Efficient sampling designs to assess biodiversity spatial autocorrelation : should we go fractal ?

F. Laroche *,[1]

[1]UMR 1201 Dynafor, Univ Toulouse, INRAE, INPT, EI PURPAN, Castanet-Tolosan, France

Abstract

1. Evaluating the autocorrelation range of species distribution in space is necessary for many applied ecological questions like implementing protected area networks or monitoring programs. The autocorrelation range can be inferred from observations, based on a spatial sampling design. However, there is a trade-off between estimating the autocorrelation range of a species distribution and estimating fixed effects affecting the mean species abundance or occupancy among sites. The random sampling design is considered as a good heuristic to estimate autocorrelation range, for it contains contrasted pairwise distances that cover a wide array of possible range values. The grid design is viewed as a better choice for estimating fixed effects, for it eliminates small pairwise distances that are more prone to pseudo-replication. Mixing random and grid ('hybrid' designs) has been presented as a way to navigate between both conflicting objectives. We postulated that fractal designs — which have

*Corresponding author: fabien.laroche@inrae.fr

a self-similarity property and well-identified scales — could make a compromise, for they preserve some regularity reminiscent of grid at each scale, but also browse a wide array of possible autocorrelation range values across scales.

2. We used maximum likelihood estimation within an optimal design of experiments approach to compare the accuracy of hybrid and fractal designs at estimating the fixed intercept and the autocorrelation range of a spatial field of values.

3. We found that hybrid designs were Pareto-optimal intermediary strategies between grid and random for small autocorrelation range values only, while classic grid design should always be preferred when autocorrelation is large. Fractal designs yielded Pareto-optimal strategies specifically good at estimating small autocorrelation ranges. However, they were generally not Pareto-optimal for higher values of autocorrelation range. At last, when the surveyed area could be changed, random designs were sufficient to reach the Pareto front in any context.

4. Fractal designs seemed relevant when specifically aiming at improving the estimation of small autocorrelation ranges in a fixed surveyed area with a limited sampling budget, which is a quite circumscribed scenario. However, they may prove more clearly advantageous to analyse biodiversity patterns when covariates are included in the analysis and ecological processes differ among spatial scales.

Keywords: beta-diversity; distance-decay; fractal; maximum likelihood; model-based inference; optimal design; sampling design; spatial autocorrelation

Introduction

Autocorrelation has a double status in the study of biodiversity patterns (Legendre, 1993). It is often seen as a nuisance, generating biases in regression models that seek to link covariates to spatial patterns of biodiversity (Lennon, 2000). Many techniques to control these undesirable effects are available, and now well popularized among ecologists (Dormann et al., 2007). However, spatial autocorrelation may also be viewed as the signature of endogeneous process driving biodiversity patterns (McGill, 2010). In particular, it is often interpreted through the prism of limited dispersal. For instance, auto-regressive modelling of species occupancy in metapopulation ecology (ter Braak et al., 1998; Bardos et al., 2015; Prugh, 2009; Ranius et al., 2010) or isolation by distance patterns on neutral markers in population genetics (Ouborg et al., 1999; Vekemans and Hardy, 2004; Manel and Holderegger, 2013) are often used to draw conclusion on species colonization or dispersal abilities. From this perspective, the accurate assessment of autocorrelation range has important implications in terms of conservation biology, especially to assess the functional connectivity of habitat networks (Tischendorf and Fahrig, 2000) or build efficient biodiversity monitoring strategies (Rhodes and Jonzén, 2011).

Few studies focused on efficient designs to accurately estimate autocorrelation range of biodiversity. In a simulation study, Bijleveld et al. (2012) showed that a grid design was the best choice to estimate spatial or temporal trends on the mean of a target field of values while random design was better at estimating autocorrelation parameters. The authors further showed that a hybrid strategy, mixing randomly chosen sites with a grid, stood as a Pareto-optimal solution on the trade-off between the conflicting objectives (i.e. changing to other designs necessarily generated performance loss on either objective). Following a distinct line of research, Marsh and Ewers (2013) suggested that fractal sam-

pling designs may be an efficient option to study the distance-decay pattern of β-diversity (Nekola and White, 1999), which can be seen as way to assess the autocorrelation range of species composition among communities. Fractal designs are characterized by a self-similar property (Mandelbrot, 1983; Falconer, 2003): sub-parts of the design look like a contraction of the total design, and natural spatial 'scales' can thus be distinguished (see Fig. 1A). Thanks to this property, a single fractal design can cover contrasted spatial scales with a low sampling effort compared to other sampling strategies, which may offer a practical way to studying autocorrelation over a broad set of possible ranges. Based on a non-parametric study of distance-decay patterns, the authors found that fractal designs lead to estimating higher values of autocorrelation range than an intensive control design, while other classic strategies (regular grid and random design) tended to yield lower values of autocorrelation range than the control. These results suggest that fractal designs have distinct properties, but precise conclusions are hindered by the fact that the model used by Marsh and Ewers (2013) to generate the data is a complex mixture of ecological scenarios that does not have a well-defined statistical formulation, which impeded a clear definition of estimation errors.

Here, we aimed at comparing fractal designs with random, grid and hybrid strategies with respect to ther ability to quantify autocorrelation range versus fixed effects on the mean of a spatial random field. We turned towards the framework of optimal design of experiments (Müller et al., 2012), which has been repeatedly used to build and compare designs of temporal (Archaux and Bergès, 2008) or spatio-temporal [see (Hooten et al., 2009) and references therein] biodiversity surveys. However, it has quite rarely been applied to the specific problem of quantifying spatial autocorrelation. A noticeable exception on that matter is the study by Müller (2007), which focused on the problem of

detecting autocorrelation with a test using Moran index (Moran, 1950). However, they did not quantify the corresponding range. Here, we focused on estimation error when simultaneously estimating the mean (i.e. a fixed intercept) and the autocorrelation range of a spatial field. We specifically considered the maximum likelihood estimation framework which offers a powerful heuristic to theoretically explore the estimation accuracy of sampling designs through the analysis of the inverse Fisher matrix (Abt and Welch, 1998). Zhu and Stein (2005) used this approach to numerically search for sampling designs able to recover autocorrelation parameters. They found that designs showing the best global performance at estimating autocorrelation parameters differ from random design, and tend to conciliate aggregated points at the center of the surveyed area with points scattered close to the frontier. Such designs might be viewed as harbouring distinct scales and might be well approached by fractal designs, hence reinforcing the interest of explicitly assessing fractal design performance.

When comparing random, grid, hybrid and fractal designs, we had two expectations grounded on the litterature previously cited: (i) hybrid designs should consitute a continuous set of intermediary Pareto-optimal designs between grid and random designs, meaning that when the proportion of random points increases from 0 (grid design) to 1 (random design), the accuracy of the mean estimate of the random field should decrease while the accuracy of the autocorrelation range estimate should increase; (ii) fractal designs should be better than other designs at estimating small autocorrelation ranges when they are built to harbour contrasted scales, hence creating new Pareto-optimal solution focused on autocorrelation range estimation.

Methods

Spatial sampling designs

All spatial sampling designs harboured $N = 27$ sampling points (the effect of larger N is discussed later on). Sampling points are spread within an area of study shaped as an equilateral triangle with a side length of $L = \sqrt{3}$ distance units. Random designs were generated by sampling each plot independently within the triangle with a homogeneous spatial density. Grid designs were obtained by first generating a triangular grid matching the area of study with mesh size equal to $L/6$ distance units, hence obtaining 28 sampling points, and then removing one point at random. Hybrid designs were characterized by a parameter p, the proportion of sites that are randomly positioned in the area of study. A hybrid design was obtained by building a grid design, selecting Np sites at random within it and resampling their new position at random in the area of study. Note that $p = 0$ yields the grid design while $p = 1$ yields the random design. Here we consider the $N + 1 = 28$ values of $p \in \{0, 1/N, 2/N, ..., 1\}$.

