BISC-869, Occupancy Models

April 8, 2018

Occupancy models

A species' *occupancy* is a measure of how many sites that species occupies.

Occupancy models

A species' *occupancy* is a measure of how many sites that species occupies.

Occupancy models extend the ideas from mark-recapture to make inferences about species' occupancy.

Occupancy models

A species' *occupancy* is a measure of how many sites that species occupies.

Occupancy models extend the ideas from mark-recapture to make inferences about species' occupancy.

Instead of re-capturing individuals (as we did with mark-recapture), here we "re-capture" the species at a site if we detect the species at that site. For example, our survey method might be listening for bird calls.

Occupancy models

A species' *occupancy* is a measure of how many sites that species occupies.

Occupancy models extend the ideas from mark-recapture to make inferences about species' occupancy.

Instead of re-capturing individuals (as we did with mark-recapture), here we "re-capture" the species at a site if we detect the species at that site. For example, our survey method might be listening for bird calls.

Collecting occupancy data is typically much less data-intensive than mark-recapture.

Single-species, single-season model

Suppose we survey 6 sites, 3 times each.

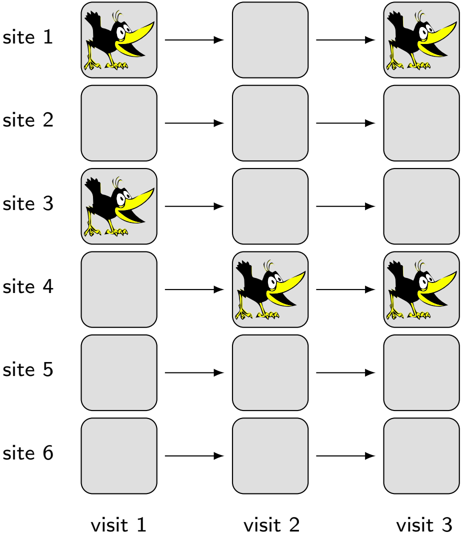

site 1

site 2

site 3

site 4

site 5

site 6

visit 1

visit 2

visit 3

Single-species, single-season model

Suppose we survey 6 sites, 3 times each.

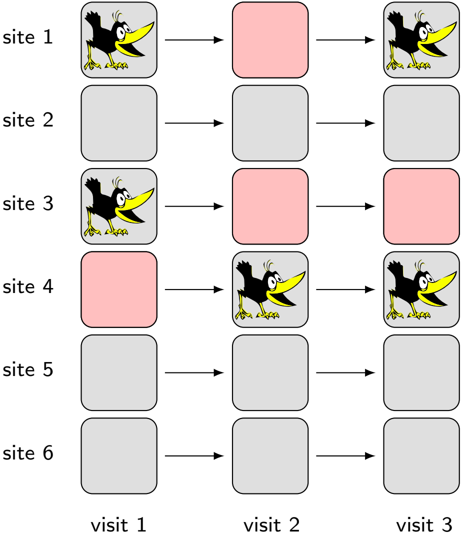

We assume “closure” of the system over the period of sampling. Thus, we know these are missed detections.

Single-species, single-season model

Suppose we survey 6 sites, 3 times each.

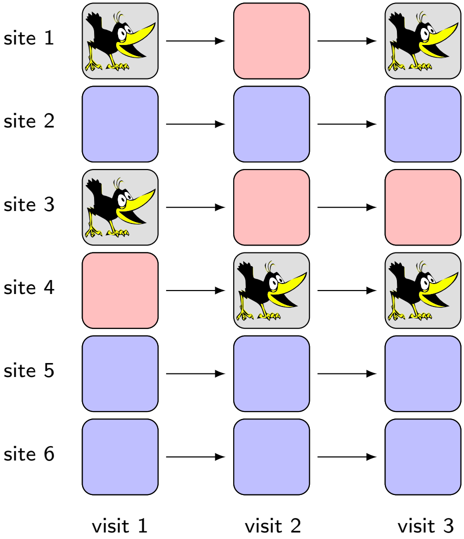

site 1

site 2

site 3

site 4

site 5

site 6

visit 1

visit 2

visit 3

We assume “closure” of the system over the period of sampling. Thus, we know these are missed detections.

These could be missed detections, or the species might really be missing from these sites.

Why do we need to do this?

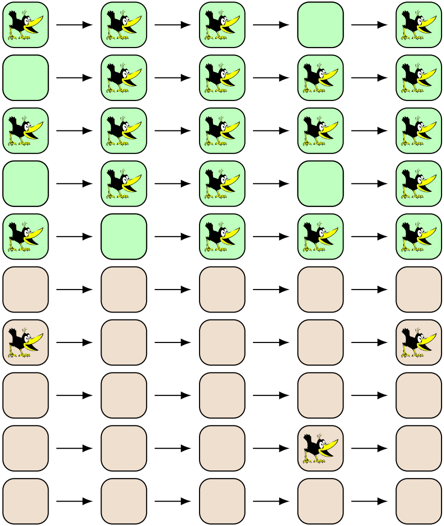

Suppose we want to know whether green versus brown sites are better habitat for our favourite cartoon bird species? We survey birds over 5 occassions across 10 sites (5 green, 5 brown).

Why do we need to do this?

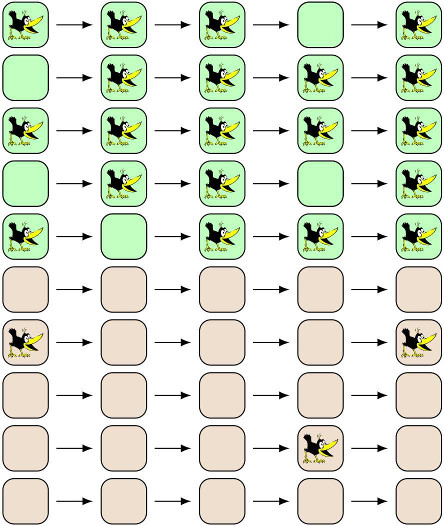

Suppose we want to know whether green versus brown sites are better habitat for our favourite cartoon bird species? We survey birds over 5 occassions across 10 sites (5 green, 5 brown).

We find that 5 green sites and two brown sites contain the bird. Based on this, we conclude that brown sites are lower quality habitat for this bird. What is potentially wrong with this?

Example

Gretchen Daily (and others) conducted point counts of bird communities across Costa Rica at 44 sites that span a land-use gradient (forest, diverse agriculture, intense agriculture). They recorded all observed species for > 10 years.

