Introduction to Geostatistics

11. Multiple regression, regression extensions

Edzer J. Pebesma

edzer.pebesma@uni-muenster.de
Institute for Geoinformatics (**ifgi**)
University of Münster

summer semester 2007/8,
June 24, 2008

The multiple linear regression model

The multiple linear regression model extends the simple regression model with one single predictor

$$y_i = \beta_0 + \beta_1 X_{i,1} + e_i$$

to two predictors

$$y_i = \beta_0 + \beta_1 X_{i,1} + \beta_2 X_{i,2} + e_i$$

or p predictors:

$$y_i = \beta_0 + \beta_1 X_{i,1} + ... + \beta_p X_{i,p} + e_i$$

Example: two groups

(Ignoring the outlier)

```
> plot(Weight ~ Gender)
```

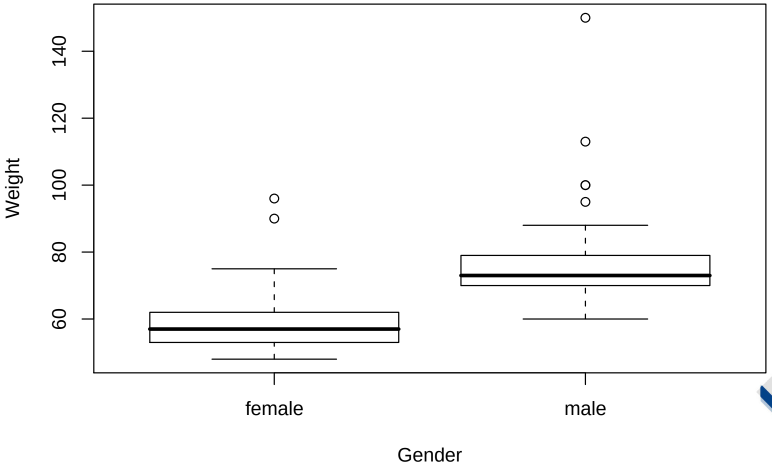

Example: ... seen through a linear regression glasses

```
> plot(Weight ~ I(as.numeric(Gender) - 1))
> abline(lm(Weight ~ Gender))
```

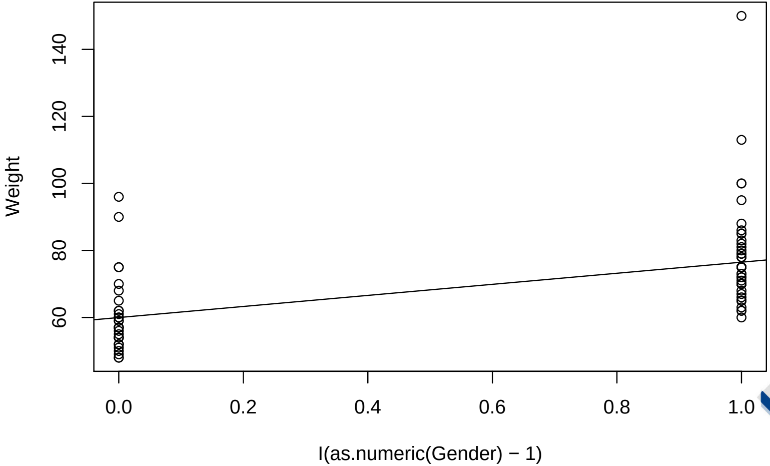

Example: simple

```
> summary(lm(Weight ~ Gender))

Call:
lm(formula = Weight ~ Gender)

Residuals:
    Min      1Q  Median      3Q     Max
-16.491  -6.491  -2.969   2.509  73.509

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept)   59.969      2.366  25.348  < 2e-16 ***
Gendermale    16.522      2.996   5.514 3.87e-07 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 13.38 on 83 degrees of freedom
Multiple R-squared: 0.2681,       Adjusted R-squared: 0.2593
F-statistic: 30.41 on 1 and 83 DF,  p-value: 3.871e-07
```

