Rank-Adapted Singular Value Decomposition Applied to NCAA Football Rankings

Dr. Kathleen Campbell*

Abstract
The practice of ranking objects, events, and people to determine relevance, importance, or competitive edge is ancient. Determining if those creating the ranks (herein called sources) are in agreement when there are more than two sources can be a daunting task. The current tests assume that all sources compare a list of the same items (herein called elements). In the case presented here (as with many other cases) each pair sources rarely choose identical elements to rank. For various reasons a new method beyond Kendall's W is needed. The use of Kendall's τ is preferable, but it can only be used when assessing association of two (and only two) sources. When there are more than two sources, Kendall's coefficient of concordance (W) is used to assess association. Notably, Kendall(1948) observed that "the most obvious procedure (to address m sources) is to average all of the possible values of τ between pairs of observers; but this is evidently very tedious when m is large." With the advance of technology, the new method proposed by this research: Rank-Adapted Singular Value Decomposition (R-A SVD) considers and addresses this ideal proposed by Kendall. Kendall's W was established for finding association among multiple sources. One issue that needs to be addressed is the possibility of masking association. Masking is when two factions (one or more sets of sources that agree with each other) agree strongly among themselves but inversely with an opposing faction. When masking exists, using W, one might conclude a lack of association where in fact two strong factions exist. The final issue, which will be addressed using R-A SVD, is that W does not provide any insight into the strength with which the sources agree. In cases where some of the sources agree moderately and others agree strongly (all in the same direction), relying on W only implies association and does not give any evidence of the structure of the data. In addressing Kendall's comments and the aforementioned weaknesses with W, R-A SVD is introduced as a conservative method that provides results comparable to averaging all the paired values of τ (herein called average τ) but also yields information beyond just whether association exists. R-A SVD also provides insight as to the structure of the data, identifies the underlying correlation, and determines whether this correlation is significant. Overall, R-A SVD will identify factions and examine the ranked elements to ascertain whether classes of good, better, and best exist among the elements being ranked (herein called equivalence classes).

Key Words: Kendall's τ, Kendall's W, rank correlation, SVD, PCA

1. R-A SVD Applied: Steps Needed to Apply the New Method

Throughout the analysis provided in this paper (and all other studies), there are four steps needed to adequately use the R-A SVD method. These steps are concisely outlined below. A more in depth discussion has been submitted to the Annals of Applied Statistics for review.

1.1 Step 1

In the first step, the correlation is assessed using Kendall's τ with the understanding that τ is near-metric. We also add placeholders for elements that do non overlap. Starting with a matrix of data ranks (P), we take the covariance of P which we call K. We factor the covariance matrix K as follows: $K = cov(P) = VDV'$. Here, V is an orthonormal basis for the rows that is useful in identifying when sources agree as to the rank order and specifically

*St. Joseph's University, 5600 City Avenue, Philadelphia, PA 19131

which sources. D is a diagonal of eigenvalues. By analogy with SVD, we then define U^* as $U^* = PV'D^{-1}$.

The diagonal matrix, D, provides the factored eigenvalues in decreasing order. Kaiser(1968) uses the first and largest eigenvalue of the population data to identify overall association. Using Hotelling's (1933) finding that a line is the first principal axis of the system and that the mean squared projection on the line is equal to the largest eigenvalue, Kaiser proposed that the average correlation can be approximated using this formula:

$$\gamma = \frac{(\lambda - 1)}{(p - 1)} \tag{1}$$

where γ represents the average correlation and λ represents the largest eigenvalue and p is the number of standardized variables. This method works well when all of the correlations are similar and positive.

We note that Kaiser did not discuss the use of his eigenvalue procedure for averaging *sample* intercorrelations and all of his equations are expressed in population parameters. Kaiser's averaging procedure is biased when applied to small samples from populations having correlations near zero, and the procedure does not handle negative values efficiently (Friedman and Weisberg (1981) and Dunlap, Silver and Phelps (1987)). Dunlap et al's (1987) modification to Kaiser's average correlation substantially reduces the bias with slightly smaller standard errors, leading to greater efficiency for small correlations. In this modification, the value of each individual paired correlation in the matrix was increased by one prior to finding the eigenvalue. Since no element in this matrix is negative, the first eigenvalue as a measure of association is valid. A back transformation is necessary to remove the added value placed into the rank when one was added to each eigenvalue. The new average for the correlation is found:

$$r = \frac{(\lambda - m - 1)}{m - 1} \tag{2}$$

where r is the average correlation, λ represents the largest eigenvalue, and m is equivalent to the rank of the matrix (representing the sum of the added ones). Dunlap et al (1987) provide evidence that the results of their Monte Carlo simulation showed substantial improvement to Kaiser's method.

