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

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

*Corresponding author: fabien.laroche@inrae.fr

Abstract

Quantifying the autocorrelation range of a species distribution in space is necessary for applied ecological questions like implementing protected area networks or monitoring programs. However, building a spatial sampling design to accurately estimate this range can conflict with other objectives such as estimating environmental effects acting upon species distribution. Mixing random sampling points and systematic grid ('hybrid' designs) may allow navigating between conflicting objectives: random sampling points induce contrasted pairwise distances that cover a wide array of possible autocorrelation range values while the grid limit small pairwise distances prone to pseudo-replication, thus improving the estimation of effects. Fractal designs (i.e. self-similar designs with well-identified scales) could also make a compromise by preserving some regularity reminiscent of grid at each scale but also browsing a wide array of possible autocorrelation range values across scales. Using maximum likelihood estimation in an optimal design of experiments approach, we compared errors of hybrid and fractal designs when estimating an effect acting upon a target spatial variable along with residual autocorrelation range. We found that Pareto-optimal sampling strategies depended on the feasible grid mesh size over the study area given the sampling budget. When the autocorrelation range was larger than the feasible mesh size, grid design was the best option on all criteria. When autocorrelation range was shorter, the choice of designs depended on the target effect. For effects driven by a monotonic environmental gradients across space, fractal designs outperformed hybrid designs on all criteria. For effects of non-linear covariates or simply the global mean of target variable over the study area, hybrid designs were more efficient. Therefore, a narrow niche for fractal designs exists. The interest of designs with a clear hierarchical structure like fractals may stand out more clearly when studying biological patterns with contrasted spatial structures across scales.

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 study 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 hybrid designs (including pure grid and pure random designs) with respect to their ability to quantify autocorrelation range and effects acting on a target variable in space. 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 an 'effect' (i.e. a fixed intercept or an effect associated to a covariate) and the autocorrelation range of a target variable. 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 differed from random design, and tended 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 designs' performance.

When comparing hybrid and fractal designs, we had two expectations based on previous litterature: (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

All the methods described below have been implemented using the R language, and made freely available online (doi : 10.5281/zenodo.7846349).

Spatial sampling designs

All spatial sampling designs harboured $N = 27$ sampling points. Sampling points are spread within an area of study shaped as an equilateral triangle with a side length of $L = \sqrt{3}$ distance units (Figs. 1, 2).

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 defined by choosing a fraction p of sites at random within a grid design and resampling their new position at random in the area of study. Here we consider the $N + 1 = 28$ values of $p \in \{0, 1/N, 2/N, ..., 1\}$. Note that $p = 1$ yields a pure random design (Fig. 1).

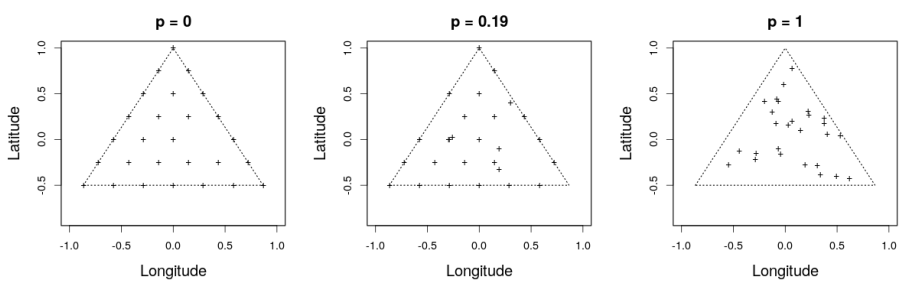

Figure 1: Effect of the proportion of random sites p on hybrid sampling designs. The triangle dashed area is the area of study. Crosses show the position of the 27 sampling points. The middle panel corresponds to $p = 5/27 \approx 0.19$.

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, auto-similar set of plots it belongs to. The values of ρ considered in the study are : $\rho = x\sqrt{3}/(2+\sqrt{3})$ with 28 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. $L/10$. Lower x yielded more irregular fractal sampling designs (Fig. 2).

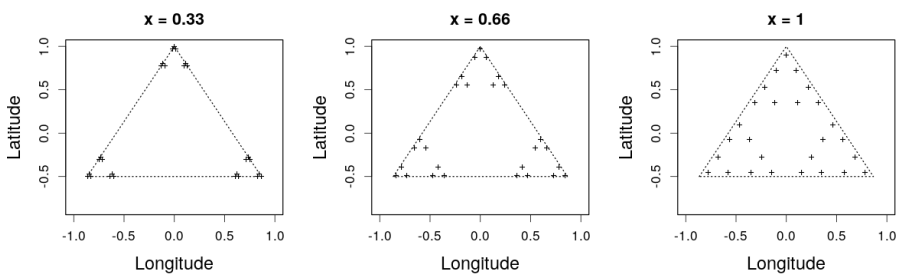

Figure 2: Effect of the contraction parameter x on fractal sampling designs. The triangle dashed area is the area of study. Crosses show the position of the 27 sampling points.

Modeling the observed variable

We assume that observations of interest are driven by an environmental gradient with auto-correlated residual variation. The vector of observations at each

sampling points $\boldsymbol{Y} = (Y_1, ..., Y_N)$ follows the model :

$$Y_i = \mu + \beta x_i + Z_i \tag{1}$$

where μ, β are parameters in $\mathbb{R}$, x_i is an environmental covariate and $\boldsymbol{Z} = (Z_1, ..., Z_N)$ is taken from a Gaussian random field with mean 0 and harbouring an exponential variogram (Cressie, 1993) without nugget effect:

$$\begin{aligned} &\forall i \in \{1, ..., N\}, Z_i \sim \mathcal{N}(0, \sigma^2) \\ &\forall i, j \in \{1, ..., N\}^2, \mathbb{E}\left[(Z_i - Z_j)^2\right] = 2\sigma^2 \left(1 - e^{-\frac{d_{ij}}{a_s}}\right) \end{aligned} \tag{2}$$

where $\sigma, a_s \in \mathbb{R}^{+*}$, 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:

$$\mathrm{Cov}\left[Z_i, Z_j\right] = \sigma^2 e^{-\frac{d_{ij}}{a_s}} \tag{3}$$

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. We considered 30 distinct a_s values evenly spaced on a log-scale between 10^{-2} and 10^2. As we will show below, the values of μ, β and σ did not affect the ranking of estimation errors among sampling designs considered in our analysis, and we could therefore set $\mu = 0$, $\beta = 1$ and $\sigma = 1$ without loss of generality.

We used two distinct models of for the environmental covariate (Fig. 3). A first model corresponded to a longitudinal gradient (x is an affine function with positive slope of longitude), a second model corresponded to a centered U-shaped (x is an affine function with positive slope of longitude^2 + latitude^2). These profiles are standardized between 0 and 1 over the triangular area of

study.

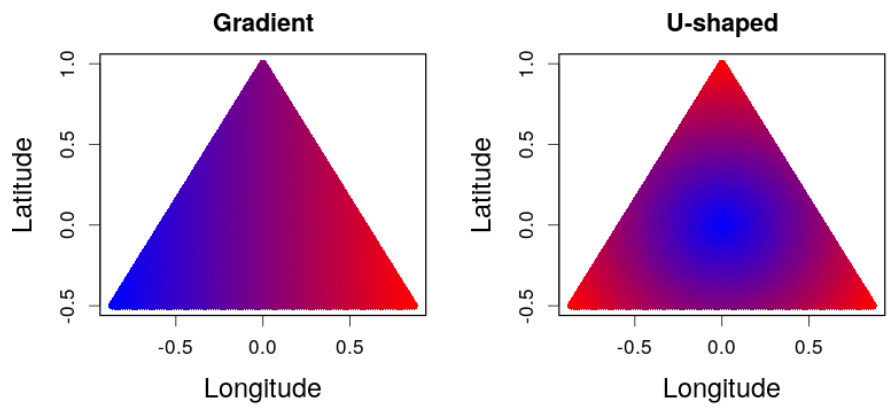

Figure 3: Two spatial profiles for the environmental variable. The darkest blue color corresponds to and environmental value of 0, the brightest red color corresponds to an environmental value of 1.