Following Marsh and Ewers (2013), we simulated fractal designs using an iterated function systems (Falconer, 2003) based on three similarities of the complex plane: $S_k(z) = \rho z + (1 - \rho)e^{\frac{2ik\pi}{3} - \frac{i\pi}{6}}$ for $k \in \{0, 1, 2\}$. A sampling design is obtained by iterating three times the system starting from a seed at the center of the area, hence yielding a sampling design with $N = 3^3 = 27$ plots. We varied the parameter ρ across designs. The parameter ρ drives the ratio between the size of a part of the design and the size of the larger, autosimilar set of plots it belongs to. The values of ρ considered in the study are : $\rho = x\sqrt{3}/(2 + \sqrt{3})$ with 240 distinct x values evenly spaced on a log-scale from $x = 10^{-1.5}$ to $x = 1$. We call x the 'contraction parameter' of fractal design below. Note that $x > 1$ would generate a sampling design with overlapping sub-components, which we considered as an undesirable property. The largest

value $x = 1$ yields a sampling design that is a subsample of the regular grid with mesh size of c.a. $L/10$.

Examples of designs are presented in figure 1.

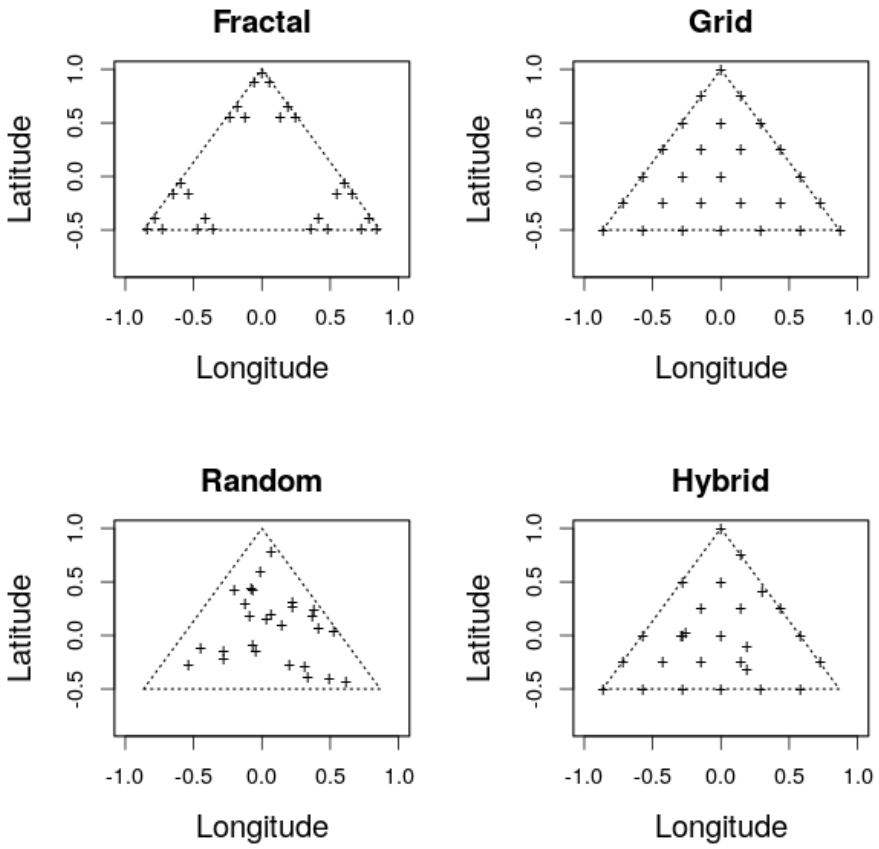

Figure 1: Examples of the four types of design considered in our study. The triangle dashed area is the area of study. Crosses show the position of the 27 sampling points. The fractal design presented here is generated with a contraction coefficient $x = 2/3$. The hybrid design is generated using a proportion of random sites $p = 5/27 \approx 0.19$.

Gaussian random field model

To study the error of estimation associated to each type of design depicted above, we assumed that the vector of observations at each sampling points $\boldsymbol{Z} = (Z_1, ..., Z_N)$ is taken from a Gaussian random field with an exponential variogram (Cressie, 1993) without nugget effect. Formally, it means that there exist $\mu \in \mathbb{R}$ and $\sigma, a_s \in \mathbb{R}^{+*}$ such that:

$$\begin{aligned} &\forall i \in \{1, ..., N\}, Z_i \sim \mathcal{N}(\mu, \sigma^2) \\ &\forall i, j \in \{1, ..., N\}^2, \mathbb{E}\left[(Z_i - Z_j)^2\right] = 2\sigma^2 \left(1 - e^{-\frac{d_{ij}}{a_s}}\right) \end{aligned} \tag{1}$$

where d_{ij} is the distance between sampling points i and j and $\mathbb{E}[.]$ denotes the expectation of a random variable. The covariance between Z_i and Z_j is $\mathrm{Cov}[Z_i, Z_j]$, is:

$$\mathrm{Cov}[Z_i, Z_j] = \sigma^2 e^{-\frac{d_{ij}}{a_s}} \tag{2}$$

which renders the simulation of $\boldsymbol{Z}$ straightforward.

Parameter a_s corresponds to a characteristic length of the autocorrelation, and we call it 'autocorrelation range' below. The parameter μ corresponds to the mean of the random field. Although one could include effects of covariates on this parameter, we followed previous studies [e.g. (Zhu and Stein, 2005)] and considered a simple fixed-intercept model, where the mean is a single parameter constant across space that one wants to accurately estimate it.

We considered 160 distinct a_s values evenly spaced on a log-scale between 10^{-3} and $10^{2.5}$. As we will show below, the values of μ and σ did not affect estimation errors considered in our analysis, and we could therefore set $\mu = 0$ and $\sigma = 1$ without loss of generality.

Estimation variance of maximum likelihood estimates

We assimilate the problem of assessing autocorrelation range to accurately estimate a_s and the problem of assessing the mean of the field of values to accurately estimating $\alpha = e^{\mu}$. We considered the exponential mean to compare estimation error of parameters that are defined on the same domain $\mathbb{R}^{+*}$. Estimation error on a parameter θ ($= \alpha$ or a_s) is quantified through the relative root mean square error:

$$RRMSE(\theta) = \sqrt{\mathbb{E}\left[(\hat{\theta} - \theta)^2\right]}/\theta$$

The statistical model used to estimate a_s and α matches the one used to generate the data (i.e. we assume no error on model specification):

$$\boldsymbol{Z} = \mu\mathbf{1} + \boldsymbol{E}$$

where $\mathbf{1}$ is a N-dimensional vector with all coordinates equal to 1, and $\boldsymbol{E}$ is a N-dimensional gaussian vector with mean $\mathbf{0}$ and variance-covariance matrix Σ following the exponential model presented in (2). Parameters of this model can be summarized in a vector $\boldsymbol{\theta} = (\alpha, \sigma, a_s)$. We focused on the maximum likelihood estimate $\hat{\boldsymbol{\theta}} = (\hat{\alpha}, \hat{\sigma}, \hat{a}_s)$ of $\boldsymbol{\theta}$.