Using the same mathematical principles though recognizing the data is not exactly metric, we create our average correlation measure ($\lambda_{\{R-A\ SVD\}}$). The equation follows the same structure as that proposed by Dunlap. However, the λ represents the largest eigenvalue found from the correlation matrix created using Kendall's τ as the underlying correlation measure (not Pearson's r as used by Dunlap et al) and m is still equivalent to the rank of the matrix.

$$\lambda_{\{R-A\ SVD\}} = \frac{(\lambda - m - 1)}{m - 1} \tag{3}$$

1.2 Step 2

In step 2 of R-A SVD, Anderson's test (1963) determines whether this correlation ($\lambda_{\{R-A\ SVD\}}$) is significant. If at least one of the eigenvalues from this matrix is significant, the correlation is significant. Anderson (1951,1963), Bartlett (1954) and Lawley (1956, 1963) found that using the covariance and variance structure of the data (assuming independent elements and the asymptotic distribution of the data) permits the construction of hypothesis tests. Within R-A SVD, we consider N independent elements being ranked by m sources, where the distribution of the m dimensional random variables is $N(\mu, \Sigma)$ and Σ has distinct eigenvalues such that $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_m \geq 0$ with the corresponding eigenvectors $\alpha_1, \ldots \alpha_m$

Our sample estimate of Σ will be S with n degrees of freedom, where $n = N - 1$. We will then use l_1 and a_1 to represent the first sample eigenvalue and eigenvector, respectively. Both the population and sample eigenvectors have a directional cosine length of 1 (Morrison:1976). Girshick (1936, 1939) and Anderson (1951, 1963) provide proof of normality and the independence structure of both the eigenvalues and eigenvectors as n becomes large without giving a lower bound of how large.

In order to determine whether some of the eigenvalues are significantly similar, Anderson (1963) created the following null hypothesis based on the covariance matrix:

$$H_0\colon \lambda_{q+1} = \ldots = \lambda_{q+r}$$

This hypothesis tests whether r of the eigenvalues of Σ are equal. The q larger and $m - q - r$ smaller are not. We specifically focus on the smallest eigenvalues of the covariance matrix, with the goal of finding which small eigenvalues are significantly small enough to be considered inconsequential. The remaining eigenvalues are then considered significant. The hypothesis test starts with the smallest two eigenvalues to test for equivalence.

$$H_0\colon \lambda_m = \lambda_{m-1}$$

If this hypothesis is found to be true, the test continues by adding the next smallest eigenvalue to be tested.

$$H_0\colon \lambda_m = \lambda_{m-1} = \lambda_{m-2}$$

This test continues until the hypothesis is rejected and only $m - q$ of the eigenvalues are found to be significantly similar and inconsequential to the data analysis. The alternative hypothesis (H_a) for all of these cases is that at least one of the roots is distinctly different than the others and hence is significantly larger.

The likelihood-ratio criterion leads to the test statistic:

$$\chi^2 = -n \sum \ln l_j + nr \ln \frac{\sum l_j}{r}$$

where $n = N-1$; r is the number of equivalent eigenvalues; and the summation extends over the values of j, where $j = m, m-1, m-2, \ldots, m-q$. To test the null hypothesis, this test statistic has a chi-squared distribution with degrees of freedom $\frac{1}{2}r(r+1) - 1$, where r increases by one each time the hypothesis is run. The a significant eigenvalues found using Anderson's test are the key to the R-A SVD method.

1.3 Steps 3 and 4

The last two steps of R-A SVD focus on the a corresponding vectors of V (step 3) and U^* (step 4), which help identify agreeing sources and factions. V is the set of orthonormal eigenvectors of the covariance of P. With one significant eigenvalue, the first eigenvector identifies which of the m sources agree (and how strongly they agree) as to the rank ordering of the data. When multiple eigenvalues are significant, the columns of eigenvectors can be used to sort the sources into groups of agreeing factions and identify other meaningful structure in the data.

