Statistical Modelling I

HES 505 Fall 2024: Session 21

Carolyn Koehn

Objectives

By the end of today you should be able to:

- Describe and implement overlay analyses
- Extend overlay analysis to statistical modeling
- Generate spatial predictions from statistical models

Overlay Analyses

Overlays

- Methods for identifying optimal site selection or suitability
- Apply a common scale to diverse or dissimilar outputs

Getting Started

1. Define the problem.
2. Break the problem into submodels.
3. Determine significant layers.
4. Reclassify or transform the data within a layer.
5. Add or combine the layers.
6. Verify

Boolean Overlays

- Successive disqualification of areas
- Series of “yes/no” questions
- “Sieve” mapping

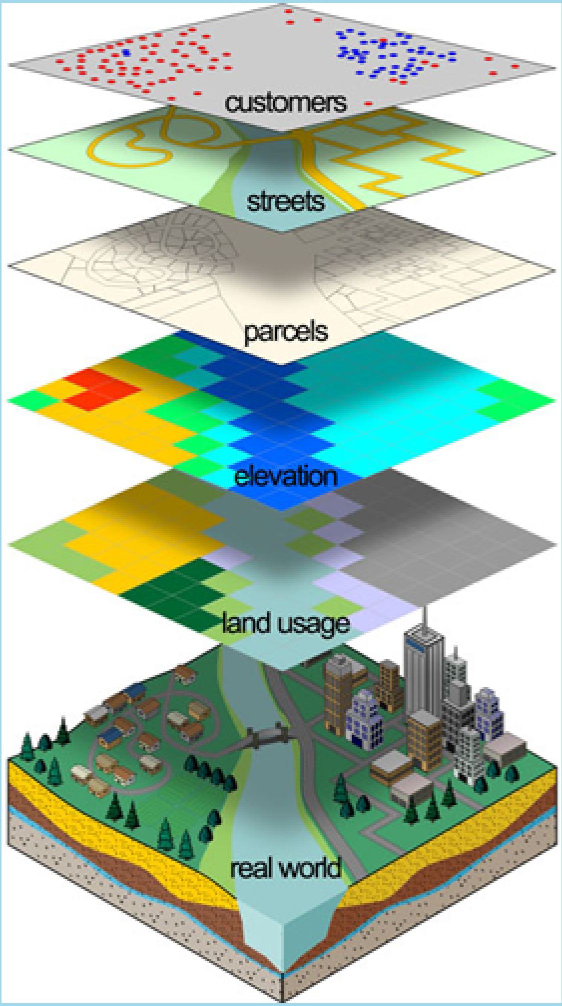

Boolean Overlays

- Reclassifying
- Which types of land are appropriate

```
nlcd <-  rast(system.file("raster/nlcd.tif", package = "spDataLarge"))
plot(nlcd)
```

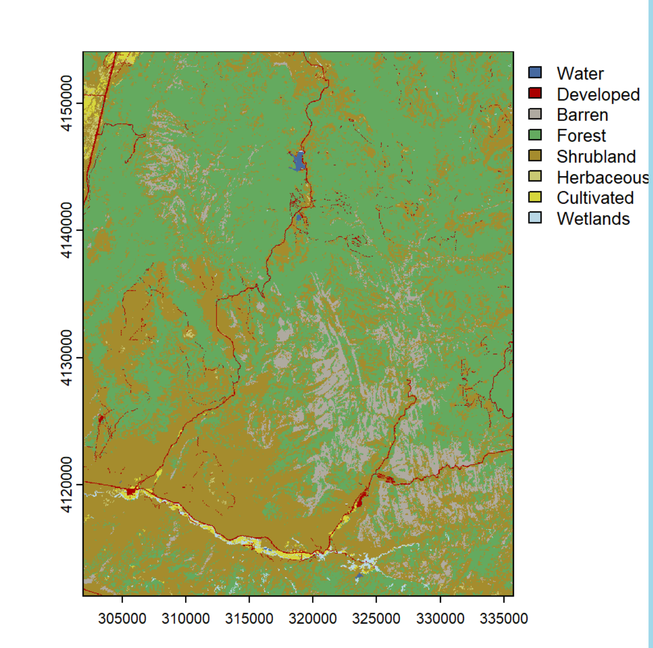

Boolean Overlays

- Which types of land are appropriate?

```r
nlcd.segments <- segregate(nlcd)
names(nlcd.segments) <- levels(nlcd)[[1]][-1,2]
plot(nlcd.segments)
```

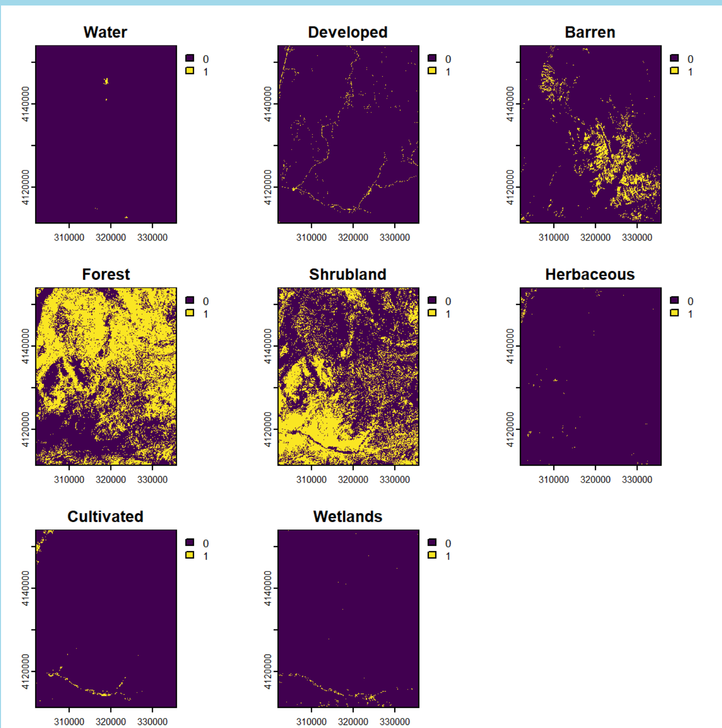

Boolean Overlays

- Which types of land are appropriate?

```
srtm <- rast(system.file("raster/srtm.tif", package = "spDataLarge"))
slope <- terrain(srtm, v = "slope")
```

