A model to account for data dependency when estimating floral cover in different land use types over a season

Charlotte Baey^{1,3} · Ullrika Sahlin¹ ·
Yann Clough¹ · Henrik G. Smith^{1,2}

Received: 21 March 2017 / Revised: 6 September 2017 / Published online: 31 October 2017
© The Author(s) 2017. This article is an open access publication

Abstract We propose a model to consider data dependencies and assess spatial and temporal variability in land use specific floral coverage across landscapes. Data dependence arising from repeated measurements across the flowering season is taken into account using hierarchical Archimedean copulas, where the correlation is assumed to be stronger within seasonal periods than between periods. For each seasonal period, a bounded probability distribution is assigned to capture spatial variability in floral cover. The model uses a Bayesian approach and can assess land-use-specific floral covers by integrating experts judgments and field data. The model is applied to assess floral covers in four land use types in southern Sweden, where seasonal variability is captured by dividing the season into two periods according to winter oilseed rape flowering. Floral cover is updated using Markov Chain Monte Carlo sampling based on data from 16 landscapes and 2 years, with repeated measures available from each of the two seasonal periods. Our results indicate that considering data dependence improved the estimation of floral cover based on data observed during a season. Different copula families specifying multivariate probability distributions were tested, and no family had a consistently higher performance in the four tested land use types. Uncertainty in both mode and variability of floral cover was higher when data dependence were accounted for. Posterior modes of floral covers in semi-natural grassland were higher

Handling Editor: Bryan F. J. Manly.

✉ Charlotte Baey
charlotte.baey@centralesupelec.fr

¹ Centre for Environmental and Climate Research, Lund University, 223 62 Lund, Sweden

² Biodiversity Unit, Department of Biology, Lund University, 223 62 Lund, Sweden

³ Present Address: Laboratoire Paul Painlevé, Université de Lille 1, Cité Scientifique, 59655 Villeneuve d'Asq, France

than in field edges, but both expert's best guesses were higher than these estimates. This confirms previous findings in expert elicitation processes that experts may fail to discriminate extreme values on a bounded range. Floral cover in flower strips were estimated to be smaller/higher than semi-natural grasslands early/late in the season. The mode of floral cover in oil seed rape was estimated to be close to 100%, and higher than estimates provided by expert judgment. Floral covers for different land use classes are key parameters when quantifying floral resources at a landscape level whose assessments rely on both expert judgment and field measurements.

Keywords Bayesian inference · Copula · Data-dependency · Floral cover data

Mathematics Subject Classification 62P10 · 62F15

1 Introduction

Several studies have produced evidence of declines in diversity and abundances of some pollinators, including managed and unmanaged bees (Biesmeijer et al. 2006; Potts et al. 2010). Multiple causes contribute to this phenomenon, but habitat loss (including floral resources) is one of the most commonly cited factor (Winfree et al. 2009). Floral resource availability is therefore seen as a crucial driver in studies of abundances of such pollinator populations and the pollination services they provide (Roulston and Goodell 2011). Assessing the floral resources of a highly structured landscape composed of several habitats is a key issue to better understand ecosystem services provided by pollinators.

Landscape assessments of floral resource availability make use of spatially explicit land use information comprising land use classes, representing administrative classes rather than climatic and ecological meaningful classes. Even for a specific climatic region, the variability of floral coverage for a given land use class can be large from one site to the other, due to the type of species growing in each site, but also to differences in management and environmental conditions like solar radiation and soil moisture. An example of variability can be seen in Fig. 1, representing different semi-natural grassland sites at different spatial locations in the world. Floral resource availability also varies over a growing season, since different plant species are blooming in spring compared to early and late summer. This variability is conceptualized in floral resource assessments by dividing a season into different floral periods. This approach is used for example in landscape assessments of pollinator abundances using spatially-explicit models (Lonsdorf et al. 2009). These models take into account the spatial structure of the landscape and more specifically the distance between nests and floral resources. Thus both spatial and temporal modelling of floral resources is important when assessing floral resources availability at landscape level.

Floral cover in a particular land use class or habitat can be estimated from field measurements of floral cover. Information is up-scaled to the landscape level by considering the frequency of habitats in each land use class and shapes of land use class information. A simple model would be to average measurements obtained from repeated measurements obtained during the seasonal period of interest. However, it is likely that floral

Fig. 1 Four different sites falling under the ‘semi-natural grassland’ category in southern Sweden. *Source:* Maj Rundlöf

cover measurements from the same site are more similar than floral cover measurements from another site. Failing to account for the dependency structure between the data from the same site may lead to a misspecified model, the main consequence being a biased estimation of the error or variability of the final parameter estimates. This can have implications when testing for differences in floral covers between habitats or periods. An under-estimation of uncertainty could lead to falsely detect a difference more often than expected.

Due to the lack of data for many land use classes and habitats, a common practice is to rely on expert judgment (see e.g. [Lonsdorf et al. 2009](#); [Polce et al. 2013](#); [Koh et al. 2015](#), among others). Expert elicitation can be carried out by asking experts to provide ‘best guess’, and in some occasions a range of possible values on parameters of interest ([O’Hagan et al. 2006](#)). However, expert elicitation are sensitive to several cognitive biases, and careful attention is needed to avoid or adjust these. Biased floral cover estimates or disregard of floral cover variability could lower the quality of assessments relying on these. Approaches combining expert judgments and available field measurements are needed to strengthen assessment of floral resources.

The objective in this study is to derive a model to estimate the spatio-temporal variation in floral coverage in a mosaic landscape consisting of several different land-use types using both expert judgments and empirical data. The intention is to build a model for floral cover parameters in pollinator abundance assessments such as [Lonsdorf et al. \(2009\)](#). Temporal within-season variability is therefore handled by dividing

the season into at least two periods. The model is further developed to (1) quantify variability in land use specific floral cover across the landscape, and to (2) consider possible dependencies in data. The model is applied on floral cover estimation using field measurements from a mosaic agricultural landscape in southern Sweden. Possible dependencies between observations are accounted for using copulas, and based on the results, we compare the performance of alternative copulas to consider data-dependencies and look for possible biases in expert judgments on floral covers.

2 Material and methods

The starting point for the assessment consists of best guesses and ranges of floral covers obtained from expert elicitation. In addition to these expert judgments, observed floral covers are also available. The possible dependency between repeated measurements obtained at a site but at different occasions during the season is then considered by specifying a model for all the seasonal periods using a hierarchy of copulas to specify the multivariate distribution of observations.