Since orthonormal vectors are normalized to have Euclidean length of 1 by definition (Johnson, 2002), the sum of the squared elements of each column of V sums to 1. When some sources agree, those sources achieve values that approach $\sqrt{1/q}$, where q is the number of sources in agreement. The remaining $m - q$ values approach zero. In PCA, the method of determining which sources are important (in our case, agreement) generally looks at the coefficients provided in the column vector space. The term *reasonably large*

Table 1: Rank Correlation Methods Compared to Empirical Data

Level of Association	Empirical	$\lambda_{R-A\ SVD}$	Avg. τ	Kendall's W
Low	0.0047	0.0017	0.0014	0.2515
Moderate	0.6637	0.5258	0.5252	0.8134
High	0.7712	0.6967	0.6963	0.9049

is used to refer to the variables of interest where larger coefficients, indicate more interest given to a specific element (Johnson, 2002). Therefore, determining the members that are part of a specific principal component is based on approximate estimates without an exact test to prove significance.

The final step of R-A SVD considers the a significant vectors of U^*. The values within the U^* matrix capture the meaningful distances between the elements that reflect the correlation information from the structure of the data. We use the significant vectors of U^* to divide ranked elements into equivalence classes by creating hierarchical clustering graphics. The initial view of the data will consider dendrograms to assess how many classes may exist and which elements belong in each class. Using these classes, other graphics established for this test are discussed. These new graphics highlight the distances of the elements both within and between classes and in many cases reveal meaningful ordering of the classes.

1.4 Simulations Providing Evidence That R-A SVD Is Less Biased Than Kendall's W

A model was created where $m = 4$ sources each ranked the same 20 elements. In each case, the sources were created to have a specified overall association: none, moderate or high. The initial 20×4 matrix was created using a continuous distribution. As a baseline measurement for association, Pearson's correlation was used to analyze the continuous variables as seen in the Empirical column (Table 1). From the continuous data matrix, the numbers were converted to ranks where the lowest number received a value of 1, next lowest received a value of 2, and so forth up to the highest number, which received a value of 20. Finally, the three rank correlation methods for association were calculated: $\lambda_{R-A\ SVD}$, average τ, and Kendall's W, respectively.

Although conservative, there is evidence that $\lambda_{R-A\ SVD}$ and average τ both assess whether there is agreement among the rankers with values slightly below that of the empirical values (Table 1). In instances where no association exists, the empirical, $\lambda_{R-A\ SVD}$, and average τ values are all approximately zero, while Kendall's W overestimates the association (0.2515). In fact, Kendall's W overestimates the value in all cases considered. Note: average τ always produces similar results (within three decimal places) to $\lambda_{R-A\ SVD}$. We simulated these tests thousands of times with the same results, leading to the conclusion that the methods of average τ and $\lambda_{R-A\ SVD}$ are comparable in addressing overall association among a set of sources.

2. NCAA Football 2011 Data Set: Weekly Analysis

R-A SVD will be applied to the NCAA rank data that compares the rankings provided by four sources. Overall, this data set ascertains the use and validity of the R-A SVD method.

Four different organizations ranking college football are considered. The Associated Press College Poll (AP Poll) provides weekly rankings of the top 25 NCAA football teams. Compiled by polling sportswriters across the nation, each writer submits a ranking of the

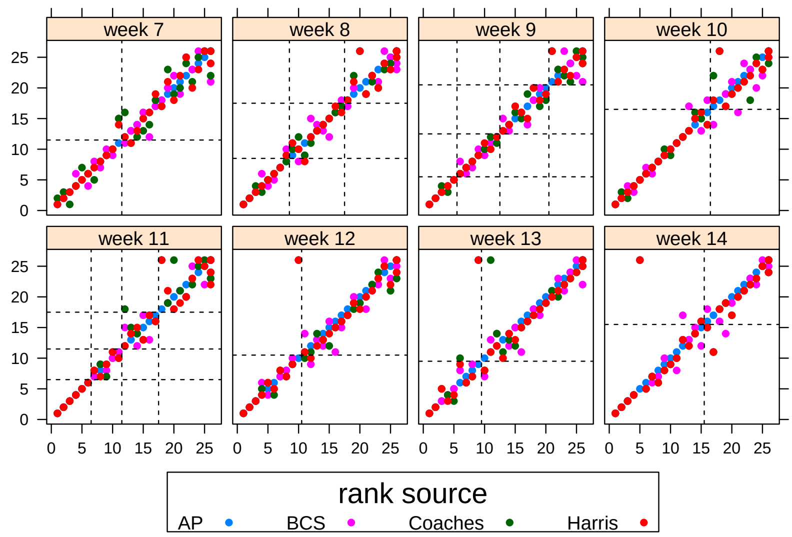

Figure 1: Modified xy plots for each week for weeks 7 through 14 that show correlation and feasible partitioning, which were chosen each week by dendrograms (not shown) created using the significant vectors of U^* as defined in each specific week discussed below.

