

Stat 425 HW7

Fritz Scholz

1. Write a function `van.der.Waerden(x, y) { ... }` that computes for given samples x (length m) and y (length n) the **exact p -value** of the 2-sample van der Waerden rank test when testing the hypothesis H_0 : no difference between treatment and control against the alternative that y -values are stochastically larger than x values.

Do this by calculating the transformed-rank vector Z_N of length $N = m + n$ according to van der Waerden's scheme (you do this just once for any given pair of samples x and y), and then (using `combn`) compute all possible sums of n transformed-ranks taken from Z_N . Get as exact p -value the proportion of such transformed-rank sums that are \geq to the observed transformed-rank sum for the given samples x and y .

Allow for ties to be present in such calculations. Apply the rank transformation to the midranks rather than averaging the tied transformed ranks. For example, when two observations are tied in rank positions 7 and 8, you take as transformed ranks $\Phi^{-1}(7.5/(N+1))$ in both positions and not $[\Phi^{-1}(7/(N+1)) + \Phi^{-1}(8/(N+1))]/2$, although there probably is not much difference between the two schemes. This aspect should only affect how Z_N is calculated, the rest of `van.der.Waerden` should stay the same.

Since any such transformed-rank sum can be viewed as the sum of a random sample of size n from the finite population of N transformed-ranks use the normal approximation for such sums that was provided in the context of sampling from a finite population. **Calculate the p -value using the normal approximation**, but do not employ a continuity correction.

Have `van.der.Waerden` also produce a **histogram** (using the option `probability=T` in the argument list of `hist`) with `breaks=seq(-10.1, 10.1, .2)` for the exact null distribution and **superimpose the approximating normal density**. On this histogram **mark the position of the observed value of the test statistic** using `abline(v=...)` and **annotate it with the exact and approximate p -values** using `text(...)`.

Let `van.der.Waerden` **output a named vector** containing the exact p -value and the one obtained by normal approximation.

Compute the result for the following two samples and comment on the two p -values and on the symmetry or lack of symmetry in the histogram:

```
x1=c(87.8, 100.1, 94.4, 99.7, 88.9, 101.5, 96.3)
y1=c(106.1, 96.4, 108.8, 119.3, 102.2, 100, 107.4, 103.6)
```

To test it for the situation with ties repeat this for the following two samples

```
x2=c(87.8, 100.1, 94.4, 99.7, 88.9, 101.8, 96.3)
y2=c(106.1, 96.4, 108.8, 119.3, 101.8, 100.1, 107.4, 103.6)
```

Give the results, plots and comments in each case, and provide the code of `van.der.Waerden`.

For extra credit: Amend the code of `van.der.Waerden` (in the spirit of problem 2 of HW 3) so that it can handle larger sample sizes m and n (where full enumeration is no longer feasible) via simulation, using `Nsim` simulated sums of size n taken from Z_N instead of the fully enumerated set of such sums. Here you would want to include `Nsim=100000` in your argument list and switch to simulation whenever the full enumeration leads to more than `Nsim` cases. When simulation is invoked your histogram will be of those `Nsim` simulated sums. As test case use

```
x3=c(102.5, 100.3, 96.8, 90, 96.4, 94.6, 100.9, 100.6, 94.6, 94.3)
y3=c(100.3, 104.5, 102.1, 97.7, 106.6, 105.3, 111.7, 97.4, 98.5, 110.2)
```

Provide the plot and commentary, and amended code instead of the code without this capability.