Maximum likelihood estimation

We consider two distinct estimation problems of ecological interest : (i) estimating the mean (μ) and the autocorrelation range (a_s) in the absence of environmental effect ($\beta = 0$); (ii) estimating the effect of the environmental covariate (β) and the autocorrelation range (a_s) when the environmental covariate has an effect on the observed variable.

Problem 1 : mean versus autocorrelation range

One supposes that the observed covariate has constant mean in space, following the simplified version of model (1):

$$Y_i = \mu + Z_i \tag{4}$$

The aim was to accurately estimating the range of spatial aurocorrelation

a_s and the intercept μ. The statistical model used to estimate parameters is model (4), i.e. we did not consider potential errors on model specification. When comparing estimation errors on intercept and autocorrelation range, we considered an exponentially-transformed intercept $\nu = e^{\mu}$ to obtain the same domain of definition $\mathbb{R}^{+*}$ as a_s. Summarizing parameters in a vector $\boldsymbol{\theta} = (\nu, \sigma, a_s)$, we focused on the corresponding maximum likelihood estimate $\hat{\boldsymbol{\theta}} = (\hat{\nu}, \hat{\sigma}, \hat{a}_s)$.

Problem 2 : environmental effect versus autocorrelation range

The observed covariate followed model (1). The aim was to accurately estimate the range of spatial aurocorrelation a_s and the slope of the environmental effect β. The statistical model used to estimate parameters was a reformulation of model (1) where we assumed that the environmental covariate has been centered in the dataset :

$$Y_i = \tilde{\mu} + \beta \tilde{x}_i + Z_i \tag{5}$$

where $\tilde{\mu} = \mu + \beta \bar{x}$, $\tilde{x}_i = x_i - \bar{x}$ and $\bar{x} = \frac{1}{N} \sum_{j=1}^{N} x_j$

We did not consider potential errors on model specification. When comparing estimation errors on intercept and autocorrelation range, we considered an exponentially-transformed slope $\gamma = e^{\beta}$ to obtain the same domain of definition $\mathbb{R}^{+*}$ as a_s. Parameters of this model are summarized in a vector $\boldsymbol{\theta} = (\tilde{\mu}, \gamma, \sigma, a_s)$. We focused on the corresponding maximum likelihood estimate $\hat{\boldsymbol{\theta}}$.

Estimation errors and Pareto fronts

Estimation error on a parameter θ (possibly ν, γ, a_s depending on the problem) is quantified through the relative root mean square error:

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

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.

When comparing designs on their ability to accurately several parameters at a time, we relied on the concept Pareto-optimality, defined as follows. Let generically denote the parameters of interest $\boldsymbol{\theta} = (\theta_1, \theta_2)$. Denote $\mathrm{RMSE}_S(\theta_i)$ the error of sampling design S at estimating parameter θ_i. Generally, a design S is said to be Pareto-optimal within the set of designs $\mathcal{S}$ at estimating the parameters in $\boldsymbol{\theta}$ if and only if:

$$\forall S' \in \mathcal{S}, \quad \begin{array}{c} \mathrm{RRMSE}_S(\theta_1) < \mathrm{RRMSE}_{S'}(\theta_1) \\ \text{or } \mathrm{RRMSE}_S(\theta_2) < \mathrm{RRMSE}_{S'}(\theta_2) \end{array}$$

Because we considered designs that could show levels of errors numerically close, we modified this definition to include an idea of 'noticeable' difference in estimation error. In our study, a design S is said to be Pareto-optimal within the set of designs $\mathcal{S}$ if and only if:

$$\forall S' \in \mathcal{S}, \quad \begin{array}{c} \rho\ \mathrm{RRMSE}_S(\theta_1) < \mathrm{RRMSE}_{S'}(\theta_1) \\ \text{or } \rho\ \mathrm{RRMSE}_S(\theta_2) < \mathrm{RRMSE}_{S'}(\theta_2) \end{array}$$

where $\rho = 0.99$ is a 1% tolerance factor on errors.

Results

Fisher information matrix and predicted errors

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

$$\mathcal{I}(\boldsymbol{\theta}) = \begin{pmatrix} \frac{1}{\nu^2}\mathbf{1}'\Sigma^{-1}\mathbf{1} & 0 & 0 \\ 0 & \frac{N}{2\sigma^4} & \frac{1}{2\sigma^2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) \\ 0 & \frac{1}{2\sigma^2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) & \frac{1}{2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) \end{pmatrix} \tag{6}$$

The relative root mean squared error associated to $\hat{\nu}$ and $\hat{a}_s$ derived from the Fisher information matrix (see Article S1 in Supporting Information, section 1) were:

$$\begin{aligned} \mathrm{RRMSE}(\nu) &= \sigma\frac{1}{\sqrt{\mathbf{1}'\tilde{\Sigma}^{-1}\mathbf{1}}} \\ \mathrm{RRMSE}(a_s) &= \frac{1}{a_s}\sqrt{\frac{2}{\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}) - \frac{1}{N}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})}} \end{aligned} \tag{7}$$

RRMSEs did not depend on ν, hence we set $\nu = 1$ without loss of generality. RRMSE(a_s) did not depend on σ, while RRMSE(ν) was proportional to σ. Because we were interested in the ranking of designs, which remains identical up to a multiplicative constant, we set $\sigma = 1$ without loss of generality in our conclusions either.

Problem 2 — The Fisher information matrix associated to parameters $\boldsymbol{\theta} = (\tilde{\mu}, \gamma, \sigma, a_s)$ in model (1) is [see Article S1 in Supporting Information, section 1; Zhu and Stein (2005); Müller (2007)]:

$$\mathcal{I}(\boldsymbol{\theta}) = \begin{pmatrix} \mathbf{1}'\Sigma^{-1}\mathbf{1} & \frac{1}{\gamma}\mathbf{1}'\Sigma^{-1}\tilde{\boldsymbol{x}} & 0 & 0 \\ \frac{1}{\gamma}\mathbf{1}'\Sigma^{-1}\tilde{\boldsymbol{x}} & \frac{1}{\gamma^2}\tilde{\boldsymbol{x}}'\Sigma^{-1}\tilde{\boldsymbol{x}} & 0 & 0 \\ 0 & 0 & \frac{N}{2\sigma^4} & \frac{1}{2\sigma^2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) \\ 0 & 0 & \frac{1}{2\sigma^2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) & \frac{1}{2}\mathrm{tr}(\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}\tilde{\Sigma}^{-1}\frac{\partial\tilde{\Sigma}}{\partial a_s}) \end{pmatrix} \tag{8}$$

The relative root mean squared error associated to $\hat{a}_s$ and $\hat{\alpha}$ derived from the Fisher information matrix (see Article S1 in Supporting Information, section 1) were:

$$\begin{aligned} \mathrm{RRMSE}(\gamma) &= \sigma\sqrt{\frac{\mathbf{1}'\tilde{\Sigma}^{-1}\mathbf{1}}{(\tilde{\boldsymbol{x}}'\tilde{\Sigma}^{-1}\tilde{\boldsymbol{x}})(\mathbf{1}'\tilde{\Sigma}^{-1}\mathbf{1})-(\mathbf{1}'\tilde{\Sigma}^{-1}\tilde{\boldsymbol{x}})^2}} \\ \mathrm{RRMSE}(a_s) &= \frac{1}{a_s}\sqrt{\frac{2}{\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s}\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})-\frac{1}{N}\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})\mathrm{tr}(\Sigma^{-1}\frac{\partial\Sigma}{\partial a_s})}} \end{aligned} \tag{9}$$

RRMSEs did not depend on $\tilde{\mu}$ and γ, hence we set $\mu = 0$ and $\gamma = 1$ throughout the studywithout loss of generality. RRMSE(a_s) did not depend on σ while RRMSE(γ) was proportional to σ. Here again, because we were interested in the ranking of designs, which remains identical up to a multiplicative constant, we set $\sigma = 1$ without loss of generality in our conclusions either.