Interpretation

So, weight depends on Gender.
But, there's also a length effect. Longer people are usually heavier, and men are usually taller than women.
Questions we could ask:

1. is there, besides a Length effect still an effect of Gender on Weight? (testing)
2. how large is the effect of Length on Weight? (estimation)
3. Does this effect depend on Gender? (testing)

Interpretation

So, weight depends on Gender.
But, there's also a length effect. Longer people are usually heavier, and men are usually taller than women.
Questions we could ask:

1. is there, besides a Length effect still an effect of Gender on Weight? (testing)
2. how large is the effect of Length on Weight? (estimation)
3. Does this effect depend on Gender? (testing)

Interpretation

So, weight depends on Gender.
But, there's also a length effect. Longer people are usually heavier, and men are usually taller than women.
Questions we could ask:

1. is there, besides a Length effect still an effect of Gender on Weight? (testing)
2. how large is the effect of Length on Weight? (estimation)
3. Does this effect depend on Gender? (testing)

Example: simple

```
> summary(lm(Weight ~ Length))

Call:
lm(formula = Weight ~ Length)

Residuals:
    Min      1Q  Median      3Q     Max
-13.776  -7.238  -2.776   3.993  75.993

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -86.9920    22.4307  -3.878  0.00021 ***
Length        0.8846     0.1259   7.024 5.49e-10 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 12.39 on 83 degrees of freedom
Multiple R-squared: 0.3728,       Adjusted R-squared: 0.3652
F-statistic: 49.33 on 1 and 83 DF,  p-value: 5.493e-10
```

ifgi

Example: simple

```
> plot(Weight ~ Length)
> abline(lm(Weight ~ Length))
```

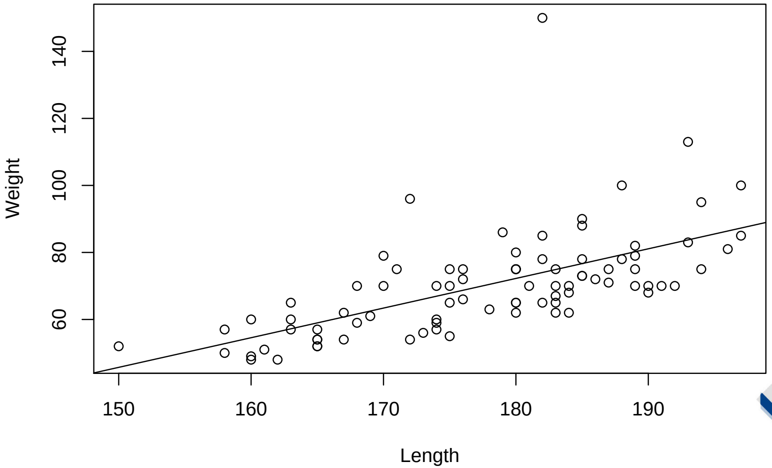

Example: simple

```
> plot(Weight ~ Length, col = Gender)
> abline(lm(Weight ~ Length))
```