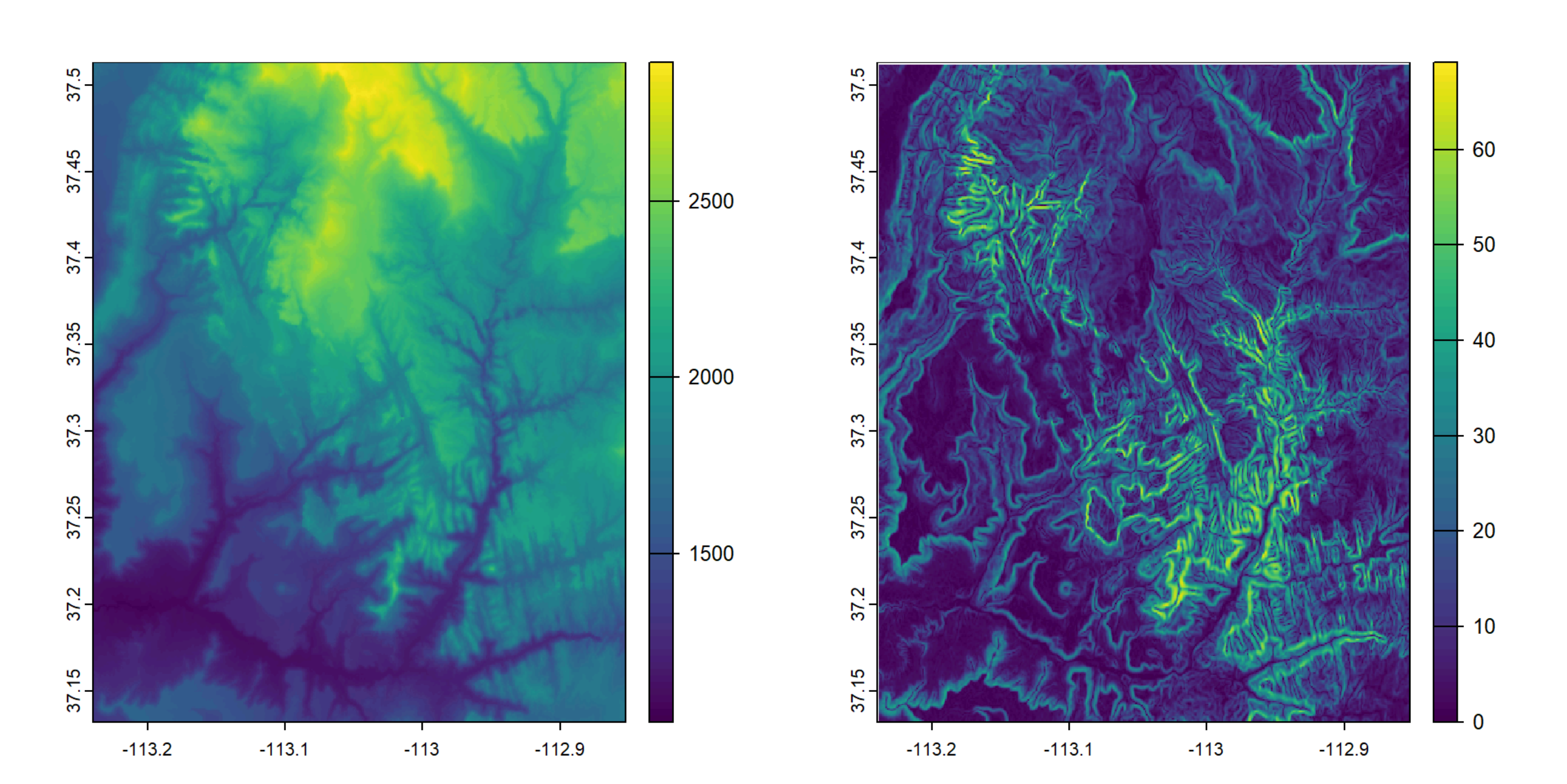

Boolean Overlays

- Make sure data is aligned!

```
suit.slope <- slope < 10
suit.landcov <- nlcd.segments["Shrubland"]
suit.slope.match <- project(suit.slope, suit.landcov)
suit <- suit.slope.match + suit.landcov
```

Boolean Overlays

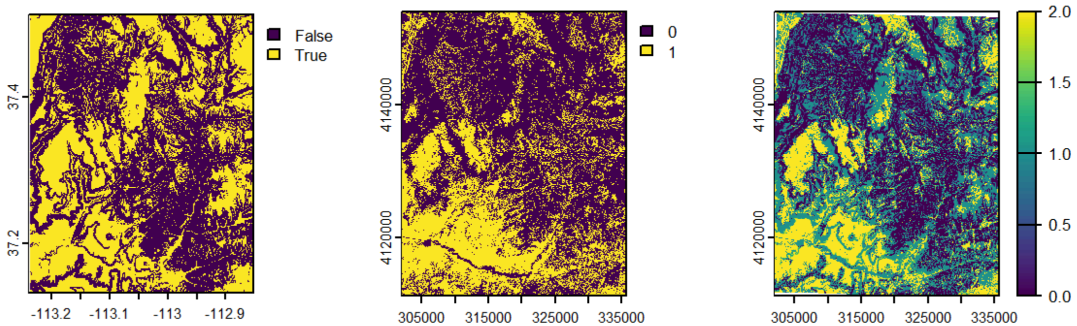

Challenges with Boolean Overlays

1. Assume relationships are really Boolean
2. No measurement error
3. Categorical measurements are known exactly
4. Boundaries are well-represented

A more general approach

- Define a *favorability* metric

- Treat as binary
- Then if all inputs () are suitable
- Then if not

Estimating favorability

- does not have to be binary (could be ordinal or continuous)
- could also be extended beyond simply 'suitable/not suitable'
- Adding weights allows incorporation of relative importance
- Other functions for combining inputs ()

Weighted Linear Combinations

- is now an index based on the values of
- can incorporate weights of evidence, uncertainty, or different participant preferences
- Dividing by normalizes by the sum of weights

Model-driven overlay

- If we estimate using data, we specify as the outcome of regression
- When is binary → logistic regression
- When is continuous → linear (gamma) regression
- When is discrete → Poisson regression
- Assumptions about matter!!

Logistic Regression and Distribution Models

Why do we create distribution models?

- To identify important correlations between predictors and the occurrence of an event
- Generate maps of the 'range' or 'niche' of events
- Understand spatial patterns of event co-occurrence
- Forecast changes in event distributions