Note that RRMSE(a_s) is identical in problems 1 and 2 (eqs. (7) and (9)).

Theoretical analysis of asymptotic errors

In the limit $a_s \to 0$ — RRMSE(a_s) increases towards $+\infty$ as a_s becomes small, irrespective of considered design (see Article S1 section 2 in Supporting Information). The increase is proportional to $a_s/d_{\min} \times \exp(d_{\min}/a_s)$ where $d_{\min}$ is the smallest distance among two distinct sampling points. Designs with smaller $d_{\min}$ should thus 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

exists a threshold on x below which fractal designs yield lower RRMSE(a_s) than any hybrid designs as $a_s \to 0$, and thus become Pareto-optimal for low a_s.

In problem 1, RRMSE(ν) should converge to $\sigma/\sqrt{N}$, which does not depend on the sampling design. When $\sigma = 1$ and $N = 27$, this yields RRMSE(ν) $\approx$ 0.19. Fractal design with low contraction parameters might therefore exclude hybrid designs from the Pareto front in the limit $a_s \to 0$, since the accuracy at estimating autocorrelation range should become the major difference among designs.

In problem 2, RRMSE(γ) should converge to $\sigma/\left[SD(\boldsymbol{x})\sqrt{N}\right]$ where $SD(\boldsymbol{x})$ is the standard deviation of the environmental covariate $\boldsymbol{x}$. Sampling designs that maximizes the variance in $\boldsymbol{x}$ should therefore yield a more accurate estimate of γ. In the linear environment, fractal designs with low contraction parameter should yield a higher variance in $\boldsymbol{x}$ than hybrid designs, hence a higher accuracy at estimating γ, for they tend to position sampling points at opposite borders of the area. They may therefore exclude hybrid designs from the Pareto front in the limit $a_s \to 0$. By contrast, if environment is hump-shaped, we expect that fractal designs with low contraction parameter yield a low variance in $\boldsymbol{x}$. The Pareto front should then be shared between fractal designs with low contraction parameter (better at estimating a_s) and hybrid designs (better at estimating γ).

In the limit $a_s \to +\infty$ — 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 problem 1 (mean versus autocorrelation range estimation), RRMSE(ν) converges to σ ($= 1$ in our example), irrespective of the sampling design. In the problem 2 (slope of environmental effect versus autocorrelation range), RRMSE(γ) converges to 0, irrespective of the sampling design. Therefore, all the sampling designs should converge to-

wards very similar performance as $a_s \to +\infty$ in either problem. Regarding RRMSE(ν), this phenomenon of convergence among designs had already been shown on simulations in a previous study (Perret et al., 2022).

Numerical analysis of Pareto fronts in problem 1

We defined below two situations : a_s being smaller than grid design mesh size ($a_s < L/6 \approx 0.29$) and a_s being larger than grid mesh size ($a_s > L/6$). We respectively called these two situations 'A' and 'B'.

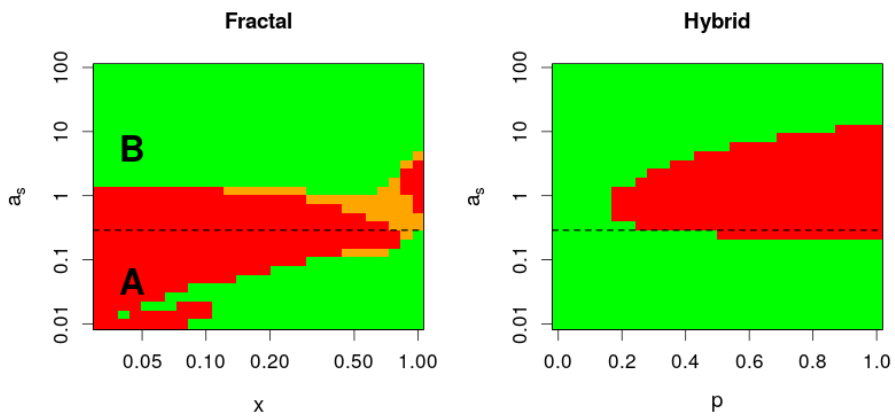

Figure 4: Pareto fronts of sampling designs in problem 1. In the left panel, colored pixels indicates, for given a_s and x values, whether the target fractal sampling is Pareto-optimal within fractal and hybrid sampling designs altogether (green pixels) or within fractal sampling designs only (orange pixels). Red pixels show, for a given a_s, fractal designs that are not Pareto optimal within the set of fractal sampling strategies. Similarly, in the right panel, colored pixels indicates, for given a_s and p values, whether the target hybrid sampling is Pareto-optimal within fractal and hybrid sampling designs altogether (green pixels) or within hybrid sampling designs only (orange pixels). Red pixels show, for a given a_s, hybrid designs that are not Pareto optimal within the set of hybrid sampling strategies. Therefore, in either panel, green and orange pixels show designs that are Pareto-optimal within their own type, and orange pixels shows the fraction of these designs that are eliminated when introducing the other type of design in the comparison. The horizontal dashed line $a_s = L/6$ shows the limit between two situations discussed in main text : 'low' autocorrelation range (A) and 'high' autocorrelation range (B).

Fractal designs alone — In situation A, the RRMSE(a_s) showed a U-shaped pattern as x increased (i.e. designs become more regular) while RRMSE(ν) decreased as x increased (Fig. S1). Therefore, sampling designs associated to an upper range of x values formed a range of Pareto optimal strategy within fractal designs (Fig. 4). This upper range of x values shrinked as a_s increased towards $L/6$. In situation B, as a_s increased above $L/6$, RRMSE(a_s) still showed a U-shaped profile as x increased (Fig. S1), while RRMSE(ν) also began to show a U-shaped profile with x. As a result, only an intermediary range of x led to Pareto-optimal sampling strategies within fractal designs (Fig. 4). This range gradually shifted towards more irregular designs (lower x) as a_s increased. When $a_s > 1.6$, the variation in errors among fractal designs became very small (Fig. S2), in line with our theoretical expectations. Designs could all be considered equivalent in practice and were included within the Pareto-front given our 1% tolerance threshold on estimation error.

Hybrid designs alone — In situation A, RRMSE(a_s) monotonically decreased while RRMSE(ν) monotonically increased as p increased (i.e. more random designs; Fig. S1). Therefore, any hybrid design along the gradient of p was a Pareto-optimal strategy (Fig. 4). As a_s crossed the frontier with situation B, the variation of RRMSE(a_s) rapidly shifted from decreasing to increasing with p while the profile of RRMSE(ν) remained unchanged. Therefore, the Pareto-front observed in situation A abruptly collapsed, only the most regular designs ($p \leq 5/27$) remained Pareto-optimal among hybrid designs. When a_s increased in situation B, the estimation errors of all hybrid designs converged towards the same value (Fig. S2) as theoretically expected. This resulted in all the designs gradually coming back to the pareto front, given ou 1% tolerance on error rates (Fig. 4).