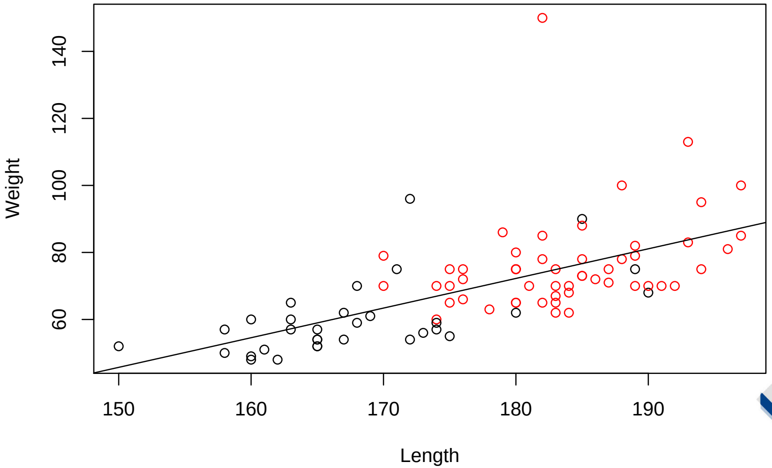

Example: the two parallel lines

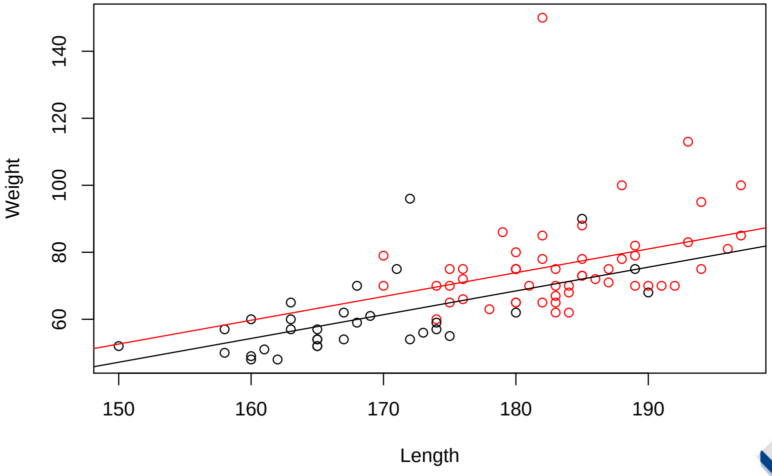

Example: the two parallel lines added

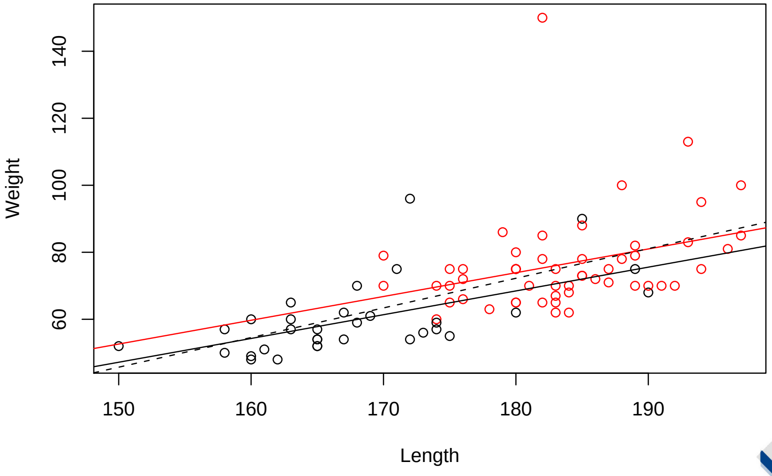

Example: corresponding model

```
> summary(lm(Weight ~ Length + Gender))

Call:
lm(formula = Weight ~ Length + Gender)

Residuals:
    Min      1Q  Median      3Q     Max
-14.732  -7.204  -2.851   3.202  74.688

Coefficients:
            Estimate Std. Error t value Pr(>|t|)
(Intercept) -59.2577    29.9528  -1.978 0.051245 .
Length        0.7095     0.1778   3.991 0.000142 ***
Gendermale    5.4322     3.9152   1.387 0.169064
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 12.32 on 82 degrees of freedom
Multiple R-squared: 0.3872,       Adjusted R-squared: 0.3722
F-statistic:  25.9 on 2 and 82 DF,  p-value: 1.909e-09
```

ifgi

3 Questions

1. is there, besides a Length effect still an effect of Gender on Weight? No, it is not significant; it can be there, but based on our data we cannot say whether it is positive or negative
2. how large is the effect of Length on Weight? Is it 0.88 or 0.71? Despite the fact that gender is not significant, assuming H_0 that the effect is zero is not very realistic. We may therefor give a preference to the 0.71 estimate .
3. Does this effect depend on Gender? See next slide.

3 Questions

1. is there, besides a Length effect still an effect of Gender on Weight? No, it is not significant; it can be there, but based on our data we cannot say whether it is positive or negative
2. how large is the effect of Length on Weight? Is it 0.88 or 0.71? Despite the fact that gender is not significant, assuming H_0 that the effect is zero is not very realistic. We may therefor give a preference to the 0.71 estimate .
3. Does this effect depend on Gender? See next slide.