Here, we'll limit our analysis to one species, the Boat-billed Flycatcher (*Megarynchus pitangua*) and one year. We'll also ignore the fact that we have multiple site types nested within multiple regions (in workshop, we'll incorporate the former of these assumptions).

```
## package data for JAGS
my.data <- list(X=X,
                nsite=nsite,
                nvisit=nvisit)
```

```
> head(my.data$X, 10)
       visit
site     1 2 3
  F-GU1 0 1 0
  F-GU2 0 0 0
  F-GU3 0 0 0
  F-LC1 0 0 0
  F-LC2 0 0 0
  F-LC3 0 0 0
  F-PV1 0 1 0
  F-PV2 0 0 0
  F-PV3 1 1 0
  F-SI1 0 0 0
> my.data$nsite
[1] 44
> my.data$nvisit
[1] 3
```

Here, we also need inits. OtherwiseJAGS will draw the initial occupancy states from the priors which will sometimes lead to a site initially being assigned a "not occupied" status, despite the species being observed there. This will crash the model.

```
## inits
my.inits <- function() {
  list(Z=Z.init)
}
```

```
> head(Z.init, 10)
F-GU1 F-GU2 F-GU3 F-LC1 F-LC2 F-LC3 F-PV1 F-PV2 F-PV3 F-SI1
    1     0     0     0     0     0     1     0     1     0
```

Example

The model:

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Example

The model:

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Flat priors.

Example

The model:

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Site is occupied with probability psi.

Example

The model:

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Species detection probability is p, provided it is present.

Example

The model:

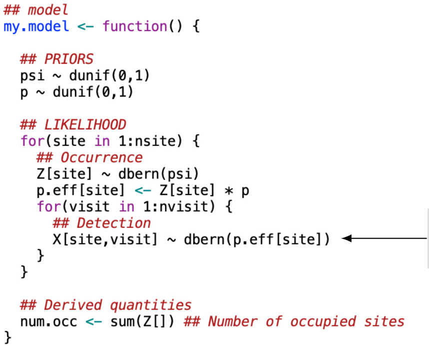

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Species is detected with probability p at sites in which it is present.

Example

The model:

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Calculate the number of occupied sites.

Example

Posterior probability distribution for occupancy, ψ

Observed occupancy.

Occupancy with highest posterior probability.

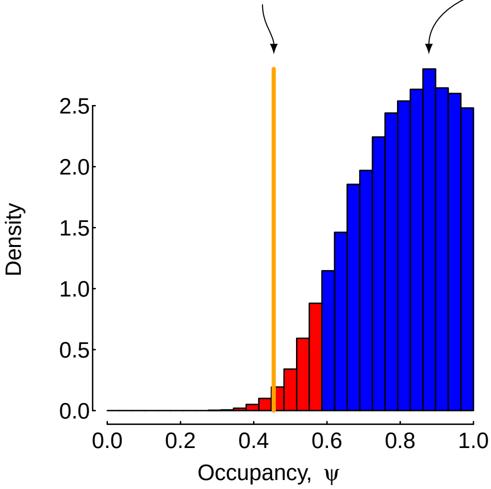

Notice that our estimated occupancy is substantially greater than our observed occupancy.

Here the blue region denotes the 95% BCI (using the "highest density region" method, due the skew in the distribution).

Why is the estimated occupancy so much higher than the observed occupancy?

Example

Posterior probability distribution for detectability, *p*

Detectability with highest posterior probability.

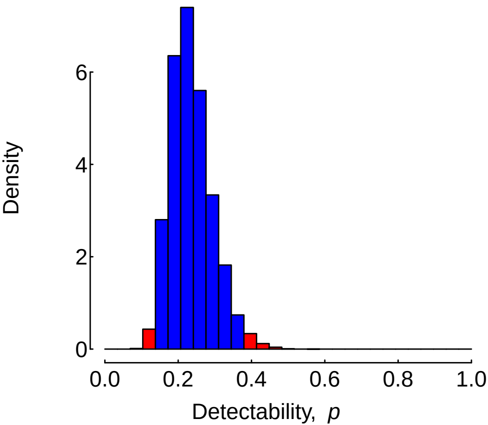

Example

Posterior probability distribution for the number of occupied sites

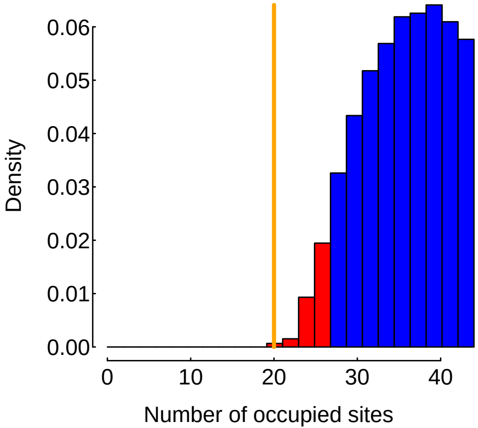

Based on this analysis, it is entirely plausible that this species is present at every one of our 44 sites, despite only having been observed at 20 of them.

Example

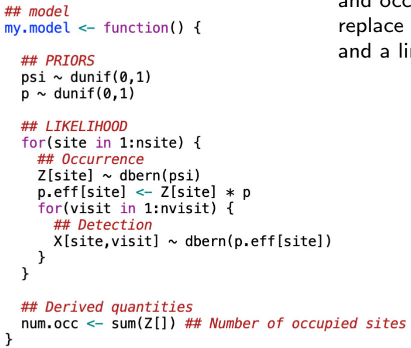

```
## model
my.model <- function() {

  ## PRIORS
  psi ~ dunif(0,1)
  p ~ dunif(0,1)

  ## LIKELIHOOD
  for(site in 1:nsite) {
    ## Occurrence
    Z[site] ~ dbern(psi)
    p.eff[site] <- Z[site] * p
    for(visit in 1:nvisit) {
      ## Detection
      X[site,visit] ~ dbern(p.eff[site])
    }
  }

  ## Derived quantities
  num.occ <- sum(Z[]) ## Number of occupied sites
}
```

Important note: we can add fixed and random effects to both detectability and occupancy, but to do so we must replace the `dunif` prior with a logit link and a linear equation.

Suppose we survey 6 sites, 3 times each, for multiple species.