Fractal and hybrid designs together — In situation A, fractal designs never excluded a hybrid design and, reciprocally, hybrid designs only rarely excluded a fractal designs (Fig. 4). Fractal designs that were Pareto-optimal within their own type either lead to larger RRMSE(ν) and lower RRMSE(a_s) than hybrid designs, or generated errors comparable to hybrid designs on both parameters. We illustrated this point for $a_s = 0.05$ (Fig. 5; see also $a_s = 0.036, 0.092$ in Fig. S2). By contrast, in situation B, hybrid designs led to the exclusion of fractal designs from the Pareto-front for all a_s values where the errors of designs still had opportunity to noticeably vary.

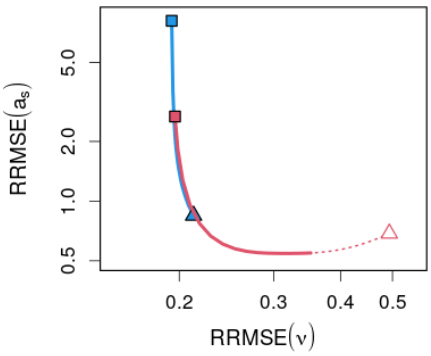

Figure 5: Joint Pareto front of fractal and hybrid sampling designs for a level of autocorrelation range in situation A (i.e. below grid mesh size; $a_s = 0.05$). Lines show the RRMSEs of fractal (red line, browsing the 28 values of x) and hybrid (blue line, browsing the 28 values of p) designs. The end of lines materialized with a square correspond to the most regular design ($x = 1$ for fractal designs; $p = 0$ for hybrid designs) while the end of lines materialized with a triangle corresponds to the most irregular design ($x = 10^{-1.5}$ for fractal designs; $p = 1$ for hybrid designs). Dashed parts of lines (and open end symbols) show designs that are not Pareto-optimal within their own type, while solid parts of lines (and solid end symbols) show those that are.

Numerical analysis of Pareto fronts in problem 2

Situations A ($a_s < L/6$) and B ($a_s > L/6$) used in problem 1 also proved useful to compare RRMSEs of designs in problem 2.

Gradient environment

The profiles of RRMSE(a_s) in all designs and situations remain identical to problem 1.

Fractal designs alone — In situation A, RRMSE(γ) harboured a U-shaped pattern as x increased (i.e. more regular designs; Fig. S3), like RRMSE(a_s). Therefore only fractal sampling designs with intermediary x, comprised between the minima for RRMSE(a_s) and RRMSE(γ), were Pareto optimal within the set of fractal designs. The intermediary range of Pareto-optimal x values gradually shifted towards higher values as a_s increased (Fig. 6). When a_s increased above $L/6$ (situation B), the RRMSE(γ) remained U-shaped (Fig. S3), like RRMSE(a_s), still yielding a intermediary range of Pareto-optimal x values. However, contrary to situation A, the range shifted towards lower x values (i.e. more irregular designs) as a_s increased. When a_s further increased in situation B, the estimation errors of all fractal designs converged towards the same value (Fig. S5) as theoretically expected. This resulted in all the designs gradually coming back to the Pareto front, given ou 1% tolerance on error rates (Fig. 6).

Hybrid designs alone — In situation A, RRMSE(γ) monotonically increased with p (i.e. more random designs; Fig. S3), while RRMSE(a_s) decreased. Therefore, any hybrid design along the gradient of p was a Pareto-optimal strategy (Fig. 6). When a_s increased above $L/6$ (situation B), the RRMSE(γ) remained monotonically increasing with p (Fig. S3), while we saw that RRMSE(a_s) abruptly shifted from decreasing to increasing. Therefore,

here again, only the most regular designs ($p \leq 5/27$) remained on the Pareto front. When a_s further increased in situation B, the estimation errors of all hybrid designs converged towards the same value (Fig. S5) as theoretically expected. This resulted in all the designs gradually coming back to the Pareto front, given ou 1% tolerance on error rates (Fig. 6).

Fractal and hybrid designs together — In situation A, fractal designs excluded all hybrid designs from the Pareto front (Fig. 6). We illustrated this point more precisely for $a_s = 0.05$ (Fig. 7). By contrast, in situation B, the introduction of regular hybrid designs ($p \leq 5/27$) led to the exclusion of fractal designs from the Pareto-front for all a_s values where the errors of designs still had opportunity to noticeably vary.

U-shaped environment

Fractal designs alone — In situation A, RRMSE(γ) harboured a decreasing pattern as x increased (i.e. towards more regular designs; Fig. S4), while we saw that RRMSE(a_s) harboured a U-shaped profile. Therefore, designs associated to an upper range of x values formed a range of Pareto optimal strategy within fractal designs (Fig. 6). This upper range of x values shrinked as a_s increased towards $L/6$. When a_s increased above $L/6$ (situation B), the RRMSE(γ) still decreased with x (Fig. S4), while RRMSE(a_s) remained U-shaped, still yielding an upper range of Pareto-optimal x values. However, contrary to situation A, the range became broader, including lower x values (i.e. more irregular designs), as a_s increased. When a_s further increased in situation B, the estimation errors of all fractal designs converged towards the same value (Fig. S5) as theoretically expected. This resulted in all the designs gradually coming back to the Pareto front, given ou 1% tolerance on error rates (Fig. 6).

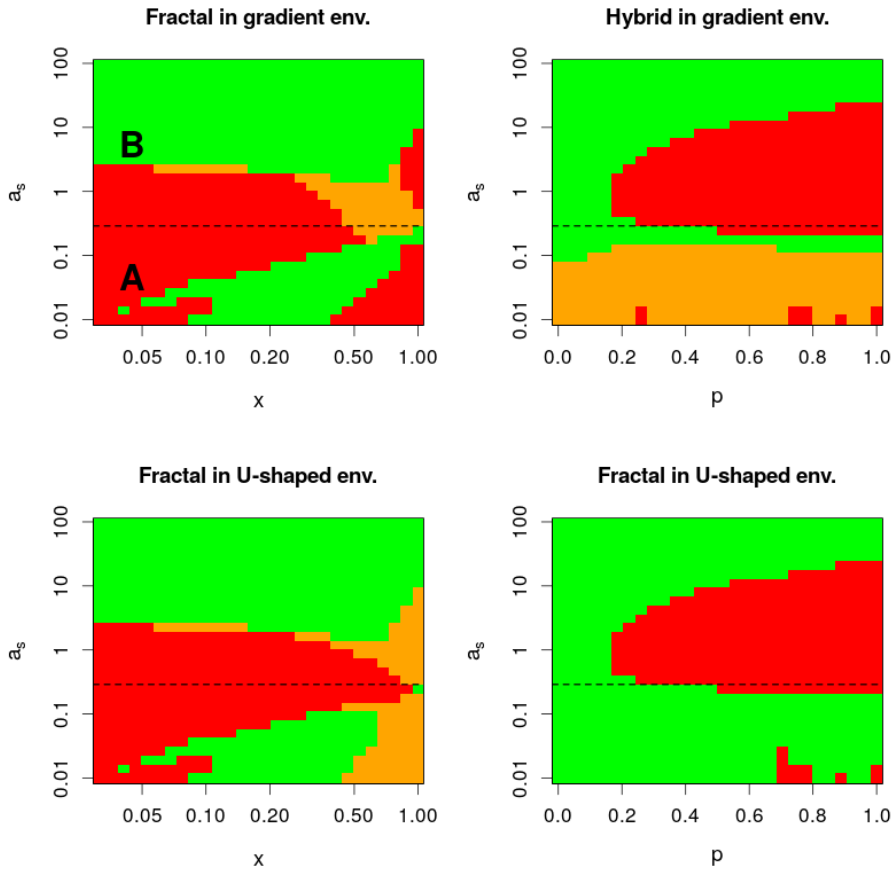

Figure 6: Pareto fronts of sampling designs in problem 2. Panels correpond to one type of design (fractal or hybrid) in one type of environment (gradient or U-shaped). In each panel, the horizontal axis presents the regularity parameter of the design (x or p for fractal or hybrid designs respectively). Colored pixels indicates, for given a_s and regularity parameter values, whether the design is Pareto-optimal within fractal and hybrid designs altogether (green pixels) or within its own type only (orange pixels). Red pixels show designs that are not Pareto optimal within their own type. Therefore, green and orange pixels show designs that are Pareto-optimal within their own type, and orange pixels shows the fraction of these designs that are eliminated when introducing the other type of design in the comparison. The horizontal dashed line $a_s = L/6$ shows the limit between two situations discussed in main text : 'low' autocorrelation range (A) and 'high' autocorrelation range (B).