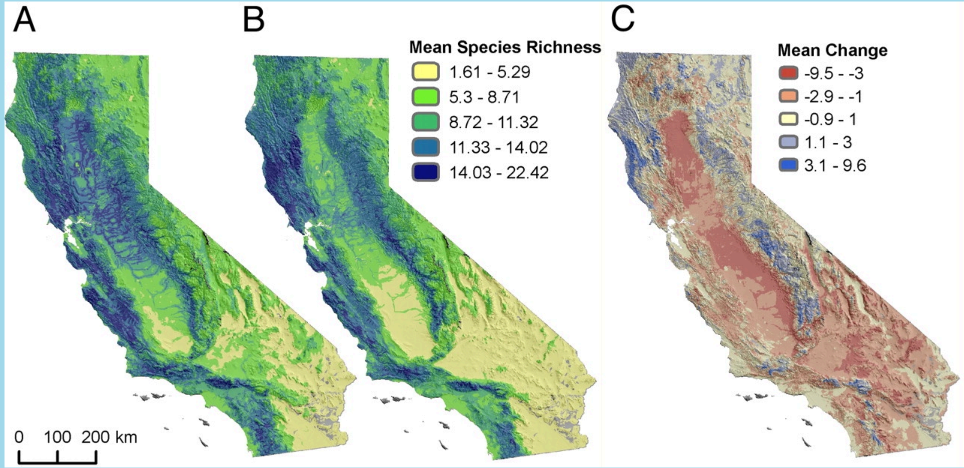

From Wiens et al. 2009

General analysis situation

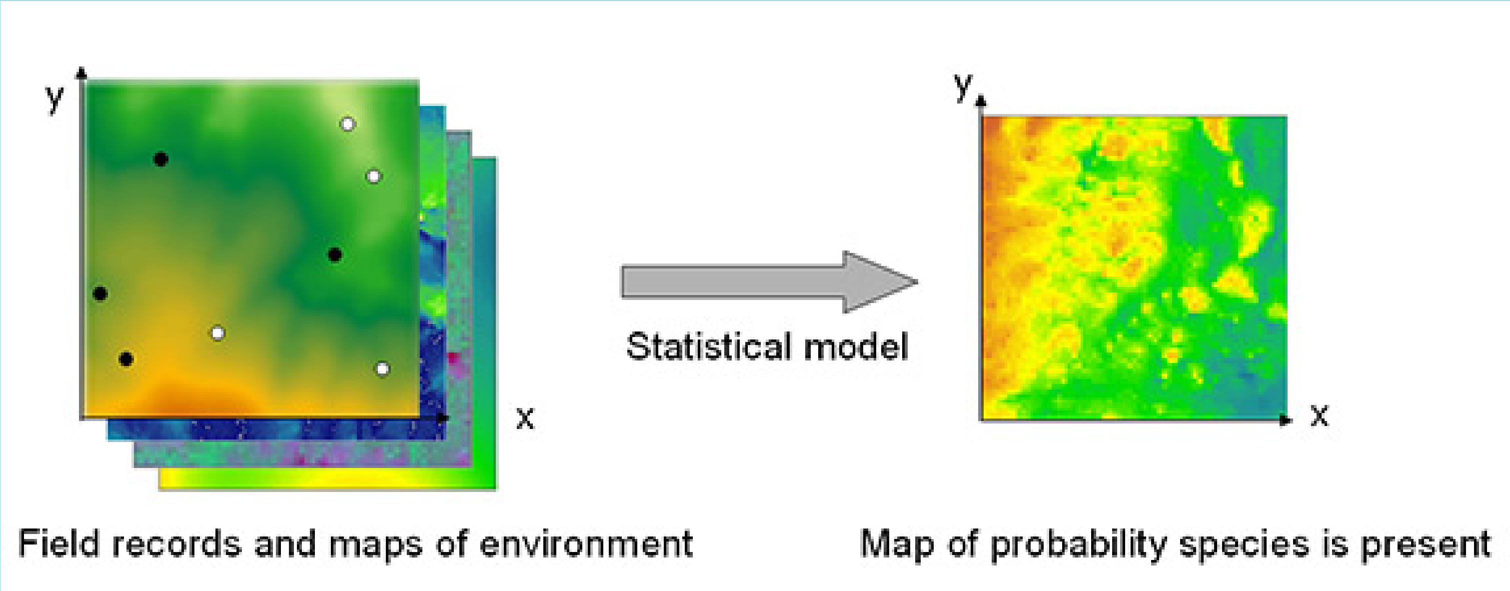

From Long

- Spatially referenced locations of events sampled from the study extent
- A matrix of predictors that can be assigned to each event based on spatial location

> **Goal**: Estimate the probability of occurrence of events across unsampled regions of the study area based on correlations with predictors

Modeling Presence-Absence Data

- Random or systematic sample of the study region
- The presence (or absence) of the event is recorded for each point
- Hypothesized predictors of occurrence are measured (or extracted) at each point

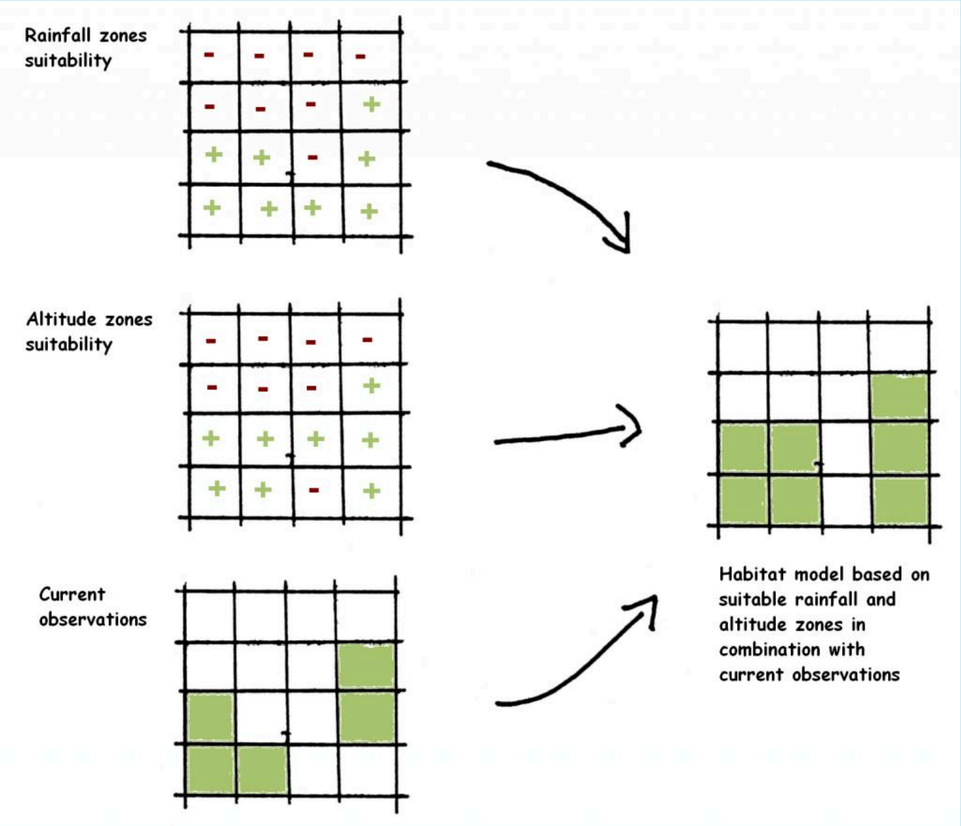

From By Ragnvald - Own work, CC BY-SA 3.0

Logistic regression

- We can model favorability as the **probability** of occurrence using a logistic regression
- A *link* function maps the linear predictor onto the support (0-1) for probabilities
- Estimates of can then be used to generate ‘wall-to-wall’ spatial predictions