2.1 Model

2.1.1 Variability in floral cover

The model assesses floral cover Φ for each land use class separately. Floral cover is defined as the proportion of a site s that is covered by flowers and is bounded between 0 and 1. Based on previous expert knowledge, a unimodal distribution was deemed appropriate to model the variability of floral coverage, and thus it was modeled as a Beta distributed random variable: $\Phi \sim \text{Beta}(a, b)$. The Beta distribution is very flexible and allows for different shapes from highly skewed to symmetric ones (see Fig. 2) and is suitable to express uncertainty in random proportions. It also includes the particular case of the uniform distribution on the $[0, 1]$ interval, when $a = b = 1$.

A season is divided into at least two seasonal periods (as in e.g. [Lonsdorf et al. 2009](#)), and floral cover for seasonal period p is assigned a period-specific Beta distribution with parameters (a_p, b_p) .

For each land use class and season there are n sites denoted by $s = 1, \dots, n$, with multiple observations of floral cover. The season is then divided into periods, and each observation made within a period is seen as a repeated measurement of the true (unobserved) floral cover in that land use class during that period (see Fig. 3 for an example where the season is divided into two periods according to the blooming period of oilseed rape).

In the example below, sites were selected by stratified randomization along a gradient of landscape heterogeneity, and within a given site, data were sampled at four occasions. This sampling scheme introduces a possible dependency between data points sampled from the same site. In the next section a model is formulated which take into account this dependency when assessing period-specific floral cover for a given land use type.

Fig. 2 Different shapes of the Beta distribution according to a and b values

Fig. 3 a) Location of the sites on the map of the province of Scania, at the southern tip of Sweden, and b) Data structure for a given land use category, where Φ_p is the true (unobserved) floral cover on period p , and y_{spr} is the r -th repeated observation of floral cover made in period p on site s

2.1.2 Dependency structure

Observations from different sites are assumed to be independent since the sampled sites were at least 7 km apart, while we allow for dependency between observations made at the same site but at different times. Further, observations are assumed to have a stronger dependency within- compared to between-seasonal periods. The latter is motivated arguing that these observations are closer in time compared to replicates from the other period. Finally, we did not keep information about the exact dates at which observations have been made, but only the corresponding periods. Since the time elapsing between two observations is different from one site to the other, this simplification allowed us to treat data from each site in a similar way. Thus a hierarchical dependence structure is assumed, where dependence is higher for data gathered during a seasonal period than between two seasonal periods (see Fig. 4).

Fig. 4 Tree-structure of an Archimedean copula: example of the $C_0(C_1(u_1, u_2), C_2(u_3, u_4))$ nested copula represented as a tree

2.1.3 A multivariate distribution with the dependency structure

A flexible way to consider the type of dependency described here is to use *copulas*. In simple terms, a copula is a multivariate probability distribution derived from the marginal distributions of each variable and a dependency structure between these variables. A key point is that a univariate marginal distribution $\mathcal{D}$ can be uniquely described by a uniform distribution on the range $[0, 1]$, using the inverse of its cumulative probability function $F_{\mathcal{D}}^{-1}$.

A copula C is a multivariate cumulative distribution function whose margins are uniform over $[0, 1]$ (see [Joe 1997](#) or [Nelsen 2006](#) for an introduction to copulas). From Sklar's theorem ([Sklar 1959](#)), we know that given a set of univariate continuous probability functions $F_i, i = 1, \dots, d$ and a copula $C : [0, 1]^d \rightarrow [0, 1]$, we can construct a (unique) multivariate probability function F with margins F_i , using:

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)). \quad (1)$$

In other words, any multivariate probability distribution can be expressed as a combination of its marginals and of the dependency structure between these marginals, given by the copula. The uniqueness of the decomposition is guaranteed whenever the margins are continuous.

Several types of copulas are available in the literature, one of the most commonly used being the class of Archimedean copulas. These copulas can be defined explicitly through the use of a generator ψ and a one- or two-dimensional parameter θ . The parameter of the generator can also be linked with Kendall's tau measure of dependence, and with tail dependence coefficients ([Joe 1997](#), see also [Hofert and Mächler \(2011\)](#) for the explicit relationships between these different quantities for some Archimedean families), which are easier to interpret.

However, the main disadvantage of this class of copulas is its exchangeability property, which implies in particular that the marginal probability distributions are identical. Hierarchical Archimedean copulas [first mentioned in [Joe \(1997\)](#), and more recently by [Savu and Trede \(2010\)](#)] are a useful tool to take into account more complex features. The basic idea is that at a given hierarchical level, copulas from the above levels are combined together via another copula. In the end, the “root” copula [using the terminology of [Hofert and Mächler \(2011\)](#)] is a copula whose arguments can be other copulas (see [Fig. 4](#) for example). If these allow for more flexibility, the resulting hierarchical structure only defines a proper copula if it verifies a sufficient nesting condition. This condition can be expressed using Kendall's tau as $\tau_i > \tau_j$, for a

Fig. 5 Pairwise distributions of a 4-dimensional random vector with Beta marginal distributions and a dependency structure defined by the hierarchical Archimedean copula of Fig. 4. Examples with Joe and Clayton families. **a** Joe family, **b** Clayton family

child copula i and its parent copula j . In other words, dependence should be stronger between elements gathered by the child copula than with the others. For example, the copula in Fig. 4 is constructed to allow dependence between u_1 and u_2 to be stronger than between u_1 and u_3 . This is the same type of structure we assumed for our data, with a stronger dependence within the periods, each of them being a child copula, than between the periods, which are then gathered at a higher level using a parent copula.

Thus, an Archimedean copula is defined through a generator function corresponding to different dependency structures. Figure 5 show simulated pairwise distributions of a four-dimensional vector whose dependency is given by the hierarchical Archimedean copula defined in Fig. 4. The choice of a copula family can thus have an influence on the model as it corresponds to different dependency structures, e.g. in tail dependence.

2.2 Parameter estimation

The model was implemented in a Bayesian framework to be able to consider expert's judgments and learn from data. The Bayesian approach quantify uncertainty in parameters and model predictions by subjective probabilities, which can be used to construct credible intervals representing uncertainty in floral cover relevant statistics such as the mean, median and spread or to account for uncertainty when sampling variability in floral cover in simulation studies at the landscape level.

We let y_{psr} denote the r -th observation of floral cover in period p and site s (see also Fig. 3). To simplify model description, the model is from now on described for two seasonal periods $p = 1, 2$ with two measurements each $r = 1, 2$.

2.2.1 Bayesian formulation

A Bayesian model is derived from a likelihood and a prior. The likelihood specify the data generating process given the model parameters. A prior is a probability dis-

Fig. 6 Directed acyclic graph (DAG) representing the structure of the model for observations made at a given site and for a given land-use category. Observations are represented inside the square, parameters of the model are represented in circles and hyperparameters in squares. Dashed rectangles represent one Archimedean copula

tribution specifying uncertainty in parameters before seeing the data. Parameters are updated using Bayes rule, resulting in posterior distributions.