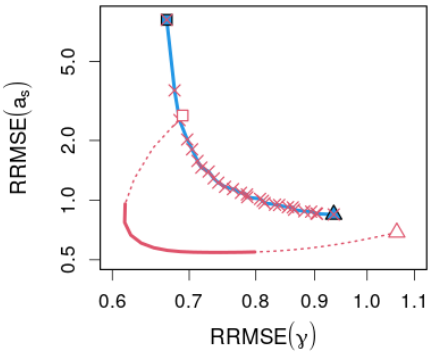

Figure 7: Joint Pareto front of fractal and hybrid sampling designs in problem 2 with gradient environment, for a level of autocorrelation range in situation A (i.e. below grid mesh size; $a_s = 0.05$). Lines show the RRMSEs of fractal (red line, browsing the 28 values of x) and hybrid (blue line, browsing the 28 values of p) designs. The end of lines materialized with a square correspond to the most regular design ($x = 1$ for fractal designs; $p = 0$ for hybrid designs) while the end of lines materialized with a triangle corresponds to the most irregular design ($x = 10^{-1.5}$ for fractal designs; $p = 1$ for hybrid designs). Dashed parts of lines (and open end symbols) show designs that are not Pareto-optimal within their own type, while solid parts of lines (and solid end symbols) show those that are. Red crosses indicate hybrid designs that are excluded from Pareto-optimal strategies when accounting for fractal designs (i.e. all the hybrid designs here).

Hybrid designs alone — Like for the case of a linear environmental covariate depicted above, RRMSE(γ) monotonically increased with p (i.e. more random designs; Fig. S3), irrespective of the value of a_s (and the situation A or B). Consequently, the situation is qualitatively similar to the one obtained with a linear environment.

Fractal and hybrid designs together — Fractal designs never excluded hybrid designs from the Pareto front (Fig. 6). By contrast, in situation A and B, hybrid designs excluded the most regular fractal designs from the Pareto front, because fractal design consistently showed markedly higher errors at esti-

mating the slope γ, but irregular ones could still ourperform hybrid designs at estimating autocorrelation range a_s. We illustrated this pattern $a_s = 0.05$ (Fig. 8).

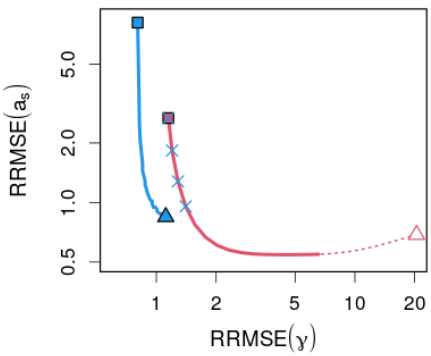

Figure 8: Joint Pareto front of fractal and hybrid sampling designs in problem 2 with U-shaped environment, for a level of autocorrelation range in situation A (i.e. below grid mesh size; $a_s = 0.05$). Lines show the RRMSEs of fractal (red line, browsing the 28 values of x) and hybrid (blue line, browsing the 28 values of p) designs. The end of lines materialized with a square correspond to the most regular design ($x = 1$ for fractal designs; $p = 0$ for hybrid designs) while the end of lines materialized with a triangle corresponds to the most irregular design ($x = 10^{-1.5}$ for fractal designs; $p = 1$ for hybrid designs). Dashed parts of lines (and open end symbols) show designs that are not Pareto-optimal within their own type, while solid parts of lines (and solid end symbols) show those that are. Blue crosses indicate fractal designs that are excluded from Pareto-optimal strategies when accounting for hybrid designs (i.e. more regular fractal designs with higher contraction parameter x here).

Discussion

We compared the estimation accuracy of hybrid and fractal sampling designs regarding two quantities: an effect acting on a response variable (fixed intercept or environmental covariate) and the spatial autocorrelation range of the response variable. These two estimation objectives are often simultaneously

present when analyzing, for instance, the spatial distribution of the occurrence or density of a target species in space at landscape scale (Ranius et al., 2010; Johansson et al., 2012). Our main findings are that: (i) hybrid designs do not necessarily spread along a trade-off between pure grid and pure random design regarding the two objectives; (ii) when the sampling budget is sufficiently large to make a grid design with mesh size lower than the autocorrelation range, grid becomes the best design among both hybrid and fractal designs for the two objectives; (iii) fractal designs can be better than hybrid designs at estimating short-ranged autocorrelation (i.e. ranges below feasible grid mesh size), but this often comes at a prohibitory cost on the ability at estimating other effects; (iv) when studying the effect of an environmental gradient spreading across the study area under short-ranged autocorrelation, fractal designs can outperform hybrid designs on both objectives. We discuss these four points below.

Hybrid designs do not necessarily browse a trade-off between pure grid and pure random design

For autocorrelation range smaller than the grid mesh size, hybrid designs spread along a Pareto-front between grid and random designs. In other words, increasing the degree of randomness (p) in the design tended to degrade the estimation accuracy on the effect but improved the estimation accuracy on the autocorrelation range. This pattern is in line with conclusions of a previous study (Bijleveld et al., 2012), and was robust to changing the target effect (mean or environmental covariate). The fact that grid design is inefficient at quantifying autocorrelation range when mesh size is above the autocorrelation range has been repeatedly emphasized, especially in the field of population genetics and isolation by distance studies (Sokal and Oden, 1978; Epperson and Li, 1997). Pairwise distances smaller than the grid mesh size are then needed to accurately estimate the autocorrelation range, and such smaller distances were provided by

the introduction of random points in our study, hence creating a Pareto front.

As the autocorrelation range increased above the threshold of grid mesh size, the above pattern vanished : adding too much randomness could become sub-optimal for both estimation objectives (effect and autocorrelation range). The maximum degree of randomness allowed to remain Pareto-optimal abruptly decreased with autocorrelation range. Here again, this pattern was robust to the particular effect considered (mean or environmental covariate). This result is not completely new: Bijleveld et al. (2012) identified that the relative performance of designs depended on the level of the underlying autocorrelation range. However, our results clearly made the point that introducing random sampling points in the design may become uninteresting in the context of large autocorrelation ranges, because accurately estimating effects and autocorrelation range did not require enriching the grid design with small pairwise distances. Adding small pairwise distances through increasing the degree of randomness led to higher pseudo-replication when estimating effects (a phenomenon already present for small autocorrelation ranges) but also proved inefficient to estimate the autocorrelation range, as distance decay of auto-correlation happened at large distances already covered by the grid design.

Our finding that Pareto front of hybrid designs collapsed for large autocorrelation ranges may seem at variance with Bijleveld et al. (2012) conclusion that, overall, there was a Pareto front of hybrid designs between grid and random. The apparent contradiction simply stems from the fact that these authors averaged the performance of designs across autocorrelation levels. Because the magnitude of errors on autocorrelation estimation rapidly increases as the autocorrelation range decreases, global averaging of errors tends to give high weights to scenarios with short autocorrelation range. This may lead to over-generalizing patterns that are in fact specific to small autocorrelation range values. The

very high magnitude of errors at the lower end autocorrelation range affects any metrics of performance integrating over an interval of autocorrelation ranges, over-emphasizing a small number of cases corresponding to small 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.

