involve

A permutation test for three-dimensional
rotation data

Daniel Bero and Melissa Bingham

msp

A permutation test for three-dimensional rotation data

Daniel Bero and Melissa Bingham

(Communicated by Mary C. Meyer)

Statistical inference procedures that require no distributional assumptions make up the area of nonparametric statistics. The permutation test is a common nonparametric test that can be used to compare measures of center for two data sets, but it is yet to be explored for three-dimensional rotation data. A permutation test for such data is developed and the statistical power of this test is considered under various scenarios. The test is then used in an application comparing movement around joints in the foot and ankle for humans, chimpanzees, and baboons.

1. Introduction

Data in the form of three-dimensional rotations are common in the study of human motion. As skeletal mammals move, the orientation of various joints can be tracked by using infrared emitting diodes attached to bones on opposite ends of the joint. Each joint orientation can be represented mathematically as a 3×3 orthogonal rotation matrix. Of interest here is comparing movement around various joints in the ankle and foot for humans, chimpanzees, and baboons by comparing the central rotation of each joint for the various species.

While other works have considered comparing sets of three-dimensional rotation data, they rely on distributional assumptions [Rancourt et al. 2000; Hendriks and Landsman 1998]. Further, existing work for studying three-dimensional rotations is often in terms of manifold considerations. As such, it is often inaccessible to practitioners outside the area. Our aim here is development of methodology for comparing central rotations that is both nonparametric and does not rely on special manifold theory, so that it can be used more broadly. The permutation test is a commonly used nonparametric test, but it has yet to be implemented for three-dimensional rotation data. We develop such a test in Section 2, explore the statistical power of the test in Section 3, and apply the test to joint data in Section 4.

2. Development of a three-dimensional permutation test

The permutation test is widely used in nonparametric statistics for determining if two data sets are different in some way (e.g., comparing means, variances, shapes). The most common example of a permutation test in one dimension is comparing population means for data sets A and B by using the difference in sample means, $\bar{x}_A - \bar{x}_B$, as a test statistic. To perform the permutation test, data sets A and B are combined and permuted so that data points are randomly reassigned to either A or B. The permuted test statistic is then calculated from this permuted data and this process is repeated a large number of times. If the means of the populations from which A and B come do in fact differ, then we expect the observed test statistic $\bar{x}_A - \bar{x}_B$ to be more extreme than the permuted test statistics. For this reason, the p-value for a permutation test is defined to be the proportion of times that the permuted test statistic is more extreme than the observed test statistic. See [Higgins 2004] for more details on permutation tests.

To translate the idea of the permutation test to three-dimensional rotation data, we first need to define a sensible test statistic that could be used for comparing two central rotations. For each set of three-dimensional rotations, we begin by finding a measure of center as follows. Compute $\bar{O} = 1/n \sum_{i=1}^{n} O_i$ for $O_1, \ldots, O_n \in SO(3)$, where $SO(3)$ represents the set of all 3×3 orthogonal rotation matrices. Next, find the matrix $T = VW$, where $\bar{O} = V \Sigma W$ is the singular value decomposition of $\bar{O}$. Using these components from the singular value decomposition is necessary since $\bar{O}$ may not be an element of $SO(3)$, but T is. This is a commonly used measure of center [León et al. 2006; Bingham et al. 2009; Khatri and Mardia 1977], which we refer to as the "mean" rotation.

Once we have found the mean rotation for each of our two data sets, a natural test statistic is the difference between these mean rotations. One way of quantifying the difference between two three-dimensional rotations is by using angles. A misorientation angle is defined as the angle needed to rotate from one three-dimensional rotation to another via a spin about some axis. For $O, P \in SO(3)$, the misorientation angle between O and P is

$$\text{mis}(O, P) = \arccos\left(\frac{\text{tr}(O'P) - 1}{2}\right), \tag{1}$$

where tr is the trace of a matrix and O' is the transpose of O. We use the misorientation angle between our two mean rotations as the test statistic for the three-dimensional permutation test of H_o: There is no difference between the population mean rotations versus H_a. There is a difference between the population mean rotations. The steps of the permutation test are given below and R code for implementing this test is provided in the Appendix.