To make it simpler to use for expert elicitation, a re-parametrized version of the Beta distribution is used, which define floral cover by its mode and the so-called sample size, defined respectively as $\mu = (a - 1)/(a + b - 2)$ and $s = a + b$, for a $\text{Beta}(a, b)$. In the sequel, we use either $\text{Beta}(a, b)$ or $\text{Beta}(\mu, s)$ to refer to the corresponding Beta distribution. Note that the use of this parameterization relies on the assumption that the mode does exist, and in particular on the assumption that $a > 1$ and $b > 1$.

Likelihood We denote by $\mathbf{y}_s = (y_{s11}, y_{s12}, y_{s21}, y_{s22})$ the set of all the observations made at site s . The likelihood of the model is given by (see also Fig. 6):

$$\begin{aligned} \mathbf{y}_s \mid \theta &\sim C_B(C_W(u_{s11}, u_{s12}), C_W(u_{s21}, u_{s22})) \\ u_{spr} &= F_p(y_{spr}) \quad \text{where } F_p = \text{cdf}(\text{Beta}(\mu_p, s_p)), \end{aligned} \quad (2)$$

where C_W and C_B are the copulas expressing respectively within- and between-period dependency with corresponding Kendall's tau τ_W and τ_B , and where F_p is the cumulative distribution function of a Beta distribution of parameters mode and sample

size (μ_p, s_p) . For a given land use category, the parameter vector is then defined as $\theta = (\mu_1, s_1, \mu_2, s_2, \tau_W, \tau_B)$.

Priors To define priors for the components of θ , we use the relationship between the parameters of a bivariate Archimedean copula and Kendall's tau measure of association. Uniform priors are assigned for all parameters.

Lower and upper bounds for the uniform priors $\pi(\mu_1)$ and $\pi(\mu_2)$ on the modes of the Beta distributions are based on bounds elicited by the experts:

$$\mu_1 \sim \mathcal{U}([m_1, M_1]), \quad \mu_2 \sim \mathcal{U}([m_2, M_2]), \quad (3)$$

where the quantities m_1, m_2, M_1 and M_2 differ for each habitat and are given in Table 1.

Lower and upper bounds for the uniform priors $\pi(s_1)$ and $\pi(s_2)$ on the sample sizes of the Beta distribution are chosen in order to ensure that the support of the prior is large enough:

$$s_1 \sim \mathcal{U}([2, 200]) \quad s_2 \sim \mathcal{U}([2, 200]). \quad (4)$$

We also assume a positive association between the observations, and to ensure that the hierarchical Archimedean copula properly defined a copula, we need to ensure that $\tau_B < \tau_W$. The latter condition is controlled by the conditional distribution of τ_B given τ_W . The last two priors are thus equivalent to assuming a uniform distribution of (τ_W, τ_B) over $D = \{(u, v) \in \mathbb{R}^2, 0 \leq v \leq u \leq 1\}$. This results in the following priors $\pi(\tau_W)$ and $\pi(\tau_B | \tau_W)$ for the copula parameters:

$$\tau_W \sim \mathcal{U}([0, 1]), \quad \tau_B | \tau_W \sim \mathcal{U}([0, \tau_W]) \quad (5)$$

Finally, the joint prior distribution for parameter θ is:

$$\pi(\theta) = \pi(\mu_1)\pi(\mu_2)\pi(s_1)\pi(s_2)\pi(\tau_B | \tau_W)\pi(\tau_W).$$

Posterior The posterior distribution of the parameters is defined as:

$$\pi(\theta | \mathbf{y}) \propto L(\mathbf{y} | \theta) \pi(\theta) = \prod_{s=1}^n L(\mathbf{y}_s | \theta) \pi(\theta), \quad (6)$$

where L is the likelihood of the data and π is the prior distribution.

2.2.2 Bayesian updating

Posterior distributions were estimated using a Markov Chain Monte Carlo (MCMC) algorithm. We used the **HAC** package in R to compute the density of the copula and the (unnormalized) posterior distribution at each iteration of the MCMC algorithm. To improve the efficiency of the algorithm, a hybrid strategy for MCMC sampling based

on a local adaptive scheme in a Robbins–Monro stochastic approximation framework, and inspired from [Andrieu and Thoms \(2008\)](#), was implemented.

More precisely, at each iteration of the algorithm, and given the current value of the chain θ^m , a candidate $\tilde{\theta}_k$ was generated for the k -th component of θ , using a univariate adaptive proposal distribution $q_m^k(\cdot|\theta^m)$. This candidate was then accepted with probability

$$\alpha_k(\theta_m, \tilde{\theta}) = \min \left(1, \frac{\pi(\tilde{\theta}_k | \mathbf{y}) q_m^k(\theta_m^k | \tilde{\theta})}{\pi(\theta_m^k | \mathbf{y}) q_m^k(\tilde{\theta}_k | \theta_m)} \right). \quad (7)$$

When the proposal distribution is symmetric, the above ratio can be simplified and expressed only using the ratio of posteriors. In our case, the proposal was chosen to be a random walk, either Gaussian or uniform. As mentioned previously, a hybrid strategy was adopted, which means that at each iteration of the algorithm, we randomly selected either the uniform or the Gaussian random walk.

The basic idea of the adaptive algorithm is to adapt the variance of this random walk in order to reach an optimal acceptance probability ([Andrieu and Thoms 2008](#)). More precisely, the variance of the random walk proposal is computed as the product of an empirical estimation of the Markov Chain's covariance matrix and a correction coefficient. This correction coefficient increases when the computed acceptance probability is too large, suggesting that the current variance of the proposal distribution is too small since too many candidates are accepted, which is often due to small moves of the chain. On the other hand, it decreases when the computed acceptance probability is too small, suggesting on the contrary that the current variance of the proposal distribution is too large, resulting in a high rejection rate.

