2 **On the use and abuse of principal component analysis in**
3 **biomechanics**

4 Daniel J Cleather^{1,2}

5 ¹ St Mary's University, Waldegrave Road, Twickenham, TW1 4SX, UK

6 ² Institute for Globally Distributed Open Research and Education (IGDORE)

7 Email: daniel.cleather@stmarys.ac.uk; Twitter: @dr_jump_uk

8 **Abstract**

9 Principal component analysis (PCA) is a data reduction technique that is increasingly popular
10 within biomechanics. However, a majority of the literature that has used PCA has employed
11 a methodology that is less strongly aligned to the philosophical basis of the technique, and
12 that may not yield the most practically interpretable results. In this technical report I
13 exemplify the two approaches to PCA that have been used in the biomechanics literature, and
14 argue that one method may be superior to the other for the majority of biomechanical
15 applications.

16 Keywords: PCA, vertical jump, knee joint moment, waveform analysis, data reduction,
17 degrees of freedom

18 Citation:

19
20 Cleather, D. J. (2019). On the use and abuse of principal component analysis in
21 biomechanics. *SportRxiv*. doi: [10.31236/osf.io/f48qk](https://doi.org/10.31236/osf.io/f48qk)
22

Introduction

The fundamental basis of principle component analysis (PCA) is to find a linear transformation (rotation) that can be used to map the original data set to a new coordinate system. The purpose of this transformation is to find a set of new axes or functional variables for the data (called the principal components; PC) that are uncorrelated with each other. Additionally, PCA will order these variables in terms of the variance in the data set that they explain. For instance, let's say that we have a data set that consists of n observations of p different variables. This data set can then be represented by the $n \times p$ matrix X . The axes of the original data set are the p variables, however these variables may be correlated with each other to some degree. The PCA will find the $p \times p$ rotation matrix U such that $Z = XU$. Z is a $n \times p$ matrix which describes the data set expressed in the new coordinate system, where each column now represents the observations for each of the new variables q , where $p = q$.

Generally, the motivation behind PCA is as a data reduction technique. In particular that if our p raw variables are correlated with each other in some way that we should be able to express our data with a reduced subset of our new variables q that still describes the majority of the variance in the data. This is achieved by setting a threshold for the desired explained variance and then only retaining the PCs that are needed to achieve this. An intuitive sense of the usefulness of PCA can be gained by looking at a trivial example. Let's say that our data matrix consists of 3 dimensional coordinates and that we have 100 data points. In addition, let's say that our data points fit closely to the line r with only minimal random deviations from the line. If we performed a PCA on the data matrix X we would expect it to find that the first PC would be an axis that closely approximates the line r and that the first PC explains a high proportion of the variance. The remaining two PCs would be orthogonal to this axis and each other, and would describe the deviations from r . In this case, we could discard the

second and third principal components. The expression of our data set in the new coordinate space would simply be a vector that gives a parameter describing the position of the data point along the first PC. In this case, the data is reduced from three degrees of freedom to one.

Use of PCA for the Comparison of Waveforms

The discussion so far has presented the use of PCA to analyse a data set that consists of time-series data for a number of different variables. For instance, in biomechanics we might have a data set that consists of hip, knee and ankle joint moments (three variables) measured across one movement trial. However, what if we want to compare the time-series data of just one variable across several different movement trials? For instance, we might want to compare the knee joint moment of an athlete across several different vertical jump trials or the knee joint moments of several athletes. There are two possible approaches to this analysis and both have been employed within the biomechanics literature. However, the two approaches yield different results and it is my contention that one is more correct than the other. The purpose of this article is to exemplify and contrast these two methods by reference to a simple example.

Methods

For the purposes of this article I will exemplify the results of the two approaches using a data set taken from 12 participants performing a vertical jump trial with arm swing (age = 23 ± 3 years, mass = 77 ± 13 kg, height = 174 ± 9 cm). Raw data measured from the participants consisted of the position of 18 retro-reflective markers placed on the lower limb (quantified with motion capture technology; Vicon MX System, Nexus 2.2 software, Vicon Motion Systems Ltd, Oxford, UK) and 3 dimensional ground reaction forces (quantified with a force

plate that was synchronised with the motion capture system; Kistler Type 9287BA, Bioware 3.24 software, Kistler Instruments Ltd, Hampshire, UK). This raw data was used to calculate the knee joint moments generated by each participant during the vertical jump using a publicly available model of the lower limb (Cleather and Bull, 2015). A detailed description of the processing of this data is provided in Cushion et al. (2019). To facilitate comparison of different time-series the data was time-normalised. The data can thus be represented by the 101×12 matrix X , where the rows represent the value of the knee joint moment at 1% intervals and each column represents a separate participant. PCA was performed using Matlab (version 9.2, R2017a).