In the context of stationary Gaussian random fields without nugget, it is known that the diagonal terms of $\mathcal{I}(\boldsymbol{\theta})^{-1}$, where $\mathcal{I}(\boldsymbol{\theta})$ is the Fisher information matrix of the model with true parameters $\boldsymbol{\theta}$, yield a qualitatively good approximation of the quadratic error on parameters in $\boldsymbol{\theta}$. By 'qualitatively', we mean that it allows to correctly rank designs according to their accuracy, even for moderate sample sizes (Abt and Welch, 1998; Zhu and Stein, 2005). We therefore use the diagonal terms of $\mathcal{I}(\boldsymbol{\theta})^{-1}$ as a theoretical approximation of quadratic error of $\hat{\boldsymbol{\theta}}$ below.

Results

Derivation of Fisher information matrix and predicted errors

The Fisher information matrix associated to parameters $\boldsymbol{\theta} = (\alpha, \sigma, a_s)$ in model (1) is [see Article S1 in Supporting Information, section 1; Zhu and Stein (2005); Müller (2007)]:

$$\mathcal{I}(\boldsymbol{\theta}) = \begin{pmatrix} \frac{1}{\alpha^2}\mathbf{1}'\Sigma^{-1}\mathbf{1} & 0 & 0 \\ 0 & \frac{N}{2\sigma^4} & \frac{1}{2\sigma^2}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}) \\ 0 & \frac{1}{2\sigma^2}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}) & \frac{1}{2}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}) \end{pmatrix} \tag{3}$$

From equation (3), one obtains the relative root mean squared error associated to $\hat{a}_s$ and $\hat{\alpha}$ (see Article S1 in Supporting Information, section 1):

$$\begin{aligned} \mathrm{RRMSE}(\alpha) &= \frac{1}{\sqrt{\mathbf{1}'\Sigma^{-1}\mathbf{1}}} \\ \mathrm{RRMSE}(a_s) &= \frac{1}{a_s}\sqrt{\frac{2}{\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}) - \frac{1}{N}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})}} \end{aligned} \tag{4}$$

Equation (4) implied in particular that RRMSEs did not depend on α (or μ), hence justifying that we set $\alpha = 1$ throughout the study without loss of generality. Recalling that we also set $\sigma = 1$, equation (4) shows that $\mathrm{RRMSE}(a_s)$ did not depend on σ but that $\mathrm{RRMSE}(\alpha)$ was proportional to σ, which implies that quantitative predictions on $\mathrm{RRMSE}(\alpha)$ should vary when changing the value of σ. However, we were mostly interested in the ranking of designs, which should remain identical up to a multiplicative constant. Therefore, setting $\sigma = 1$ did not imply any loss of generality on our results either.

Theoretical analysis of asymptotic errors

When $a_s \to 0$ Our theoretical analysis (see Article S1 in Supporting Information, section 2) yielded that RRMSE(a_s) should increase towards $+\infty$ as a_s becomes small, irrespective of considered design. The increase is quite abrupt, proportional to $a_s/d_{\min} \times \exp(d_{\min}/a_s)$ where $d_{\min}$ is the smallest distance among two distinct sampling points. Although the proportionality constant depends on the sampling design, the feature of designs with strongest effect on RRMSE(a_s) when a_s becomes arbitrarily small is $d_{\min}$: designs with smaller $d_{\min}$ should yield markedly smaller RRMSE(a_s). The grid design tends to maximize $d_{\min}$ for a given sampling effort N (see Article S1 in Supporting Information, section 3) and should thus yield consistently higher RRMSE(a_s) than other designs as $a_s \to 0$. Fractal design can harbour arbitrarily small $d_{\min}$ values by decreasing contraction parameter x. As a result, there should exist a threshold on x below which fractal designs yield lower RRMSE(a_s) than hybrid sampling designs, and thus become Pareto-optimal.

In the meantime, RRMSE(α) should converge to $\sigma/\sqrt{N}$ irrespective of the sampling design. This corresponds to the expected standard error on the mean when sampling points are independent. When $\sigma^2 = 1$ and $N = 27$, this yields RRMSE(α) ≈ 0.19.

Both results suggest that fractal design with low contraction parameters may exclude hybrid designs from the Pareto front, since the accuracy at estimating autocorrelation range should become the major difference among designs.

When $a_s \to +\infty$ RRMSE(a_s) converges towards $\sqrt{2N/(N-1)}$ (≈ 1.44 when $N = 27$), irrespective of the sampling design (see Article S1 in Supporting Information, section 2). In the meantime, RRMSE(α) converges to σ ($= 1$ in our example), irrespective of the sampling design. This is the expected result for a single observation, hence illustrating the fact that all the sampling points are

perfectly correlated. Both results suggest that all the sampling designs should converge towards very similar performance as $a_s \to +\infty$, hence rendering their ordination impossible.

Numerical analysis of Pareto fronts for hybrid designs

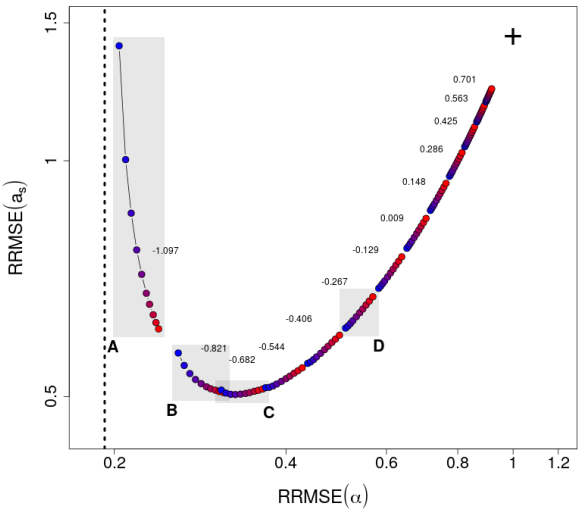

Figure 2: Relative root mean square error of estimation of exponential mean (α) and autocorrelation range (a_s) for hybrid designs along a gradient of a_s values. For each a_s value, we present a line of dots showing the RRMSEs for the 28 hybrid designs. Dots color indicates the value of p, increasing from blue ($p = 0$; grid design) to red ($p = 1$; random design).The lines of dots shift towards the right as a_s increases, following a U-shaped global pattern. We presented results for 13 values a_s out of the 160 values explored, for readibility. Those values are reported above each line of dots. The vertical dotted line shows the predicted asymptote when $a_s \to 0$, the cross shows the predicted limit when $a_s \to +\infty$. We obtained the same qualitative pattern when assessing the quadratic error of estimation through a Monte-Carlo approach (see Article S1 in Supporting Information, section 4). The four grey rectangles show the lines of hybrid designs that are further considered in fig. 3.

In line with the asymptotic study when $a_s \to 0$ detailed above, the predicted RRMSE(a_s) of hybrid designs rapidly increased at the lower margin of explored

a_s values. It exceeded standard numerical precision of computers and software when considering designs with high $d_{\min}$ (e.g. grid design). Below, we focused our analysis on the range of a_s values for which RRMSE(a_s) could be computed for all designs. This led us to ignore a_s values smaller than $10^{-2.101}$ (26 values out of the 160 initially considered).

