative to their peers than a player averaging 27 points in a slower 2010s season.

Prior work in basketball analytics has addressed this in various ways. Hollinger’s Player Efficiency Rating (PER) Hollinger [2004] aggregates positive and negative contributions into a per-minute rating. Value Over Replacement Player (VORP) and Box Plus/Minus (BPM) estimate a player’s contribution relative to a league-average or replacement-level player Myers [2012]. Win Shares (WS) attempt to credit players with a portion of team wins Oliver [2004]. While informative, these composite metrics embed specific weighting choices that may not generalize across eras, and they do not explicitly account for the changing statistical landscape of the league.

A more principled approach is to ask: *relative to the other players in the same season, how statistically unusual was this player’s performance?* The Mahalanobis distance Mahalanobis [1936] is a natural tool for this question. By incorporating the covariance structure of the data, it measures how many “multivariate standard deviations” a point lies from the center of a distribution, accounting for correlations between variables.

However, standard Mahalanobis distance has a fundamental limitation for our purposes: it is *unsigned*. It measures statistical unusualness without distinguishing the direction of that unusualness. A player who blocks an extraordinary number of shots but contributes little else (e.g., Mark Eaton) may register

a higher Mahalanobis distance than a player who is excellent across all dimensions (e.g., Michael Jordan), simply because the specialist’s profile is more *surprising* relative to the league distribution.

In this paper, we propose **Signed Mahalanobis Distance**, which addresses this limitation by projecting each player’s Mahalanobis-whitened displacement vector onto the first principal component (PC1) of the era-normalized feature space. Since PC1 empirically captures the axis of overall playing quality—loading heavily on PER, WS/48, and BPM—this projection yields a signed score that distinguishes excellence from mere oddity. We apply this method to NBA seasons from 1981 to 2017, producing era-adjusted rankings of individual player-seasons and careers, and validate the metric through a downstream predictive task.

Contributions. Our main contributions are:

1. A **Signed Mahalanobis Distance** metric that corrects the directional ambiguity of standard Mahalanobis distance by projecting onto an empirically derived quality axis (Section 3).
2. **Era-adjusted season and career rankings** of NBA players from 1981–2017 that surface Jordan’s 1988 season as the most dominant and Jordan as the career GOAT, with LeBron James as a close second (Section 4).
3. A **predictive validation** showing that Signed Mahalanobis features outperform traditional VORP for predicting team playoff qualification (Section 4).

2 Related Work

Composite basketball metrics. The most widely used single-number player evaluation metrics include PER Hollinger [2004], which sums positive and negative contributions per minute of play; Win Shares Oliver [2004], which estimates win credit; BPM and VORP Myers [2012], which estimate a player’s per-possession contribution above league average or replacement level; and True Shooting Percentage, which adjusts for free throws and three-pointers. Each metric embeds specific modeling assumptions and weighting choices. Our approach differs by learning the quality direction from the data via PCA rather than prescribing it.

Mahalanobis distance in statistics and sports. The Mahalanobis distance Mahalanobis [1936] has been widely used for outlier detection and classification. In sports analytics, multidimensional statistical methods including PCA and clustering have been applied to NBA player evaluation Sarlis and Tjortjis [2020], Page et al. [2013]. Our work adapts the Mahalanobis distance specifically for player ranking by adding a directional correction via PCA.

Covariance estimation. Computing Mahalanobis distance requires inverting the covariance matrix, which is problematic when the number of features is comparable to the sample size. Ledoit and Wolf Ledoit and Wolf [2004] proposed a shrinkage estimator that convexly combines the sample covariance with a structured target (the identity matrix), producing a well-conditioned estimate. We adopt this estimator in our pipeline.

3 Methods

3.1 Data

We use the “NBA Players Stats Since 1950” dataset Kaggle [2018], which compiles per-season statistics from Basketball Reference. We restrict our analysis to the **three-point era** (1981–2017), since the introduction of the three-point line in 1979–80 fundamentally altered offensive strategy and statistical distributions. We apply a minimum playing-time filter of **500 minutes per season** to exclude small-sample-size artifacts. Players traded mid-season are represented by their combined “TOT” row. After filtering, the dataset contains 5,674 player-seasons from 1,676 unique players.