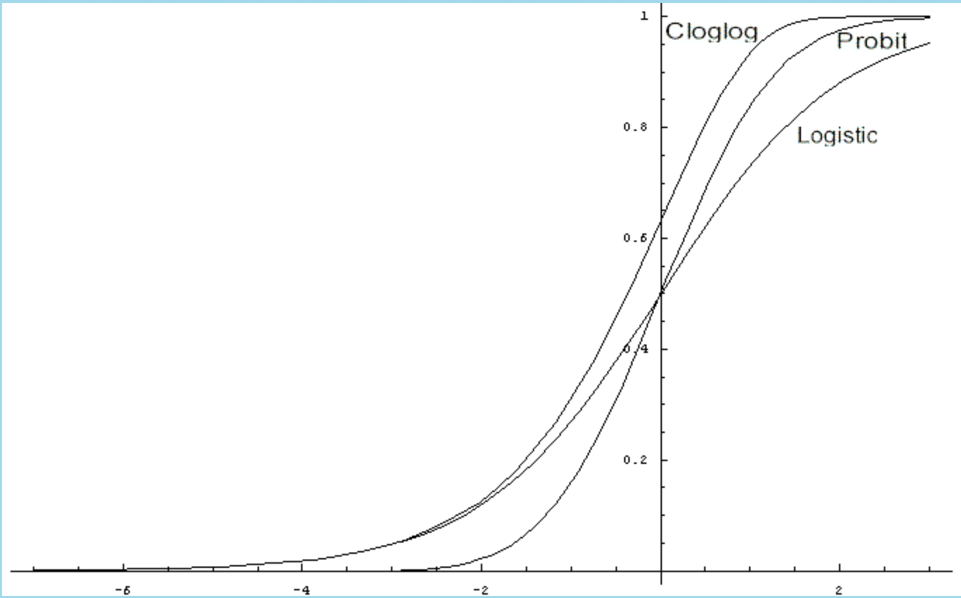

From Mendoza

An Example

Inputs from the `dismo` package

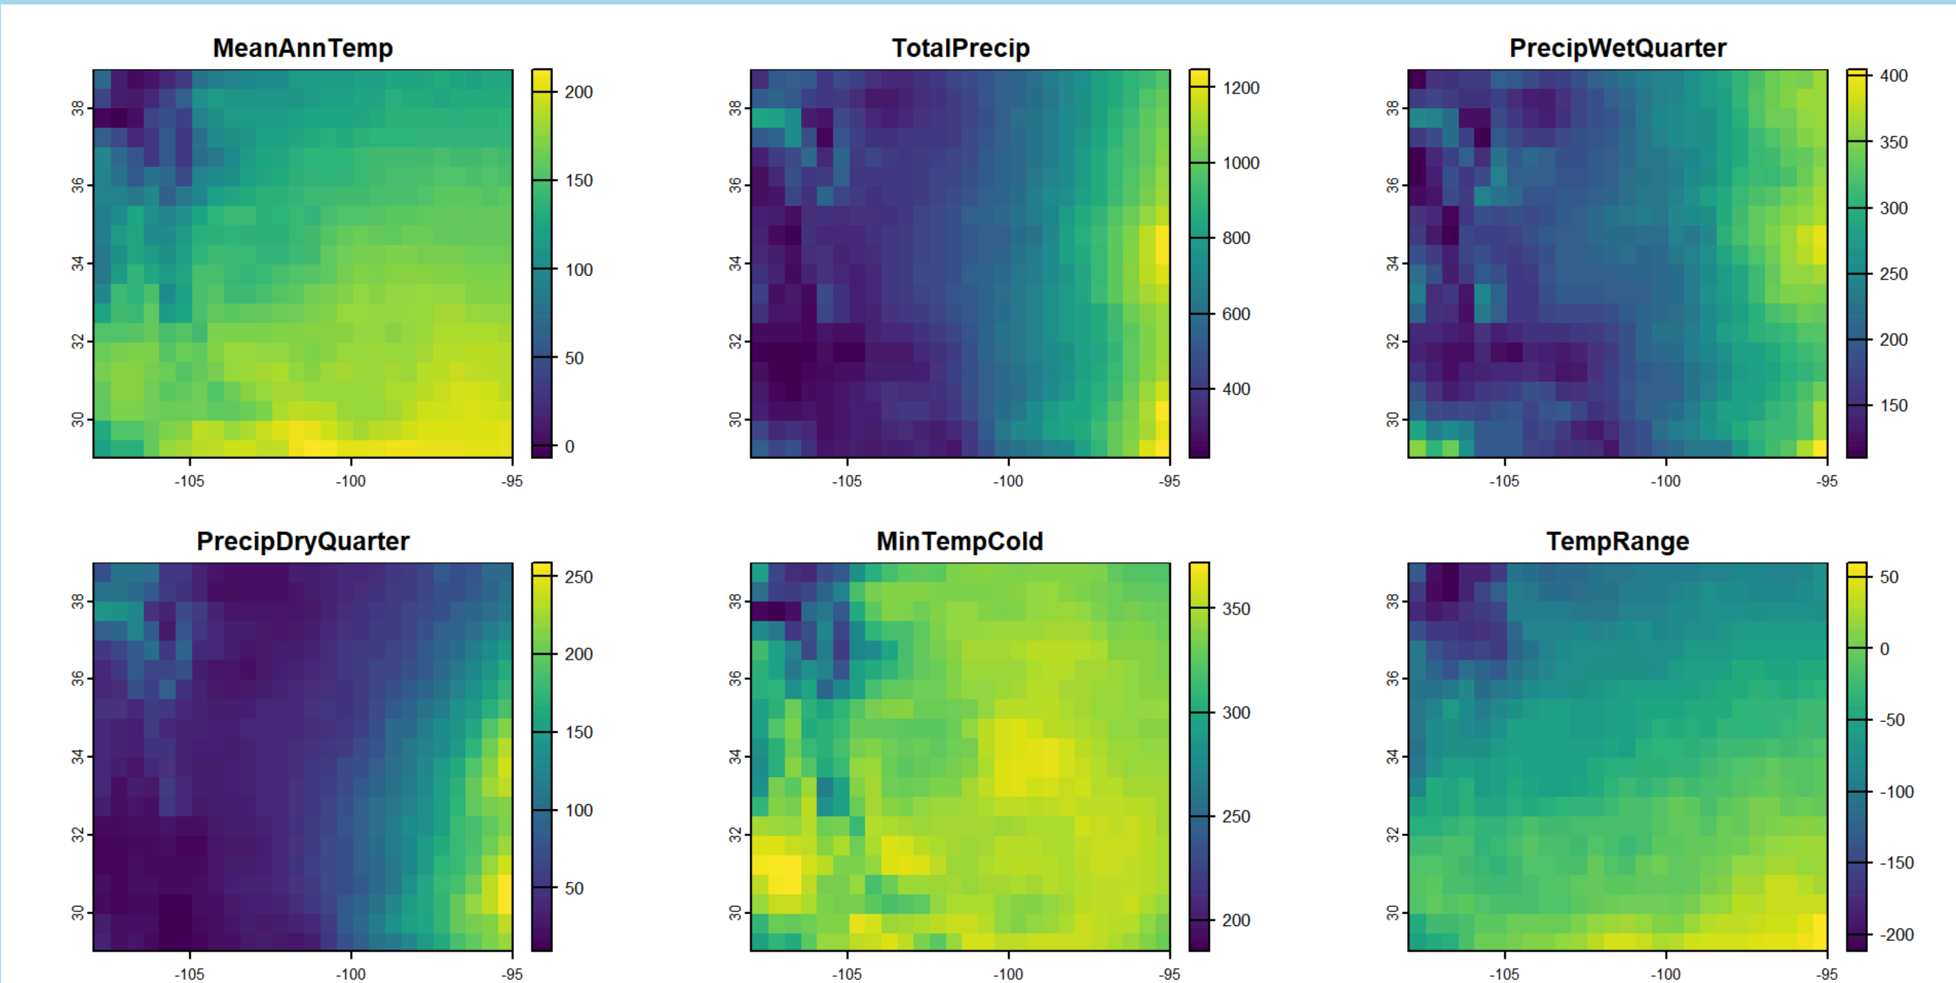

An Example

The sample data

An Example

Building our dataframe

```
pts.df <- terra::extract(pred.stack, vect(pres.abs), df=TRUE)
head(pts.df)
```

```
  ID MeanAnnTemp TotalPrecip PrecipWetQuarter PrecipDryQuarter MinTempCold
1  1         155         667              253               71         350
2  2         147         678              266               66         351
3  3         123         261              117               40         329
4  4         181         533              198               69         348
5  5         127         589              257               48         338
6  6          83         438              213               38         278
  TempRange
1       -45
2       -58
3       -64
4        -5
5       -81
6      -107
```

