

Statistical Analysis of Performance Data – Euler Group

Thursday 16 January, 2020

Methods and Software

Statistical analysis was based on the Generalized Estimating Equations method (Liang and Zeger 1986) as this is implemented in the R (R Core Team 2019) package `geepack` (Højsgaard, Halekoh, and Yan 2005). In addition, the function `ComparisonStats` was developed to evaluate the statistical significance of the desired comparisons for the accuracy and time data.

```
> ComparisonStats <- function(FittedModel, Lmatrix, alpha = 0.05) {  
+   Lmatrix <- matrix(Lmatrix, nrow = 1)  
+   ModelBetas <- FittedModel$geese$beta  
+   ModelVCov <- FittedModel$geese$vbeta  
+   Estimate <- drop(Lmatrix %*% ModelBetas)  
+   SdError <- sqrt(drop(Lmatrix %*% ModelVCov %*% t(Lmatrix)))  
+   CBs <- Estimate + qnorm(c(alpha/2, 1 - alpha/2)) * SdError  
+   pvalue <- 2 * pnorm(-abs(Estimate/SdError))  
+   ans <- c(exp(c(Estimate, CBs)), round(pvalue, 4))  
+   names(ans) <- c("Estimate", paste0((1 - alpha) * 100, "% LB"), paste0((1 - alpha) *  
+     100, "% UB"), "p-value")  
+   ans  
+ }
```

Import Data

The performance data for the Euler group were obtained as follows:

```
> library("readxl")  
> performance_data <- read_excel("data/mainAllPerformance_GS.xlsx")  
> names(performance_data) <- gsub("[.]", "_", make.names(names(performance_data), unique = TRUE))  
> names(performance_data)[4] <- "Treatment"  
> euler_performance_data <- performance_data[performance_data$Group == "Euler", ]
```

Analysis of Accuracy Data

First, the following regression model was fitted to the accuracy data

$$\log \left[\frac{\Pr(Y_{ij} = 1)}{1 - \Pr(Y_{ij} = 1)} \right] = \beta_0 + \beta_1 x_{ij1} + \beta_2 x_{ij2} + \beta_3 x_{ij3} + \beta_4 x_{ij4} + \beta_5 x_{ij1} x_{ij4} + \beta_6 x_{ij2} x_{ij4} + \beta_7 x_{ij3} x_{ij4}$$

where

- $\Pr(Y_{ij} = 1)$ is the probability for participant i to answer question j correctly,
- x_{ij1} is the indicator for the *NWMA* treatment,
- x_{ij2} is the indicator for the *WM* treatment,
- x_{ij3} is the indicator for the *WMA* treatment,
- x_{ij4} is the indicator for the *subset* task type,

for $i = 1, \dots, 101$, corresponding to the individual participants and $j = 1, \dots, 8$ corresponding to questions. This model assumes that the success probability depends on the treatment, on the task type and on their interaction.

```
> library(geepack)
> fitted_model <- geeglm(formula = Correct ~ Treatment * Question_Type, id = Study_Id,
+   data = euler_performance_data, family = binomial(link = "logit"))
```

By comparing this model with one with no interaction terms

```
> reduced_model <- geeglm(formula = Correct ~ Treatment + Question_Type, id = Study_Id,
+   data = euler_performance_data, family = binomial(link = "logit"))
> Waldts <- anova(fitted_model, reduced_model)
> Waldts
Analysis of 'Wald statistic' Table

Model 1 Correct ~ Treatment * Question_Type
Model 2 Correct ~ Treatment + Question_Type
  Df      X2 P(>|Chi|)
1  3 4.9695   0.174
> fitted_model <- reduced_model
```

we can conclude that at least one of the interplay of *treatment* and *task type* is not significant (p - value = 0.174). This means that the effect of each treatment is constant across the two tasks and the effect of each task is constant across the four treatments.