Feature set. We select nine advanced statistics that capture complementary aspects of player performance:

Feature	Description
PER	Player Efficiency Rating (per-minute composite)
TS%	True Shooting Percentage (scoring efficiency)
AST%	Assist Percentage (passing creation)
TRB%	Total Rebound Percentage (rebounding)
STL%	Steal Percentage (perimeter defense)
BLK%	Block Percentage (rim protection)
USG%	Usage Rate (offensive load)
WS/48	Win Shares per 48 minutes
BPM	Box Plus/Minus

These features span offensive efficiency (TS%, USG%), playmaking (AST%), defense (STL%, BLK%), rebounding (TRB%), and composite value (PER, WS/48, BPM). We additionally retain Win Shares (WS) and Value Over Replacement Player (VORP) as external validation metrics, but these are *not* included in the PCA or Mahalanobis computation.

3.2 Era-Specific Z-Score Normalization (Step A)

To account for era effects, we standardize each feature independently within each season. Let $x_{i,j}^{(t)}$ denote the value of feature j for player i in season t , and let $\mu_j^{(t)}$ and $\sigma_j^{(t)}$ denote the mean and standard deviation of feature j across all qualifying players in season t . The era-normalized value is:

$$z_{i,j}^{(t)} = \frac{x_{i,j}^{(t)} - \mu_j^{(t)}}{\sigma_j^{(t)}}. \quad (1)$$

After this transformation, a z -score of $+2.5$ means the player was 2.5 standard deviations above the league average *for that specific season*, regardless of how the raw statistics compare across eras.

3.3 PCA Dimensionality Reduction (Step B)

Within each season, we perform PCA on the 9-dimensional z -scored feature matrix. Given the era-normalized data matrix $\mathbf{Z}^{(t)} \in \mathbb{R}^{n_t \times 9}$ for season t (with n_t qualifying players), PCA finds the orthogonal directions of maximum variance by computing the eigendecomposition of the sample covariance matrix:

$$\mathbf{C}^{(t)} = \frac{1}{n_t - 1} (\mathbf{Z}^{(t)})^\top \mathbf{Z}^{(t)} = \mathbf{V}^{(t)} \mathbf{\Lambda}^{(t)} (\mathbf{V}^{(t)})^\top, \quad (2)$$

where $\mathbf{V}^{(t)}$ contains the eigenvectors (principal component directions) and $\mathbf{\Lambda}^{(t)}$ is the diagonal matrix of eigenvalues. Empirically, the first principal component (PC1) consistently captures approximately 40–45% of the total variance across seasons, and a scree analysis confirms that 4 components explain roughly 90% of the variance (see Section 4).

Crucially, PC1 loads heavily on PER, WS/48, and BPM—the features most associated with overall player quality. This means that PC1 empirically captures a “quality axis” in the feature space. We exploit this observation in our signed distance computation.

3.4 Mahalanobis Distance (Step C)

The standard Mahalanobis distance of a point \mathbf{z}_i from the centroid $\boldsymbol{\mu}$ of a distribution with covariance $\boldsymbol{\Sigma}$ is defined as:

$$D_M(\mathbf{z}_i) = \sqrt{(\mathbf{z}_i - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{z}_i - \boldsymbol{\mu})}. \quad (3)$$

This measures how many “multivariate standard deviations” a point lies from the center, accounting for the correlation structure of the data. Unlike Euclidean distance, which treats all directions equally, Mahalanobis distance weights directions by the inverse of their variance—a point that lies along a direction of high variance is less unusual than one at the same Euclidean distance along a direction of low variance.

Ledoit-Wolf covariance estimation. A challenge in computing the Mahalanobis distance is that the sample covariance matrix may be poorly conditioned when the number of features ($p = 9$) is not negligible relative to the number of observations per season ($n_t \approx 130\text{--}200$). While the ratio p/n_t is small enough that the sample covariance is invertible, the resulting inverse can amplify estimation error. We therefore use the Ledoit-Wolf shrinkage estimator Ledoit and Wolf [2004], which replaces the sample covariance \mathbf{S} with:

$$\hat{\Sigma}_{\text{LW}} = \alpha \cdot \mu_{\text{target}} \mathbf{I} + (1 - \alpha) \cdot \mathbf{S}, \quad (4)$$

where $\mu_{\text{target}} = \text{tr}(\mathbf{S})/p$ is the average eigenvalue, \mathbf{I} is the identity matrix, and $\alpha \in [0, 1]$ is the analytically optimal shrinkage intensity that minimizes expected quadratic loss. This estimator “shrinks” extreme eigenvalues toward their grand mean, producing a well-conditioned covariance matrix that is more robust to finite-sample noise.