3 Questions

1. is there, besides a Length effect still an effect of Gender on Weight? No, it is not significant; it can be there, but based on our data we cannot say whether it is positive or negative
2. how large is the effect of Length on Weight? Is it 0.88 or 0.71? Despite the fact that gender is not significant, assuming H_0 that the effect is zero is not very realistic. We may therefor give a preference to the 0.71 estimate .
3. Does this effect depend on Gender? See next slide.

Does the effect depend on Gender?

- Both models (simple linear, and multiple linear) give a *single* dependence (slope of the line) for Weight on Length.
- The question whether this effect (the slope) depends on Gender, is the following: does the slope (Weight ~ Length) differ for male persons from that of female persons?

Let $X_{i,1}$ be Length, and let $X_{i,2}$ be zero for female, and one for male persons. Then

$$y_i = \beta_0 + \beta_1 X_{i,1} + \beta_2 X_{i,2} + \beta_3 X_{i,1} X_{i,2} + e_i$$

is a single regression model that reduces for female persons to

$$y_i = \beta_0 + \beta_1 X_{i,1} + e_i$$

and for male persons to

$$y_i = (\beta_0 + \beta_2) + (\beta_1 + \beta_3) X_{i,1} + e_i$$

so, we have two completely free regression lines, each with a unique slope and intercept.

Does the effect depend on Gender?

- Both models (simple linear, and multiple linear) give a *single* dependence (slope of the line) for Weight on Length.
- The question whether this effect (the slope) depends on Gender, is the following: does the slope (Weight ~ Length) differ for male persons from that of female persons?

Let $X_{i,1}$ be Length, and let $X_{i,2}$ be zero for female, and one for male persons. Then

$$y_i = \beta_0 + \beta_1 X_{i,1} + \beta_2 X_{i,2} + \beta_3 X_{i,1} X_{i,2} + e_i$$

is a single regression model that reduces for female persons to

$$y_i = \beta_0 + \beta_1 X_{i,1} + e_i$$

and for male persons to

$$y_i = (\beta_0 + \beta_2) + (\beta_1 + \beta_3) X_{i,1} + e_i$$

so, we have two completely free regression lines, each with a unique slope and intercept.

ifgi

The R model

```
> summary(lm(Weight ~ Length * Gender))

Call:
lm(formula = Weight ~ Length * Gender)

Residuals:
    Min      1Q  Median      3Q     Max
-14.726  -7.301  -2.688   2.967  74.659

Coefficients:
                   Estimate Std. Error t value Pr(>|t|)
(Intercept)       -61.99199   41.97702  -1.477  0.14360
Length              0.72582    0.24948   2.909  0.00467 **
Gendermale         11.32815   63.13357   0.179  0.85805
Length:Gendermale  -0.03349    0.35788  -0.094  0.92568
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 12.4 on 81 degrees of freedom
Multiple R-squared: 0.3872,       Adjusted R-squared: 0.3645
F-statistic: 17.06 on 3 and 81 DF,  p-value: 1.115e-08
```

Multiple linear regression with two variables.