When budget allows setting a mesh size lower than the autocorrelation range, grid outperforms other explored designs

As we just mentioned, when autocorrelation range was above the grid mesh size (and not too high to ensure that designs still differed in accuracy), adding a proportion of randomness in hybrid designs became sub-optimal: it increased simultaneously errors on the effect and the autocorrelation range estimates. Only hybrid designs with a small degree of randomness that generated less than a 1% increase of error compared to grid design remained on the Pareto front. Importantly, fractal designs were also suboptimal compared to grid, because their main contribution was, like random sampling points, to generate small pairwise distances.

These findings were all robust to the effect considered and emphasized the need of comparing expected autocorrelation range to feasible grid mesh size when choosing a sampling design. This is all the more important that the relative magnitudes of autocorrelation range and feasible grid mesh size have been shown to determine other statistical properties of sampling designs, beyond those considered here. For instance, a recent simulation study (Perret et al., 2022) focusing on the problem of estimating the mean density of a population (analogous to ν in problem 1 here) showed that when grid mesh size matches

the autocorrelation range of the density, the contrast in error estimation of the intercept was maximized between grid and random design, the two extreme hybrid strategies.

Before choosing the spatial design, empiricists interested in assessing both an effect and the autocorrelation range of a variable of interest should therefore bear in mind three important quantities : the expected autocorrelation range (a) of the target variable, the sampling budget (B, the feasible number of sampling points) and the study site area (A).

A priori knowledge about expected autocorrelation range can stem from previous studies of spatial patterns on similar systems or from mechanistic models relating patterns to measurable mechanisms at a lower scale. The latter reductionist approach has been used on closely-related topics like making prediction of landscape 'scale of effect' on species distribution based on species dispersal traits (Jackson and Fahrig, 2012; Laroche et al., 2020). However, it is not always easier to generate knowledge about dispersal traits than to set up a pilot study on distribution patterns, and empirical studies often rather use patterns to better understand dispersal than dispersal to anticipate spatial patterns (especially in population and landscape genetics; Manel and Holderegger (2013)).

Combining sampling budget with study site area can lead to feasible grid mesh size. Generally, we expect that this mesh size will be about $\sqrt{A/B}$. Thus we suggest as a rule of thumb that when $a > \sqrt{A/B}$ grid design should be preferred. Consider the practical case where one wants to define a strategy to position 27 sampling plots in a forest stand of about 16 ha in order to study the distribution of saproxylic beetles species. The mesh size of a regular grid spread over the stand would then be of c.a. 75m (with fluctuations depending on the geometry of both stand 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 within the range $100 - 1000$m. One thus expects the mesh size to be smaller than autocorrelation range in this example, and that grid design is a robust choice. By contrast, if the study area is a whole forest of say 3600 ha, the mesh size of a regular grid is likely to be of c.a. 1200m, and the sampling design choice may differ as we discuss below.

The previous example shows the critical interplay between the surveyed area and sampling budget. Keeping the studied area limited could be considered as an option to make grid design an unambiguously robust choice. However, this may often be at variance with scientific objectives. For instance, the area study cannot be freely reduced when one aims at assessing environmental covariates effects on a target response variable. Then, it is crucial to keep environmental gradients within the area of study, which may result in constraints on the extent of the area. Such need of covering environmental gradients clearly stands out from our results : the accuracy of the estimate of a covariate effect increases with the standard deviation of the covariate across the area of study (see e.g. our asymptotic analysis of error on γ for short autocorrelation range). It has also been well developped in the litterature (Field et al., 2009; Albert et al., 2010).

Fractal designs are better than others at estimating short-ranged autocorrelation, but the cost on estimating other effects is prohibitory

When confronted to a practical situation where budget is not sufficient to make grid mesh size lower than expected autocorrelation range, choosing the design becomes less straightforward. We already discussed the fact that, in those situations, hybrid designs are spread along a trade-off between grid and random designs : increasing the degree of randomness generates small pairwise distances useful to estimate short autocorrelation ranges, but increases the estimation er-

ror on other effects because of pseudo-replication.

Fractal designs seems to offer a way to deterministically generate arbitrarily small pairwise distances. When considering autocorrelation ranges for which even the typical pairwise distances generated by random designs are still too large, this feature allows fractals to reach the lowest estimation error of autocorrelation range among all explored designs, and thus to constitute Pareto-optimal strategies.

However, this theoretical advantage has to be nuanced. First, 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 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. In words, one cannot aribtrarily increase the contrast among scales of the design.

Beyond these practical constraints, our results showed that decreasing the contraction parameter to estimate short autocorrelation ranges also generated a marked agregation of points in space and a very sparse design (i.e. large empty spaces in the middle of the area). This degraded a lot the ability to

estimate the mean of response variable across the study area (because of pseudo replication) as well as the effect of a U-shaped environmental covariate (because of both pseudo-replication and sparsity). Choosing such fractal designs thus amounted to focus nearly exclusively on estimating the autocorrelation range. But if the study mainly focuses on estimating autocorrelation range, there is no real necessity of keeping a large study area anymore, and reducing the area of study may be a better option than implementing a fractal design. It should be noted that the problem of sparsity that prevents capturing non linear trends in space could be adressed by using a more complex geometrical shape than a triangle to generate the fractal design (e.g. a triangle with one additional site at the center). Because using more points in the elementary shape that is self-replicated has an exponential impact on the total number of points in the design, this increase in complexity should probably be combined with subsampling, as suggested by Marsh and Ewers (2013).

Fractal design is the best choice when studying the effect of a linear environmental gradient under short-ranged autocorrelation.

There are many empirical questions related to species distributions that involves monotonic environmental gradient in space : the effect of latitudinal gradient and associated climatic trends across a biogeographic region, the effect of altitude on a mountaine slope, the effect of salinity in an estuarine river etc. When the target effect is associated to an environmental covariate regularly changing in space, the sparsity of fractal designs turned out to be advantageous. Indeed, moderate contraction parameter values yielded a higher standard deviation of the covariate in the design compared to hybrid designs, hence leading to lower error of the estimated effect. In the meantime, aggregating sampling points in space lead to better estimates of short autocorrelation ranges. Thus fractal designs with intermediate contraction outperformed hybrid designs on all ob-

jectives and constituted the only Pareto-optimal strategies. However, too much contrast between scales of the fractal design ultimately lead to similar pseudo-replication problems as those presented in previous section, which emphasizes the need of determining bounds of appropriate values for the contraction parameter. In our results, the range of interesting contraction parameter values shifted towards higher values (more regular design, with milder contrast between scales) as autocorrelation range increased, but precise quantitative bounds should probably be explored through simulations on a case-by-case basis, depending on the expected order of magnitude of autocorrelation range.

Although monotonic environmental gradients in space are frequent, they are often mixed with more rugged covariates in empirical studies (e.g. soil cover in agricultural landscapes). For covariates that do not steadily vary in space, fractal designs are not optimal, as already suggested by our study of a simple U-shaped covariate. Indeed, the strength of fractal design to estimate the effect of a monotonic environmental gradient in space in space comes from the fact that sampling effort is concentrated at extreme points of the study area, which are ideal to detect a steady change across the area but have little power to depict complex variations within the area. Non-monotonicity of target covariates should thus bring operators to consider hybrid designs, and use the degree of randomness to tune the priority between estimating effects and autocorrelation range.

Conclusions

A key contribution of our work has been to pinpoint that adequate sampling strategies to jointly estimate environmental effects and autocorrelation range of a target variable depends on the sampling budget and associated feasible grid mesh size over the study area. This mesh size defines a boundary between con-

trasted situation leading to distinct decision for operators. Strikingly, when the autocorrelation range is large (i.e. above feasible grid mesh size), grid designs stands out as the best design on the two criteria explored here, and may thus be a robust choice in practical application. By contrast, when autocorrelation range is short, the choice of designs depended on the effect of interest. Importantly, fractal designs outperformed hybrid designs when studying monotonous environmental gradients across the study area. In other situations involving non-linear rugged covariates or simply estimating the global mean of target variable over the study area, it seemed more efficient and less complicated to implement hybrid designs.