3.5 Signed Mahalanobis Distance (Step D)

The standard Mahalanobis distance D_M is non-negative by construction and thus *unsigned*: it measures how unusual a player is, but not in which direction. To assign directionality, we project the Mahalanobis-whitened displacement vector onto the first principal component direction.

For player i in season t , define the displacement vector $\delta_i = \mathbf{z}_i - \boldsymbol{\mu}^{(t)}$ and the whitened displacement $\mathbf{w}_i = \hat{\Sigma}_{\text{LW}}^{-1} \delta_i$. Let $\mathbf{v}_1^{(t)}$ denote the first principal component direction from the PCA of season t . The **Signed Mahalanobis Distance** is:

$$S_M(\mathbf{z}_i) = \mathbf{w}_i^\top \mathbf{v}_1^{(t)} = \left(\hat{\Sigma}_{\text{LW}}^{-1} (\mathbf{z}_i - \boldsymbol{\mu}^{(t)}) \right)^\top \mathbf{v}_1^{(t)}. \quad (5)$$

Orientation convention. PC1 is defined only up to a sign flip. To ensure a consistent interpretation across seasons, we orient $\mathbf{v}_1^{(t)}$ so that its loadings on PER, WS/48, and BPM are positive on average. That is, if the mean of the loadings on these three features is negative, we negate $\mathbf{v}_1^{(t)}$. After this orientation step, positive values of S_M correspond to players who are outliers in the direction of *excellence*, while negative values correspond to outliers in “bad” or irrelevant directions.

Normalization. We standardize S_M within each season to have unit variance (but not zero mean), so that scores are comparable across years.

3.6 Career Rankings

To produce career-level rankings, we compute each player’s average Signed Mahalanobis score over their top 5 seasons (by S_M), requiring a minimum of 5 qualifying seasons. This balances peak performance with sustained excellence: a player with a single extraordinary season will not outrank a player with many dominant seasons.

3.7 Predictive Validation

As an external validation, we test whether our dominance metric predicts team-level playoff qualification. For each team-season, we construct features from the Signed Mahalanobis scores of the team’s players: the *star* score (highest S_M on the roster) and the *depth* score (mean S_M of the 2nd–5th ranked players).

Since we do not have a direct playoff qualification column in the dataset, we use a proxy: for each season, teams are ranked by total Win Shares, and the top 16 are labeled as having “made the playoffs.” This proxy is imperfect—actual playoff qualification depends on conference standings and division records—but it provides a reasonable approximation that is fully derivable from the available data.

We compare four feature sets using 5-fold stratified cross-validation with both logistic regression and random forest classifiers:

1. **Raw Mahalanobis:** star and depth features using unsigned D_M .
2. **Signed Mahalanobis:** star and depth features using S_M .
3. **Baseline (VORP):** era-adjusted VORP of the best player.
4. **Combined:** Signed Mahalanobis star + VORP + Signed Mahalanobis depth.

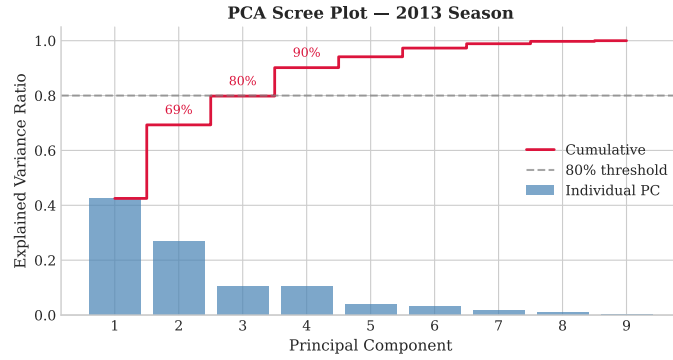


Figure 1: PCA scree plot for the 2013 NBA season, showing individual and cumulative explained variance. Four components explain approximately 90% of the total variance. The first component, which we use as the quality axis in our Signed Mahalanobis computation, captures 42.5% of the variance.