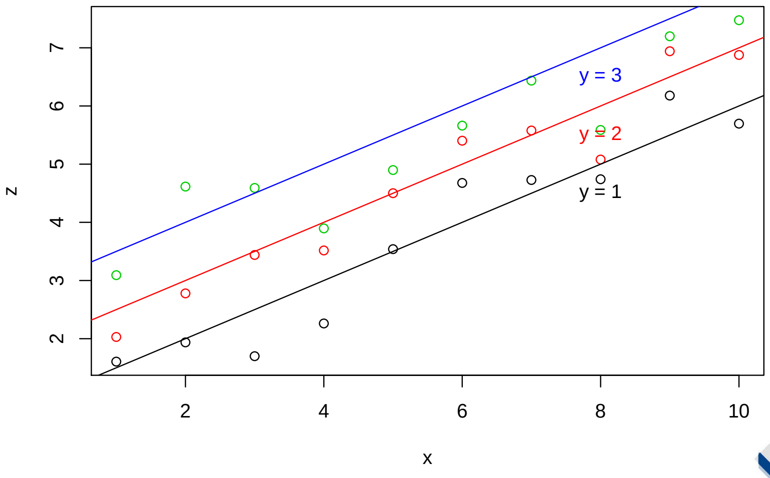

```
> summary(lm(z ~ x))

Call:
lm(formula = z ~ x)

Residuals:
     Min       1Q   Median       3Q      Max 
-1.59744 -0.65875  0.02085  0.61690  1.82123 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)  1.78683    0.33593   5.319 1.16e-05 ***
x            0.50339    0.05414   9.298 4.68e-10 ***
---
Signif. codes:  0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 ‘.’ 0.1 ‘ ’ 1

Residual standard error: 0.8517 on 28 degrees of freedom
Multiple R-squared: 0.7553,	Adjusted R-squared: 0.7466 
F-statistic: 86.45 on 1 and 28 DF,  p-value: 4.682e-10
```

```
> summary(lm(z ~ x + y))

Call:
lm(formula = z ~ x + y)

Residuals:
     Min       1Q   Median       3Q      Max 
-1.04537 -0.25856  0.04558  0.25999  1.00188 

Coefficients:
            Estimate Std. Error t value Pr(>|t|)    
(Intercept)   0.1481     0.3011   0.492    0.627    
x             0.5034     0.0321  15.681 4.37e-15 ***
y             0.8194     0.1129   7.256 8.37e-08 ***
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1 

Residual standard error: 0.505 on 27 degrees of freedom
Multiple R-squared: 0.9171,	Adjusted R-squared: 0.9109 
F-statistic: 149.3 on 2 and 27 DF,  p-value: 2.531e-15
```

Why using multiple regression?

1. There is a difference in interpretation for slopes
   - when (some of) the predictors X are correlated, the slopes differ from eachother.
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + e$ is simply *the expected change in y as a function of X_1, ignoring everything else*
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + e$ is simply *the expected change in y as a function of X_1, everything else (meaning: X_2) held constant.*
   - in the first model, the slope may be partly due to X_2.
2. Their power is often larger (smaller residual standard error).

Why using multiple regression?

1. There is a difference in interpretation for slopes
   - when (some of) the predictors X are correlated, the slopes differ from eachother.
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + e$ is simply *the expected change in y as a function of X_1, ignoring everything else*
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + e$ is simply *the expected change in y as a function of X_1, everything else (meaning: X_2) held constant.*
   - in the first model, the slope may be partly due to X_2.
2. Their power is often larger (smaller residual standard error).

Why using multiple regression?

1. There is a difference in interpretation for slopes
   - when (some of) the predictors X are correlated, the slopes differ from eachother.
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + e$ is simply *the expected change in* y *as a function of* X_1*, ignoring everything else*
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + e$ is simply *the expected change in* y *as a function of* X_1*, everything else (meaning:* X_2*) held constant.*
   - in the first model, the slope may be partly due to X_2.
2. Their power is often larger (smaller residual standard error).

Why using multiple regression?

1. There is a difference in interpretation for slopes
   - when (some of) the predictors X are correlated, the slopes differ from eachother.
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + e$ is simply *the expected change in y as a function of X_1, ignoring everything else*
   - the slope for the model $y = \beta_0 + \beta_1 X_1 + \beta_2 X_2 + e$ is simply *the expected change in y as a function of X_1, everything else (meaning: X_2) held constant.*
   - in the first model, the slope may be partly due to X_2.
2. Their power is often larger (smaller residual standard error).

Correlated errors

When observations are correlated, and cannot be considered independent (e.g. by the random sampling argument), regression can be applied under a more general model that addresses these correlations.