More precisely, following [Andrieu and Thoms \(2008\)](#), a stochastic approximation procedure was adopted to update the algorithm's parameters, using a time step γ_m . If we denote by λ_m^k and ν_m^k respectively the current mean and variance of the Markov Chain for the k -th component of θ at iteration m of the algorithm, and by κ_m^k the corresponding correction coefficient, we have:

$$\lambda_{m+1}^k = \lambda_m^k + \gamma_{m+1}(\theta_k^{m+1} - \lambda_m^k) \quad (8)$$

$$\nu_{m+1}^k = \nu_m^k + \gamma_{m+1}((\theta_k^{m+1} - \lambda_m^k)^2 - \nu_m^k) \quad (9)$$

$$\log \kappa_{m+1}^k = \log \kappa_m^k + \gamma_{m+1}(\alpha_k(\theta_m, \tilde{\theta}) - \alpha^*). \quad (10)$$

The support of the uniform proposal at iteration m of the MCMC algorithm is then defined through its mean and variance as $[\theta_m^k - \sqrt{3\kappa_m^k \nu_m^k}, \theta_m^k + \sqrt{3\kappa_m^k \nu_m^k}]$. The choice of the time step γ_m is crucial to ensure the convergence of the stochastic approximation, and should verify $\sum_m \gamma_m = \infty$ and $\sum_m \gamma_m^2 < \infty$, which is the case for sequences of the type $\gamma_m = 1/m^a$, for some $a \in (1/2, 1]$.

Several approaches can then be used to take into account the strong constrained supports of parameters μ_1, μ_2, τ_W and τ_B . A first one is to use unconstrained proposals, resulting in the rejection of candidates falling outside of the support thanks to the prior

Fig. 7 Accounting for the constraints on τ_B when using a uniform proposal distribution to generate a candidate $\tilde{\tau}_B$ at iteration m of the MCMC algorithm. Depending on the adaptive variance at iteration m , $\kappa_m^B \nu_m^B$, and on the candidate value $\tilde{\tau}_W$, the lower and upper bounds might need to be truncated. Case (a): the current adaptive variance lead to a uniform distribution which allow for the sufficient nesting condition to be fulfilled, so that no truncation is needed. Cases (b): the current adaptive variance lead to a uniform distribution with an upper bound which is too high and could lead to candidates $\tilde{\tau}_B$ that would not fulfilled the nesting condition: in this case, the upper bound of the proposal distribution is truncated. **a** No truncation is needed, **b** The upper bound is truncated to $\tau_B^m = \tilde{\tau}_W$

distributions densities in (7). However, several issues may arise in this case. First, it is somehow inefficient since we allow the algorithm to generate candidates for which we already know that they will be rejected, without the need to compute the acceptance probability. In our context, since the algorithm is already time consuming, it seems wiser to make a better use of the resources by generating candidates within the support of the target distribution. Second, when using an adaptive scheme, rejecting candidates which fall outside of the targeted support will have an influence on the variance of the proposal, which is adapted at each step of the algorithm according to the current acceptance rate.

Another approach, which was adopted here, is to include the constraints in the proposal. The adaptive proposal is a bit more complex than in the unconstrained case, and in particular it is no longer symmetric which means that the ratio in (7) can no longer be simplified. When q_m^k is chosen to be a Gaussian distribution, the variance of the distribution is adapted at each iteration of the MCMC algorithm as described above in (8), and then constraints on the parameters are taken into account using truncated versions of the Gaussian distribution. When q_m^k is chosen to be a uniform distribution, the variance of the distribution is also adapted as in (8), and then the constrained are accounted for by truncating the upper and lower bounds when necessary (see Fig. 7).

Moreover, to ensure that the sufficient nesting condition is verified, τ_B was sampled conditionally on the value of $\tilde{\tau}_W$ and two cases were accounted for: (i) if $\tau_B^m < \tilde{\tau}_W$, a random walk (either Gaussian or uniform) was used, taking into account the constraint that $\tilde{\tau}_B$ should remain smaller than $\tilde{\tau}_W$ (an illustration in the case of the uniform proposal is given in Fig. 7), and (ii) if $\tau_B^m \geq \tilde{\tau}_W$, then we set $\tilde{\tau}_B = \tilde{\tau}_W$. The algorithm is described in more details in Algorithm 1.

2.3 Convergence in the Bayesian updating

Convergence and properties of the chains were assessed by performance measures derived from the **coda** package in R. In particular, acceptance rates, autocorrelation and cumulative quantile plots were surveyed. Convergence diagnoses were made using

Algorithm 1 MCMC algorithm for the Bayesian updating

Initialize the parameters $\theta^0 = (\mu_1^0, \mu_2^0, s_1^0, s_2^0, \tau_B^0, \tau_W^0)^t$

Initialize the algorithm variables $\lambda_1^0, \dots, \lambda_K^0, v_1^0, \dots, v_K^0$ and $\kappa_1^0, \dots, \kappa_K^0$, where $K = 6$ is the number of components of θ .

for $m = 1, \dots, M$ **do**

1. Randomly select a type of random walk, either uniform or Gaussian
2. Generate a candidate for each component of θ .
 - If the Gaussian random walk was selected at the previous step, let:

$$\begin{aligned}\tilde{\mu}_1 &\sim \mathcal{N}(\mu_1^m, \kappa_1^m v_1^m, [0, 1]) \\ \tilde{\mu}_2 &\sim \mathcal{N}(\mu_2^m, \kappa_2^m v_2^m, [0, 1]) \\ \tilde{s}_1 &\sim \mathcal{N}(s_1^m, \kappa_3^m v_3^m) \\ \tilde{s}_2 &\sim \mathcal{N}(s_2^m, \kappa_4^m v_4^m) \\ \tilde{\tau}_W &\sim \mathcal{N}(\tau_W^m, \kappa_5^m v_5^m, [0, 1]) \\ \tilde{\tau}_B \mid \tilde{\tau}_W &\sim \begin{cases} \mathcal{N}(\tau_B^m, \kappa_6^m v_6^m, [0, \tilde{\tau}_W]) & \text{if } \tau_B^m \leq \tilde{\tau}_W \\ \tilde{\tau}_W & \text{otherwise} \end{cases}\end{aligned}$$

where $\mathcal{N}(\mu, \sigma)$ is the Gaussian distribution of mean μ and variance σ , and $\mathcal{N}(\mu, \sigma, [a, b])$ its truncated counterpart on the interval $[a, b]$.