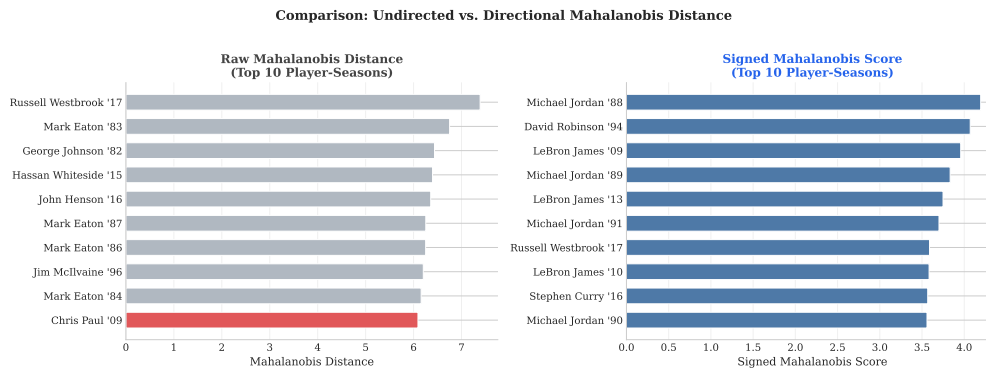


Figure 2: Comparison of top 10 player-seasons under Raw Mahalanobis Distance (left) and Signed Mahalanobis Distance (right). Raw Mahalanobis is dominated by extreme defensive specialists and statistical anomalies (Mark Eaton, Hassan Whiteside, Westbrook’s triple-double season), while Signed Mahalanobis correctly surfaces consensus all-time greats.

4 Results

4.1 PCA Structure

Figure 1 shows a representative scree plot (2013 season). The first principal component explains approximately 42.5% of the variance, with 4 components reaching roughly 90%. The first component loads most heavily on PER (+0.49), WS/48 (+0.48), and BPM (+0.47), confirming its interpretation as a “quality axis.” The second component tends to separate defensive specialists (high BLK%, TRB%) from playmakers (high AST%), capturing a positional or stylistic dimension.

4.2 Signed vs. Raw Mahalanobis Rankings

Table 1 compares the top 10 player-seasons under raw Mahalanobis distance and Signed Mahalanobis distance. The contrast is striking: raw Mahalanobis ranks Russell Westbrook’s 2017 triple-double season first—an extraordinary but positionally anomalous season—followed by several extreme shot-blocking specialists (Mark Eaton, George Johnson, Hassan Whiteside). By contrast, Signed Mahalanobis ranks Michael Jordan’s 1988 season first, followed by David Robinson’s 1994 MVP campaign and LeBron James’s 2009 season.

Figure 2 visualizes this comparison as side-by-side horizontal bar charts. The raw Mahalanobis ranking is dominated by players whose statistical profiles are extreme along a single dimension

Table 1: Top 10 player-seasons ranked by Raw Mahalanobis Distance (left) and Signed Mahalanobis Distance (right). Raw Mahalanobis finds the most statistically unusual players, while Signed Mahalanobis identifies the most dominant. Notable GOAT candidates are bolded.

Rank	Player (Year)	Raw Mahal.	VORP
1	Russell Westbrook (2017)	7.39	12.4
2	Mark Eaton (1983)	6.75	0.5
3	George Johnson (1982)	6.44	1.8
4	Hassan Whiteside (2015)	6.40	0.3
5	John Henson (2016)	6.36	0.9
6	Mark Eaton (1987)	6.26	2.2
7	Mark Eaton (1986)	6.25	2.1
8	Jim McIlvaine (1996)	6.20	0.6
9	Mark Eaton (1984)	6.16	2.1
10	Chris Paul (2009)	6.09	10.0

Rank	Player (Year)	Signed Mahal.	VORP
1	Michael Jordan (1988)	4.19	11.8
2	David Robinson (1994)	4.07	10.6
3	LeBron James (2009)	3.96	11.6
4	Michael Jordan (1989)	3.83	12.0
5	LeBron James (2013)	3.75	9.8
6	Michael Jordan (1991)	3.70	9.8
7	Russell Westbrook (2017)	3.59	12.4
8	LeBron James (2010)	3.58	10.9
9	Stephen Curry (2016)	3.57	9.8
10	Michael Jordan (1990)	3.56	10.1