top 25 teams. A team receives 25 points for a first place vote, 24 points for a second place vote, and so forth to 1 point for a 25^{th} place vote. To produce the national ranking, the sum of the scores is considered. The team with the most points receives a rank of 1, the next highest a rank of 2, until the top 25 have been ranked. The *USA Today*'s Coaches Poll (Coaches) compiles rankings in a similar way but relies on a board of 59 head coaches from Division I football bowl subdivision institutions who vote for the top 25 teams. The Bowl Championship Series (BCS), which is only released for weeks 7 through 14, relies on a combination of polls and computer selection methods to determine relative rankings. Of note, this source can include data from the previous two rank methods and therefore has an inherently dependent relationship. Finally, the Harris Interactive College Football Poll (Harris), only available for weeks 6 to 14, is compiled by Harris Interactive, a market research company that specializes in Internet research.

Although there are 15 weeks in a standard college football season, the purpose of examining this data is to compare multiple sources ranking the same set of elements, with the assumption that the data will be highly correlated sets of ranks produced with significant overlap of elements. For the analysis of weekly data, only weeks 7 through 14 were considered since it is the only timeframe during which all four sources released rankings. We examined the 2011 season because the data was recent and readily available. Of note, the rankings change during the season due to the nature of the game, resulting in teams either joining or leaving the list and other teams exchanging places. Although the ranking is of the top 25 teams by week, there are 43 teams ranked within the polls between weeks 7 and 14 in 2011. Per Fagin et al's (2003) imputation method, a placeholder value of 26 is used when data is missing.

In order to concisely see that using R-A SVD to assess association is a better method than W as well as provide a snapshot of the rest of the findings for each week, a modified xy plot of the AP Poll data (Figure 1) is provided. Table 2 provides the correlation averages by

Table 2: Average Correlation Methods by Week NCAA 2011 Football Rankings

Week	Average τ	$\lambda_{R-A\ SVD}$	Kendall's W	# of Eigenvalues
7	0.8759	0.8796	0.9786	1
8	0.9004	0.9006	0.9810	2
9	0.8971	0.8972	0.9814	1
10	0.8786	0.8788	0.9702	3
11	0.8575	0.8576	0.9607	1
12	0.8733	0.8736	0.9553	3
13	0.8421	0.8424	0.9267	1
14	0.8600	0.8607	0.9318	3

Table 3: First Vector of V: Week 7

Source	Vector 1
AP	0.5031
BCS	0.4915
Coaches	0.4978
Harris	0.5075
Eigenvalue	547.3

method for each of the eight weeks studied in the first three columns: average τ (column 1) to reinforce its equivalent results to $\lambda_{R-A\ SVD}$ (column 2). Kendall's W is in column 3. In the fourth column, the number of significant eigenvalues found ($\alpha = 0.05$) for each week is listed. When examining the correlation values from Table 2 and Figure 1 together, there is overwhelming evidence that Kendall's W overestimates the association among the sources as judged by the patterns on Figure 1. If the association was as high as the values attained by Kendall's W, one would expect to see a nearly perfect straight line along the diagonals of each plot. For this article, we discuss the cases of 1, 2, and 3 significant eigenvalues in detail (weeks 7, 8 and 10, respectively).

2.1 Week 7

For week 7, the overall association for $\lambda_{R-A\ SVD}$ is 0.8796 and for W is 0.9786 and one significant eigenvalue is found (Table 2) . V_1 is considered to determine feasible factions. The V_1 column of Table 3 indicates that all four sources strongly agree as to the top 25 rankings in week 7 of the 2011 season, which is evidenced by each value closely approaching $\sqrt{1/4}$.

Finally, we consider the first column of the U^* matrix to identify feasible equivalence classes. The dendrogram (see Figure 2) identifies three classes using a height between 0.35 and 0.55 on the y-axis. The xy plot (Figure 3) is created by mapping U_1^* to itself, with color coding identified from the dendrogram (see Figure 2). Both figures identify these classes, although the larger spaces between points (along the y-axis) show the distance both within and among classes on the xy plot. Moreover, the xy plot orders the initial data (not evident on the dendrogram) where the top 6 ranked teams (blue) appear at the top. The same teams appear in the center of the dendrogram (Figure **??**). The second class (middle 10 ranked teams) appear in red on xy plot while the lowest 11 ranked teams (green) appear at the bottom.