Some patterns were common to all hybrid designs (fig. 2). The RRMSE(α) increased with the autocorrelation range, starting from the expected value of 0.19 towards the predicted value upper limit of 1. This increase was quite expected : stronger autocorrelation increases pseudo-replication and makes the mean of the field harder to estimate. The RRMSE(a_s) showed a non monotonic profile first decreasing from infinity, then increasing again towards the expected limit of 1.44.

Increasing the degree of randomness p within hybrid designs consistently increased the RRMSE(α) along the gradient of a_s (see fig. 2). By contrast, the ordination of RRMSE(a_s) among hybrid designs with various degree of randomness p changed as a_s increased. For small a_s values ($a_s \leq 10^{-0.786}$), increasing p decreased the RRMSE(a_s). Therefore, any hybrid design along the gradient of p was a Pareto-optimal strategy (see figs. 3A, 3B, 4A). For intermediate a_s values ($10^{-0.752} \leq a_s \leq 10^{-0.579}$), the RRMSE($a_s$) harboured a U-shaped pattern as p increased. Therefore, there existed a threshold on a_s above which increasing p too much did not lead to Pareto-optimal strategies anymore (see figs. 3C, 4A). For larger a_s values ($a_s \geq 10^{-0.544}$), the RRMSE(a_s) increased with p, making grid design ($p = 0$) the only Pareto-optimal strategy among hybrid designs (see also figs. 3D, 4A). We retrieved those three types of patterns for small, intermediary and large values of a_s when estimating RRMSEs from simulations in a Monte-Carlo approach (see Article S1 in Supporting Information, section 4).

Considering the intermediate range of a_s values where the effect of p grad-

ually changes from all hybrid designs being Pareto-optimal to grid design only, we observed that it contained the a_s value corresponding to the mesh size of the grid design ($\sqrt{3}/6 \approx 10^{-0.540}$). In practice, the transition might therefore occur when the autocorrelation range reach values close to the mesh size of the grid design. For simulated RRMSEs, the three ranges of a_s values associated to distinct patterns seemed to be positioned later on the autocorrelation gradient (see Article S1 in Supporting Information, section 4), but the rule of thumb that transition occurs for autocorrelation range values close to mesh size was not rejected.

Numerical comparison of fractal designs to the Pareto front of hybrid designs

For small a_s values ($a_s \leq 10^{-1.893}$), fractal designs with intermediate to high contraction parameter ($10^{-0.969} \leq x \leq 1$; fig. 4B) excluded all the hybrid designs from the Pareto front except the pure grid design ($p = 0$; fig. 4A), which remained the most efficient design to estimate the mean of the field. We had the theoretical conjecture — derived from our theoretical analysis of asymptotic errors above — that fractal designs with low contraction parameters x could become unilaterally better that hybrid designs at small autocorrelation range, because the performance of all designs at estimating α should become similar while fractal design with low x should be better at estimating a_s. The observed exclusion of most hybrid designs can be seen as a result of this process. However, the conjecture was not fully verified over the range of a_s values explored: first fractal designs with very low contraction parameters ($x < 10^{-0.969}$) were not Pareto-optimal among fractal designs, second grid design still persisted as a Pareto-optimal option. Maybe smaller a_s values would have matched the theoretical conjecture better, but as explained above, they could not be explored

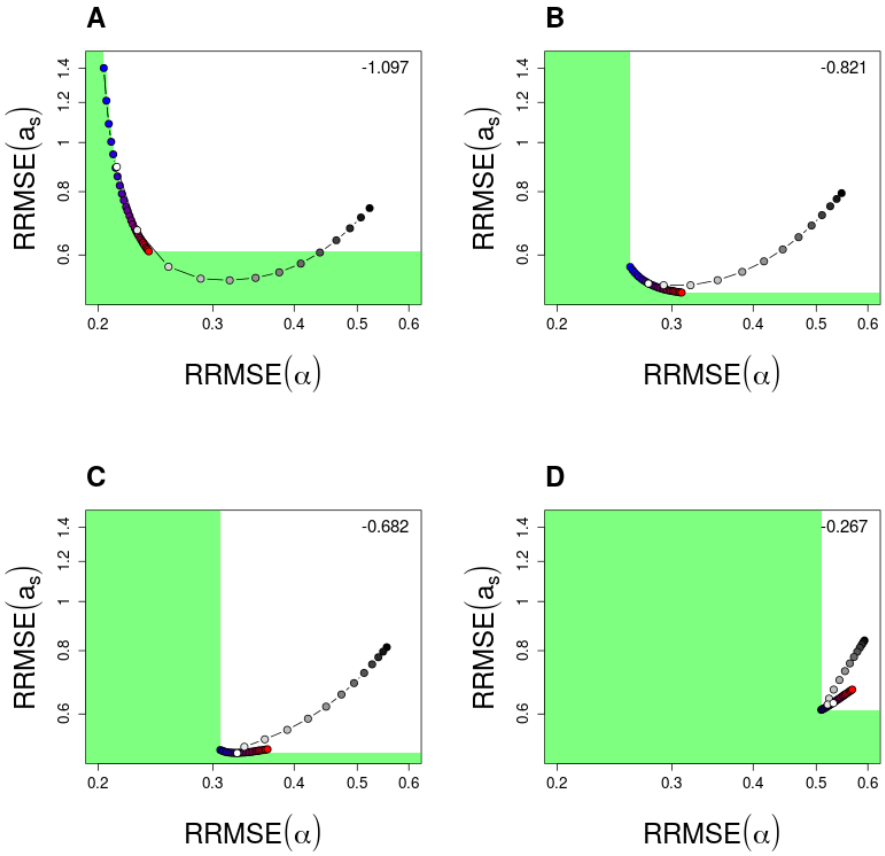

Figure 3: Comparing the relative root mean square errors of fractal designs to the Pareto front of hybrid designs in four typical situations. Panels A, B, C and D correspond to increasing values of a_s (values of $\log_{10}(a_s)$ are reported at the upper right corner). They were identified as grey rectangles in figure 2. In each panel, the RRMSEs of fractal designs are presented as a line of grey dots. The grey level of dots indicates the value of the contraction parameter x, increasing from black ($x = 10^{-1.5}$) to white ($x = 1$). Only a sub-sample of the 240 explored values on x were presented, to improve readibility. RRMSEs of hybrid strategies are presented using the same caption as in figure 2 with blue-to-red gradient. The Pareto front associated to hybrid designs is presented as a green polygon. When fractal designs reach the green area, they offer a new Pareto-optimal design compared to hybrid strategies.

because the theoretical prediction of grid design error on a_s parameter exceeded computer limits.

Then we observed a narrow range of a_s values ($10^{-1.858} \leq a_s \leq 10^{-1.513}$) where all the hybrid designs gradually came back to the Pareto front as a_s increased, starting from pure random design ($p = 1$; fig. 4A). Nearly simultaneously, as a_s increased above $10^{-1.789}$, fractal designs within a range of intermediary contraction parameter values x became excluded from the Pareto front by hybrid designs. The range of excluded x values initiated at $x = 10^{-0.306}$ and expanded while shifting towards high x values, until encompassing the higher end of the range ($x = 1$; fig. 4B). Fractal designs with high contraction parameters ($x > 10^{-0.1}$) could sporadically become Pareto optimal again at larger a_s values but, in those cases, they were quantitatively very close to hybrid design in terms of error (see for instance the fractal design with $x = 1$ in fig. 3C).