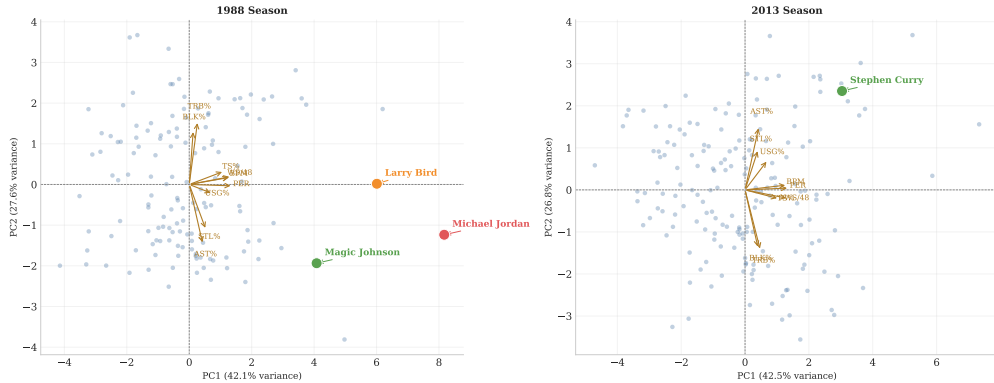


Figure 3: PCA biplots for the 1988 season (left, highlighting Jordan, Bird, Magic) and the 2013 season (right, highlighting LeBron, Durant, Curry). Feature loading arrows show that PC1 aligns with the quality metrics (PER, WS/48, BPM). The dominant players in each era are positioned far in the positive PC1 direction.

(shot-blocking), while the signed version produces rankings that align closely with expert consensus on the greatest NBA players.

4.3 PCA Biplot Analysis

Figure 3 shows PCA biplots for the 1988 and 2013 seasons, highlighting the positions of key players in the reduced two-dimensional space. In both cases, the GOAT candidates (Jordan in 1988, LeBron in 2013) are located far along the positive PC1 direction—the direction our method uses to sign the Mahalanobis distance. Feature loading arrows show that PC1 aligns with PER, WS/48, and BPM, while PC2 captures positional differences (e.g., BLK% vs. AST%).

Table 2: Top 15 career rankings by 5-season peak average Signed Mahalanobis score (minimum 5 qualifying seasons). “Elite seasons” counts the number of seasons with $S_M > 2.5$.

Rank	Player	Peak 5yr Avg	Best Season	Elite Seasons	Years Active
1	Michael Jordan	3.77	4.19	10	1985–2003
2	LeBron James	3.61	3.96	11	2004–2017
3	David Robinson	3.42	4.07	9	1990–2003
4	Karl Malone	3.04	3.26	8	1986–2004
5	Charles Barkley	3.01	3.18	7	1985–2000
6	Larry Bird	2.96	3.21	5	1981–1992
7	Chris Paul	2.94	3.55	6	2006–2017
8	Tim Duncan	2.93	3.24	7	1998–2016
9	Kevin Durant	2.92	3.37	5	2008–2017
10	Magic Johnson	2.90	3.29	5	1981–1996
11	Dwyane Wade	2.79	3.35	5	2004–2017
12	Hakeem Olajuwon	2.67	3.37	2	1985–2002
13	Russell Westbrook	2.61	3.59	3	2009–2017
14	John Stockton	2.58	2.78	4	1985–2003
15	Dirk Nowitzki	2.56	2.69	4	1999–2017

Table 3: ROC-AUC for predicting playoff qualification (5-fold stratified CV). Signed Mahalanobis outperforms both raw Mahalanobis and the VORP baseline. The combined model performs best.

Feature Set	Logistic Regression	Random Forest
Raw Mahalanobis	0.708 ± 0.051	0.683 ± 0.010
Signed Mahalanobis	0.917 ± 0.024	0.895 ± 0.018
Baseline (VORP)	0.881 ± 0.014	0.811 ± 0.024
Combined (Signed + VORP)	0.926 ± 0.018	0.914 ± 0.016

4.4 Career Rankings

Table 2 presents the top 15 career rankings by 5-season peak average Signed Mahalanobis score. Michael Jordan leads with a peak average of 3.77, followed by LeBron James at 3.61. Jordan has 10 “elite” seasons (defined as $S_M > 2.5$), while LeBron has 11—the most of any player in the dataset. The rankings align well with conventional wisdom about the greatest NBA players, lending face validity to the method.

4.5 Predictive Validation

Table 3 reports the ROC-AUC scores for predicting playoff qualification under each feature set and model. The Signed Mahalanobis features achieve AUC = 0.917 with logistic regression and 0.895 with random forest, substantially outperforming both the raw Mahalanobis features (AUC = 0.708/0.683) and the VORP baseline (AUC = 0.881/0.811). The combined model achieves the best performance (AUC = 0.926/0.914).

The improvement of Signed Mahalanobis over raw Mahalanobis is dramatic (AUC gap of +0.21 for logistic regression), demonstrating the critical importance of directional correction: knowing that a team’s best player is unusual *in the direction of excellence* is far more predictive than simply knowing they are unusual. The further improvement of the combined model suggests that Signed Mahalanobis and VORP capture partially complementary information.