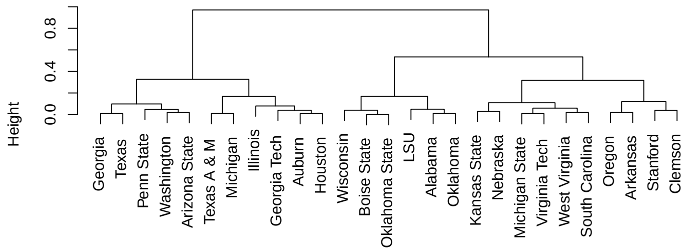

Figure 2: Classic dendrogram using the rankings from week 7 to identify proper partitioning. Three equivalence classes are considered using a height on the $y-$axis between 0.35 and 0.55.

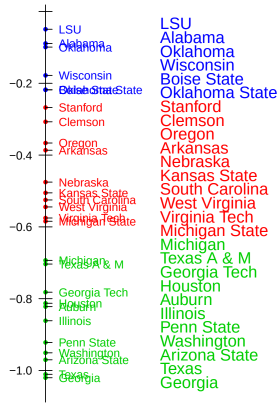

Figure 3: An *xyplot* of week 7 created by mapping the first vector of U^* to itself. The top class of teams (1 through 6) is in blue, the next best class of teams (7 through 16) is in red and the bottom class of teams (17 through 26) is in green. The figure (left) represents the actual positioning of the points, which leads to overprinting. The names are reprinted in the column (right) in the same order for ease of reading.

Table 4: First Two Vectors of V from the Week 8 College Football Rankings

Source	V_1	V_2
AP	−0.498	−0.418
BCS	−0.488	0.841
Coaches	−0.505	−0.337
Harris	−0.509	−0.063
Eigenvalues	2776.7	121.9

2.2 Week 8 Analyzed Using R-A SVD

In week 8, $\lambda_{R-A\ SVD}$ is 0.9006 and W is 0.9810; two significant eigenvalues are identified. The first two vectors of V are examined to identify feasible factions and anomalies (Table 4). V_1 again implies that all four sources exist within one faction, with values approaching $\sqrt{1/4}$. These sources strongly agree as to the top 25 rankings. V_2 provides evidence that the second source, BCS (with a strong positive value), may have some issues in regards to the strength with which it agrees to the other sources. V_2 demonstrates that two of the sources (AP Poll and Coaches) have inverse signs (indicating opposition) to BCS and moderate agreement with one another. Harris approaches zero (indicative of noise). Since the first eigenvalue is much larger than the second, the first relationship is more meaningful than the second.

To further investigate this potential anomaly, we review the modified xy plot of each source versus the AP Poll (Figure 6). This graphic reveals multiple locations where BCS agrees to the inclusion of football teams within a specific range of rankings but orders the teams inversely to every other rank source or, more succinctly, the magenta points representing BCS often form a perpendicular hash mark across the diagonal as the line increases from 0 to 26. This is more obvious at two specific locations, between rankings 11 and 15 and from 23 to 26. In these ranges there is evidence that the other three sources agree in their opposition of BCS.

Finally, feasible equivalence classes are identified in both the dendrogram (see Figure 5) and xy plot (Figure 6) which are created using the first two vectors of U^*. Using a height between 0.45 and 0.60 on the $y-$axis, three equivalence classes are identified on the dendrogram. The color coded classes from the dendrogram appear on the xy plot (Figure 6) revealing class partitioning as well as the ordering of these classes. The top eight teams (black) are at the top of the figure, the next best nine teams (red) in the middle, and the bottom eleven teams (green) at the bottom of the figure.

2.3 Week 10 Analyzed Using R-A SVD

When week 10 is factored using R-A SVD, $\lambda_{R-A\ SVD}$ is 0.8788 and W is 0.9702. Since three significant eigenvalues are found, factions of sources and/or equivalence classes may exist. In interpreting the data, we noted that a feasible influential observation was identified when looking at the graphics. Specifically, USC was ranked 18 by the AP Poll and not considered by any other source, leading us to posit that this datum has an effect on the analysis. Therefore, we consider week 10 a second time eliminating USC from the model. Without USC, the correlation value is 0.879 for $\lambda_{R-A\ SVD}$ and is 0.97 for W (no change); there are still three significant eigenvalues.