Overall, the niche for fractal designs exists but it may appear quite narrow, limited to jointly estimating short autocorrelation ranges and the effect of monotonic environmental gradients across space. It should nonetheless be noted that we evaluated designs on a simple scenario with a parsimonious autocorrelation structure (only positive autocorrelation, steadily decreasing in space with a single well-defined range). However, biological patterns often stem from heterogeneous drivers acting at different scales (Thuiller et al., 2015; Ricklefs, 2008). For instance, when studying the distribution of a species in space, competition among conspecifics may drive negative autocorrelation (i.e. underdispersion) at a given spatial scale, while limited dispersal may drive positive autocorrelation at a coarser grain (see Bolker and Pacala (1997) for a theoretical analysis of these effects in plant populations). In the meantime, non-measured environmental variables may also leave various autocorrelation signatures across scales (Legendre, 1993). 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 hierarchy of the design.

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.

This research was funded in whole 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 author 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, N. E., and Thuiller, W. (2010). Sampling in ecology and evolution - bridging the gap between theory and practice. *Ecography*, 33(6):1028–1037.

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

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

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

Bolker, B. and Pacala, S. W. (1997). Using moment equations to understand stochastically driven spatial pattern formation in ecological systems. *Theoretical Population Biology*, 52(3):179–197.

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

Dormann, C. F., McPherson, J. M., Araújo, M. B., Bivand, R., Bolliger, J., Carl, G., Davies, R. G., Hirzel, A., Jetz, W., Kissling, D. W., Kühn, I., 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.

Epperson, B. K. and Li, T.-Q. (1997). Gene Dispersal and Spatial Genetic Structure. *Evolution*, 51(3):672.

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.

Jackson, H. B. and Fahrig, L. (2012). What size is a biologically relevant landscape? *Landscape Ecology*, 27(7):929–941.

Johansson, V., Ranius, T., and Snäll, T. (2012). Epiphyte metapopulation dynamics are explained by species traits, connectivity, and patch dynamics. *Ecology*, 93(2):235–241.

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.

Laroche, F., Balbi, M., Grébert, T., Jabot, F., and Archaux, F. (2020). Three points of consideration before testing the effect of patch connectivity on local species richness: patch delineation, scaling and variability of metrics. *bioRxiv*, page 640995.

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.

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

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

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

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

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

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

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

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

Perret, J., Charpentier, A., Pradel, R., Papuga, G., and Besnard, A. (2022). Spatially balanced sampling methods are always more precise than random ones for estimating the size of aggregated populations. *Methods in Ecology and Evolution*, 13(12):2743–2756.

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

Ranius, T., Johansson, V., and Fahrig, L. (2010). A comparison of patch connectivity measures using data on invertebrates in hollow oaks. *Ecography*, 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.

Sokal, R. R. and Oden, N. L. (1978). Spatial autocorrelation in biology: 1. Methodology. *Biological Journal of the Linnean Society*, 10(2):199–228.

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.

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.

Supplementary figures

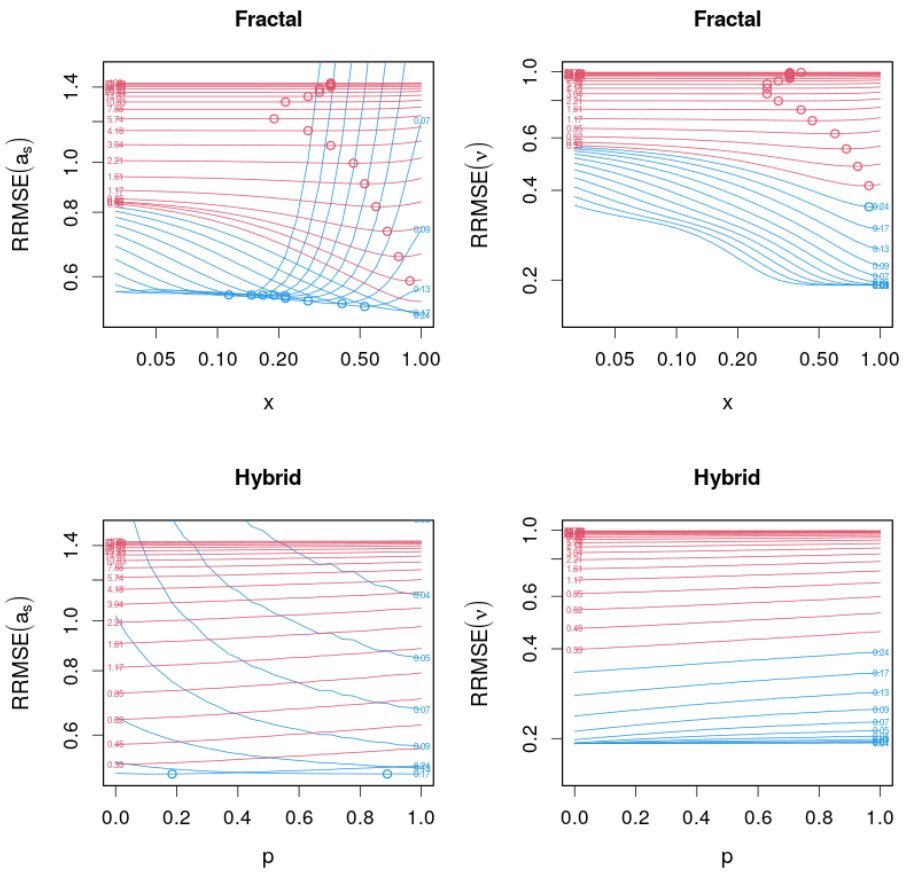

Figure S1: Relative root mean square error of estimation (RRMSE) of parameters in problem 1 for various sampling designs. Each panel corresponds to one parameter (a_s or ν) for one type of design (fractal or hybrid). The horizontal axis presents the parameter of the design that is varied (x or p for fractal or hybrid designs respectively). Each curve on a panel correspond to a level of a_s. The value of a_s is reported at one end of the curve. Blue and red curves correspond to situation A and B respectively. When a curve has an inner minimum over the range of design parameter values explored, it is presented with a dot.

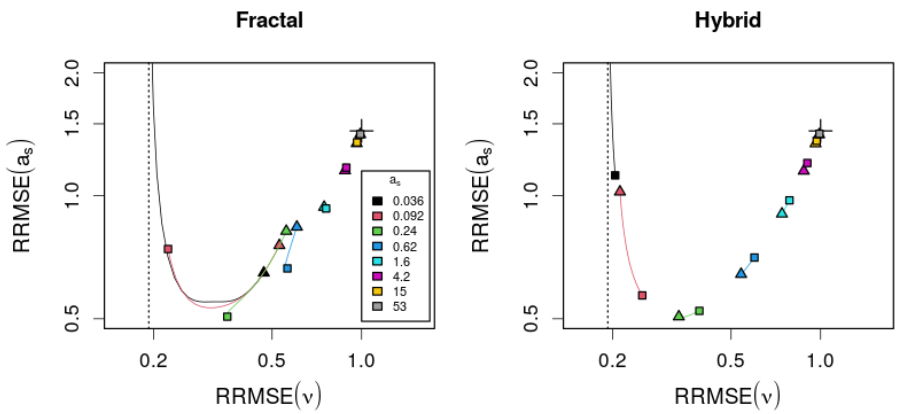

Figure S2: Relative root mean square error of estimation of exponential mean (ν) and autocorrelation range (a_s) for fractal (left panel) and hybrid (right panel) designs along a gradient of a_s values. We selected 8 values of a_s among the 30 values explored for readibility. For each selected value, we present a colored line of dots showing the RRMSEs of the various designs (28 values of x for fractal designs in the left panel, 28 values of p for hybrid designs in the right panel; see methods). For each line (i.e. each a_s value), the end of the line materialized with a square correspond to the most regular design ($x = 1$ for fractal designs; $p = 0$ for hybrid designs) while the end of the line materialized with a triangle corresponds to the most irregular design ($x = 10^{-1.5}$ for fractal designs; $p = 1$ for hybrid designs). The vertical dotted line shows the predicted asymptote when $a_s \to 0$, the cross shows the predicted limit when $a_s \to +\infty$.