An Example

Building our dataframe

```r
pts.df[,2:7] <- scale(pts.df[,2:7])
summary(pts.df)
```

```
       ID           MeanAnnTemp        TotalPrecip       PrecipWetQuarter
 Min.   :  1.00   Min.   :-3.3729   Min.   :-1.3377   Min.   :-1.6926
 1st Qu.: 25.75   1st Qu.:-0.4594   1st Qu.:-0.7980   1st Qu.:-0.6895
 Median : 50.50   Median : 0.2282   Median :-0.2373   Median :-0.2224
 Mean   : 50.50   Mean   : 0.0000   Mean   : 0.0000   Mean   : 0.0000
 3rd Qu.: 75.25   3rd Qu.: 0.7118   3rd Qu.: 0.7140   3rd Qu.: 0.6508
 Max.   :100.00   Max.   : 1.4285   Max.   : 2.4843   Max.   : 2.2713
 PrecipDryQuarter   MinTempCold        TempRange
 Min.   :-1.0828   Min.   :-3.9919   Min.   :-2.7924
 1st Qu.:-0.7013   1st Qu.:-0.0598   1st Qu.:-0.5216
 Median :-0.3770   Median : 0.3582   Median : 0.2075
 Mean   : 0.0000   Mean   : 0.0000   Mean   : 0.0000
 3rd Qu.: 0.4290   3rd Qu.: 0.5495   3rd Qu.: 0.6450
 Max.   : 3.1713   Max.   : 1.1092   Max.   : 2.0407
```