By contrast, fractal designs with lower x values (e.g. $x < 10^{-0.306}$) were not excluded from the Pareto front when a_s increased above $10^{-1.789}$ (fig. 4B), and remained Pareto-optimal over a larger range of a_s values. These designs came as an extension of — rather than in competition with — the Pareto front associated to hybrid designs. They were associated to lower error on a_s but higher error on α (as illustrated in fig. 3A). However, when a_s values increased above $10^{-0.924}$, this type of Pareto-optimal fractal strategies based on accurate a_s estimation were excluded by hybrid designs, as illustrated by the transition between fig. 3A and fig. 3B.

For larger a_s values ($a_s > 10^{-0.682}$), fractal designs became excluded from the Pareto front by hybrid designs, irrespective of x value (figs. 3D, 4B).

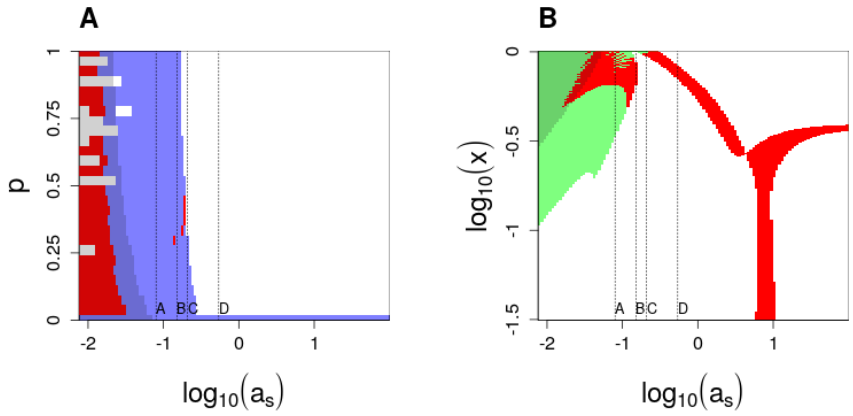

Figure 4: Intersection between Pareto fronts associated hybrid and fractal designs. Panel A presents the impact of fractal designs on the Pareto front of hybrid designs. Colored pixels (either blue or red) show, for each explored a_s value, the values of the proportion of random observations (p) that lead to a Pareto-optimal hybrid design. White 'holes' on the left side of the graph are artifacts due to limits in the numerical precision at very high errors on a_s. Among those Pareto-optimal values of p, red pixels show which values are not Pareto-optimal anymore when adding fractal designs, while blue pixels show the p values that remain Pareto optimal. Panel B presents the impact of hybrid designs on the Pareto front of fractal designs. Colored pixels (either green or red) show, for each explored a_s value, the values of the contraction parameter (x) that lead to a Pareto-optimal fractal design. Among those Pareto-optimal values of x, red pixels show which values are not Pareto-optimal any more when adding hybrid designs, while green pixels show the x values that remain Pareto optimal. In both panels, the shaded area (left on panel A, upper-left on panel B) shows designs where the RRMSE(a_s) is above 2. Vertical dotted lines show the positions of examples detailed in figure 3.

Theoretical analysis of changing the size of the surveyed area or the sampling effort

In this section, we used the shorthand notations $E_a(a_s, L, N)$ [resp. $E_\alpha(a_s, L, N)$] for the RRMSE(a_s) [resp. RRMSE(α)] when true autocorrelation range is a_s, surveyed area side length is L and sample size is N.

Surveyed area Until now, we considered the problem of sampling within a fixed triangular area constraining designs to cover the whole surface. We relaxed this assumption and allowed the side length L of the surveyed area to vary as a free parameter. A re-scaling argument (see Article S1 in Supporting Information, section 5) yielded that:

$$E_a(a_s, \lambda L, N) = E_a(a_s/\lambda, L, N)$$
$$E_\alpha(a_s, \lambda L, N) = E_\alpha(a_s/\lambda, L, N)$$

where $\lambda > 0$ is the dilatation factor applied to side length. In words, changing the size of the area, through dilatation or contraction, is exactly equivalent to changing the value of a_s while keeping the size of the area to its original value. Therefore, when the surveyed area can freely change, Pareto fronts of sampling strategies can be obtained directly through merging RRMSEs previously obtained at distinct a_s values, and computing the Pareto front of the pooled dataset. By doing so, we can compare hybrid and fractal sampling strategies and see their respective contributions to a global Pareto front (fig. 5A).

Grid design was never a Pareto-optimal design (fig. 5A). It was consistently excluded from the front by other hybrid designs including some degree of randomness and by fractal designs. By contrast pure random designs alone reached a Pareto front very similar if not identical to the Pareto front of all hybrid designs, suggesting that when the size of the area can be adapted it may not be useful to add regular elements within the random design. Fractal and random designs showed quantitatively very close Pareto fronts, which both contributed to the global Pareto front. Fractal designs seemed slightly more performant when seeking intermediary error levels on the mean and the autocorrelation range, but we reckoned that the magnitude of the difference was too reduced to justify a deep interpretation. Focusing on random designs ($p = 1$), the transition from designs oriented towards estimating the mean to designs oriented

towards estimating the autocorrelation range occured when L decreased from $L = 100a_s$ to $L = 10a_s$ (fig. 5B).

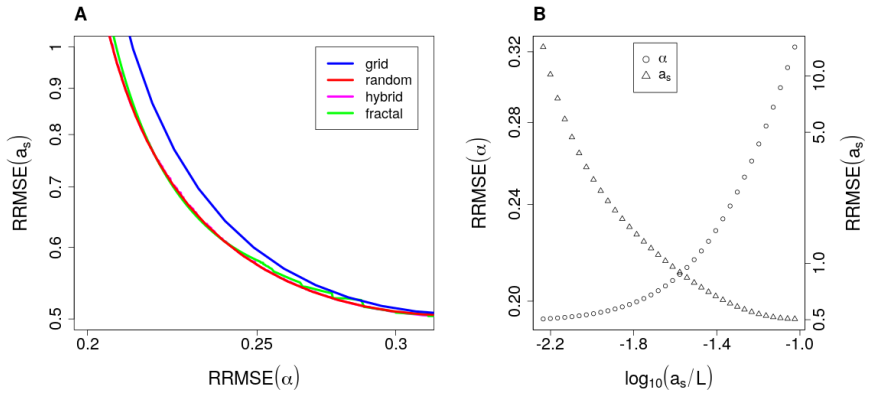

Figure 5: Pareto front of relative root mean square error of estimation of exponential mean RRMSE(α) and autocorrelation range RRMSE(a_s) for four sampling strategies when allowing to change surveyed area, hence removing the dependence on a_s. The 'hybrid' Pareto front corresponds to considering all the values of p simultaneouly (hence encompassing grid and random designs as particular cases). The 'fractal' Pareto front corresponds to considering all the values of x simultaneouly. Hybrid (pink) and random-only (red) Pareto fronts are nearly exactly super-imposed. Panel A shows the Pareto fronts for the different types of design. Panel B relates the position on the random-only Pareto front to the value of $\log_{10}(a_s/L)$, where L is the length of the side of the triangular area.