Results and Discussion

The use of PCA (method 1)

One approach to the PCA of this type of waveform data is to treat the individual trials as different variables and to treat the time-points as observations (Daffertshofer et al., 2004). For the purposes of the example data here, this would mean that the PCA is simply performed on our matrix X , and U will be the 12×12 rotation matrix that defines the linear transformation to the new coordinate system. The new coordinate system will have 12 orthogonal axes (PCs) that are a linear combination of the 12 trials, and the PCA will give us the variance explained by each axis. Finally, Z will be the 101×12 matrix which describes the values of the data set given in the new coordinate system, and the columns will be ordered according to the variance explained.

Typically when using PCA as a data reduction technique, the next step would be to reduce the number of axes to those that describe the majority of the variance in the data. For the purposes of this example, we will take the PCs that describe 90% of the variance in the data,

which in this case is just the first (PC1) and second (PC2) PCs. We can then plot the data set in the new coordinate system, focussing on just the new variables (PC1 and PC2) that explain most of the variance in the data set (Figure 1).

Figure 1. Principal component (PC) scores for the first 2 PCs of the data matrix X (percentages indicate the variance explained by each PC).

The abuse of PCA (method 2)

The method described above is not the most common approach to the PCA of waveform data that has been taken in biomechanics however. Instead, as described in an article by Deluzio and Astephen (2007), it has been more usual to treat the normalised time-points as the ‘variables’ within the PCA whereas the trials are treated as the different observations. Thus the matrix that is analysed within the PCA is actually the transpose of our data matrix, X^T . In this scenario, the rotation matrix V is a 101×11 matrix and the matrix expressing the data in the new coordinate system W is a 12×11 matrix. (Note that V is not a 101×101 matrix and W is not a 12×101 matrix as would be suggested by the general statement of PCA given in

108 the introduction simply because the number of trials is smaller than the number of variables.
109 There are only 11 degrees of freedom, and thus we can only have variation along a maximum
110 of 11 axes.)

111 What needs to be understood here is that by treating the time-points as variables we are
112 actually defining each axis to be the value of the knee joint moment at a particular time-point.
113 That is, column 1 of X^T is the value at time $t = 0\%$ and so on, whereas the rows of the matrix
114 give the value for each trial. Our data set is thus described within a coordinate system that
115 has 101 different time axes, where each axis represents a discrete (non-continuous) moment
116 in time. The transformation that is given by the PCA then maps the data set to a new
117 coordinate system that has 101 uncorrelated time axes (the PCs). If we then reduce the
118 dimensionality of the data set by just taking the PCs that describe 90% of the variance, what
119 we are actually doing is trying to find a limited set of time-points in the new coordinate
120 system that describe most of the variance in the data. In our example, this can be achieved
121 with 4 PCs. We can then plot the data set within the new coordinate system (Figure 2).

Figure 2. Principal component (PC) scores for the first 4 PCs of the data matrix X^T (percentages indicate the variance explained by each PC).

124

Figure 2 is clearly hard to interpret. Instead, when authors have used this approach to PCA, they have more commonly presented the columns of V , the so-called loading vectors. For instance, column 1 of V is a unit vector in the original coordinate space in the direction of the first PC. Figure 3 presents the loading vectors that describe the first 4 PCs for this example. It is then probably easiest to understand Figure 2 by observing that $X^T = WV^T$. That is, we can reconstruct our original data as a linear combination of the loading vectors, and the PC scores that we plotted in Figure 2 give the relative weight of each loading vector in each particular trial.

Figure 3. Loading vectors (LV) for the first 4 principal components of the data matrix X^T (percentages indicate the variance explained by each principal component).

Use or abuse?

So we have seen that there are two possible approaches to the use of PCA with waveform data, and that both have been employed within the biomechanics literature. It is also apparent that the two methods can yield markedly different results. So which approach should we use?

Clearly method 1 is more clearly aligned to the traditional, theoretical description of PCA, whereas method 2 is based on some unusual assumptions which seem harder to justify philosophically. In particular, the treatment of time as a discrete variable with multiple time-like axes is a non-standard approach, and it is strange to try and reduce time-series data to a limited number of time-points (especially given that the new time-points will not be ordered temporally).