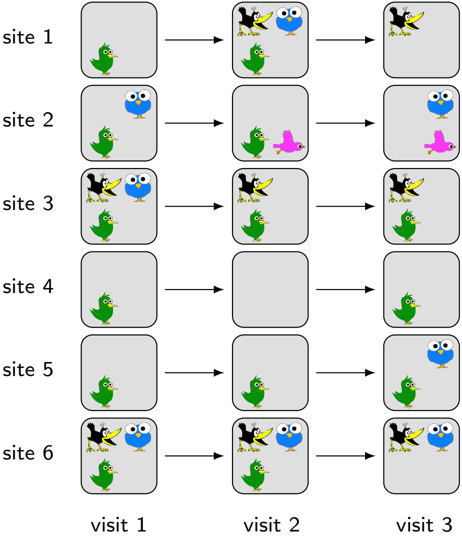

This is a powerful extension, because now we can look at how attributes of species affect occupancy.

A strength of this approach is that you can use data from all species, even ones only detected a few times.

Implementation is straight-forward; we just add a random effect of species.

Multi-species occupancy model

Here, I like to use multi-dimensional arrays for data. These are more work to construct, but make indexing within the model more intuitive.

Multi-species occupancy model

Here, I like to use multi-dimensional arrays for data. These are more work to construct, but make indexing within the model more intuitive.

A multidimensional array is similar to a data-frame, but with more than 2 dimensions.

Multi-species occupancy model

Here, I like to use multi-dimensional arrays for data. These are more work to construct, but make indexing within the model more intuitive.

A multidimensional array is similar to a data-frame, but with more than 2 dimensions.

Here, I will use a 3-dimensional array, `X`, with dimensions corresponding to "site", "visit", and "species", in that order.

Multi-species occupancy model

Here, I like to use multi-dimensional arrays for data. These are more work to construct, but make indexing within the model more intuitive.

A multidimensional array is similar to a data-frame, but with more than 2 dimensions.

Here, I will use a 3-dimensional array, `X`, with dimensions corresponding to "site", "visit", and "species", in that order.

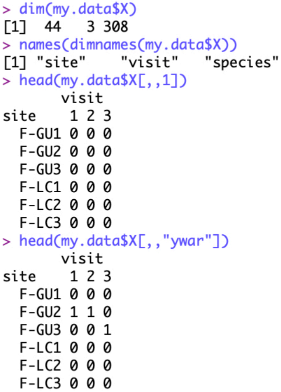

```
> dim(my.data$X)
[1]  44   3 308
> names(dimnames(my.data$X))
[1] "site"    "visit"   "species"
> head(my.data$X[,,1])
       visit
site    1 2 3
  F-GU1 0 0 0
  F-GU2 0 0 0
  F-GU3 0 0 0
  F-LC1 0 0 0
  F-LC2 0 0 0
  F-LC3 0 0 0
> head(my.data$X[,,"ywar"])
       visit
site    1 2 3
  F-GU1 0 0 0
  F-GU2 1 1 0
  F-GU3 0 0 1
  F-LC1 0 0 0
  F-LC2 0 0 0
  F-LC3 0 0 0
```

In addition to this more complex structure for our observations, our inits for `Z` must also now contain a species dimension:

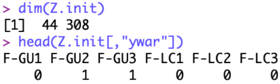

```
> dim(Z.init)
[1]  44 308
> head(Z.init[,"ywar"])
F-GU1 F-GU2 F-GU3 F-LC1 F-LC2 F-LC3
    0     1     1     0     0     0
```

We must also have a parameter indicating the number of species:

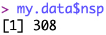

```
> my.data$nsp
[1] 308
```

Multi-species occupancy model

The model:

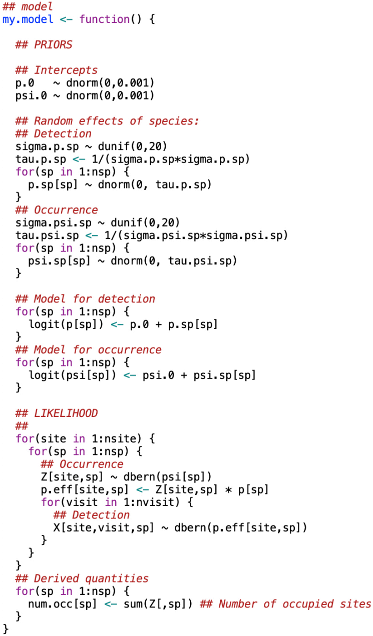

```
## model
my.model <- function() {

  ## PRIORS

  ## Intercepts
  p.0   ~ dnorm(0,0.001)
  psi.0 ~ dnorm(0,0.001)

  ## Random effects of species:
  ## Detection
  sigma.p.sp ~ dunif(0,20)
  tau.p.sp <- 1/(sigma.p.sp*sigma.p.sp)
  for(sp in 1:nsp) {
    p.sp[sp] ~ dnorm(0, tau.p.sp)
  }
  ## Occurrence
  sigma.psi.sp ~ dunif(0,20)
  tau.psi.sp <- 1/(sigma.psi.sp*sigma.psi.sp)
  for(sp in 1:nsp) {
    psi.sp[sp] ~ dnorm(0, tau.psi.sp)
  }

  ## Model for detection
  for(sp in 1:nsp) {
    logit(p[sp]) <- p.0 + p.sp[sp]
  }
  ## Model for occurrence
  for(sp in 1:nsp) {
    logit(psi[sp]) <- psi.0 + psi.sp[sp]
  }

  ## LIKELIHOOD
  ##
  for(site in 1:nsite) {
    for(sp in 1:nsp) {
      ## Occurrence
      Z[site,sp] ~ dbern(psi[sp])
      p.eff[site,sp] <- Z[site,sp] * p[sp]
      for(visit in 1:nvisit) {
        ## Detection
        X[site,visit,sp] ~ dbern(p.eff[site,sp])
      }
    }
  }
  ## Derived quantities
  for(sp in 1:nsp) {
    num.occ[sp] <- sum(Z[,sp]) ## Number of occupied sites
  }
}
```

Multi-species occupancy model

The model:

```
## model
my.model <- function() {

  ## PRIORS

  ## Intercepts
  p.0   ~ dnorm(0,0.001)
  psi.0 ~ dnorm(0,0.001)

  ## Random effects of species:
  ## Detection
  sigma.p.sp ~ dunif(0,20)
  tau.p.sp <- 1/(sigma.p.sp*sigma.p.sp)
  for(sp in 1:nsp) {
    p.sp[sp] ~ dnorm(0, tau.p.sp)
  }
  ## Occurrence
  sigma.psi.sp ~ dunif(0,20)
  tau.psi.sp <- 1/(sigma.psi.sp*sigma.psi.sp)
  for(sp in 1:nsp) {
    psi.sp[sp] ~ dnorm(0, tau.psi.sp)
  }

  ## Model for detection
  for(sp in 1:nsp) {
    logit(p[sp]) <- p.0 + p.sp[sp]
  }
  ## Model for occurrence
  for(sp in 1:nsp) {
    logit(psi[sp]) <- psi.0 + psi.sp[sp]
  }

  ## LIKELIHOOD
  ##
  for(site in 1:nsite) {
    for(sp in 1:nsp) {
      ## Occurrence
      Z[site,sp] ~ dbern(psi[sp])
      p.eff[site,sp] <- Z[site,sp] * p[sp]
      for(visit in 1:nvisit) {
        ## Detection
        X[site,visit,sp] ~ dbern(p.eff[site,sp])
      }
    }
  }
  ## Derived quantities
  for(sp in 1:nsp) {
    num.occ[sp] <- sum(Z[,sp]) ## Number of occupied sites
  }
}
```

Our equations for detectability and occurrence now contain both an intercept and a random effect of species. Thus, we must use the logit link in order to ensure that each remains between 0 and 1.