(a)   (b)

Figure 1. Plots of two simulated three-dimensional rotation data
sets (each with $n = 50$) with mean rotations that (a) are not signifi-
cantly different and (b) are significantly different.

(1) Calculate the mean rotation for each data set and then find the misorientation
angle between these means. This serves as the observed test statistic, θ_{obs}.

(2) Permute the data a large number (say 10,000) of times, storing the misorienta-
tion angle between the permuted mean rotations, θ_{perm}, each time.

(3) Let the p-value be the fraction of times that the permuted misorientation angle
is greater than the observed misorientation angle; that is,

$$p\text{-value} = \frac{\# \text{ of times } \theta_{\text{perm}} > \theta_{\text{obs}}}{\# \text{ of permutations}}.$$

The three-dimensional permutation test outlined above is briefly illustrated in
two different examples. Figure 1 shows three-dimensional data sets plotted as
points on the sphere, with one observation represented by three points that would
correspond to three orthogonal axes. In Figure 1(a), the two simulated data sets (in
white and black, each of size 50) show considerable overlap. Under the permutation
test, these data sets resulted in a test statistic of 0.0546 and a p-value of 0.3101.
In Figure 1(b), the simulated data sets are more separated. These data sets gave a
test statistic of 0.6102 and a p-value of 0, indicating a significant difference in the
population mean rotations. These examples suggest that the p-value decreases as
expected when the data sets have mean rotations that increase in distance.

3. Power: a simulation study

To examine the effectiveness of the three-dimensional permutation test developed
in Section 2, we perform a simulation study to investigate statistical power. Power
is the probability of correctly rejecting a false null hypothesis. We simulate data
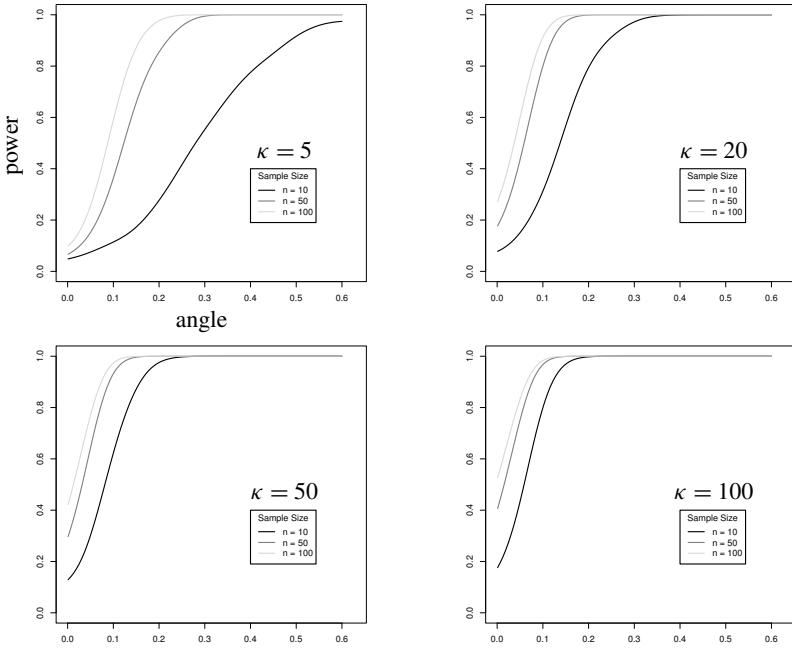sets with centers that differ by a known misorientation angle, ϕ, (i.e., there is a