- If the uniform random walk was selected at the previous step, let:

$$\begin{aligned}\tilde{\theta}_k &\sim \mathcal{U}(l_k^m, u_k^m), \quad k = 1, \dots, 5 \\ \tilde{\tau}_B \mid \tilde{\tau}_W &\sim \begin{cases} \mathcal{U}(l_6^m, u_6^m) & \text{if } \tau_B^m \leq \tilde{\tau}_W \\ \tilde{\tau}_W & \text{otherwise} \end{cases}\end{aligned}$$

where the lower and upper bounds of the uniform proposal l_k^m and u_k^m are given by:

$$\begin{aligned}l_1^m &= \max(0, \mu_1^m - \sqrt{3\kappa_1^m v_1^m}), & u_1^m &= \min(1, \mu_1^m + \sqrt{3\kappa_1^m v_1^m}) \\ l_2^m &= \max(0, \mu_2^m - \sqrt{3\kappa_2^m v_2^m}), & u_2^m &= \min(1, \mu_2^m + \sqrt{3\kappa_2^m v_2^m}) \\ l_3^m &= s_1^m - \sqrt{3\kappa_3^m v_3^m}, & u_3^m &= s_1^m + \sqrt{3\kappa_3^m v_3^m} \\ l_4^m &= s_2^m - \sqrt{3\kappa_4^m v_4^m}, & u_4^m &= s_2^m + \sqrt{3\kappa_4^m v_4^m} \\ l_5^m &= \max(0, \tau_W^m - \sqrt{3\kappa_5^m v_5^m}), & u_5^m &= \min(1, \tau_W^m + \sqrt{3\kappa_5^m v_5^m}) \\ l_6^m &= \max(0, \tau_W^m - \sqrt{3\kappa_6^m v_6^m}), & u_6^m &= \min(\tilde{\tau}_W - \tau_B^m, \tau_W^m + \sqrt{3\kappa_6^m v_6^m}),\end{aligned}$$

3. Set:

$$\theta_k^{m+1} = \begin{cases} \tilde{\theta}_k & \text{with probability } \alpha(\theta^m, \tilde{\theta}) \\ \theta_k^m & \text{with probability } 1 - \alpha(\theta^m, \tilde{\theta}), \end{cases} \quad (11)$$

where $\alpha(\theta^m, \tilde{\theta})$ is defined in equation (7).

4. Update

$$\begin{aligned}\lambda_{m+1}^k &= \lambda_m^k + \gamma_{m+1}(\theta_k^{m+1} - \lambda_m^k) \\ v_{m+1}^k &= v_m^k + \gamma_{m+1}((\theta_k^{m+1} - \lambda_m^k)^2 - v_m^k) \\ \log \kappa_{m+1}^k &= \log \kappa_m^k + \gamma_{m+1}(\alpha(\theta^m, \tilde{\theta}) - \alpha^*),\end{aligned}$$

end for

Heidelberger and Welch’s criterion (Heidelberger and Welch 1983) implemented in the `heidel.diag` function of the `coda` package.

2.4 The model applied on floral cover assessment

2.4.1 Land use information

Floral cover was estimated on four pollinator relevant land use types in southern Sweden: oilseed rape, which is a mass-flowering crop and attractive for foraging bees, semi-natural grassland, uncultivated field borders, which are important nesting habitats with a low density of floral resources, and flower strips, which are created with the purpose of providing floral resources to pollinators in intensive agricultural landscapes. The season is divided into two periods, an early period corresponding to the flowering of oilseed rape, and a late period after this flowering period to the end of the growing season.

By definition of the two floral periods, oilseed rape is only blooming in the first period. Therefore floral cover of oilseed rape was only assessed during the first seasonal period, and the model was simplified to only take into account within-period dependency for this land use class.

2.4.2 Floral cover from expert elicitation

A group of experts were asked to assess the floral cover of different habitat types in Sweden. The experts had been asked to provide a best guess value, along with lower and upper bounds representing a range of variability or their uncertainty about which values floral cover may take.

Prior model of floral covers, i.e. the parameter μ of the land use class and period specific Beta distributions (see Table 1), were partly based on the expert’s judgments. The lower bound of priors were set to zero for the land use semi-natural grassland, flower strips and field edges, while the upper bounds were set to the experts’ judgments on upper bounds. For oil seed rape, the experts judged floral cover to be high, and therefore the prior was set to the range 0.5–1. The best guess on floral cover provided by the experts were not used here to justify the comparison between the posterior mode and the best guess provided by experts. However, it is possible to include the best guess when specifying the prior.

Table 1 Lower and upper bounds for the uniform priors on the modes of the Beta distribution for each floral period

Habitat	Period 1	Period 2
Semi-natural grassland	[0–0.20]	[0–0.15]
Field edge	[0–0.20]	[0–0.20]
Flower strips	[0–0.30]	[0–0.30]
Oilseed rape field	[0.5–1]	–

2.5 Floral cover data

Data on floral coverage were obtained from a large field study in 2011 and 2012 in 16 sites in Scania (southernmost part of Sweden) (Holzschuh et al. in press). The study consisted in 16 landscape sectors with either a high or a low cover of oilseed rape and an orthogonal gradient of semi-natural habitat. Within each sector, different habitats which were assumed to contain flowers were sampled. In 2012, sown wildflower strips were also planted and surveyed in half of the 16 sites. Observations obtained in 2011 and 2012 were considered as independent replicates, i.e. as if coming from different sites.

The percentage of flowers was recorded on two occasions during the flowering period of oilseed rape, i.e. between middle May and beginning of June, and on two other occasions after OSR flowering period, i.e. between late June and beginning of July. For each habitat, observations were made along $150 \times 1 \text{ m}^2$ transects, and then averaged at the square meter level.

Finally, for semi-natural grassland, field edges and oilseed rape fields, we have 32 vectors of 2 repeated observations for each period, and for flower strips we have 8 vector of 2 repeated observations for each period. In other words, $n = 32$ for semi-natural grasslands, field edges and oilseed rape fields, and $n = 8$ for flower strips.

2.6 Model selection

We compared the results obtained with different copula families, in order to identify the one that best suits the data, but also to check the robustness of the results concerning the Beta distribution parameters. The latter are of main interest, and it is important to be sure that the choice of the copula family does not have a too strong influence on the marginal distribution parameters. The DIC and WAIC criteria were used to compare the different competing models, i.e. those corresponding to different copulas family, but also the model where we assumed independence between the observations.

The DIC criterion [Deviance Information Criterion, Spiegelhalter et al. (2002)] is defined as:

$$DIC = -2 \left(\log L(\mathbf{y} \mid \hat{\theta}) - p_{DIC} \right), \quad (12)$$

where $L(\mathbf{y} \mid \hat{\theta})$ is the likelihood function evaluated at $\hat{\theta}$, a Bayesian point-estimate of θ , and where p_{DIC} is the effective number of parameters, measuring the model complexity. This quantity can be defined in two different ways (see Gelman et al. 2014), either as:

$$p_{DIC} = 2 \left(\log L(\mathbf{y} \mid \hat{\theta}) - \mathbb{E}_{\theta}(\log L(\mathbf{y} \mid \theta)) \right),$$

or as:

$$p_{DIC,2} = 2 \text{Var}_{\theta} \left(\log L(\mathbf{y} \mid \hat{\theta}) \right),$$

where the expectations are taken over the posterior distribution of θ . Using the mode a posteriori as the point estimate for θ , p_{DIC} can be negative if the posterior mode

is far from the posterior mean. On the other hand, the definition of $p_{DIC,2}$ ensures that is always positive, but results in a quantity that is less stable numerically. For this reason, we used the p_{DIC} to compute the DIC criterion. $\hat{\theta}$ is usually taken to be the posterior mean, but as suggested in Spiegelhalter et al. (2002), it can also be the mode or the median of the posterior distribution. Here, DIC was derived from the median of the MCMC sample.