Multi-species occupancy model

The model:

```
## model
my.model <- function() {

  ## PRIORS

  ## Intercepts
  p.0   ~ dnorm(0,0.001)
  psi.0 ~ dnorm(0,0.001)

  ## Random effects of species:
  ## Detection
  sigma.p.sp ~ dunif(0,20)
  tau.p.sp <- 1/(sigma.p.sp*sigma.p.sp)
  for(sp in 1:nsp) {
    p.sp[sp] ~ dnorm(0, tau.p.sp)
  }
  ## Occurrence
  sigma.psi.sp ~ dunif(0,20)
  tau.psi.sp <- 1/(sigma.psi.sp*sigma.psi.sp)
  for(sp in 1:nsp) {
    psi.sp[sp] ~ dnorm(0, tau.psi.sp)
  }

  ## Model for detection
  for(sp in 1:nsp) {
    logit(p[sp]) <- p.0 + p.sp[sp]
  }
  ## Model for occurrence
  for(sp in 1:nsp) {
    logit(psi[sp]) <- psi.0 + psi.sp[sp]
  }

  ## LIKELIHOOD
  ##
  for(site in 1:nsite) {
    for(sp in 1:nsp) {
      ## Occurrence
      Z[site,sp] ~ dbern(psi[sp])
      p.eff[site,sp] <- Z[site,sp] * p[sp]
      for(visit in 1:nvisit) {
        ## Detection
        X[site,visit,sp] ~ dbern(p.eff[site,sp])
      }
    }
  }
  ## Derived quantities
  for(sp in 1:nsp) {
    num.occ[sp] <- sum(Z[,sp]) ## Number of occupied sites
  }
}
```

Likelihood remains the same, except now terms are indexed by species.

Posterior probability distributions for occupancy for two of the species

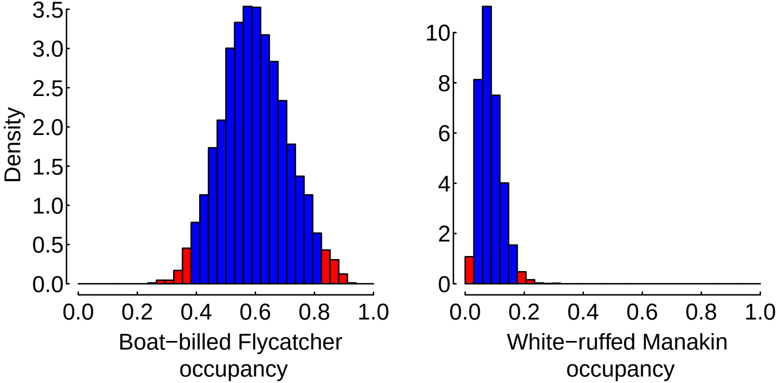

Posterior probability distributions for detectability for the same two species

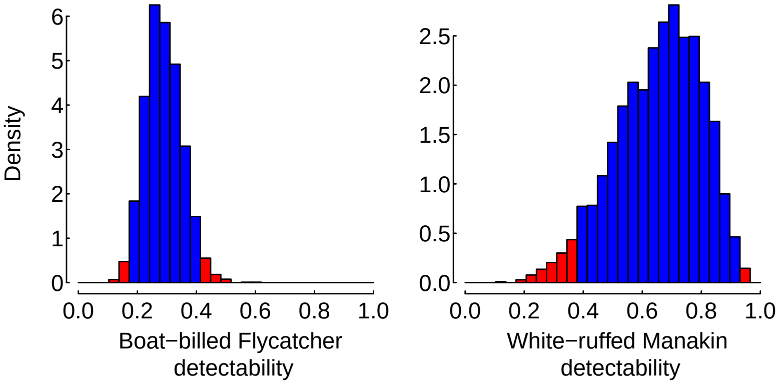

Next let's look at all the species' rates of occurence/detectability.

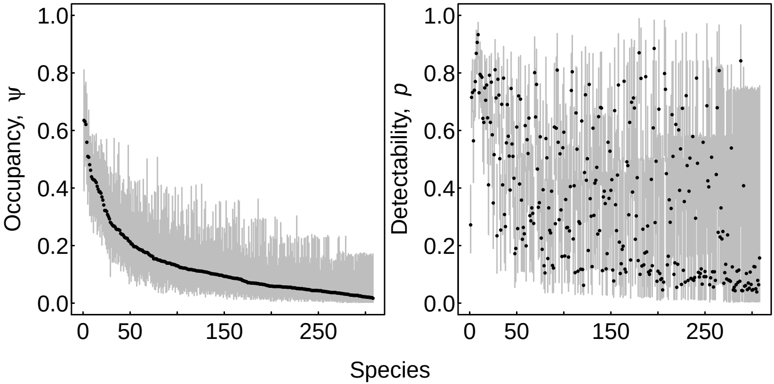

Just as we did for mark-recapture models, we could add predictors to either or both the occurrence or detectability components of the model. We will do this in workshop.