Figure 2. Plots of power versus misorientation angle for the von Mises version of the UARS distributions with $\kappa = 5, 20, 50, 100$.
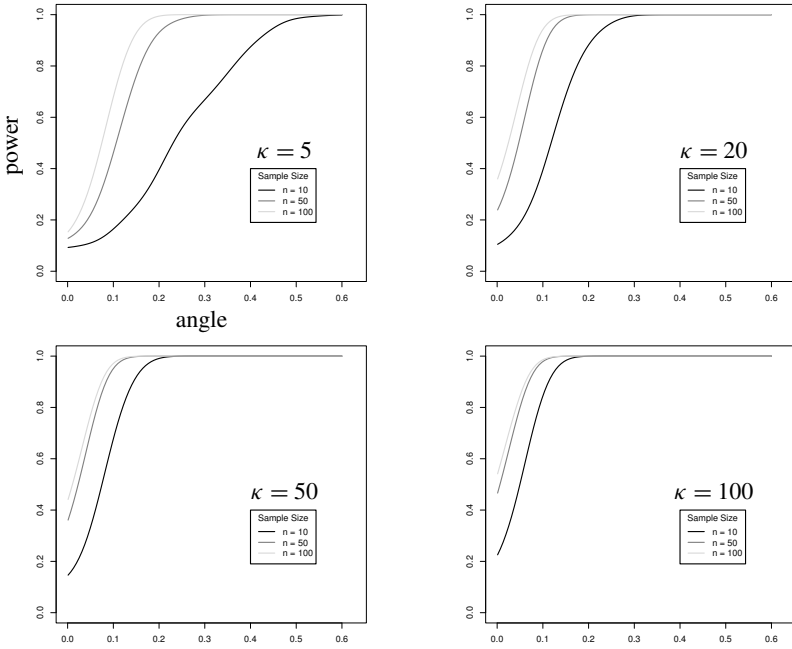

Figure 3. Plots of power versus misorientation angle for the symmetric matrix von Mises–Fisher distribution with $\kappa = 5, 20, 50, 100$.
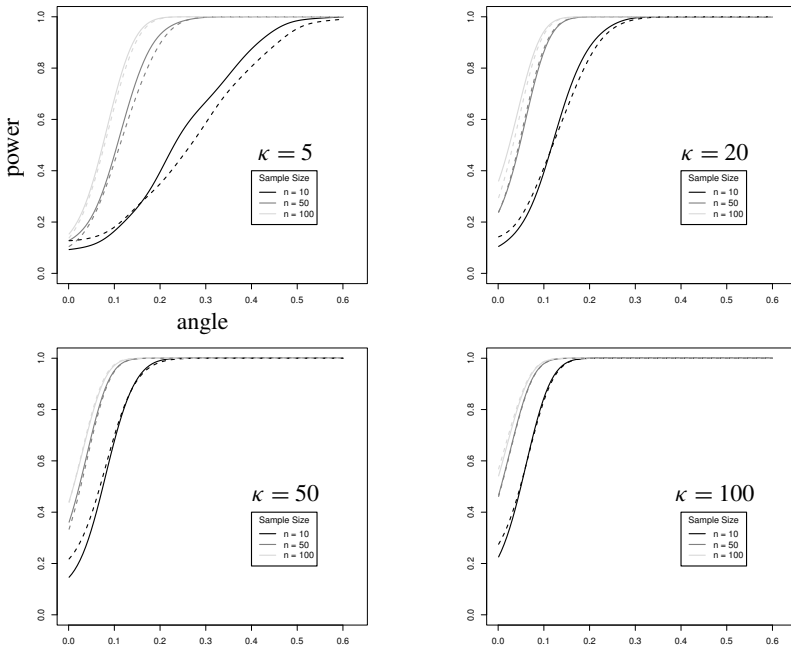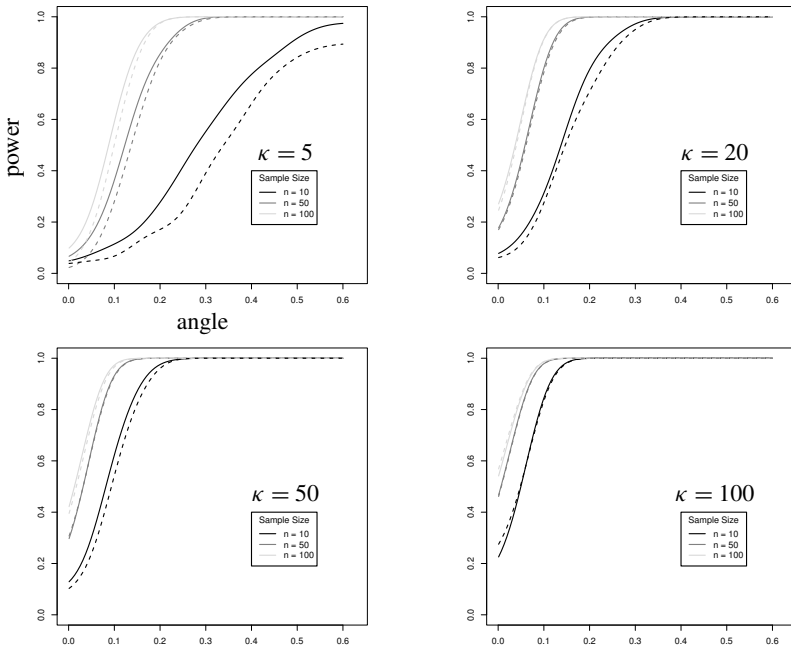
Figure 4. Plots of power versus misorientation angle for the vM-F
distribution with solid lines representing the permutation test and
dashed lines representing the parametric approach for $\kappa =$
5, 20, 50, 100.