- the structure of the correlation needs to be assessed
  - correlation in space: a function of spatial distance?
  - correlation over time: a function of time separation?
  - within-item correlation: e.g. longitudinal studies.
- the magnitude of the correlations needs to be assessed

Correlated errors

When observations are correlated, and cannot be considered independent (e.g. by the random sampling argument), regression can be applied under a more general model that addresses these correlations.

- the structure of the correlation needs to be assessed
  - correlation in space: a function of spatial distance?
  - correlation over time: a function of time separation?
  - within-item correlation: e.g. longitudinal studies.
- the magnitude of the correlations needs to be assessed

Correlated errors

When observations are correlated, and cannot be considered independent (e.g. by the random sampling argument), regression can be applied under a more general model that addresses these correlations.

- the structure of the correlation needs to be assessed
  - correlation in space: a function of spatial distance?
  - correlation over time: a function of time separation?
  - within-item correlation: e.g. longitudinal studies.
- the magnitude of the correlations needs to be assessed

Correlated errors

When observations are correlated, and cannot be considered independent (e.g. by the random sampling argument), regression can be applied under a more general model that addresses these correlations.

- the structure of the correlation needs to be assessed
  - correlation in space: a function of spatial distance?
  - correlation over time: a function of time separation?
  - within-item correlation: e.g. longitudinal studies.
- the magnitude of the correlations needs to be assessed

Correlated errors

When observations are correlated, and cannot be considered independent (e.g. by the random sampling argument), regression can be applied under a more general model that addresses these correlations.

- the structure of the correlation needs to be assessed
  - correlation in space: a function of spatial distance?
  - correlation over time: a function of time separation?
  - within-item correlation: e.g. longitudinal studies.
- the magnitude of the correlations needs to be assessed

Generalized linear models

Generalized linear models extend the (multiple) linear regression models by

- not assuming a (free) continuous variable as dependent
- not assuming a Gaussian distribution for the residuals

Examples:

- logistic regression: dependent variable is 0/1 (absence/presence)
- log-linear models: dependent variable is a count (Poisson)
- regression on log-transforms: the logarithm of y is taken instead of y

These models are very common in ecology.

Generalized linear models

Generalized linear models extend the (multiple) linear regression models by

- not assuming a (free) continuous variable as dependent
- not assuming a Gaussian distribution for the residuals

Examples:

- logistic regression: dependent variable is 0/1 (absence/presence)
- log-linear models: dependent variable is a count (Poisson)
- regression on log-transforms: the logarithm of y is taken instead of y

These models are very common in ecology.

Generalized linear models

Generalized linear models extend the (multiple) linear regression models by

- not assuming a (free) continuous variable as dependent
- not assuming a Gaussian distribution for the residuals

Examples:

- logistic regression: dependent variable is 0/1 (absence/presence)
- log-linear models: dependent variable is a count (Poisson)
- regression on log-transforms: the logarithm of y is taken instead of y

These models are very common in ecology.

Generalized linear models

Generalized linear models extend the (multiple) linear regression models by

- not assuming a (free) continuous variable as dependent
- not assuming a Gaussian distribution for the residuals

Examples:

- logistic regression: dependent variable is 0/1 (absence/presence)
- log-linear models: dependent variable is a count (Poisson)
- regression on log-transforms: the logarithm of y is taken instead of y

These models are very common in ecology.

Generalized linear models

Generalized linear models extend the (multiple) linear regression models by

- not assuming a (free) continuous variable as dependent
- not assuming a Gaussian distribution for the residuals

Examples:

- logistic regression: dependent variable is 0/1 (absence/presence)
- log-linear models: dependent variable is a count (Poisson)
- regression on log-transforms: the logarithm of y is taken instead of y

These models are very common in ecology.

R-squared and adjusted R-squared

Coefficient of multiple correlation:

$$R^2 = \frac{\sum_{i=1}^{n}(\hat{Y}_i - \bar{Y})^2}{\sum_{i=1}^{n}(Y_i - \bar{Y})^2}$$

Adjusted R^2:

$$\bar{R}^2 = \frac{(n-1)R^2 - k}{n-k-1}$$

with n the number of observations, and k the number of parameters fitted.