Multi-season (aka "dynamic") occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

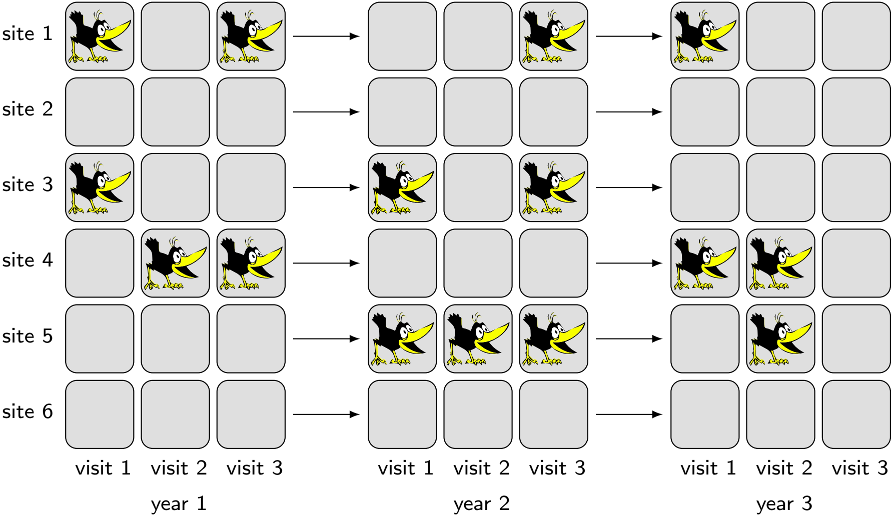

Multi-season (aka "dynamic") occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

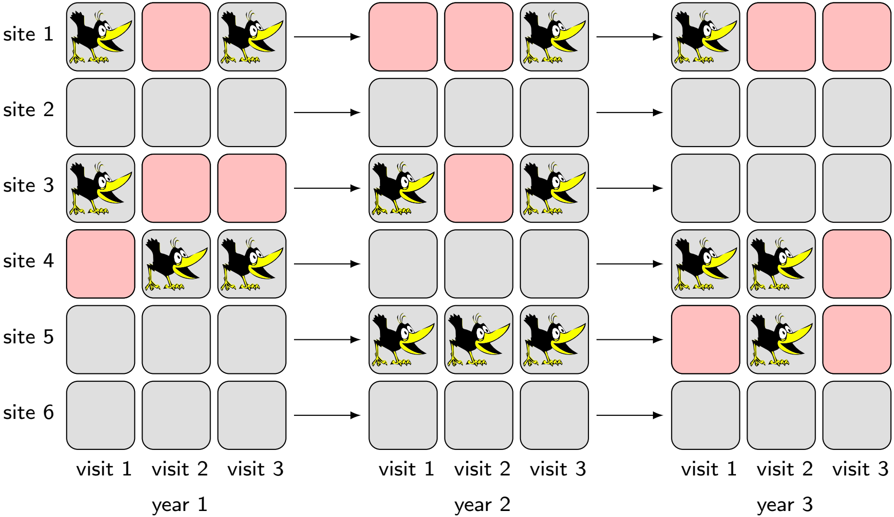

We know these are missed detections.

Multi-season (aka ”dynamic”) occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

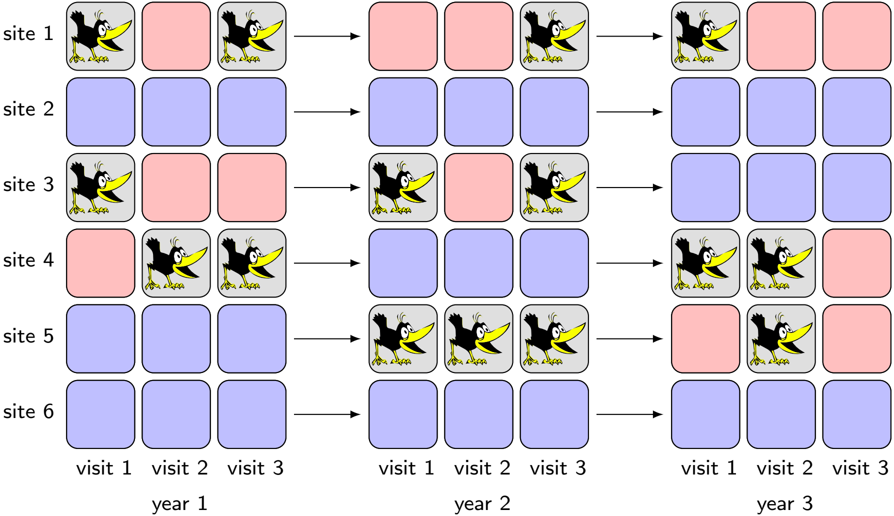

We know these are missed detections.

These could be missed detections, or the species might really be absent at these sites in these years.

Multi-season (aka ”dynamic”) occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

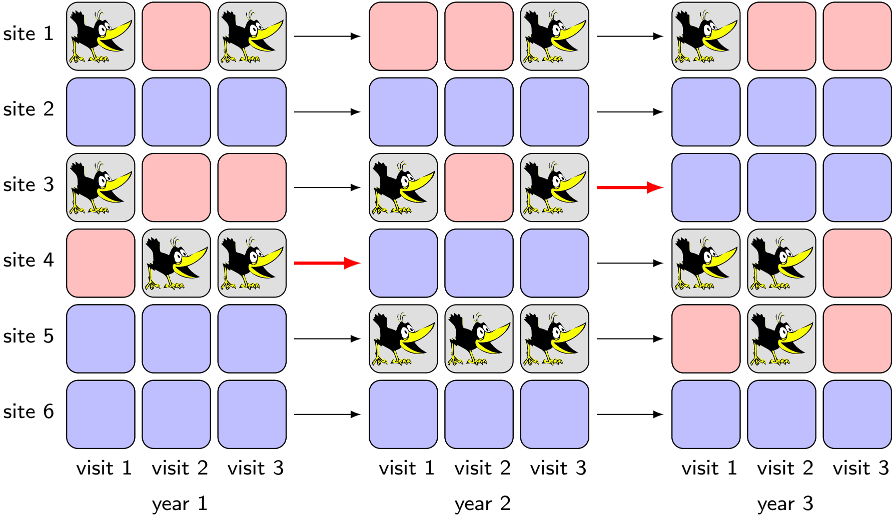

→ These **could be** instances where the species went extinct at a site.

Multi-season (aka "dynamic") occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

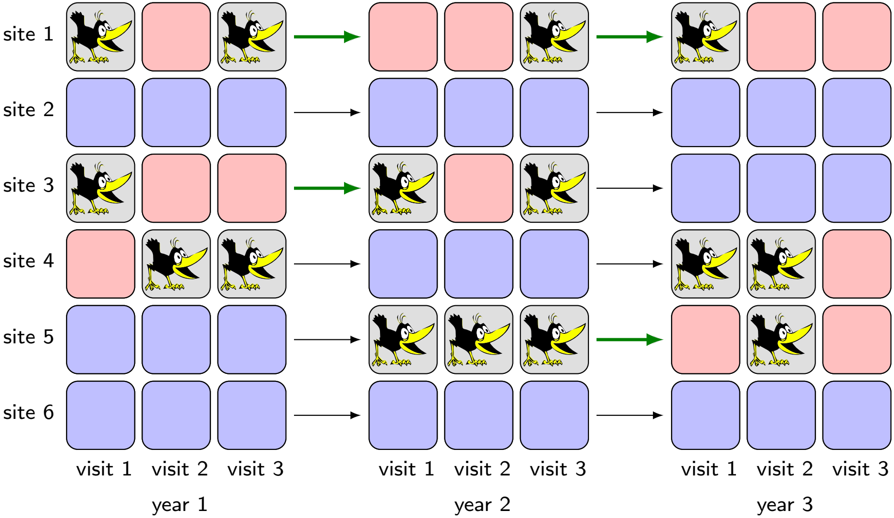

→ These **are** instances where the species "persisted" at a site (the opposite of extinction).