An Example

Looking at correlations

```
pairs(pts.df[,2:7])
```

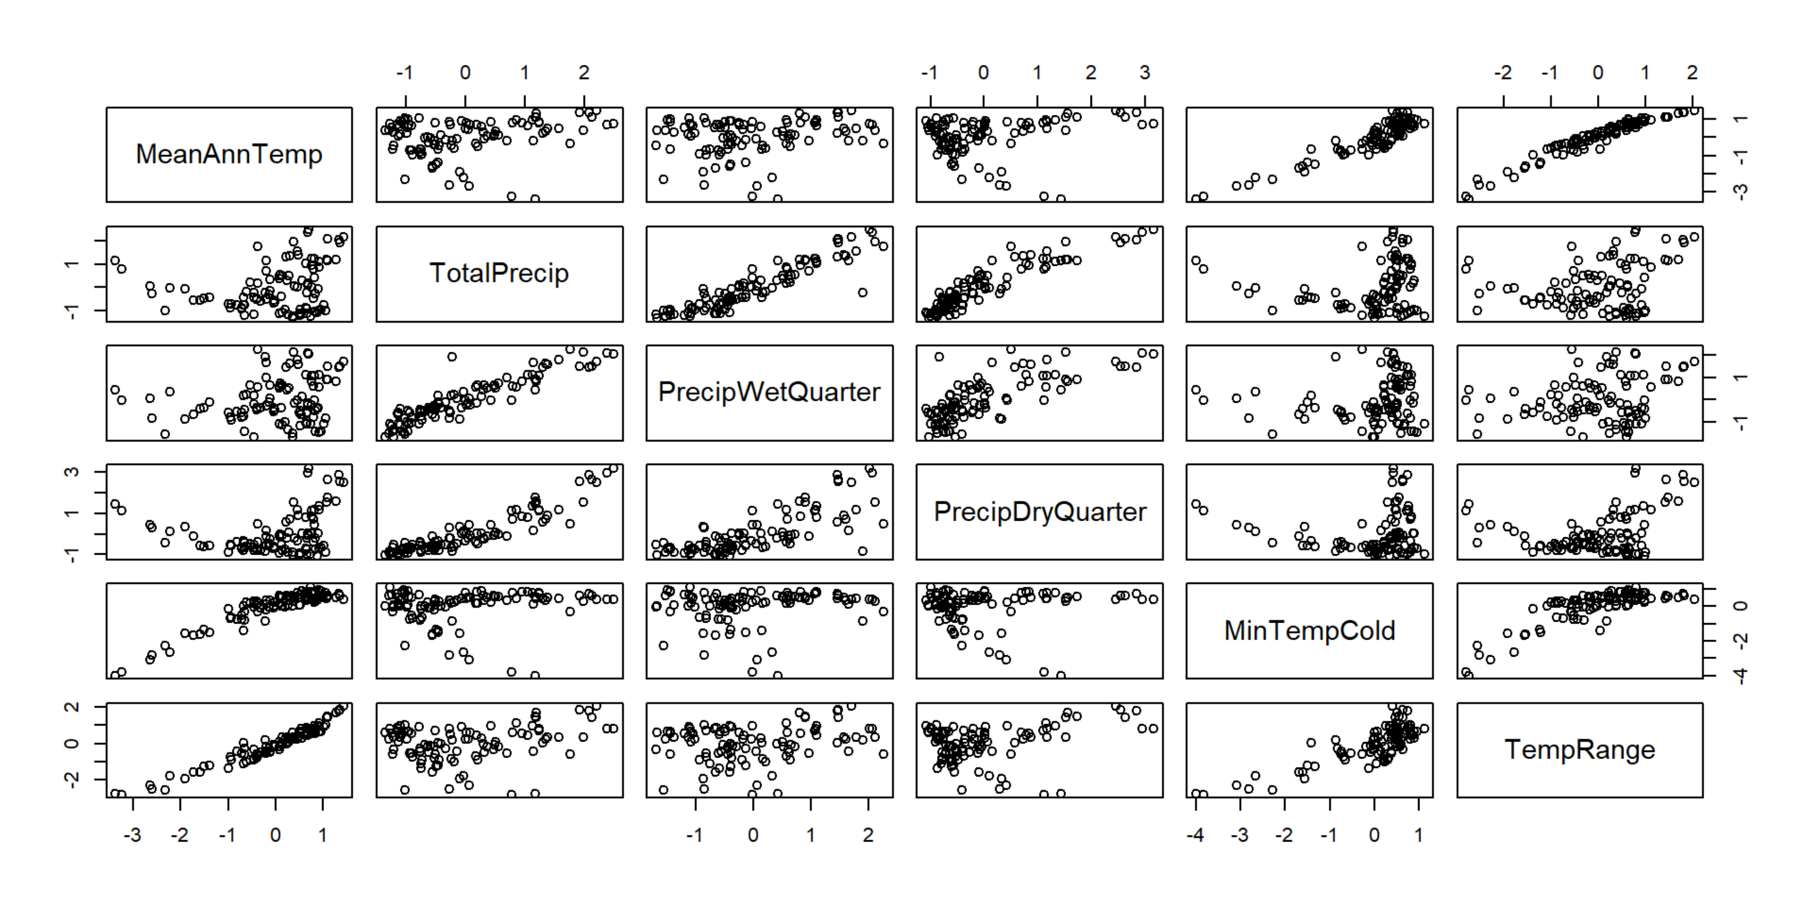

An Example

Looking at correlations

```r
corrplot(cor(pts.df[,2:7]), method = "number")
```

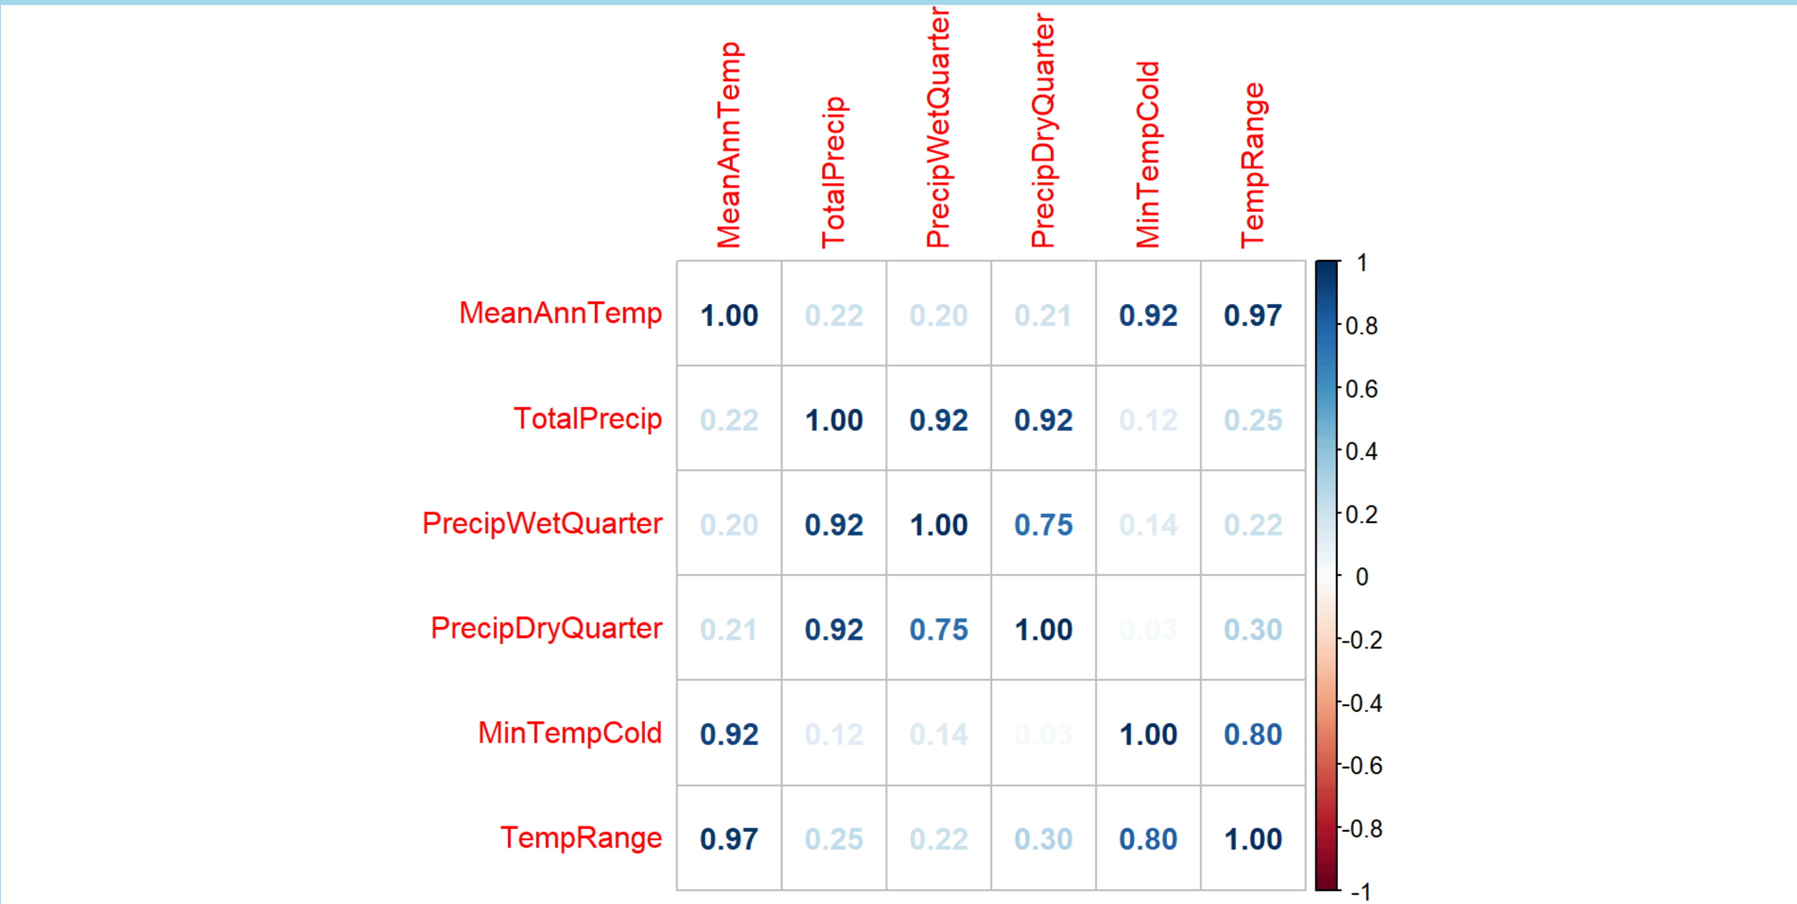

An Example

Fitting some models

```
pts.df <- cbind(pts.df, pres.abs$y)
colnames(pts.df)[8] <- "y"
logistic.global <- glm(y~., family=binomial(link="logit"), data=pts.df[,2:8
logistic.simple <- glm(y ~ MeanAnnTemp + TotalPrecip, family=binomial(link=
logistic.rich <- glm(y ~ MeanAnnTemp + PrecipWetQuarter + PrecipDryQuarter,
```

An Example

Checking out the results

```r
summary(logistic.global)
```

```

Call:
glm(formula = y ~ ., family = binomial(link = "logit"), data = pts.df[,
    2:8])

Coefficients:
                 Estimate Std. Error z value Pr(>|z|)
(Intercept)       -1.4461     0.5096  -2.837  0.00455 **
MeanAnnTemp       -6.3578     6.1645  -1.031  0.30237
TotalPrecip        7.1453     4.5577   1.568  0.11694
PrecipWetQuarter  -5.4207     3.0432  -1.781  0.07487 .
PrecipDryQuarter  -1.3110     2.2482  -0.583  0.55981
MinTempCold        3.0890     2.6334   1.173  0.24080
TempRange         -0.6213     4.5470  -0.137  0.89131
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

An Example

Checking out the results

```r
summary(logistic.simple)
```

```
Call:
glm(formula = y ~ MeanAnnTemp + TotalPrecip, family = binomial(link =
"logit"),
    data = pts.df[, 2:8])

Coefficients:
            Estimate Std. Error z value Pr(>|z|)
(Intercept)  -0.9880     0.3145  -3.141  0.00168 **
MeanAnnTemp  -2.9990     0.6647  -4.512 6.42e-06 ***
TotalPrecip   0.3924     0.3827   1.025  0.30517
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)
```

An Example

Checking out the results

```r
summary(logistic.rich)
```

```
Call:
glm(formula = y ~ MeanAnnTemp + PrecipWetQuarter + PrecipDryQuarter,
    family = binomial(link = "logit"), data = pts.df[, 2:8])

Coefficients:
                 Estimate Std. Error z value Pr(>|z|)    
(Intercept)      -0.96504    0.35650  -2.707  0.00679 ** 
MeanAnnTemp      -2.85446    0.66142  -4.316 1.59e-05 ***
PrecipWetQuarter  0.03212    0.43102   0.075  0.94060    
PrecipDryQuarter  0.16759    0.64935   0.258  0.79634    
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

(Dispersion parameter for binomial family taken to be 1)
```

An Example

Comparing models

```
AIC(logistic.global, logistic.simple, logistic.rich)
```

```
                df      AIC
logistic.global  7 65.76394
logistic.simple  3 74.10760
logistic.rich    4 77.00622
```

An Example

Generating predictions

