NAG Library Function Document

nag_mann_whitney (g08amc)

1   Purpose

nag_mann_whitney (g08amc) performs the Mann–Whitney U test on two independent samples of possibly unequal size and calculates the exact probability for the Mann–Whitney rank sum test statistic for the case where there are either ties or no ties in the samples pooled together.

2   Specification

```
#include <nag.h>
#include <nagg08.h>
```

```
void nag_mann_whitney (Integer n1, const double x[], Integer n2,
    const double y[], Nag_TailProbability tail, Nag_CompProb exact,
    double *u, double *z, double *p, NagError *fail)
```

3   Description

The Mann–Whitney U test investigates the difference between two populations defined by the distribution functions $F(x)$ and $G(y)$ respectively. The data consist of two independent samples of size n_1 and n_2, denoted by $x_1, x_2, \ldots, x_{n_1}$ and $y_1, y_2, \ldots, y_{n_2}$, taken from the two populations.

The hypothesis under test, H_0, often called the null hypothesis, is that the two distributions are the same, that is $F(x) = G(x)$, and this is to be tested against an alternative hypothesis H_1 which is

$H_1 : F(x) \neq G(y)$; or

$H_1 : F(x) < G(y)$, i.e., the x's tend to be greater than the y's; or

$H_1 : F(x) > G(y)$, i.e., the x's tend to be less than the y's,

using a two tailed, upper tailed or lower tailed probability respectively. You select the alternative hypothesis by choosing the appropriate tail probability to be computed (see the description of argument **tail** in Section 5).

Note that when using this test to test for differences in the distributions one is primarily detecting differences in the location of the two distributions. That is to say, if we reject the null hypothesis H_0 in favour of the alternative hypothesis H_1: $F(x) > G(y)$ we have evidence to suggest that the location, of the distribution defined by $F(x)$, is less than the location, of the distribution defined by $G(y)$.

The Mann–Whitney U test differs from the Median test (see nag_median_test (g08acc)) in that the ranking of the individual scores within the pooled sample is taken into account, rather than simply the position of a score relative to the median of the pooled sample. It is therefore a more powerful test if score differences are meaningful.

The test procedure involves ranking the pooled sample, average ranks being used for ties. Let r_{1i} be the rank assigned to x_i, for $i = 1, 2, \ldots, n_1$ and r_{2j} the rank assigned to y_j, for $j = 1, 2, \ldots, n_2$. Then the test statistic U is defined as follows;

$$U = \sum_{i=1}^{n_1} r_{1i} - \frac{n_1(n_1 + 1)}{2}$$

U is also the number of times a score in the second sample precedes a score in the first sample (where we only count a half if a score in the second sample actually equals a score in the first sample).

nag_mann_whitney (g08amc) returns:

(a) the test statistic U;

(b) the approximate Normal test statistic,

$$z = \frac{U - \text{mean}(U) \pm \frac{1}{2}}{\sqrt{\text{var}\,(U)}}$$

where

$$\text{mean}(U) = \frac{n_1 n_2}{2}$$

and

$$\text{var}\,(U) = \frac{n_1 n_2 (n_1 + n_2 + 1)}{12} - \frac{n_1 n_2}{(n_1 + n_2)(n_1 + n_2 - 1)} \times TS$$

where

$$TS = \sum_{j=1}^{\tau} \frac{(t_j)(t_j - 1)(t_j + 1)}{12}$$

τ is the number of groups of ties in the sample and t_j is the number of ties in the jth group.

Note that if no ties are present the variance of U reduces to $\frac{n_1 n_2}{12}(n_1 + n_2 + 1)$.

(c) An indicator as to whether ties were present in the pooled sample or not.

(d) The tail probability, p, corresponding to U, depending on the choice of **tail**, i.e., the choice of alternative hypothesis, H_1. An exact probability or a normal approximation may be selected using the **exact** option. For large values of n_1 and n_2 the normal approximation should be adequate while the exact computation may require extensive calculation, particularly in cases where there are ties in the observations. For small samples the exact probability should be used.

The value of p can be used to perform a significance test on the null hypothesis H_0 against the alternative hypothesis H_1. Let α be the size of the significance test (that is, α is the probability of rejecting H_0 when H_0 is true). If $p < \alpha$ then the null hypothesis is rejected. Typically α might be 0.05 or 0.01.