Multi-season (aka "dynamic") occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

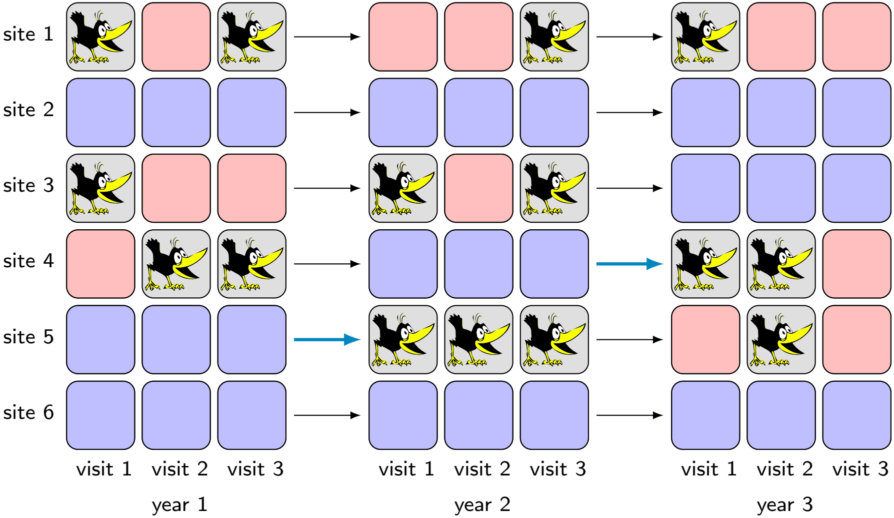

→ These **could be** instances where the species colonized a previously unoccupied site.

Multi-season (aka "dynamic") occupancy model

Suppose we survey 6 sites, 3 times per year, for 3 years.

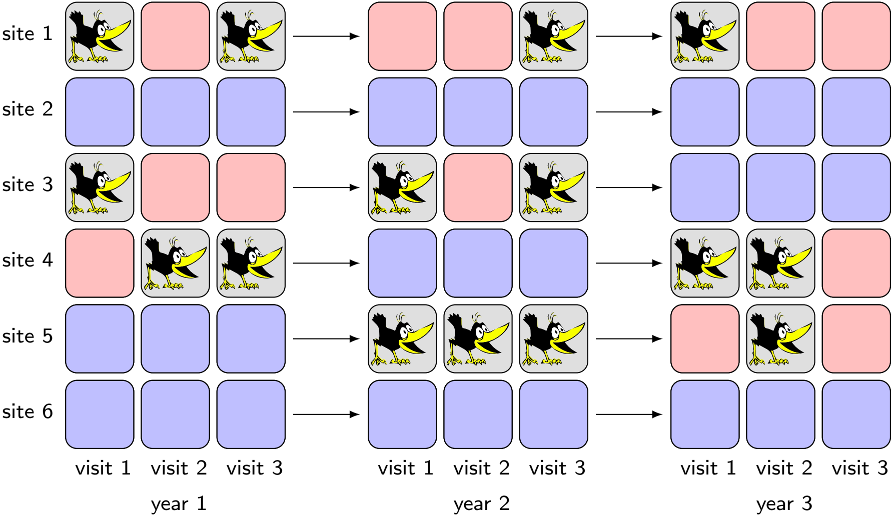

Here, we will parameterize 1) the rates at which species **persist** in patches in which they already occur and 2) the rates at which species **colonize** patches in which they don't occur. **Occupancy** then emerges as the net outcome of these two processes.

Example: Phylogenies and land-use change

Luke Frishkoff and Daniel Karp were interested in asking whether a species' phylogenetic history predicts its response to land-use change. To do this, they used the same long-term data-set that we have already examined (birds in Costa Rica, sampled across a land-use gradient). Frishkoff et al, 2014. *Science*. 345:1343-1346.

Bird surveys were conducted for 12 years at 44 sites (3 site types).

Forest reserves

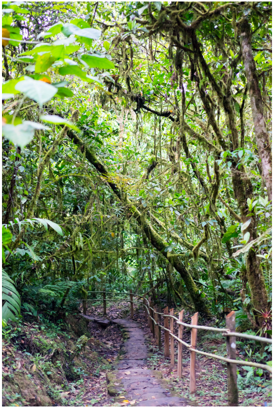

Diversified agriculture

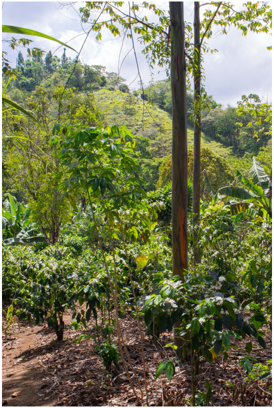

Intensive monocultures

Example: Phylogenies and land-use change

Diversified agriculture appears to support the same number of species as forests.

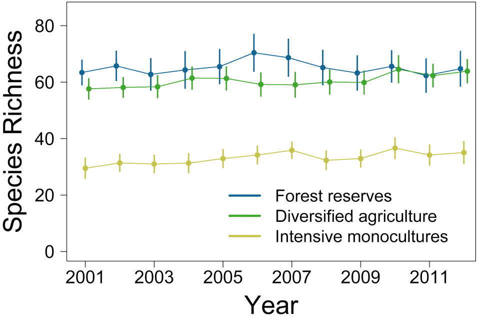

Example: Phylogenies and land-use change

The authors coded birds according to their evolutionary uniqueness (the length of that species' branch in the phylogeny) and put all this into a multi-season, multi-species occupancy model (which simply combines multi-species models with multi-season models). We'll go through the key steps.