```
preds <- predict(object=pred.stack, model=logistic.simple)
plot(preds)
plot(pres.pts$geometry, add=TRUE, pch=3, col="blue")
plot(abs.pts$geometry, add=TRUE, pch ="-", col="red")
```

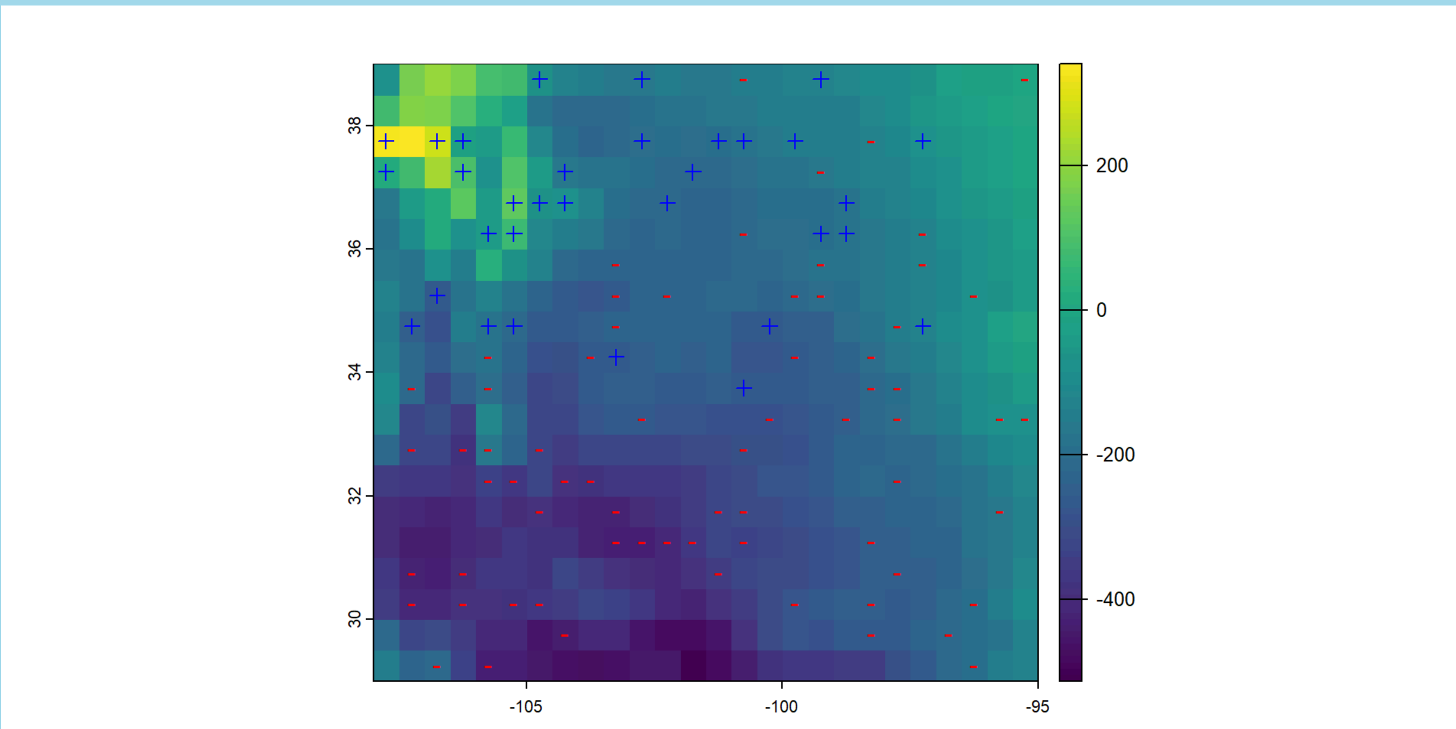

An Example

Generating predictions

```
preds <- predict(object=pred.stack, model=logistic.simple, type="response")
plot(preds)
plot(pres.pts$geometry, add=TRUE, pch=3, col="blue")
plot(abs.pts$geometry, add=TRUE, pch ="-", col="red")
```

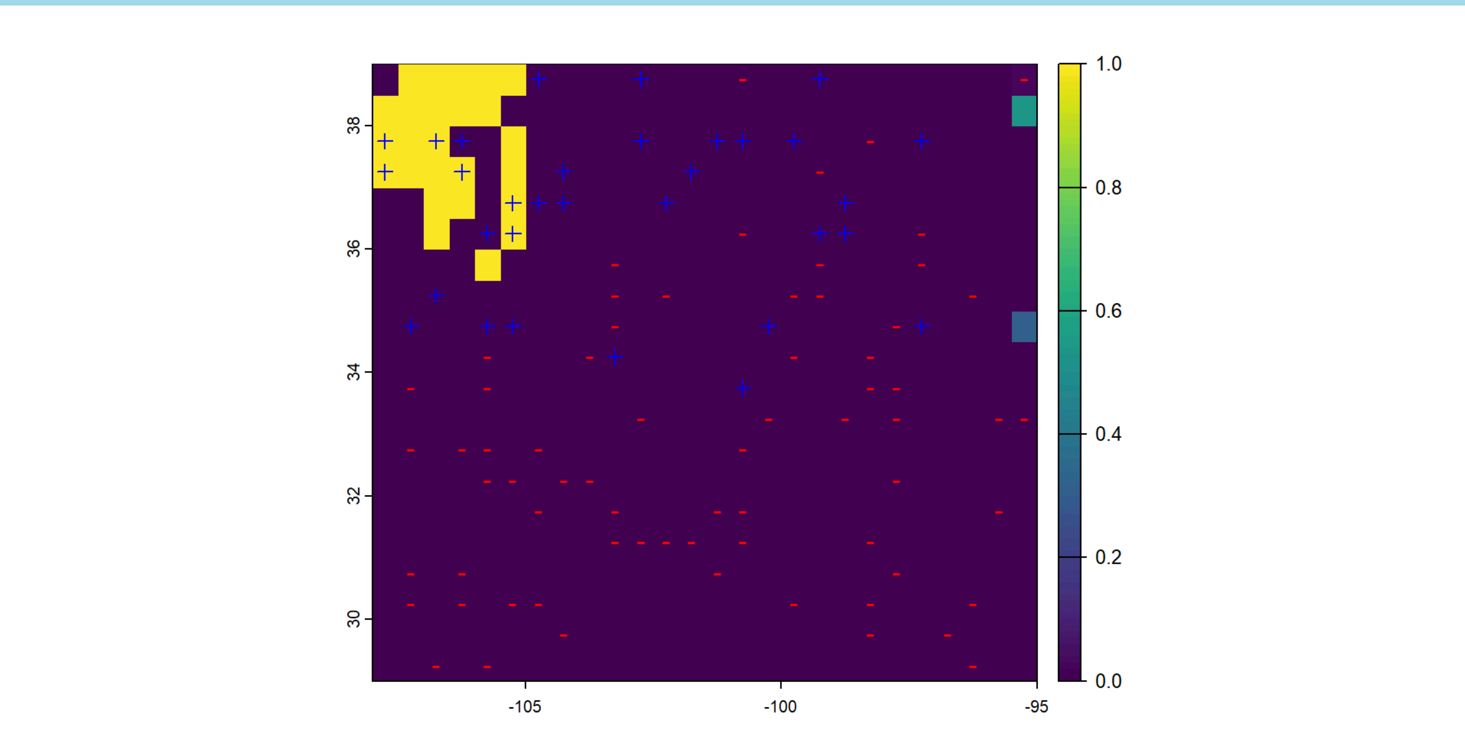

An Example

Generating predictions

