Knee Injuries - Proportional Odds Models

December 4, 2025

At the beginning the knee dataset is loaded:

```
rm(list=ls(all=TRUE))
library(catdata)
data(knee)
attach(knee)
```

First of all a simple χ^2 – test of independence between the therapy (Th) and the pain level (R4) is applied.

```
suppressWarnings(chisq.test(knee$Th,knee$R4))
```

In the following the variable Age is centered around 30 years and the quadratic variable Age2 is created.

```
Age <- Age - 30
Age2 <- Age^2
```

The response pain (R4) has to be an ordered factor, the covariates therapy (Th) and sex (Sex) need to be factors.

```
R4 <- as.ordered(R4)
Th <- as.factor(Th)
Sex <- as.factor(Sex)
```

A proportional odds model can be fitted by the function "polr" from the "MASS" – library. Attention has to be paid to the algebraic signs of the co-efficients. These are inverse to the usual interpretation in porportional odds models.

```
library(MASS)
```

The first model only uses therapy as covariate, to achieve a proportional odds model the option "method" needs to use the logistic link function.

```
polr1 <- polr(R4 ~ Th, method="logistic")
```

```r
summary(polr1)
```

The corresponding odds–ratio can be recieved by the following command (consider the inverse sign!):

```r
exp(-coef(polr1))
```

Now a model with the covariates therapy, sex and age is fitted.

```r
polr2 <- polr(R4 ~ Th + Sex + Age, method="logistic")
```

```r
summary(polr2)
```

Odds–ratios for the second model:

```r
exp(-coef(polr2))
```

To get the Wald–statistic, the standard errors have to be extracted from the summary. Afterwards the Wald–statistic and the corresponding p–values are easily recieved.

```r
se <- summary(polr2)[1][[1]][1:3,2]
(wald2 <- -coef(polr2)/se)
```

P–values for the second model:

```r
1-pchisq(wald2^2, df=1)
```

Finally the quadratic age–effect is added to the previous model.

```r
polr3 <- update(polr2, ~. + Age2)
```

```r
summary(polr3)
```

Odds–ratios for the final model:

```r
exp(-coef(polr3))
```

Wald–statistic for the final model:

```r
se <- summary(polr3)[1][[1]][1:4,2]
(wald3 <- -coef(polr3)/se)
```

P–values for the final model:

```
1-pchisq(wald3^2, df=1)
```

As the proportional odds–model is the most popular model for ordinal data, there are several different ways to fit such models. Now the final model is additionally fitted with function "vglm" from the "VGAM"–library and with function "lrm" from the "rms"–library.

Model fitted with "vglm":

```
library(VGAM)
```

```
m.vglm <- vglm(R4 ~ Th + Sex + Age + Age2, family = cumulative (link="logitlink",
parallel=TRUE))
summary(m.vglm)
```

The resulting coefficients are very similar to the coefficients in the model above fitted with function "polr", but they have inverse signs. Therefore they can be interpreted in the usual way.

Model fitted with "lrm":

```
library(rms)
```

```
m.lrm <- lrm(R4 ~ Th + Sex + Age + Age2)
m.lrm
```

Again no big differences are found concerning the coefficients. Here the signs are the same as with the function "polr".