difference between the population mean rotations and the null is false) from both
the von Mises version of the uniform axis-random spin (vM-UARS) distribution
[Bingham et al. 2009] and the symmetric version of the matrix von Mises–Fisher
(vM-F) distribution [Khatri and Mardia 1977] . A vM-UARS or vM-F distribution
can be specified by a central rotation $S \in SO(3)$ and a spread parameter $\kappa \in (0, \infty)$,
where κ is best termed as a concentration parameter since larger values of κ indicate
rotations that are less spread about the center at S. Two samples, each of size n,
are generated from vM-UARS(S_1, κ) and vM-UARS(S_2, κ) distributions, where
$\phi = \text{mis}(S_1, S_2)$ as in (1). We consider κ values of 5, 20, 50, and 100, set n at 10,
50, and 100, and let the misorientation angle, ϕ, vary between 0 and $\pi/5$. The same
is done for the vM-F distribution.

   For each combination of κ, n, and ϕ, the permutation test was conducted 1,000
times with 1,000 permutations per test. The power was then found as the proportion
of times (out of 1,000) that the test correctly rejected the null hypothesis of equal
means. Plots of the power against the misorientation angle, ϕ, for the various choices
of n and κ are provided in Figure 2 for the vM-UARS distribution and in Figure 3
for the vM-F distribution. It can be seen from all plots that as sample size increases,

Figure 5. Plots of power versus misorientation angle for the vM-UARS distribution with solid lines representing the permutation test and dashed lines representing the parametric approach for $\kappa = 5, 20, 50, 100$.

the power of the test increases. In addition, as the concentration parameter, κ, increases (i.e., data sets become more clustered around their mean rotation), the power increases. Finally, as the misorientation angle increases and the true centers become farther apart, the power increases. This mimics properties of power for traditional hypothesis tests for differences in means (for nonrotational data), giving evidence that the three-dimensional permutation test performs as desired.