5 Conclusion

We have proposed Signed Mahalanobis Distance as a principled method for quantifying era-relative player dominance in the NBA. By projecting the Mahalanobis-whitened displacement onto PC1, we obtain a directed measure that distinguishes excellence from mere statistical peculiarity. Applied to 1981–2017 NBA data, the method produces rankings that align with expert consensus—Michael Jordan’s 1988 season emerges as the most dominant, with Jordan as the career GOAT and LeBron

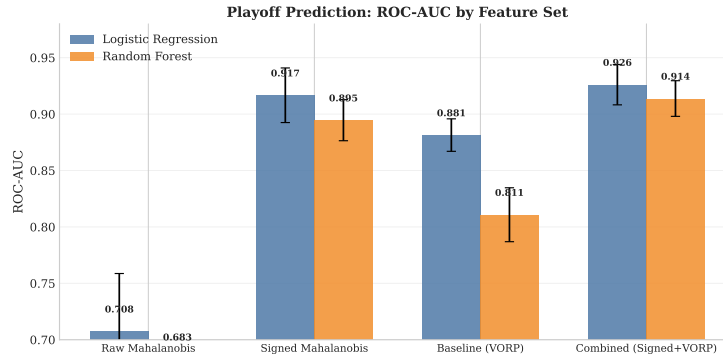


Figure 4: ROC-AUC comparison across feature sets and models for predicting playoff qualification. Error bars indicate standard deviation across 5 cross-validation folds. The combined Signed Mahalanobis + VORP model achieves the highest AUC.

James as a close second—and the derived features outperform traditional metrics for predicting playoff qualification.

Limitations. Several limitations should be noted. First, our quality axis is determined by PC1, which we assume captures “overall quality” based on its empirical loadings. This assumption, while empirically supported, is not rigorously justified—there is no theoretical guarantee that the direction of maximum variance in era-normalized advanced statistics corresponds to playing excellence. In seasons with unusual statistical distributions, PC1 could in principle capture a different dimension. Second, the orientation convention (flipping PC1 based on PER, WS/48, and BPM loadings) introduces a mild circularity: while these features are not used as a direct ranking criterion, they influence the direction of the quality axis. Third, the choice of 5 seasons for career averaging is somewhat arbitrary and could be varied. Fourth, our predictive validation uses a proxy for playoff qualification (top 16 teams by Win Shares) rather than actual playoff data, which introduces additional noise. Fifth, the dataset ends in 2017 and thus does not include recent seasons. Finally, the method relies entirely on box-score statistics and does not incorporate contextual information such as team quality, opponent strength, playoff performance, or intangible leadership qualities.

Extensions. Future work could explore several directions: (1) replacing the PC1 assumption with a supervised quality axis learned from All-NBA selections or MVP votes; (2) extending the dataset to include recent seasons; (3) incorporating tracking data (e.g., player movement, shot quality) as additional features; (4) applying the method to other sports where cross-era comparison is challenging; and (5) investigating the sensitivity of rankings to the choice of features and filtering thresholds.

Acknowledgments

This project was completed for CS 141 at Stanford University. We thank the course staff for their guidance and feedback.

References

- John Hollinger. *Pro Basketball Forecast, 2004-05*. Potomac Books, 2004.
- Omri Goldstein. NBA Players Stats Since 1950. Kaggle Dataset, 2018. <https://www.kaggle.com/datasets/drgilermo/nba-players-stats>
- Justin Kubatko, Dean Oliver, Kevin Pelton, and Dan T. Rosenbaum. A starting point for analyzing basketball statistics. *Journal of Quantitative Analysis in Sports*, 3(1), 2007.
- Olivier Ledoit and Michael Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2):365–411, 2004.

- Prasanta Chandra Mahalanobis. On the generalized distance in statistics. *Proceedings of the National Institute of Sciences of India*, 2(1):49–55, 1936.
- Daniel Myers. About Box Plus/Minus (BPM). Basketball Reference, 2012. <https://www.basketball-reference.com/about/bpm2.html>
- Dean Oliver. *Basketball on Paper: Rules and Tools for Performance Analysis*. Potomac Books, 2004.
- Garritt L. Page, Bradley J. Barney, and Aaron T. McGuire. Effect of position, usage rate, and per game minutes played on NBA player production curves. *Journal of Quantitative Analysis in Sports*, 9(4):337–345, 2013.
- Karl Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(11):559–572, 1901.
- Vangelis Sarlis and Christos Tjortjis. Sports analytics—evaluation of basketball players and team performance. *Information Systems*, 93:101562, 2020.