146 Fundamentally for method 1, PCA is analogous to finding the eigenvectors of the covariance
147 matrix $X^T X$. This matrix is constructed by multiplying the columns of X with each other, and
148 thus the PCA is based on making the comparison of the complete time-series of one trial to
149 another (for all possible combinations). In contrast, method 2 is based on the analysis of the
150 matrix XX^T – the multiplication of the rows of X with one another. It is thus based on
151 comparing the values for all trials at one particular time-point to those at another time-point
152 (for all possible combinations of time-points). I would argue that the theoretical basis for
153 method 1 is thus much stronger – we are comparing complete waveforms to one another
154 which we might expect to be similar. In contrast, method 2 involves the comparison of the
155 trial values at different time-points despite the fact that we have little reason for suspecting
156 that the values at $t = 0\%$ bear any relation to the values at $t = 50\%$ (say).

157 Figure 4 presents the results of our two analyses compared to the mean values of the raw
158 data. It is clear that the PC1 that is a result of method 1 is qualitatively much more similar to
159 the raw data than the LV1 that was produced by method 2. Similarly, consideration of
160 Figures 1-3 shows that method 1 is able to explain 94.4% of the variance in the data set using
161 only 2 PCs where method 2 requires 4 PCs to explain approximately the same amount of
162 variance (94.6%). One factor that is likely to contribute significantly to the lower amount of
163 variance described by method 2 is that the coordinate space is underdetermined by the data
164 set – we have 101 axes, but only 12 trials.

Figure 4. Comparison of the mean knee joint moment for the data set with the first principal components derived from the analysis of the data matrices X and X^T (dotted grey lines indicate 95% confidence interval of the mean values).

It is much easier to interpret the output of method 1. In this example the PCs are ‘moment-like’ – they are a linear combination of the knee joint moment data from each participant. PC1 explains 84% of the variance across subjects, and is thus a fundamental pattern that is clearly the basis of (nearly?) all trials. PC2 likely describes the variation of PC1 between subjects. In addition, the data reduction implies that we can describe the data set with just 2 moment-like axes, suggesting that there are effectively just 2 functional degrees of freedom in the data set. In contrast, the PCs for method 2 are ‘time-like’, and thus harder to relate to the mechanics of the movement.

It therefore seems hard to understand the popularity of method 2 in the biomechanics literature. The method is less philosophically aligned with the fundamental nature of PCA, while at the same time being less readily interpretable in terms of the underlying mechanics

of movement. Similarly, in practical terms, method 1 is likely to be more appropriate for use with typical biomechanical data sets. Firstly, the fact that method 1 involves the comparison of complete waveforms is likely to result in fewer PCs that explain more of the variance. Secondly, as biomechanical data sets often consist of limited number of trials, in many applications there may not be enough trials to fully specify all of the variance in the coordinate space for method 2.

It is important to note that if we consider the two methods purely pragmatically there is no doubt that both provide a reliable and consistent way to analyse a set of waveforms. When considering the two methods in this context, both will yield a subset of waveforms that can be compared in linear combination to reconstruct the original data set. The pragmatist would therefore argue that the most useful method will be the one which facilitates the most useful and interpretable outputs. I do have sympathy for this position. However, I would suggest that in the absence of a compelling reason to use method 2, that method 1 should be the default approach. Similarly, I am not sure that it is correct to refer to method 2 as a principal component analysis, given that the transformation is time-like.

Acknowledgments

I'd like to thank Emily Cushion for providing the data and analysis that is presented here, and for a number of critical discussions in which I developed the understanding presented here. I'd also like to thank her for her wisdom in advising me to moderate the strength of my claims – advice that I have probably not taken seriously enough. Any errors in statistical or philosophical interpretations of PCA are my own.

203

204 **References**

205 Cleather DJ, Bull AMJ. The development of a segment-based musculoskeletal model of the
206 lower limb: introducing FreeBody. *R Soc Open Sci*, 2015; 2(6): 140449

207 Cushion EJ, Warmenhoven J, North J, Cleather DJ. Principal component analysis reveals the
208 proximal to distal pattern in vertical jumping is governed by two functional degrees of
209 freedom. *Front Bioeng Biotechnol*, 2019; in press

210 Daffertshofer A, Lamoth CJC, Meijer OG, Beek PJ. PCA in studying coordination and
211 variability: a tutorial. *Clin Biomech Bristol Avon*, 2004; 19(4): 415–28

212 Deluzio KJ, Astephen JL. Biomechanical features of gait waveform data associated with knee
213 osteoarthritis: an application of principal component analysis. *Gait Posture*, 2007;
214 25(1): 86–93

215