The power of the three-dimensional permutation test was also compared to that of the parametric approach presented in [Rancourt et al. 2000], which requires the observations be distributed according to the matrix von Mises–Fisher distribution. The plots in Figure 4 show power versus misorientation angle for the various choices of n and κ using the matrix von Mises–Fisher distribution. The solid lines represent power for the permutation test, with the dashed lines representing power for the parametric approach. We see that the power of the permutation test is comparable to the power of parametric approach in all cases. The permutation test was also compared to the parametric approach for the vM-UARS distribution, with power plots given in Figure 5. We see that the permutation test outperforms the parametric approach in terms of resulting in a larger power, with this fact more visible when

Figure 6. Bones in the ankle and foot (image taken from
http://www.ceuarmy.com/BSFAFpdf.pdf).

we have smaller sample sizes or data that is more spread (small κ). Thus, the
three-dimensional permutation test is comparable to the parametric approach when
the assumptions of the parametric test are met, and it performs better than the
parametric approach when the assumptions are not met.

4. Application to ankle joint rotation data

Now that we have verified that the three-dimensional permutation test performs as
expected with regard to power, we apply the test to ankle/foot joint rotation data
collected by Prof. Thomas Greiner of the Department of Physical Therapy at the
University of Wisconsin-La Crosse. Data was collected from humans, baboons, and
chimps during circumduction, which is the movement characterized by the foot being
placed flat on the floor and the leg rotating in a circular motion around it. Infrared
emitting diodes attached to bones on each side of a joint give the orientation of each
bone as the movement occurs. If the orientation of the first bone is represented as $\boldsymbol{F}$
and the orientation of second bone is represented as $\boldsymbol{G}$, then the resulting orientation
of the joint is defined as $\boldsymbol{F}'\boldsymbol{G}$. Because markers may not have been placed identically
on all subjects, the orientations of all joints under consideration were measured
with the tibia-talus joint as the reference to allow for comparison of species. Joints
considered were the cuboid-calcaneus, navicular-cuboid, navicular-talus, talus -
calcaneus, and fifth metatarsal-cuboid. (See Figure 6 for a diagram of the bones
in the foot and ankle region.) Orientations were collected for six human subjects,
four chimpanzee subjects, and seven baboon subjects, and the base alignment matrix

corresponding to the primary rotational axis (see [Ball and Greiner 2012]) was used in the three-dimensional permutation test to compare species.

Species were compared pairwise (human versus chimpanzee, human versus baboon, and chimpanzee versus baboon) for each of the joints mentioned above, and each test was done using 1,000 permutations. Out of all tests, there were four significant differences found. There was significant evidence to suggest that the orientation of the navicular-talus joint differs between the humans and chimpanzees (p-value = 0.001) and humans and baboons (p-value ≈ 0). The orientation of the talus-calcaneus joint was found to be significantly different between humans and chimpanzees (p-value = 0.019) and humans and baboons (p-value = 0.001). Therefore, it appears that movement for humans differs from baboons and chimps when considering two specific joints.

5. Conclusion

The analysis of joint rotation data provided here is just one of many applications that the three-dimensional permutation test could be used for. Given the abundance of three-dimensional rotation data in the study of human motion, as well as in the other fields like materials science, having methodology for comparing measures of center for three-dimensional data is important. The three-dimensional permutation test developed here provides that methodology without the need for any distributional assumptions on where the data sets come from. It also does not require any theory on special manifolds, making the three-dimensional permutation test an important addition to the field of statistics, as well as to practitioners who collect data in this form.