When the first three vectors of V are considered to identify which sources belong within each faction no difference is noticed with the absence of USC. The results listed in Table 5 indicate very strong association among all four sources, with values from both V_1 columns

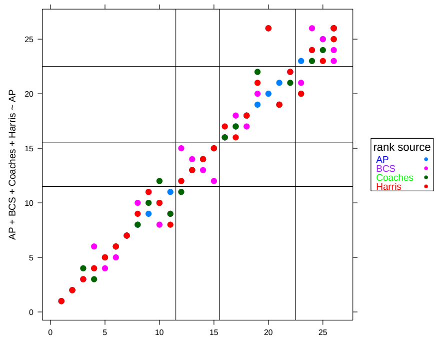

Figure 4: The modified xy plot of all sources plotted against the AP Poll during week 8. Notice that the perpendicular lines (magenta) along the diagonal indicate the disagreement with specific ordering by BCS. Color code key: blue = AP Poll; magenta = BCS; green = Coaches; red = Harris.

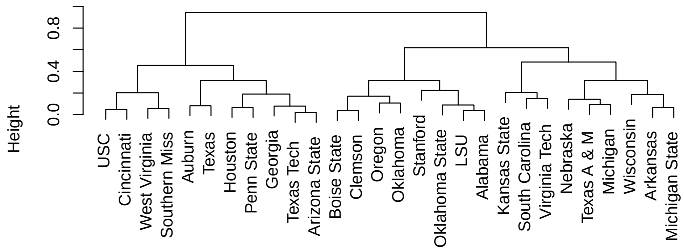

Figure 5: Classic dendrogram using the rankings from week 8 to identify proper partitioning. Three equivalence classes are considered using a height on the y-axis between 0.45 and 0.60.

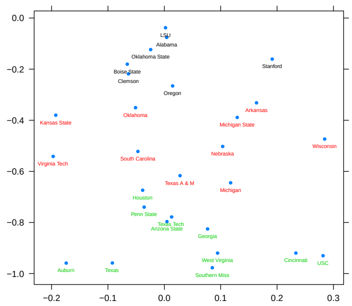

Figure 6: An *xyplot* of week 8 created by mapping the second vector of U^* to the first. The top class of 8 teams (black), the next best class of 9 teams (red) and the bottom class of 11 teams (green).

Table 5: First Three Vectors of V for the Week 10 College Football Rankings First with Then Without USC Included.

	With	USC		Without	USC	
Source	V_1	V_2	V_3	V_1	V_2	V_3
AP	−0.5010	0.0848	0.7841	−0.5105	0.0374	0.5435
BCS	−0.4861	−0.8395	−0.2392	−0.4822	−0.8252	−0.2777
Coaches	−0.4991	0.4746	−0.5725	−0.4957	0.5038	−0.6902
Harris	−0.5134	0.2507	0.0178	−0.5110	0.2527	0.3887
Eigenvalues	2545.9	132.3	91.1	2395.8	129.5	59.6

approaching $\sqrt{1/4}$ (and all signs are the same). Using 0.30 as a cutoff, both V_2 identify BCS strongly. BCS has the opposite sign to the Coaches Poll, showing disagreement between these sources (the other sources approach zero). With USC included, V_3, identifies the AP Poll strongly, which also is in contention with the Coaches poll (again, the other rank sources approach zero). One slight difference observed that when USC is dropped is that V_3 shows Harris and Coaches in agreement with one another and in contention with Coaches (using 0.30 as a cutoff). Again, the first vector corresponds to the largest eigenvalue by far (see Table 5), providing most of the useful information in both cases.

Finally, we will identify feasible equivalence classes using the corresponding significant eigenvectors of U^* to create three graphics: a dendrogram, a scatterplot matrix, and a three-dimensional scatterplot. In all cases, we show these graphics side by side where the left represents the data including USC and the right is the data eliminating USC. On the dendrogram (Figure 9), using a height between 0.95 and 1.05 on the y-axis reveals two classes (one weaker and one more cohesive) and the singular straying point (USC). USC was ranked only by the AP Poll. When this fact is combined with the behavior of the graphic, it suggests that the results of R-A SVD may be affected when an element is only considered by a single source, which is more evident in the U^* graphics. The dendrogram (Figure 10) when USC is dropped reveals that three classes are easily identified at heights ranging between 0.65 and 0.95 on the y-axis.