While the DIC relies on a point estimate for θ , the WAIC criterion [Watanabe-Akaike Information Criterion, also called the Widely Applicable Information criterion, Watanabe (2010)] is a fully Bayesian criterion, since it relies on the whole posterior distribution of θ . Its expression is given by:

$$WAIC = -2 \left(\sum_{i=1}^n \log \int L(y_i | \theta) \pi(\theta | \mathbf{y}) d\theta - p_{WAIC} \right), \quad (13)$$

where the first term between the parentheses is the log pointwise predictive density, and p_{WAIC} is the effective number of parameters, accounting for the complexity of the model. Here again, two alternative definitions are possible for this term (Gelman et al. 2014):

$$p_{WAIC} = 2 \sum_{i=1}^n [\log \mathbb{E}_{\theta}(L(y_i | \theta)) - \mathbb{E}(\log L(y_i | \theta))] \quad (14)$$

$$p_{WAIC,2} = \sum_{i=1}^n \text{Var}(\log L(y_i | \theta)). \quad (15)$$

(Gelman et al. 2014) recommend to use $p_{WAIC,2}$ in practice, due to the fact that it gives results which are closer to what one would get using leave-one-out cross validation. Both were calculated in our case.

As for the AIC and BIC criteria, one should choose the model with the smallest DIC or WAIC values.

3 Results

3.1 Convergence of the Markov Chains

We obtained similar results when running the algorithm from different starting points, suggesting that the posterior distributions are independent of the initialization of the algorithm. Acceptance rates were all comprised between 0.26 and 0.29, reflecting a good mixing of the chains. Figure 8 gives an example of the trace plot obtained when running the algorithm on the semi-natural grassland data and with the Clayton copula family. The auto correlation plots were satisfying for all variables except for τ_W , due to the strongest dependency of this parameter on the history of the chain. After 100 000 iterations, the Heidelberger and Welch's criterion indicated that every chain has reached its stationary distribution. Thus, the convergence was judged acceptable for the forthcoming analysis.

Fig. 8 Trace plots of the parameters, on the semi-natural grassland data and with the Clayton copula family

3.2 Model selection

Model comparison were performed using DIC and WAIC values obtained in each land use type and for each copula family (Table 2). Both criteria lead to the selection of the same copula families. Model comparison showed that taking into account the dependence allowed for a better fit of the data (Table 2), except for flower strips. The sample size is much smaller in this land use category, since it existed in only half of the sites and only in 2012. In this case it is more difficult to estimate the dependency structure, as shown for example by [Huard et al. \(2006\)](#). For the semi-natural grasslands, the DIC and the WAIC computed with the effective number of parameters p_{WAIC} (see Sect. 2.6 for the definition of these quantities) gave very similar values for the Clayton copula and the independent model. The difference is bigger between these two models when using the WAIC criterion computed with $p_{WAIC,2}$, and tends to favor the independent model.

When Kendall's tau is low, identifying the right copula is difficult, particularly for families sharing similar characteristics under small values of the Kendall's tau ([Huard et al. 2006](#)).

3.3 Floral cover estimates

Floral cover parameters for each land use type and each period were chosen from the copula identified by model selection (Table 3). The corresponding assessments of land

Table 2 DIC and WAIC criteria and their associated effective number of parameters, for each habitat type and each copula family. $WAIC$ is calculated using the effective number of parameters p_{WAIC} and $WAIC_2$ using $p_{WAIC,2}$. The values obtained with the different runs of the algorithm were very similar, and we give here the smallest criteria values

Habitat	Copula family	DIC	p_{DIC}	$WAIC$	p_{WAIC}	$WAIC_2$	$p_{WAIC,2}$
Semi-natural grassland	Clayton	-761.2	9.41	-759.3	11.30	-753.2	14.36
	Joe	-741.3	10.71	-739.0	12.97	-731.6	16.68
	Frank	-746.6	9.38	-743.8	12.11	-736.9	15.59
	Gumbel	-748.1	9.48	-745.7	11.81	-739.7	14.81
	Independence	-761.1	2.98	-760.0	4.06	-759.1	4.50
Flower strip	Clayton	-157.7	10.43	-159.3	8.83	-151.4	12.79
	Joe	-157.8	9.58	-157.9	9.51	-149.9	13.51
	Frank	-160.3	8.98	-161.6	7.74	-155.8	10.62
	Gumbel	-158.3	9.75	-159.0	9.0	-151.0	13.0
	Independence	-179.0	2.88	-179.0	2.83	-178.2	3.24
Field edge	Clayton	-740.1	9.40	-728.9	20.62	-711.9	29.13
	Joe	-660.2	13.44	-659.7	26.47	-642.3	35.16
	Frank	-708.4	8.70	-694.0	23.12	-677.3	31.48
	Gumbel	-675.1	9.76	-673.5	21.37	-660.7	27.78
	Independence	-697.6	3.00	-688.8	11.81	-685.5	13.49
Oilseed rape field	Clayton	-88.3	-1.51	-80.7	2.48	-80.5	2.60
	Joe	-157.2	-28.42	-95.1	2.85	-94.8	2.97
	Frank	-97.9	-5.76	-81.4	2.92	-81.1	3.05
	Gumbel	-153.3	-29.33	-89.6	2.65	-89.3	2.77
	Independence	-84.0	1.58	-83.1	2.48	-82.9	2.59

Table 3 Median and 95% credible interval of floral cover for four land use types

Parameter	Land use			
	Semi-natural grassland	Flower strip	Field edge	Oilseed rape field
μ_1 (%)	0.32 [0.04–0.78]	0.18 [0.02–0.51]	0.16 [0.02–0.52]	99.5 [98.3–99.9]
μ_2 (%)	0.17 [0.02–0.43]	1.13 [0.13–3.11]	0.06 [0.01–0.22]	–
s_1	49.5 [32.9–71.7]	120 [52.2–228]	21.2 [13.7–31.7]	4.70 [3.90–5.64]
s_2	89.4 [58.9–130]	24 [9.31–51.4]	46.7 [29.6–68.5]	–
τ_W	0.15 [0.03–0.26]	0.15 [0.02–0.35]	0.17 [0.08–0.26]	0.12 [0.07–0.19]
τ_B	0.026 [0.003–0.07]	0.01 [0.002–0.05]	0.09 [0.03–0.16]	–
Copula	Clayton	Frank	Clayton	Joe
Chain size	100,000	100,000	100,000	100,000

Fig. 9 Comparison of the spatial variability of floral cover across a landscape for four land use types and two seasonal periods estimated by our model, and provided by expert judgment when available. The colored regions are 95% prediction intervals based on the posterior distributions of the parameters, and the shape of these regions reflect the different shapes obtained for the Beta distribution, for different values of μ_1 , s_1 , μ_2 and s_2 . The solid dark green line is the prediction obtained with the posterior median, and the solid light green line is the triangle distribution obtained using expert knowledge. The ticks on the x-axis mark data on floral cover

use and seasonal period specific floral cover are plotted in Fig. 9, together with expert judgments.