The interaction terms can be removed and the regression model fitted to the accuracy data can be simplified to the reduced mode

$$\log \left[\frac{\Pr(Y_{ij} = 1)}{1 - \Pr(Y_{ij} = 1)} \right] = \beta_0 + \beta_1 x_{ij1} + \beta_2 x_{ij2} + \beta_3 x_{ij3} + \beta_4 x_{ij4}.$$

Comparison of treatments

Ranking: WM > WMA > NWM > NWMA

```
> ## NWMA vs. NWM
> ComparisonStats(fitted_model, c(0, 1, 0, 0, 0))
Estimate   95% LB   95% UB  p-value
0.7006712 0.5096690 0.9632529 0.0285000
> ## WM vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 1, 0, 0))
Estimate   95% LB   95% UB  p-value
8.163696  4.972179 13.403769 0.000000
> ## WMA vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 0, 1, 0))
Estimate   95% LB   95% UB  p-value
```

```

1.547445 1.000263 2.393956 0.049900
> ## NWMA vs. WM
> ComparisonStats(fitted_model, c(0, 1, -1, 0, 0))
  Estimate    95% LB    95% UB  p-value
0.08582769 0.05272851 0.13970415 0.00000000
> ## NWMA vs. WMA
> ComparisonStats(fitted_model, c(0, 1, 0, -1, 0))
  Estimate    95% LB    95% UB  p-value
0.4527923 0.3053805 0.6713620 0.0001000
> ## WM vs. WMA
> ComparisonStats(fitted_model, c(0, 0, 1, -1, 0))
  Estimate    95% LB    95% UB  p-value
5.275597 3.113593 8.938846 0.000000

```

Comparison of task type

Ranking: subset < disjoint

```

> ## subset vs. disjoint
> ComparisonStats(fitted_model, c(0, 0, 0, 0, 1))
  Estimate    95% LB    95% UB  p-value
0.1307947 0.0881310 0.1941117 0.0000000

```

Time Data

First, the following regression model was fitted to the time data

$$\log(Z_{ij}) = \gamma_0 + \gamma_1 x_{ij1} + \gamma_2 x_{ij2} + \gamma_3 x_{ij3} + \gamma_4 x_{ij4} + \gamma_5 x_{ij1} x_{ij4} + \gamma_6 x_{ij2} x_{ij4} + \gamma_7 x_{ij3} x_{ij4}$$

where Z_{ij} is the time required for participant i to answer question j correctly for $i = 1, \dots, 100$, corresponding to the individual participants and $j = 1, \dots, 8$ corresponding to questions. This model assumes that the time to answer a question correctly depends on the treatment, on the question type and on their interaction.

```

> euler_time <- euler_performance_data[euler_performance_data$Correct == 1, ]
> fitted_model <- geeglm(formula = log(Time) ~ Treatment * Question_Type, id = Study_Id,
+   data = euler_time, family = gaussian(link = "identity"))

```

Testing for the interaction term,

```

> anova(fitted_model)
Analysis of 'Wald statistic' Table
Model: gaussian, link: identity
Response: log(Time)
Terms added sequentially (first to last)

              Df      X2 P(>|Chi|)
Treatment      3 118.330 < 2.2e-16 ***
Question_Type  1  18.675 1.55e-05 ***
Treatment:Question_Type 3  11.181 0.01079 *
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

```

implies that the interplay of *treatment* and *task type* is significant (p -value = 0.01079).