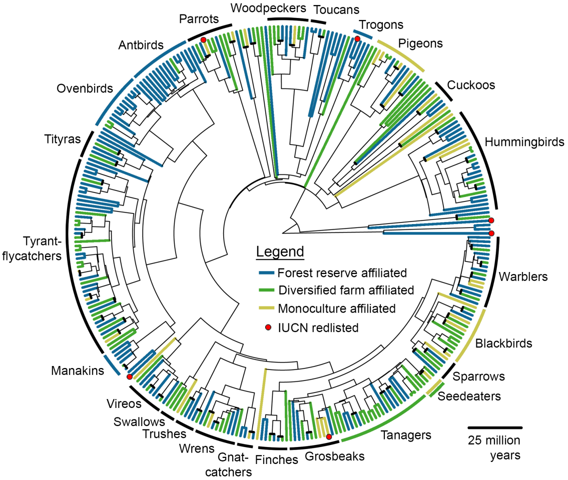

Example: Phylogenies and land-use change

Here, our occurrence data could now be stored in a 4-dimensional array, `X`, with dimensions corresponding to “site”, “year”, “visit”, and “species”.

Example: Phylogenies and land-use change

Here, our occurrence data could now be stored in a 4-dimensional array, `X`, with dimensions corresponding to “site”, “year”, “visit”, and “species”.

In addition, our inits for `Z` must also now contain a species dimension *and* a year dimension.

Example: Phylogenies and land-use change

Part of the model:

```
## Occurrence in first year
mu.psi.1[site,1,sp] ~ dunif(0,1)
psi[site,1,sp] <- mu.psi.1[site,1,sp] * region.presence[sp,region[site]]
Z[site,1,sp] ~ dbern(psi[site,1,sp])

## Subsequent years
for(yr in 1:(nyear-1)) {

  ## Rate of persistence
  logit(phi[site,yr,sp]) <-
    phi.0[yr,sp] +
      phi.intensity[intensity[site]] +
        phi.region[region[site]] +
          phi.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Rate of colonization
  logit(gam[site,yr,sp]) <-
    gam.0[yr,sp] +
      gam.intensity[intensity[site]] +
        gam.region[region[site]] +
          gam.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Occurrence
  psi[site,yr+1,sp] <-
    (Z[site,yr,sp] * phi[site,yr,sp] +
     (1-Z[site,yr,sp]) * gam[site,yr,sp]) *
      region.presence[sp,region[site]]

  Z[site,yr+1,sp] ~ dbern(psi[site,yr+1,sp])

  ## Detection
  mu.p[site,yr+1,sp] <- Z[site,yr+1,sp]*p[sp]
  for(rep in 1:nrep[site,yr+1,sp]) {
    X[site,yr+1,rep,sp] ~ dbern(mu.p[site,yr+1,sp])
  }
}
```

Example: Phylogenies and land-use change

Part of the model:

```
## Occurrence in first year
mu.psi.1[site,1,sp] ~ dunif(0,1)
psi[site,1,sp] <- mu.psi.1[site,1,sp] * region.presence[sp,region[site]]
Z[site,1,sp] ~ dbern(psi[site,1,sp])

## Subsequent years
for(yr in 1:(nyear-1)) {

  ## Rate of persistence
  logit(phi[site,yr,sp]) <-
    phi.0[yr,sp] +
      phi.intensity[intensity[site]] +
        phi.region[region[site]] +
          phi.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Rate of colonization
  logit(gam[site,yr,sp]) <-
    gam.0[yr,sp] +
      gam.intensity[intensity[site]] +
        gam.region[region[site]] +
          gam.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Occurrence
  psi[site,yr+1,sp] <-
    (Z[site,yr,sp] * phi[site,yr,sp] +
     (1-Z[site,yr,sp]) * gam[site,yr,sp]) *
      region.presence[sp,region[site]]

  Z[site,yr+1,sp] ~ dbern(psi[site,yr+1,sp])

  ## Detection
  mu.p[site,yr+1,sp] <- Z[site,yr+1,sp]*p[sp]
  for(rep in 1:nrep[site,yr+1,sp]) {
    X[site,yr+1,rep,sp] ~ dbern(mu.p[site,yr+1,sp])
  }
}
```

Our equations for persistence and colonization should look familiar. We use the logit link, because these are rates so must be between 0 and 1.

Example: Phylogenies and land-use change

Part of the model:

We must specify a prior for occurrence in the first year.

```
## Occurrence in first year
mu.psi.1[site,1,sp] ~ dunif(0,1)
psi[site,1,sp] <- mu.psi.1[site,1,sp] * region.presence[sp,region[site]]
Z[site,1,sp] ~ dbern(psi[site,1,sp])

## Subsequent years
for(yr in 1:(nyear-1)) {

  ## Rate of persistence
  logit(phi[site,yr,sp]) <-
    phi.0[yr,sp] +
      phi.intensity[intensity[site]] +
        phi.region[region[site]] +
          phi.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Rate of colonization
  logit(gam[site,yr,sp]) <-
    gam.0[yr,sp] +
      gam.intensity[intensity[site]] +
        gam.region[region[site]] +
          gam.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Occurrence
  psi[site,yr+1,sp] <-
    (Z[site,yr,sp] * phi[site,yr,sp] +
     (1-Z[site,yr,sp]) * gam[site,yr,sp]) *
       region.presence[sp,region[site]]

  Z[site,yr+1,sp] ~ dbern(psi[site,yr+1,sp])

  ## Detection
  mu.p[site,yr+1,sp] <- Z[site,yr+1,sp]*p[sp]
  for(rep in 1:nrep[site,yr+1,sp]) {
    X[site,yr+1,rep,sp] ~ dbern(mu.p[site,yr+1,sp])
  }
}
```

Example: Phylogenies and land-use change

Part of the model:

```
## Occurrence in first year
mu.psi.1[site,1,sp] ~ dunif(0,1)
psi[site,1,sp] <- mu.psi.1[site,1,sp] * region.presence[sp,region[site]]
Z[site,1,sp] ~ dbern(psi[site,1,sp])

## Subsequent years
for(yr in 1:(nyear-1)) {

  ## Rate of persistence
  logit(phi[site,yr,sp]) <-
    phi.0[yr,sp] +
      phi.intensity[intensity[site]] +
        phi.region[region[site]] +
          phi.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Rate of colonization
  logit(gam[site,yr,sp]) <-
    gam.0[yr,sp] +
      gam.intensity[intensity[site]] +
        gam.region[region[site]] +
          gam.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Occurrence
  psi[site,yr+1,sp] <-
    (Z[site,yr,sp] * phi[site,yr,sp] +
     (1-Z[site,yr,sp]) * gam[site,yr,sp]) *
      region.presence[sp,region[site]]

  Z[site,yr+1,sp] ~ dbern(psi[site,yr+1,sp])

  ## Detection
  mu.p[site,yr+1,sp] <- Z[site,yr+1,sp]*p[sp]
  for(rep in 1:nrep[site,yr+1,sp]) {
    X[site,yr+1,rep,sp] ~ dbern(mu.p[site,yr+1,sp])
  }
}
```

Occurrence each year after the first depends on rates of persistence and colonization.