```
preds <- predict(object=pred.stack, model=logistic.global, type="response")
plot(preds)
plot(pres.pts$geometry, add=TRUE, pch=3, col="blue")
plot(abs.pts$geometry, add=TRUE, pch ="-", col="red")
```

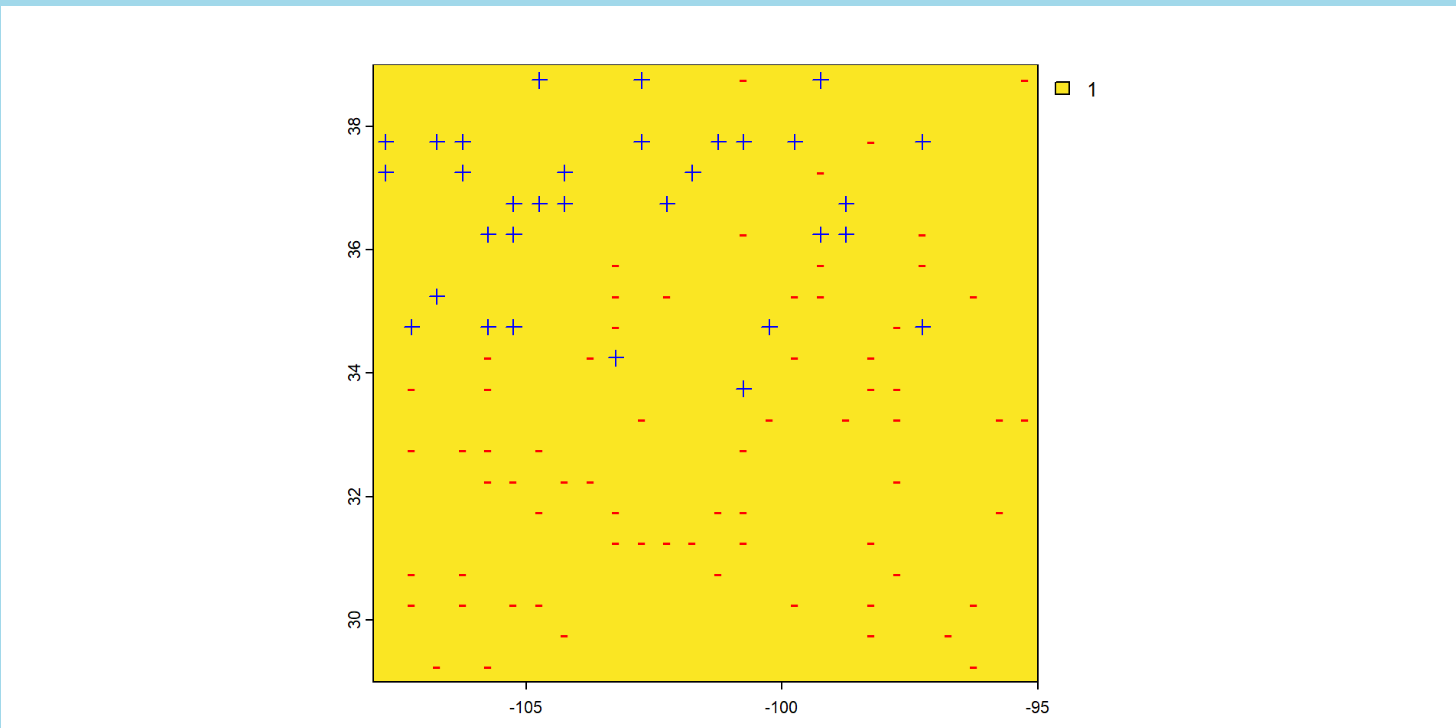

An Example

Generating predictions

```
preds <- predict(object=pred.stack, model=logistic.rich, type="response")
plot(preds)
plot(pres.pts$geometry, add=TRUE, pch=3, col="blue")
plot(abs.pts$geometry, add=TRUE, pch ="-", col="red")
```

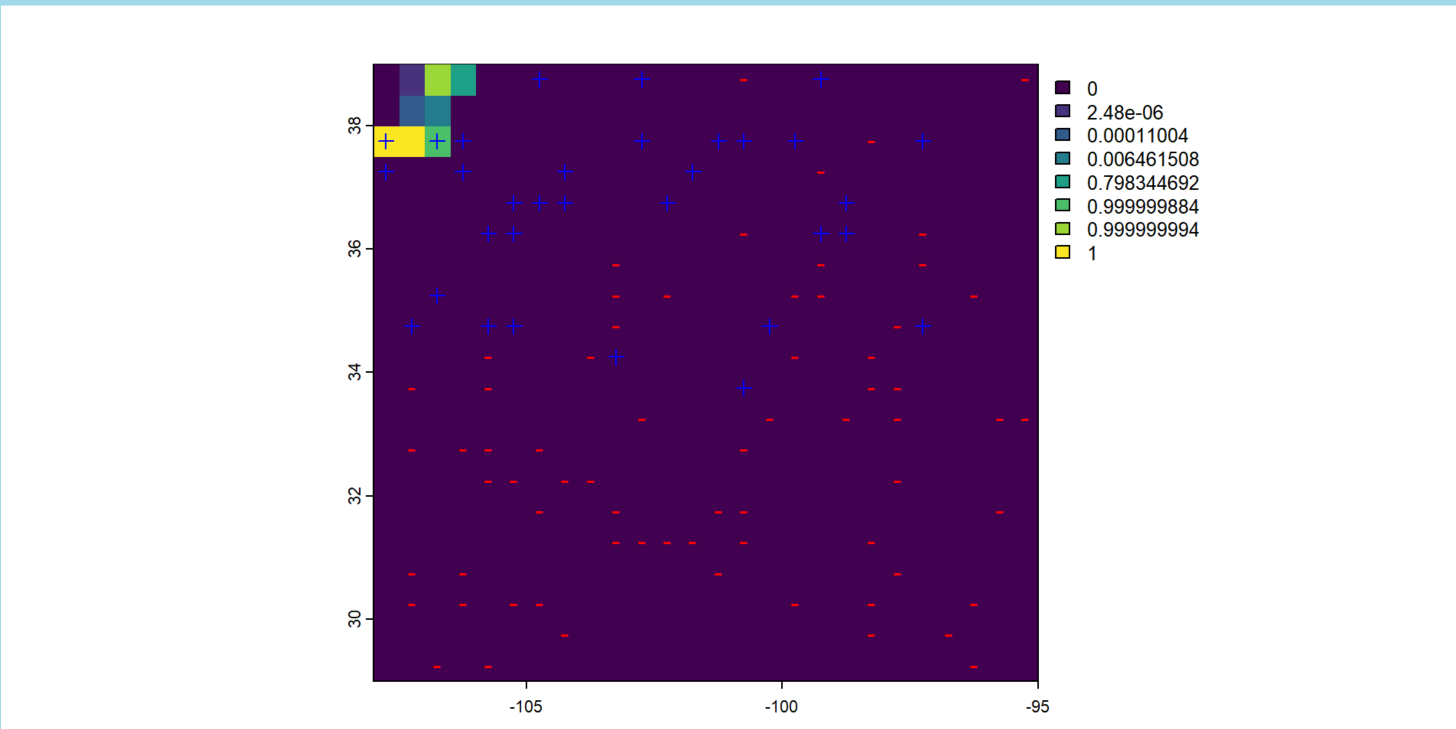

Key assumptions of logistic regression

- Dependent variable must be binary
- Observations must be independent (important for spatial analyses)
- Predictors should not be collinear
- Predictors should be linearly related to the log-odds
- **Sample Size**