Sampling effort We now theoretically explore the implication of increasing sample size N by a factor $\eta = 3^q$ with $q \in \mathbb{N}^*$. For fractal designs, the increase of sample size is done by further iterating q times the iterating function system depicted in methods. For hybrid designs, the increase of sample size is done by increasing the density of sampling points by a η factor. We propose the following approximation for a design i (see Article S1 in Supporting Information, section

5 for a justification):

$$E_a^{(i)}(a_s, L, \eta N) \approx \eta^{-\frac{1}{2}} E_a^{(i)}(a_s \times \eta^{\frac{1}{\delta_i}}, L, N)$$
$$E_\alpha^{(i)}(a_s, L, \eta N) \approx \eta^{-\frac{1}{2}} E_\alpha^{(i)}(a_s \times \eta^{\frac{1}{\delta_i}}, L, N)$$

where δ_i is known as the 'fractal dimension' of the design i, equal to 2 for hybrid designs and to $-\log(3)/\log(\rho)$ for fractal designs (always strictly lower than 2). The first effect of increasing sample size should thus be to decrease RRMSEs by a factor $\eta^{-\frac{1}{2}}$, irrespective of the design, which should not change the ordination of designs, and suggest considering rescaled RRMSEs to discuss the question of ordination :

$$\eta^{\frac{1}{2}} E_a^{(i)}(a_s, L, \eta N) \approx E_a^{(i)}(a_s \times \eta^{\frac{1}{\delta_i}}, L, N)$$
$$\eta^{\frac{1}{2}} E_\alpha^{(i)}(a_s, L, \eta N) \approx E_\alpha^{(i)}(a_s \times \eta^{\frac{1}{\delta_i}}, L, N)$$

These rescaled RRMSEs suggest that the effect of increasing sampling effort on designs ordination is equivalent to increasing a_s. The equivalent increase on a_s depend on the fractal dimension, it is larger for fractal designs than for hybrid designs.

We previously observed that RRMSE(α) tended to increase with the degree of randomness p of hybrid designs (fig. 2) irrespective of a_s value. Because hybrid designs all have the same fractal dimension, this pattern should persist as sampling effort increases. The effect of increasing sampling effort on the ordination of RRMSE(a_s) among hybrid design is harder to predict since the variation of RRMSE(a_s) along the a_s gradient is non-monotonic. However, using previous results (fig. 2), one expects that when sampling effort has increased enough to ensure that the mesh size of grid sampling design has become lower than autocorrelation range a_s, the grid design would become the best hybrid design with respect to RRMSE(a_s), and therefore the unique Pareto-optimal design among hybrid designs.

Combining the facts that (i) grid design consistently outperformed fractal designs on RRMSE(α) at the same a_s value (e.g. fig. 4A), (ii) fractal designs have higher 'equivalent' a_s than hybrid designs when sampling effort increases and (iii) RRMSE(α) of fractal designs increased with a_s (fig. 6A), one can expect that increasing sampling effort preserves grid design as the best design among all with respect to RRMSE(α). When the autocorrelation range a_s is higher than the mesh size of grid sampling design, RRMSE(a_s) of fractal designs increases with a_s (fig. 6B). This tends to suggest that when sampling effort has increased enough to ensure a mesh size of grid sampling design lower than a_s, the grid design may also outperform fractal designs with $x > 10^{-1.5}$ in terms of RRMSE(a_s), and thus be the unique Pareto-optimal design among all designs.

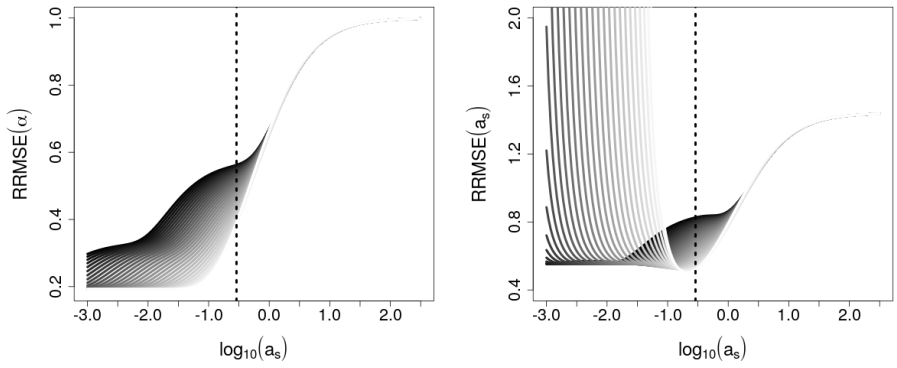

Figure 6: RRMSEs of fractal designs as a function of autocorrelation range a_s. Panel A shows RRMSE(α), panel B shows RRMSE($a[s]$). In both panels, the vertical dotted line shows the grid design mesh size. The grey level of curves indicates the value of the contraction parameter x, increasing from black ($x = 10^{-1.5}$) to white ($x = 1$). Only a sub-sample of the 240 explored values on x were presented, to improve readibility.

Discussion

Within a fixed surveyed area, hybrid designs are not always intermediary Pareto-optimal strategies between grid and random designs For autocorrelation range smaller than the grid mesh size, we retrieved the expected continuum of Pareto-optimal hybrid designs between grid and random designs. In this context, pairwise distances among sampling points smaller than the grid mesh size were needed to accurately estimate the autocorrelation range, and such smaller distances were provided by the introduction of random points. Increasing the degree of randomness in designs thus lead to a gradual shift in the accuracy from estimating the mean of the field to estimating the autocorrelation range. For larger autocorrelation ranges, we obtained less expected results: adding too much randomness could depart from the Pareto front of designs and become sub-optimal. The upper threshold of acceptable randomness decreased with autocorrelation range and, for large autocorrelation ranges, the grid design stood as the unique best strategy among hybrid designs to estimate both the mean and the autocorrelation range of the field.

In practice, choosing among hybrid designs thus relies on *a priori* knowledge about the order of magnitude of the autocorrelation range for the quantity of interest. Let us consider the practical case where one wants to define a strategy to position 27 sampling plots in a forest of about 3600ha in order to study the distribution of saproxylic beetles species living in hollow trees. The mesh size of a regular grid spread over the forest would then be of c.a. 1500m (with fluctuations depending on the geometry of both forest and the chosen shape of the mesh). A previous studies based on auto-regressive occupancy models [Ranius et al. (2010), fig. 1] suggested that cavicolous beetles often harbour a spatial autocorrelation with range below 1000m. One thus expects the mesh size to be larger than autocorrelation range in this example. If estimating the autocorre-

lation range were a strong priority of the study, random sampling should thus be preferred. If one rather looked for a compromise between mean and autocorrelation range estimation, truly hybrid strategies should be preferred. In the latter case, the shape of the Pareto front seems to be convex when autocorrelation is smaller than mesh size (fig. 2), suggesting that the pay-off of adding randomness decreases as the proportion of random points increase. Therefore, a choice for a low degree of randomness [e.g. $p = 0.1$; Bijleveld et al. (2012)] could be appropriate.

Bijleveld et al. (2012) had already identified that the relative performance of designs depended on the level of the underlying autocorrelation range. For instance, they found that the bias when estimating autocorrelation range was minimized by random design for small autocorrelation ranges, but minimized by hybrid strategy with $p = 0.1$ at higher autocorrelation level. However, because they averaged the performance of designs out across autocorrelation levels and intersample distances explored in their analysis, the authors further concluded that, overall, there was a Pareto front of hybrid designs between grid and random. Our findings discourage averaging across autocorrelation ranges, because the magnitude of errors on autocorrelation estimation rapidly increases as the autocorrelation range decreases. Global averaging thus tends to give too much of weight to scenarios with small autocorrelation range compared to intersample distance, and may lead to over-generalizing patterns that are in fact specific to small autocorrelation range values. Diverging magnitude of error at the lower end autocorrelation range raises the same problem for all metrics of perfomance integrating over an interval of autocorrelation ranges. For instance, Zhu and Stein (2005) mentioned that minimax or average metrics of estimation error across the autocorrelation range considered in their study were very unstable and hard to optimize, probably due this phenomenon. For this reason, we chose

not to derive global metrics in our study but focused on the qualitative analysis of Pareto fronts.