Appendix

The following gives an R function called `PermTest` for performing the three-dimensional permutation test on data sets A (of size n_A) and B (of size n_B). The argument A must be an array of dimension $3 \times 3 \times n_A$ and B must be an array of dimension $3 \times 3 \times n_B$. The argument `nspec` specifies the number of times the data should be permuted. The function `PermTest` outputs the test statistic (misorientation angle between the two sample mean rotations) and p-value.

```
PermTest=function(A,B,nspec){
  ##Loads functions needed for test
  trace=function(M){sum(diag(M))}
  Mis.Ang=function(C,D){acos((trace(t(C)%*%D)-1)/2)}

  ##Finds mean matrices for both sets of data
  na=dim(A)[3]
  Abar=matrix(rep(0,9),nrow=3)
  for(i in 1:na){Abar=Abar+A[,,i]}
```

```
  Abar=Abar/na
  M.A=svd(Abar)$u%*%t(svd(Abar)$v)
  nb=dim(B)[3]
  Bbar=matrix(rep(0,9),nrow=3)
  for(i in 1:nb){Bbar=Bbar+B[,,i]}
  Bbar=Bbar/nb
  M.B=svd(Bbar)$u%*%t(svd(Bbar)$v)

  ##Finds the test statistic
  Test.Stat=Mis.Ang(M.A,M.B)

  ##Puts data into one array
  T=array(c(A,B),dim=c(3,3,(na+nb)))

  ##Performs the permutation test
  nsim=nspec
  ang=rep(0,nsim)
  for(i in 1:nsim){
    samp=sample(1:(na+nb))
    O=T[,,samp[1:na]]
    P=T[,,samp[(na+1):(na+nb)]]
    Obar=matrix(rep(0,9),nrow=3)
    for(j in 1:na){Obar=Obar+O[,,j]}
    Obar=Obar/na
    M.O=svd(Obar)$u%*%t(svd(Obar)$v)
    Pbar=matrix(rep(0,9),nrow=3)
    for(k in 1:nb){Pbar=Pbar+P[,,k]}
    Pbar=Pbar/nb
    M.P=svd(Pbar)$u%*%t(svd(Pbar)$v)
    ang[i]=Mis.Ang(M.O,M.P)
    }
  p.value=sum(ang>Test.Stat)/nsim
  list(Test.Statistic=Test.Stat,P.Value=p.value)
}
```

References

[Ball and Greiner 2012] K. A. Ball and T. M. Greiner, "A procedure to refine joint kinematic assessments: functional alignment", *Comput. Methods Biomech. Biomed. Eng.* **15**:5 (2012), 487–500.

[Bingham et al. 2009] M. A. Bingham, D. J. Nordman, and S. B. Vardeman, "Modeling and inference for measured crystal orientations and a tractable class of symmetric distributions for rotations in three dimensions", *J. Amer. Statist. Assoc.* **104**:488 (2009), 1385–1397. MR 2011a:62189 Zbl 1205.62215

[Hendriks and Landsman 1998] H. Hendriks and Z. Landsman, "Mean location and sample mean location on manifolds: asymptotics, tests, confidence regions", *J. Multivariate Anal.* **67**:2 (1998), 227–243. MR 2000a:62125 Zbl 0941.62069

[Higgins 2004] J. J. Higgins, *Introduction to modern nonparametric statistics*, Brooks/Cole, Pacific Grove, CA, 2004.

[Khatri and Mardia 1977] C. G. Khatri and K. V. Mardia, "The von Mises–Fisher matrix distribution in orientation statistics", *J. Roy. Statist. Soc. (B) Stat. Methodol.* **39**:1 (1977), 95–106. MR 58 #13506 Zbl 0356.62044

[León et al. 2006] C. A. León, J.-C. Massé, and L.-P. Rivest, "A statistical model for random rotations", *J. Multivariate Anal.* **97**:2 (2006), 412–430. MR 2234030 Zbl 1085.62066

[Rancourt et al. 2000] D. Rancourt, L.-P. Rivest, and J. Asselin, "Using orientation statistics to investigate variations in human kinematics", *J. Roy. Statist. Soc. (C) Appl. Stat.* **49**:1 (2000), 81–94. MR 1817876 Zbl 0974.62107

dbero@iastate.edu   *Colony Brands, Monroe, WI 53566, United States*

mbingham@uwlax.edu   *Mathematics Department, University of Wisconsin-La Crosse, 1725 State Street, La Crosse, WI 54601, United States*

involve

msp.org/involve

involve