Color coding for the feasible equivalence classes is identified from each dendrogram, establishing the coloring for each scatterplot matrix (Figure 8) where the significant U^* vectors are mapped to one other. Figure 8 reveals the two equivalence classes and the singular point (USC) which reveals that removing USC created three equivalence classes. In both cases, the true partitioning is most evident when comparing U_1 by U_2. When USC is included, the single green dot (USC) looks as though it could belong in the center of red class. The mapping of each vector by U_3 reveals overlaps among the data partitioning. Finally, the three-dimensional graphic including USC (Figure 11) clearly shows the partitioning of the classes as well as the single straying point identified by the other graphics (Figures 8 and 9). When seen live, the axis can be rotated to show that these classes exist in most views of the data. The snapshot seen in Figure 11 only captures the third dimension slightly. The blue, red, and green points continue to show two separated classes and a stray point (USC). The three-dimensional plot excluding USC (Figure 12) allows for a better view of the absolute differentiation of the classes. From the initial data, the top 12 teams (identified in dark blue points) are clearly the best class. The other two classes are partitioned properly, although the ordering is not as evident.

We believe that USC may be an influential observation similar to those found in regression analysis. An influential observation can completely change the results of an analysis and is identified when Cook's distance metric is applied and a value over 1 is attained. We recognize that a similar test may be needed when elements are only ranked by a single

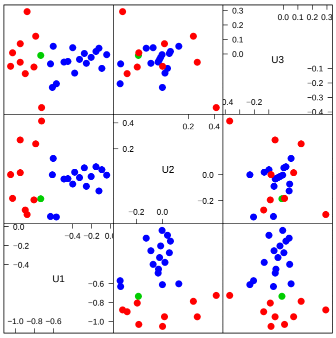

Figure 7: A scatterplot matrix representing the first three vectors of U^* that assists in identifying two classes and a stray point in the week 10 data. The partitioning is best identified when U_1 is mapped to U_2. The outlier (green dot), USC, is mostly embedded in the other classes. The top class of 16 teams (dark blue) and the lower class of teams (17, 19 through 26) are in red, excluding USC in green (18).

source influencing the data and will explore this in future work. In this case, the influential observation masks the true partitioning of the equivalence classes. Removing USC does not affect either the correlation or the factions but does affect the proper partitioning of the classes.

2.4 Conclusions of Weekly In-Depth Analysis

Overall, each week (7, 8, and 10) show that there is a great deal of agreement among all four sources. In all cases the agreement is not perfect as evidenced by the modified xy plots in Figure 1; hence, W overestimates the association every time. For week 7, the singular dimensional xy plot reveals the ordering and distances within and among equivalence classes. For week 8, the second eigenvalue identifies unexpected results, possibly influenced by the BCS poll where the BCS poll provides inverse orderings within smaller groups when compared to the remaining three sources. As such, R-A SVD may be affected by the slightly rebellious BCS (hence, the second significant eigenvalue). In week 10, three significant eigenvalues are found where the anomaly of USC was discovered.

In all cases, using the corresponding significant eigenvectors of U^* provides evidence of feasible equivalence classes. In week 8, the two dimensions provide a better view of the equivalence classes. In week 10, using three significant eigenvalues allows for an in-depth look into the data, which identifies a single faction and two equivalence classes with a single straying point. In this case, removing USC only affects the equivalence classes. It should be noted that the partitioned equivalence classes become more visible as more dimensions are available to view them. In future work, we will consider eliminating points that are only ranked by a single source.

REFERENCES

Anderson, T. W., (1951), "The Asymptotic Distribution of Certain Characteristic Roots and Vectors," *In proceedings Second Berkeley Symposium on Mathematical Statistics and Probability*, 103–130.

Anderson, T. W. (1963), "Asymptotic Theory for Principle Component Analysis" *Annals of Mathematical Statistics*, 34, 122-148.

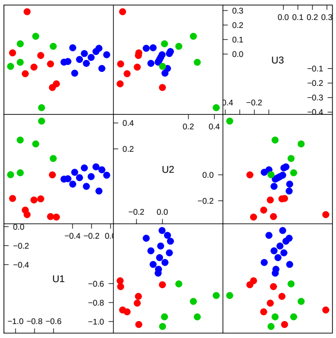

Figure 8: Scatterplot of the first three vectors of U^* mapped to one another after USC is removed from the data set. Each plot shows three classes mostly partitioned. The top 12 teams (dark blue) and clearly separated, but the ordering of the remaining classes is not as evident (red and green).