4   References

Conover W J (1980) *Practical Nonparametric Statistics* Wiley

Neumann N (1988) Some procedures for calculating the distributions of elementary nonparametric teststatistics *Statistical Software Newsletter* **14(3)** 120–126

Siegel S (1956) *Non-parametric Statistics for the Behavioral Sciences* McGraw–Hill

5   Arguments

1:    **n1** – Integer   *Input*

   On entry: the number of non-tied pairs, n_1.

   Constraint: **n1** ≥ 1.

2:    **x**[**n1**] – const double   *Input*

   On entry: the first vector of observations. $x_1, x_2, \ldots, x_{n_1}$.

3:    **n2** – Integer   *Input*

   On entry: the size of the second sample, n_2.

   Constraint: **n2** ≥ 1.

4:     **y**[**n2**] – const double   *Input*

On entry: the second vector of observations. $y_1, y_2, \ldots, y_{n_2}$.

5:     **tail** – Nag_TailProbability   *Input*

On entry: indicates the choice of tail probability, and hence the alternative hypothesis.

tail = Nag_TwoTail
   A two tailed probability is calculated and the alternative hypothesis is $H_1 : F(x) \neq G(y)$.

tail = Nag_UpperTail
   An upper tailed probability is calculated and the alternative hypothesis $H_1 : F(x) < G(y)$, i.e., the x's tend to be greater than the y's.

tail = Nag_LowerTail
   A lower tailed probability is calculated and the alternative hypothesis $H_1 : F(x) > G(y)$, i. e., the x's tend to be less than the y's.

Constraint: **tail** = Nag_TwoTail, Nag_UpperTail or Nag_LowerTail.

6:     **exact** – Nag_CompProb   *Input*

On entry: indicates if exact probability, p, is to be computed.

exact = Nag_CompProbExact
   Compute the exact probability.

exact = Nag_CompProbApprox
   Compute the approximate probability.

Constraint: **exact** = Nag_CompProbExact or Nag_CompProbApprox.

7:     **u** – double *   *Output*

On exit: the Mann–Whitney rank sum statistic, U.

8:     **z** – double *   *Output*

On exit: the approximate Normal test statistic, z, as described in Section 3.

9:     **p** – double *   *Output*

On exit: the exact tail probability, p, as specified by the argument **tail**.

10:    **fail** – NagError *   *Input/Output*

The NAG error argument (see Section 2.7 in How to Use the NAG Library and its Documentation).

6   Error Indicators and Warnings

NE_ALLOC_FAIL

Dynamic memory allocation failed.

NE_BAD_PARAM

On entry, argument **exact** had an illegal value.

On entry, argument **tail** had an illegal value.

NE_G08AH_SAMP_IDEN

The pooled samples are all the same, that is the variance of $U = 0.0$.

NE_INT_ARG_LT

On entry, **n1** must not be less than 1: **n1** = $\langle value \rangle$.

On entry, **n2** must not be less than 1: **n2** = $\langle value \rangle$.

NE_INTERNAL_ERROR

An internal error has occurred in this function. Check the function call and any array sizes. If the call is correct then please contact NAG for assistance.

NE_MANN_WHITNEY_STAT

The computed Mann−Whitney statistic, $U = \langle value \rangle$. If the exact tail probability is to be computed, then $U \geq 0$.

7 Accuracy

The approximate tail probability, p, returned by nag_mann_whitney (g08amc) is a good approximation to the exact probability for cases where $\max(n_1, n_2) \geq 30$ and $(n_1 + n_2) \geq 40$. The relative error of the approximation should be less than 10 percent, for most cases falling in this range.

8 Parallelism and Performance

nag_mann_whitney (g08amc) is not threaded in any implementation.

9 Further Comments

The time taken by nag_mann_whitney (g08amc) increases with n_1 and n_2.

10 Example

The example program performs the Mann−Whitney test on two independent samples of sizes 16 and 23 respectively. This is used to test the null hypothesis that the distributions of the two populations from which the samples were taken are the same against the alternative hypothesis that the distributions are different. The test statistic, the approximate Normal statistic and the approximate two-tail probability are printed. An exact tail probability is also calculated and printed depending on whether ties were found in the pooled sample or not.