In semi-natural grassland and field edges, both the mode and the variance of floral coverage are higher in the first period. This may reflect a loss of floral resources with grazing and/or cutting, as well as to the ending of the main flowering period of species

Fig. 10 Comparison of the posterior densities obtained with each copula family in semi-natural grasslands

Table 4 Median and 95% credible interval of floral cover when assuming independence between the observations

Parameter	Land use			
	Semi-natural grassland	Flower strip	Field edge	Oilseed rape field
μ_1 (%)	0.16 [0.017–0.44]	0.07 [0.007–0.23]	0.05 [0.004–0.18]	99.7 [98.9–100]
μ_2 (%)	0.09 [0.01–0.25]	0.53 [0.053–1.45]	0.02 [0.002–0.07]	–
s_1	48.3 [38.7–60.4]	115 [74.1–167.8]	29.4 [23.9–35.7]	5.33 [4.50–6.28]
s_2	88.3 [70.3–110.3]	27.0 [17.6–40.1]	63.9 [51.9–77.4]	–

that account for a high cover early in the season, such as dandelions. Conversely, the mode and the variance of floral coverage in flower strips is higher during the second period. Floral cover in flower strips were estimated to be smaller/higher than semi-natural grasslands early/late in the season. Floral cover in oil seed rape were estimated to be close to 100%. There is also a very large variability in the floral cover of oilseed rape fields, which might be due to differences in varieties and in sowing dates from one field to the other.

Results were similar across copula families, especially for the main parameters of interest, i.e. the parameters of the Beta distribution, which means that our results are robust with respect to the choice of the copula family. Figure 10 illustrates the posterior distributions obtained for semi-natural grasslands with each copula family.

While the results are consistent across the different copula families, floral cover estimates were different when no data dependencies were considered (Table 4).

As illustrated in Fig. 10, uncertainty associated to each parameter is lower in the independent case, due to the fact that we artificially increase the quantity of information of the sample when assuming that data are independent. We observed a larger difference between the modes of the posterior distributions for parameters μ_1 and μ_2 , whether we accounted for the dependence or not.

3.4 Expert bias

Expert judgment on the mode of floral cover are represented by triangular distributions using the minimum, maximum and best guess values provided by the experts (Fig. 9). For small floral coverage values, parameters provided by expert judgment are higher than those obtained with our model. For semi-natural grasslands in the first period the minimum value provided by the experts is even higher than the mode of the Beta distribution estimated with our approach. This indicates that the expert elicitation suffer from low discriminative power, i.e. that low proportions are judged as higher than they are (O'Hagan et al. 2006). In semi natural grassland and field edges, experts tended to overestimate variability, i.e. they believe there are more flowers than there actually are. On the contrary in the case of oilseed rape fields, the expert's judgements were less biased but more narrow compared to the posterior variability (Fig. 7).

4 Discussion

In this paper, we proposed a model to estimate the floral cover in different land use types by taking into account the dependency between data observed at repeated times in the season. Our results showed that failing to account for the dependency between observations can have an impact on parameter estimates.

The four land use categories represented in our study correspond to very low (semi-natural grassland, flower strips, field edges) or very high (oilseed rape fields) flower coverage, and our results tend to show that in such extreme situations it might be easier for experts to provide over estimated or under estimated values. For these types of habitats, using empirical evidence and accounting for the potential dependency between the data is therefore crucial in order to provide accurate estimates of the corresponding floral coverage.

The model is based on a partition of the season into two periods, one early and one late season, in order to capture the within-season dynamics. The methodology can be easily generalized to more than two periods, for example by considering other nodes on the hierarchical copula. The framework of hierarchical Archimedean copulas is very flexible, and one can also think of more complex hierarchical structures to account for other types of data dependency. In particular, if in our case the sampling sites were far enough from each others to assume spatial independence, this might not always be the case. Copulas can also be used to model spatial dependence, for example using the natural graph representation of neighboring points to construct a vine copula (Gräler and Pebesma 2011). Other time-dependent structures can also be implemented, for example conditional copula taking into account the time elapsing between two observations.

The focus has been here on floral cover to assess floral resources at the landscape scale. The variability of floral coverage in each period was modeled using a Beta distribution, while the dependency structure was accounted for using a nested Archimedean copula. The method is also applicable on other types of data on floral cover, e.g. using more trait-based approaches to quantify floral resources (Hicks et al. 2016). In our case, the marginal distributions were fixed as Beta distributions, but one can also think of more complex models where not only different copula families but also different marginal distributions are compared [see for example Silva and Lopes (2008)]. The model can then be modified to estimate land-use associated nectar and pollen resources across the season or species abundances. Another interesting perspective would be to consider a non-parametric model, where both the marginal distributions and the copula are estimated empirically.

To simplify the model, we assumed that the association between observations was stronger within than between periods, so that a hierarchical structure as described in Fig. 4 could be used. We also assumed that the parameters of the two copulas linking within-period observations were identical. This is a strong assumption, and the small sample sizes make a proper validation of this assumption difficult using empirical estimates of Kendall's taus. However, our main interest here was not to accurately estimate the copulas parameters, but rather the Beta distributions parameters. Moreover, we obtained consistent and robust results with the different copula families, corresponding to different dependency structures. This enhance the confidence on the results.

Compared to the results obtained by expert elicitation, we are able to distinguish both the variability of the floral cover in the landscapes and uncertainty in model parameters and predictions of land use specific floral resources. Any bias discovered in the expert elicitation can be due to the procedure of elicitation rather than to the judgment of the experts itself. There are several reasons why expert judgments can be biased and there are ways to consider these. One issue is the differentiation between asking for uncertainty in the mode of floral cover, and asking for the full range of variability in floral cover. Experts are actually asked to provide a range of a mixture of aleatory and epistemic uncertainty, which makes the elicitation harder (O'Hagan et al. 2006). A difference between expert judgments and data can also be explained by observation biases. For example, the floral cover data in semi-natural habitats had actually been sampled from the parts of selected landscapes where floral cover was, from visual inspection, judged as relatively high. This sampling strategy deliberately introduced an over estimation of floral cover in semi-natural habitats, which can be motivated to get an upper estimate of floral cover in habitats with in general, a low amount of flowers. In this study, this bias in floral data from semi-natural habitats does not change the conclusions, since the floral cover updated from data were still lower than the judgment provided by experts (Section 2.2).