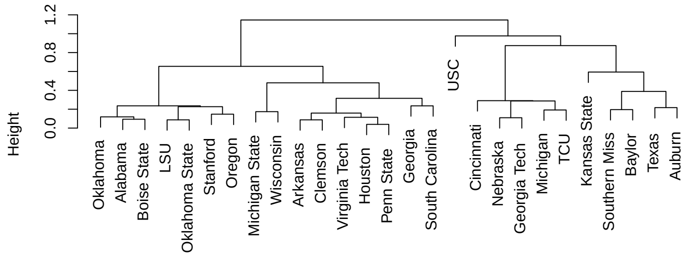

Figure 9: A dendrogram showing two classes and a singular straying point (USC) when looking at the rankings for week 10.

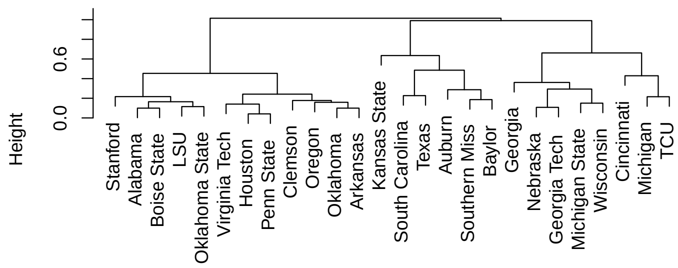

Figure 10: New dendrogram of the week 10 data after USC is dropped from the data. Three classes are evident using a height on the y-axis between 0.65 and 0.95.

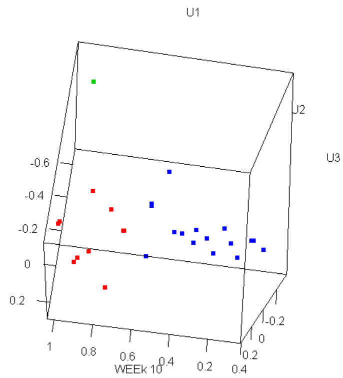

Figure 11: A still frame of the three-dimensional plot generated using the same color scheme as the scatterplot matrix and dendrogram to show the identified classes. The top class of teams (1 through 16) is in blue, the next best class of teams (17 through 26) are in red, excluding USC in green.

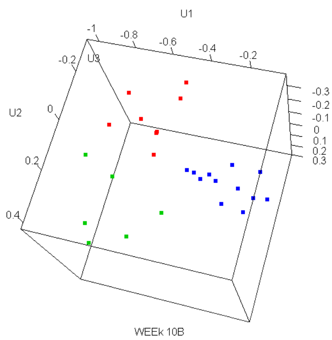

Figure 12: A still frame of the three-dimensional plot generated using the color scheme identified by the dendrogram which shows three classes clearly partitioned after USC was removed. The top 12 teams are dark blue. The next two classes of teams (17 through 25) are in red and green. There is no clarity on which is the better class.

Bartlett, M. S. (1954), "A note on multiplying factors for various chi-squared approximations," *Journal of the Royal Statistical Socitey*, 16, 296–298.

Dunlap, W. and Silver, C. and Phelps, G., (1987), "A Monte Carlo Study of Using the First Eigenvalue for Averaging Intercorrelations," *Educational and Psychological Measurement*, 47, 917–923.

Freidman, S. A. and Weisberg, H. (1981), "Interpretting the First Eigenvalue of a Correlation Matrix," *Educational and Psychological Measurement*, 41, 11–21.

Girshick, M. A. (1936), "Principal Components," *Journal of the American Statistical Association*, 31, 519–528.

Hotelling, H.(1933), "Analysis of a Complex of Statistical Variables into Principal Components," *Journal of Educational Psychology*, 24, 498–519.

Kaiser, H. (1968), "A Measure of the average Intercorrelation," *Educational and Psychological Measurement*, 28, 245–247.

Kendall, M. G. (1948), *Rank Correlation Methods*, Charles Griffin and Company Limited.

Lawley, D. N. (1956), "A Modified Method of Estimation in Factor Analysis and Some Large Sample Results," *Uppsala Symposium of Psychological Factor Analysis*, 35–42.

Lawley, D. N. (1963), "On Testing a Set of Correlation Coefficients for Equality," *Annals of Mathematical Statistics*, 34, 149–151.

Morrison, D. F. (1976), *Multivariate Statistical Methods*, McGraw-Hill.