10.1 Program Text

```
/* nag_mann_whitney (g08amc) Example Program.
 *
 * NAGPRODCODE Version.
 *
 * Copyright 2016 Numerical Algorithms Group.
 *
 * Mark 26, 2016.
 */

#include <stdio.h>
#include <nag.h>
#include <nag_stdlib.h>
#include <nagg08.h>

int main(void)
{

  Integer exit_status = 0, i, n1, n2;
  NagError fail;
  double p, u, *x = 0, *y = 0, z;

  INIT_FAIL(fail);
```

```
  printf("nag_mann_whitney (g08amc) Example Program Results\n\n");

  /* Skip heading in data file */
#ifdef _WIN32
  scanf_s("%*[^\n]");
#else
  scanf("%*[^\n]");
#endif

#ifdef _WIN32
  scanf_s("%" NAG_IFMT " %" NAG_IFMT " ", &n1, &n2);
#else
  scanf("%" NAG_IFMT " %" NAG_IFMT " ", &n1, &n2);
#endif
  printf("%s%5" NAG_IFMT "\n", "Sample size of group 1 = ", n1);
  printf("%s%5" NAG_IFMT "\n", "Sample size of group 2 = ", n2);
  if (!(x = NAG_ALLOC(n1, double))
      || !(y = NAG_ALLOC(n2, double)))
  {
    printf("Allocation failure\n");
    exit_status = -1;
    goto END;
  }
  printf("\n");
  for (i = 1; i <= n1; ++i)
#ifdef _WIN32
    scanf_s("%lf", &x[i - 1]);
#else
    scanf("%lf", &x[i - 1]);
#endif
  printf("%s\n", "Mann-Whitney U test");
  printf("\n");
  printf("%s\n", "Data values");
  printf("\n");
  printf("%s", "    Group 1  ");
  for (i = 1; i <= n1; ++i)
    printf("%5.1f%s", x[i - 1], i % 8 ? "" : "\n             ");
  for (i = 1; i <= n2; ++i)
#ifdef _WIN32
    scanf_s("%lf", &y[i - 1]);
#else
    scanf("%lf", &y[i - 1]);
#endif
  printf("\n");
  printf("%s", "    Group 2  ");
  for (i = 1; i <= n2; ++i)
    printf("%5.1f%s", y[i - 1], i % 8 ? "" : "\n             ");

  /* nag_mann_whitney (g08amc).
   * Performs the Mann-Whitney U test on two independent
   * samples
   */
  nag_mann_whitney(n1, x, n2, y, Nag_LowerTail, Nag_CompProbApprox,
                   &u, &z, &p, &fail);
  if (fail.code != NE_NOERROR) {
    printf("Error from nag_mann_whitney (g08amc).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
  }
  printf("\n\n");
  printf("%s%8.4f\n", "Test statistic                = ", u);
  printf("%s%8.4f\n", "Normal Statistic              = ", z);
  printf("%s%8.4f\n", "Approximate tail probability   = ", p);
  /* nag_mann_whitney (g08amc), see above. */
  nag_mann_whitney(n1, x, n2, y, Nag_LowerTail, Nag_CompProbExact,
                   &u, &z, &p, &fail);
  if (fail.code != NE_NOERROR) {
    printf("Error from nag_mann_whitney (g08amc).\n%s\n", fail.message);
    exit_status = 1;
    goto END;
  }
```

```
    printf("%s%8.4f\n", "Exact tail probability         = ", p);
END:
  NAG_FREE(x);
  NAG_FREE(y);
  return exit_status;
}
```

10.2 Program Data

```
nag_mann_whitney (g08amc) Example Program Data
 16 23
 13.0   6.0 12.0   7.0 12.0   7.0 10.0   7.0
 10.0   7.0 16.0   7.0 10.0   8.0  9.0   8.0
 17.0   6.0 10.0   8.0 15.0   8.0 15.0 10.0 15.0 10.0 14.0 10.0
 14.0 11.0 14.0 11.0 13.0 12.0 13.0 12.0 13.0 12.0 12.0
```

10.3 Program Results

```
nag_mann_whitney (g08amc) Example Program Results

Sample size of group 1 =    16
Sample size of group 2 =    23

Mann-Whitney U test

Data values

    Group 1   13.0   6.0 12.0   7.0 12.0   7.0 10.0   7.0
              10.0   7.0 16.0   7.0 10.0   8.0  9.0   8.0

    Group 2   17.0   6.0 10.0   8.0 15.0   8.0 15.0 10.0
              15.0 10.0 14.0 10.0 14.0 11.0 14.0 11.0
              13.0 12.0 13.0 12.0 13.0 12.0 12.0

Test statistic           =   86.0000
Normal Statistic         =   -2.8039
Approximate tail probability  =    0.0025
Exact tail probability        =    0.0020
```