Comparison of treatments

Subset

Ranking: WM > NWMA \approx NWM \approx WMA

```
> ## NWMA vs. NWM
> ComparisonStats(fitted_model, c(0, 1, 0, 0, 0, 1, 0, 0))
Estimate   95% LB   95% UB  p-value
0.9448820 0.7684551 1.1618143 0.5908000
> ## WM vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 1, 0, 0, 0, 1, 0))
Estimate   95% LB   95% UB  p-value
0.5983785 0.5050249 0.7089884 0.0000000
> ## WMA vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 0, 1, 0, 0, 0, 1))
Estimate   95% LB   95% UB  p-value
0.9657124 0.7798228 1.1959132 0.7491000
> ## NWMA vs. WM
> ComparisonStats(fitted_model, c(0, 1, -1, 0, 0, 1, -1, 0))
Estimate   95% LB   95% UB  p-value
1.579071 1.249260 1.995953 0.000100
> ## NWMA vs. WMA
> ComparisonStats(fitted_model, c(0, 1, 0, -1, 0, 1, 0, -1))
Estimate   95% LB   95% UB  p-value
0.9784301 0.7510786 1.2746009 0.8716000
> ## WM vs. WMA
> ComparisonStats(fitted_model, c(0, 0, 1, -1, 0, 0, 1, -1))
Estimate   95% LB   95% UB  p-value
0.6196240 0.5073292 0.7567745 0.0000000
```

Disjoint

Ranking: WM \approx WMA > NWM \approx NWMA

```
> ## NWMA vs. NWM
> ComparisonStats(fitted_model, c(0, 1, 0, 0, 0, 0, 0, 0))
Estimate   95% LB   95% UB  p-value
1.0260187 0.9188046 1.1457434 0.6483000
> ## WM vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 1, 0, 0, 0, 0, 0))
Estimate   95% LB   95% UB  p-value
0.6479520 0.5869747 0.7152639 0.0000000
> ## WMA vs. NWM
> ComparisonStats(fitted_model, c(0, 0, 0, 1, 0, 0, 0, 0))
Estimate   95% LB   95% UB  p-value
0.7060551 0.6203149 0.8036464 0.0000000
> ## NWMA vs. WM
> ComparisonStats(fitted_model, c(0, 1, -1, 0, 0, 0, 0, 0))
Estimate   95% LB   95% UB  p-value
1.583480 1.424612 1.760063 0.000000
> ## NWMA vs. WMA
> ComparisonStats(fitted_model, c(0, 1, 0, -1, 0, 0, 0, 0))
```

```

Estimate  95% LB  95% UB  p-value
1.453171 1.277985 1.652372 0.000000
> ## WM vs. WMA
> ComparisonStats(fitted_model, c(0, 0, 1, -1, 0, 0, 0, 0))
Estimate  95% LB  95% UB  p-value
0.9177074 0.8214411 1.0252553 0.1288000

```

Comparison of task type

NWMA and WM: subset \approx disjoint

NWM and WMA: subset $<$ disjoint

```

> ## (subset & NWMA) vs. (disjoint & NWMA)
> ComparisonStats(fitted_model, c(0, 0, 0, 0, 1, 1, 0, 0))
Estimate  95% LB  95% UB  p-value
1.103037 0.875334 1.389973 0.405800
> ## (subset & WM) vs. (disjoint & WM)
> ComparisonStats(fitted_model, c(0, 0, 0, 0, 1, 0, 1, 0))
Estimate  95% LB  95% UB  p-value
1.1061166 0.9829637 1.2446991 0.0940000
> ## (subset & WMA) vs. (disjoint & WMA)
> ComparisonStats(fitted_model, c(0, 0, 0, 0, 1, 0, 0, 1))
Estimate  95% LB  95% UB  p-value
1.638238 1.338628 2.004907 0.000000
> ## (subset & NWM) vs. (disjoint & NWM)
> ComparisonStats(fitted_model, c(0, 0, 0, 0, 1, 0, 0, 0))
Estimate  95% LB  95% UB  p-value
1.197754 1.007704 1.423648 0.040700

```

References

- Højsgaard, S., U. Halekoh, and J. Yan. 2005. “The R Package geepack for Generalized Estimating Equations.” *Journal of Statistical Software* 15: 1–11.
- Liang, K. Y., and S. L. Zeger. 1986. “Longitudinal Data Analysis Using Generalized Linear Models.” *Biometrika* 73 (1): 13–22.
- R Core Team. 2019. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. <http://www.R-project.org/>.