Evaluating the availability of floral resources is a key issue when studying the dynamics of pollinators populations in the landscape. This is particularly true when dealing with spatially explicit pollination models, which use land use type and seasonal period-specific floral resources as inputs. An accurate estimation of these quantities is then crucial in order to enhance the model predictions, in particular for the main land use types encountered in the landscapes of interest.

Acknowledgements The research leading to these results has received funding from the European Community's Seventh Framework Programme under Grant Agreement No. 311781, the LIBERATION Project (www.fp7liberation.eu) (to HS), the FORMAS-funded Swedish research environment SAPES (to HS), the FORMAS-funded project Scaling up uncertain environmental evidence(219-2013-1217) (to US) and the 2013–2014 Biodiv-ER-sA/FACCE-JPI joint call for research proposals, with the national funders ANR, BMBF, FORMAS, FWF, MINECO, NWO and PT-DLR, the ECODEAL project (to YC).

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

- Andrieu C, Thoms J (2008) A tutorial on adaptive MCMC. *Stat Comput* 18:343–373
- Biesmeijer JC, Roberts SPM, Reemer M, Ohlemüller R, Edwards M, Peeters T, Schaffers AP, Potts SG, Kleukers R, Thomas CD, Settele J, Kunin WE (2006) Parallel declines in pollinators and insect-pollinated plants in Britain and The Netherlands. *Science* 313(5785):351–354
- Gelman A, Hwang J, Vehtari A (2014) Understanding predictive information criteria for Bayesian models. *Stat Comput* 24(6):997–1016
- Gräler B, Pebesma E (2011) The pair-copula construction for spatial data: a new approach to model spatial dependency. *Procedia Environ Sci* 7:206–211
- Heidelberger P, Welch PD (1983) Simulation run length control in the presence of an initial transient. *Oper Res* 31(6):1109–1144
- Hicks DM, Ouvrard P, Baldock KCR, Baude M, Goddard MA, Kunin WE, Mitschunas N, Memmott J, Morse H, Nikolitsi M, Osgathorpe LM, Potts SG, Robertson KM, Scott AV, Sinclair F, Westbury DB, Stone GN (2016) Food for pollinators: quantifying the nectar and pollen resources of urban flower meadows. *PLoS ONE* 11(6):1–37
- Hofert M, Mächler M (2011) Nested archimedean copulas meet R: The nacopula package. *J Stat Softw* 39(9):1–20
- Holzschuh A, Dainese M, Gonzalez-Varo JP, Mudri-Stojni S, Riedinger V, Rundlf M, Scheper J, Wickens JB, Wickens VJ, Bommarco R, Kleijn D, Potts SG, Roberts SP, Smith HG, Vil M, Vuji A, Steffan-Dewenter I (2016) Mass-flowering crops dilute pollinator abundance in agricultural landscapes across Europe. *Ecol Lett* 19(10):1228–1236
- Huard D, Évin G, Favre AC (2006) Bayesian copula selection. *Comput Stat Data Anal* 51(2):809–822
- Joe H (1997) Multivariate models and multivariate dependence concepts. CRC Press, Boca Raton
- Koh I, Lonsdorf EV, Williams NM, Brittain C, Isaacs R, Gibbs J, Ricketts TH (2015) Modeling the status, trends, and impacts of wild bee abundance in the United States. *Proc Natl Acad Sci* 113(1):140–145
- Lonsdorf EV, Kremen C, Ricketts TH, Winfree R, Williams N, Greenleaf S (2009) Modelling pollination services across agricultural landscapes. *Ann Bot* 103(9):1589–1600
- Nelsen R (2006) An introduction to copulas. Lecture notes in statistics, 2nd edn. Springer, Berlin
- O'Hagan A, Buck C, Daneshkhah A, Eiser J, Garthwaite P, Jenkinson D, Oakley J, Rakow T (2006) Uncertain judgements: eliciting experts' probabilities. Statistics in practice. Wiley, New York
- Polce C, Termansen M, Aguirre-Gutiérrez J, Boatman ND, Budge GE, Crowe A, Garratt MP, Pietravalle S, Potts SG, Ja Ramirez, Somerville KE, Biesmeijer JC (2013) Species Distribution models for crop pollination: a modelling framework applied to Great Britain. *PLoS ONE* 8(10):e76308
- Potts SG, Biesmeijer JC, Kremen C, Neumann P, Schweiger O, Kunin WE (2010) Global pollinator declines: trends, impacts and drivers. *Trends Ecol Evol* 25(6):345–353
- Roulston TH, Goodell K (2011) The role of resources and risks in regulating wild bee populations. *Annu Rev Entomol* 56(1):293–312
- Savu C, Tiede M (2010) Hierarchies of archimedean copulas. *Quant Finance* 10(3):295–304
- Silva RdS, Lopes HF (2008) Copula, marginal distributions and model selection: a bayesian note. *Stat Comput* 18(3):313–320
- Sklar A (1959) Fonctions de répartition à n dimensions et leurs marges. Publications de l'Institut de Statistique de l'Université de Paris 8:229–231

- Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A (2002) Bayesian measures of model complexity and fit. *J R Stat Soc Ser B (Stat Methodol)* 64(4):583–639
- Watanabe S (2010) Asymptotic equivalence of bayes cross validation and widely applicable information criterion in singular learning theory. *J Mach Learn Res* 11:3571–3594
- Winfree R, Aguilar R, Vázquez DP, LeBuhn G, Aizen MA (2009) A meta-analysis of bees' responses to anthropogenic disturbance. *Ecology* 90(8):2068–2076

Charlotte Baey received her Ph.D. in Applied Mathematics from Ecole Centrale Paris, France in 2014. She is an assistant professor in University of Lille, France, and is mainly interested in the application of statistics to biological science.

Ullrika Sahlin received her Ph.D. in Ecology from Lund University, Sweden in 2010. Her research focus on the treatment of uncertainty in risk assessment.

Yann Clough is a researcher with the Centre for Environmental and Climate Research at Lund University. He uses empirical, statistical and process-based modelling approaches to understand and predict the relationship between landscape-scale land use by humans, species communities and ecosystem services such as pollination and pest control.

Henrik G. Smith is professor in Animal Ecology at the Department of Biology, Lund University. His research focuses on effects of agricultural intensification and climate change on meta-communities of interacting organisms and the resulting consequences for ecosystem services.