Example: Phylogenies and land-use change

Part of the model:

```
## Occurrence in first year
mu.psi.1[site,1,sp] ~ dunif(0,1)
psi[site,1,sp] <- mu.psi.1[site,1,sp] * region.presence[sp,region[site]]
Z[site,1,sp] ~ dbern(psi[site,1,sp])

## Subsequent years
for(yr in 1:(nyear-1)) {

  ## Rate of persistence
  logit(phi[site,yr,sp]) <-
    phi.0[yr,sp] +
      phi.intensity[intensity[site]] +
        phi.region[region[site]] +
          phi.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Rate of colonization
  logit(gam[site,yr,sp]) <-
    gam.0[yr,sp] +
      gam.intensity[intensity[site]] +
        gam.region[region[site]] +
          gam.sp.uniq[intensity[site]] * sp.uniq[sp]

  ## Occurrence
  psi[site,yr+1,sp] <-
    (Z[site,yr,sp] * phi[site,yr,sp] +
     (1-Z[site,yr,sp]) * gam[site,yr,sp]) *
      region.presence[sp,region[site]]

  Z[site,yr+1,sp] ~ dbern(psi[site,yr+1,sp])

  ## Detection
  mu.p[site,yr+1,sp] <- Z[site,yr+1,sp]*p[sp]
  for(rep in 1:nrep[site,yr+1,sp]) {
    X[site,yr+1,rep,sp] ~ dbern(mu.p[site,yr+1,sp])
  }
}
```

Detection is as before.

As we have done for the earlier models, we would need to first look at posteriors for individual parameters, checking convergence, etc.

Here, we'll skip to the results.

Example: Phylogenies and land-use change

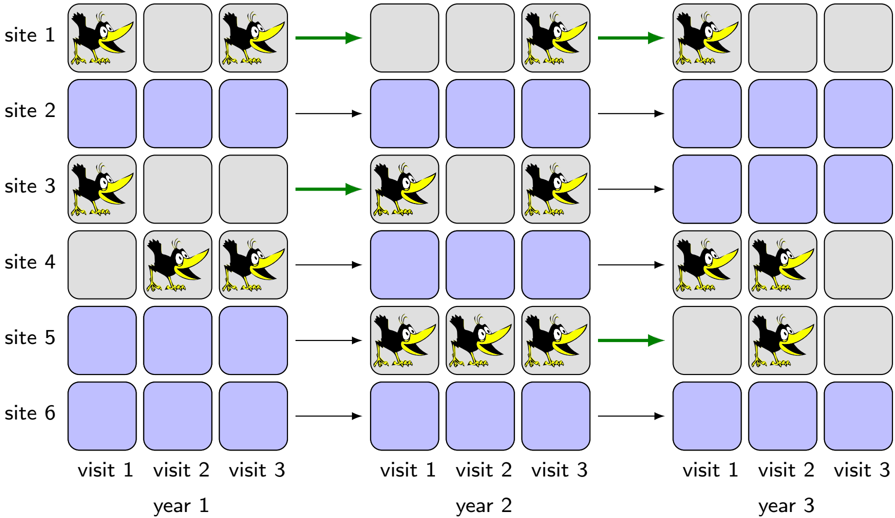

site 1

site 2

site 3

site 4

site 5

site 6

visit 1 visit 2 visit 3

year 1

visit 1 visit 2 visit 3

year 2

visit 1 visit 2 visit 3

year 3

Let's remind ourselves what **persistence** is.

Some of these other arrows may also have been cases where the species persisted; the model takes that into account.

Example: Phylogenies and land-use change

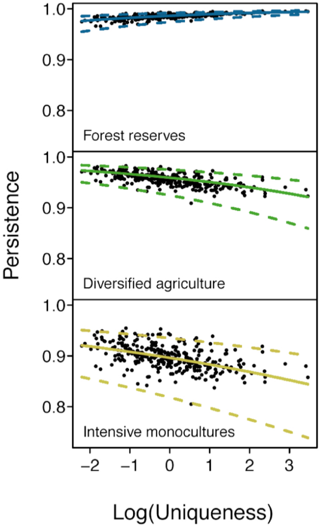

Both forms of agriculture have negative effects on persistence, particularly for evolutionarily unique birds.

Example: Phylogenies and land-use change

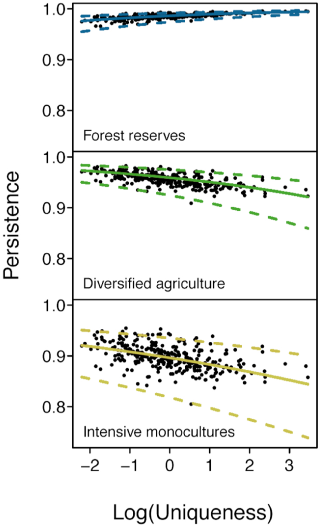

Both forms of agriculture have negative effects on persistence, particularly for evolutionarily unique birds.

Does this change how you interpret the below figure?

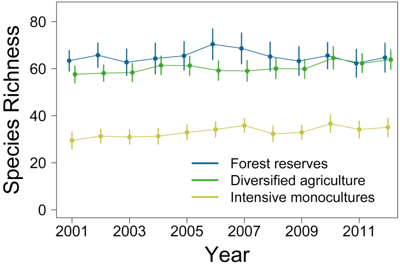

Conclusions

Occupancy models are a powerful tool that let you account for imperfect detection.

Conclusions

Occupancy models are a powerful tool that let you account for imperfect detection.

Conclusions can differ dramatically between models that do and do not account for imperfect detection.

Conclusions

Occupancy models are a powerful tool that let you account for imperfect detection.

Conclusions can differ dramatically between models that do and do not account for imperfect detection.

There are many different types of occupancy models; which to use will depend on the data you have.

Conclusions

Occupancy models are a powerful tool that let you account for imperfect detection.

Conclusions can differ dramatically between models that do and do not account for imperfect detection.

There are many different types of occupancy models; which to use will depend on the data you have.

Sometimes collecting data that is appropriate for an occupancy model doesn't take much extra effort; the key is that you have multiple visits to each site (but such "visits" could be observations by different observers or different time periods within the same trip).