Within a fixed surveyed area, fractal designs can be Pareto-optimal strategy to estimate small autocorrelation ranges At very small autocorrelation ranges, all hybrid designs except grid were excluded by fractal designs, i.e. the latter were more efficient at estimating both the autocorrelation range and the mean. Fractal design with intermediate contraction parameter seemed particularly interesting because the associated absolute level of error on autocorrelation range remained moderate (see non-shaded area on fig. 4A). These designs remained Pareto-optimal when autocorrelation range increased up to values close to grid mesh size, because they extended the hybrid Pareto front towards estimating autocorrelation range more accurately. In other words, they offered a way to go further than the random design towards the aim of accurately estimating the autocorrelation range while paying a cost on the estimation of the mean. Coming back to the example of saproxylic beetles mentioned above, if estimating the autocorrelation range were a strong priority of the study, fractal designs with intermediate x may be even more interesting than random design. Under the assumptions of our study, the figure 4B seems to suggest that taking $x = 10^{-0.4} \approx 0.4$ is quite a robust choice. One may object that when accurately estimating small autocorrelation ranges is a strong priority of a survey, it may be relevant to combine the optimization of sampling design with the reduction of the area of study. This specific point is discussed in the next section.

It should be noted that, in practice, the choice of contraction parameter to build sampling design comes with sterical contraints when sampling units cannot be too close one from another. This may happen when sampling units have a large size (see our example below) or if sampling induces a disturbance than would alter the outcome of sampling nearby (e.g. because of organisms have

large home range or because they are sensitive to the presence of observers), a phenomenon akin to 'interference among sampling unit' in causal theory (Kimmel et al., 2021). Considering our example about saproxylic beetles, sampling units could be circular plots of 1ha (a radius of c.a. 57m). Then the minimal distance between two sampling units would have to be of at least 114m to avoid overlapping. If one assumes that the forest under study ($\approx$ 3600ha) has a diameter of c.a. 7km, building a triangular fractal design with 27 plots implies that the largest distance between plots in the design is $114/\rho^2$ where $\rho = x\sqrt{3}/(2+\sqrt{3})$. The constraint that this distance must be below 7km implies that x cannot be lower than 0.27. Similarly, it is straightforward to show that there could not be more than five scales in the triangular fractal sampling design without generating overlapping of sampling units (i.e. sampling effort must be lower than $N = 3^5 = 243$ plots). However, this threshold on sampling size could be overcome by considering a more complex geometrical shape of the fractal. More generally, the number of sampling sites can be modulated by combining the choice of the geometrical shape with subsampling (Marsh and Ewers, 2013).

If the size of surveyed area can be adapted or sampling effort increased, fractal designs are outperformed by more classic options Assuming that the size of the area of study is not predetermined by external constraints, random design was sufficient to reach — or get very close to — the global Pareto front of designs explored in our study through adjusting the size of the surveyed area. In this case, moving towards hybrid or random designs seems adding complexity without subsequent payoff. The problem of knowing the order of magnitude of autocorrelation range *a priori* is still present though, for the size of the area has to be adapted to this quantity. For 27 sampling point, the typical dimension — side length in our case — of surveyed area should be comprised between ten and a hundred times the target autocor-

relation depending on whether the main goal is autocorrelation range or mean estimation, respectively. Our results about the effect of sampling effort suggest that if the number of sampling points is increased e.g. fourfold, then the range of side length values to consider for the surveyed area should be shifted upwards, and approximately comprised between twenty and two hundred times the target autocorrelation range.

However, there are several reasons in practice for which the area of study may not be a real degree of freedom when building the study design. First, the area open to sampling may be limited in space either for biological reasons (e.g. a spatially-limited habitat, like a lake) or practical reasons (restricted access, time of travel, etc.). This would prevent extending at will the area of study and potentially limit the opportunities for improving the estimation of the mean that way. Conversely, the area study cannot be freely reduced when one aims at relating environmental covariates to target biodiversity patterns (especially non-linear ones; Albert et al. (2010)), because the range of covariate values has to be appropriately covered. This implies e.g. stratifying among various type of soil cover (Yoccoz et al., 2001), or to span the full extent of an environmental gradient (Field et al., 2009; Albert et al., 2010). Our study does not include these constraints, for we did not consider a third criterion that would be accurately estimating patterns along an environmental covariate. By taking this simplified regression framework, we could easily adress the question of the trade-off between estimating a fixed effect and estimating the spatial structure of residuals. At that stage, we showed that fractal and random designs showed very similar Pareto-fronts when freely adjusting the size of surveyed area (fig. 5A). Consequently, if fractal designs happened to better estimate the effect of gradients by forcing the presence of large pairwise distances, they may exclude random design from the Pareto front when including this third axis of evaluation.

Our choice of sampling effort $N = 27$ was done to reflect realistic settings that one can observe in many resarch projects on biodiversity, ours included. However, when data acquisition at a sampling point is not very demanding, it is also frequent to observe larger designs. Sticking with the hollow trees example, if one simply aims at describing features of the trees like tree-related microhabitats (Larrieu et al., 2018), which are proxys for saproxylic beetle biodiversity (Bouget et al., 2014, 2013), but does not aim at sampling and identifying beetles themselves, then the sampling budget can considerably increase. In this case, our theoretical results tend to suggest that when the sampling budget is sufficient for the mesh size of a grid design to become equivalent to or lower than the anticipated magnitude of autocorrelation range, a grid design should be preferred among the other strategies and fractal designs are excluded from the Pareto-front.

Conclusions

In the context of our study, the main advantage of fractal designs occured when aiming at estimating short autocorrelation ranges while constrained on covering a large area of survey with a limited sampling budget. In other situations, it seemed more efficient and less complicated to implement more classic designs. The niche for fractal designs may thus appear quite limited. It should nonetheless be noted that we evaluated designs on a simple scenario with a parsimonious autocorrelation structure and no effect of covariates. The question of jointly estimating the effects of covariates and the autocorrelation range should now be further adressed, for it adds new axes to the trade-off among designs. In particular, biological patterns often stem from heterogeneous drivers acting at different scales (Thuiller et al., 2015; Ricklefs, 2008). Designs that harbour a clear hierarchical structure — like fractal designs — may be particularly adapted

to capture such heterogeneity (Simpson and Pearse, 2021), provided that the scales of variation induced by the hypothesized processed match the geometrical constraint of self-similarity inherent to fractals.

Acknowledgements

The author thanks C. Bouget and A. Brin for discussion about the practical implementation of fractal sampling designs, J. Crabot for pointing useful bibliography, B. Laroche for orienting towards the framework of optimal design, C. Sirami for emphasizing the interest of random designs as a benchmark and A. Tortosa for pointing the need of exploring sampling effort.

This research was funded, in whole or in part, by the Agence Nationale de la Recherche (ANR), through the BloBiForM project grant ANR-19-CE32-0002-01. A CC-BY public copyright license has been applied by the authors to the present document and will be applied to all subsequent versions up to the Author Accepted Manuscript arising from this submission, in accordance with the grant's open access conditions.

Conflicts of interest

The author has no conflict of interest to declare.

References

Abt, M. and Welch, W. J. (1998). Fisher information and maximum-likelihood estimation of covariance parameters in Gaussian stochastic processes. *Canadian Journal of Statistics*, 26(1):127–137.