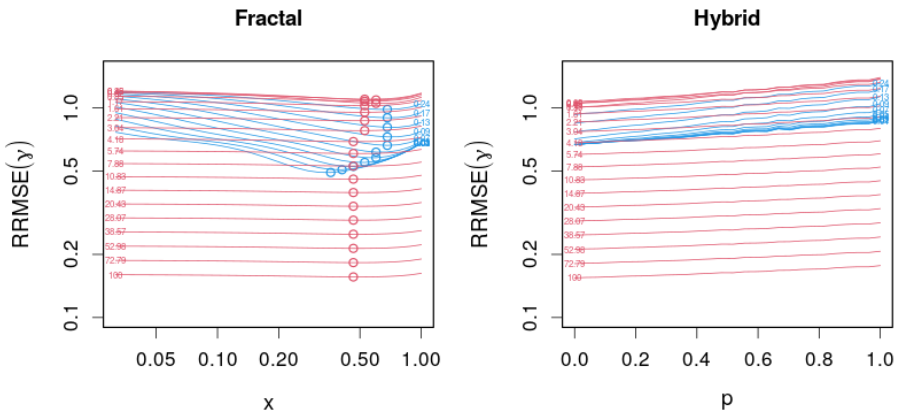

Figure S3: Relative root mean square error of estimation (RRMSE) of environmental effect parameter γ in problem 2 with a linear environment for various sampling designs. Each panel corresponds to one type of design (fractal or hybrid). The horizontal axis presents the regularity parameter of the design used (x or p for fractal or hybrid designs respectively). Each curve on a panel correspond to a level of a_s. The value of a_s is reported at one end of the curve. Blue and red curves correspond to situation A and B respectively. When a curve has an inner minimum over the range of design parameter values explored, it is presented with a dot.

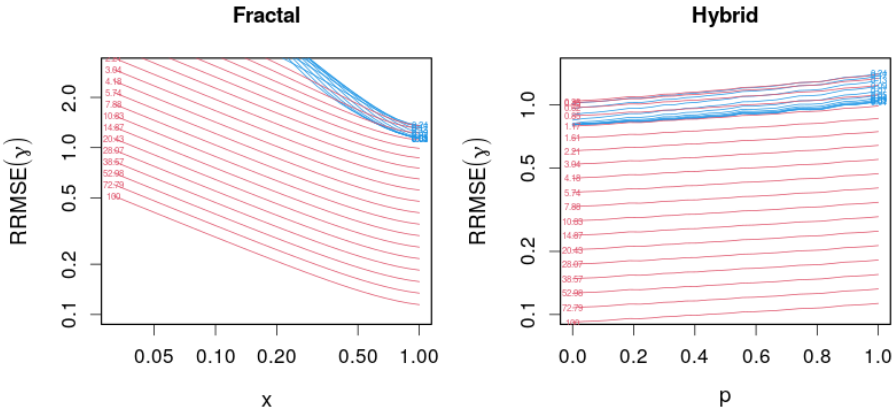

Figure S4: Relative root mean square error of estimation (RRMSE) of environmental effect parameter γ in problem 2 with a U-shaped environment for various sampling designs. Each panel corresponds to one type of design (fractal or hybrid). The horizontal axis presents the regularity parameter of the design used (x or p for fractal or hybrid designs respectively). Each curve on a panel correspond to a level of a_s. The value of a_s is reported at one end of the curve. Blue and red curves correspond to situation A and B respectively. Mind the difference in ordinate axis graduation between the panels.

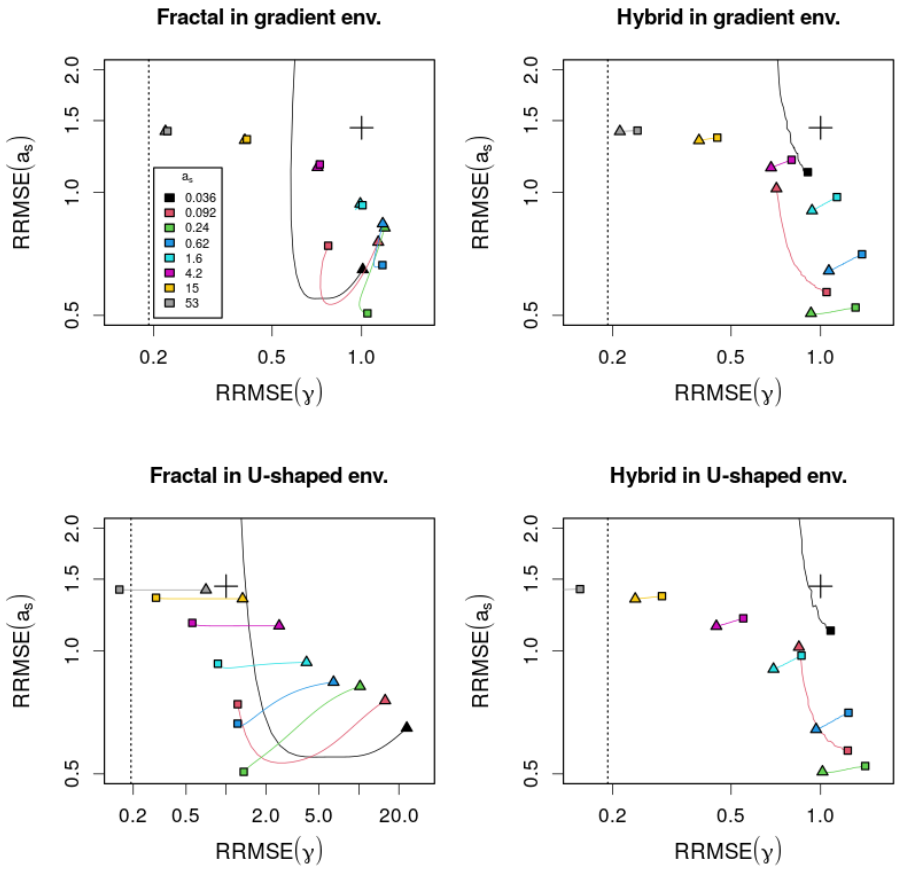

Figure S5: Relative root mean square error of estimation of exponential slope (γ) for various sampling designs and environmental profiles, along a gradient of a_s values. Left (resp. right) panels show results for fractal (resp. hybrid) designs. Upper (resp. lower) panels show results for a West-East gradient (resp. a U-shaped environmental profile). We selected 8 values of a_s among the 30 values explored for readibility. For each selected value, a colored line of dots shows the RRMSEs of the various designs (28 values of x for fractal designs in the left panels, 28 values of p for hybrid designs in the right panels; see methods). For each line (i.e. each a_s value), the end of the line materialized with a square correspond to the most regular design ($x = 1$ for fractal designs; $p = 0$ for hybrid designs) while the end of the line materialized with a triangle corresponds to the most irregular design ($x = 10^{-1.5}$ for fractal designs; $p = 1$ for hybrid designs). The horizontal dotted line shows the predicted asymptote when $a_s \to +\infty$ on a log-scale.