Albert, C. H., Yoccoz, N. G., Edwards, T. C., Graham, C. H., Zimmermann,

577 N. E., and Thuiller, W. (2010). Sampling in ecology and evolution - bridging
578 the gap between theory and practice. *Ecography*, 33(6):1028–1037.

579 Archaux, F. and Bergès, L. (2008). Optimising vegetation monitoring. A case
580 study in A French lowland forest. *Environmental Monitoring and Assessment*,
581 141(1-3):19–25.

582 Bardos, D., Guillera-Arroita, G., and Wintle, B. A. (2015). Valid auto-models
583 for spatially autocorrelated occupancy and abundance data. *Methods in Ecol-*
584 *ogy and Evolution*, 6(10):1137–1149. Publisher: John Wiley & Sons, Ltd.

585 Bijleveld, A. I., van Gils, J. A., van der Meer, J., Dekinga, A., Kraan, C.,
586 van der Veer, H. W., and Piersma, T. (2012). Designing a benthic monitor-
587 ing programme with multiple conflicting objectives. *Methods in Ecology and*
588 *Evolution*, 3(3):526–536.

589 Bouget, C., Larrieu, L., and Brin, A. (2014). Key features for saproxylic beetle
590 diversity derived from rapid habitat assessment in temperate forests. *Ecolog-*
591 *ical Indicators*, 36:656–664.

592 Bouget, C., Larrieu, L., Nusillard, B., and Parmain, G. (2013). In search of the
593 best local habitat drivers for saproxylic beetle diversity in temperate decidu-
594 ous forests. *Biodiversity and Conservation*, 22(9):2111–2130.

595 Cressie, N. A. C. (1993). Geostatistics. In *Statistics for Spatial*
596 *Data*, pages 27–104. John Wiley & Sons, Ltd. Section: 2 _eprint:
597 https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119115151.ch2.

598 Dormann, C. F., McPherson, J. M., Araújo, M. B., Bivand, R., Bolliger, J.,
599 Carl, G., Davies, R. G., Hirzel, A., Jetz, W., Kissling, D. W., Kühn, I.,
600 Ohlemüller, R., Peres-Neto, P. R., Reineking, B., Schröder, B., Schurr, F. M.,

and Wilson, R. (2007). Methods to account for spatial autocorrelation in the analysis of species distributional data: a review. *Ecography*, 30(5):609–628.

Falconer, K. J. (2003). *Fractal geometry: mathematical foundations and applications*. Wiley, Chichester, England, 2nd ed edition.

Field, R., Hawkins, B. A., Cornell, H. V., Currie, D. J., Diniz-Filho, J. A. F., Guégan, J.-F., Kaufman, D. M., Kerr, J. T., Mittelbach, G. G., Oberdorff, T., O'Brien, E. M., and Turner, J. R. G. (2009). Spatial species-richness gradients across scales: a meta-analysis. *Journal of Biogeography*, 36(1):132–147.

Hooten, M. B., Wikle, C. K., Sheriff, S. L., and Rushin, J. W. (2009). Optimal spatio-temporal hybrid sampling designs for ecological monitoring. *Journal of Vegetation Science*, 20(4):639–649.

Kimmel, K., Dee, L. E., Avolio, M. L., and Ferraro, P. J. (2021). Causal assumptions and causal inference in ecological experiments. *Trends in Ecology & Evolution*, 36(12):1141–1152.

Larrieu, L., Paillet, Y., Winter, S., Bütler, R., Kraus, D., Krumm, F., Lachat, T., Michel, A. K., Regnery, B., and Vandekerkhove, K. (2018). Tree related microhabitats in temperate and Mediterranean European forests: A hierarchical typology for inventory standardization. *Ecological Indicators*, 84:194–207.

Legendre, P. (1993). Spatial Autocorrelation: Trouble or New Paradigm? *Ecology*, 74(6):1659–1673.

Lennon, J. J. (2000). Red-shifts and red herrings in geographical ecology. *Ecography*, 23(1):101–113.

Mandelbrot, B. (1983). *The fractal geometry of nature*, volume 173. WH freeman New York.

625 Manel, S. and Holderegger, R. (2013). Ten years of landscape genetics. *Trends*
626 *in Ecology and Evolution*, 28(10):614–621.

627 Marsh, C. J. and Ewers, R. M. (2013). A fractal-based sampling design for eco-
628 logical surveys quantifying beta-diversity. *Methods in Ecology and Evolution*,
629 4(1):63–72.

630 McGill, B. J. (2010). Towards a unification of unified theories of biodiversity.
631 *Ecology Letters*, 13:627–642.

632 Moran, P. A. P. (1950). Notes on continuous stochastic phenomena. *Biometrika*,
633 37(1-2):17–23.

634 Müller, W. G. (2007). *Collecting spatial data: optimum design of experiments*
635 *for random fields*. Springer, Berlin ; New York, 3rd rev. and extended ed
636 edition.

637 Müller, W. G., Rodríguez-Díaz, J. M., and Rivas López, M. J. (2012). Opti-
638 mal design for detecting dependencies with an application in spatial ecology.
639 *Environmetrics*, 23(1):37–45.

640 Nekola, J. C. and White, P. S. (1999). The distance decay of similarity in
641 biogeography and ecology. *Journal of Biogeography*, 26(4):867–878.

642 Ouborg, N. J., Piquot, Y., and Van Groenendael, J. M. (1999). Population
643 genetics, molecular markers and the study of dispersal in plants. *Journal of*
644 *Ecology*, 87(4):551–568. Publisher: John Wiley & Sons, Ltd.

645 Prugh, L. R. (2009). An evaluation of patch connectivity measures. *Ecological*
646 *Applications*, 19(5):1300–1310.

647 Ranius, T., Johansson, V., and Fahrig, L. (2010). A comparison of patch con-
648 nectivity measures using data on invertebrates in hollow oaks. *Ecography*,
649 33(5):971–978.

Rhodes, J. R. and Jonzén, N. (2011). Monitoring temporal trends in spatially structured populations: how should sampling effort be allocated between space and time? *Ecography*, 34(6):1040–1048. Publisher: John Wiley & Sons, Ltd.

Ricklefs, R. (2008). Disintegration of the ecological community. *The American Naturalist*, 172(6):741–750.

Simpson, E. G. and Pearse, W. D. (2021). Fractal triads efficiently sample ecological diversity and processes across spatial scales. *Oikos*, 130(12):2136–2147.

ter Braak, C. J., Hanski, I., and Verboom, J. (1998). The incidence function approach to modeling of metapopulation dynamics. In *Modeling spatiotemporal dynamics in ecology*, pages 167–188. Berlin, springer-verlag edition.

Thuiller, W., Pollock, L. J., Gueguen, M., and Münkemüller, T. (2015). From species distributions to meta-communities. *Ecology Letters*, 18(12):1321–1328.

Tischendorf, L. and Fahrig, L. (2000). On the usage and measurement of landscape connectivity. *Oikos*, 90(1):7–19.

Vekemans, X. and Hardy, O. J. (2004). New insights from fine-scale spatial genetic structure analyses in plant populations. *Molecular Ecology*, 13(4):921–935. Publisher: John Wiley & Sons, Ltd.

Yoccoz, N. G., Nichols, J. D., and Boulinier, T. (2001). Monitoring of biological diversity in space and time. *Trends in Ecology & Evolution*, 16(8):446–453.

Zhu, Z. and Stein, M. L. (2005). Spatial sampling design for parameter estimation of the covariance function. *Journal of Statistical Planning and Inference*, 134(2):583–603.