

Bergische Universität Wuppertal

Fachbereich Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational
Mathematics (IMACM)

Preprint BUW-IMACM 20/08

Hanno Gottschalk und Karsten Kahl

**Coarsening in Algebraic Multigrid using
Gaussian Processes**

April 25, 2020

<http://www.math.uni-wuppertal.de>

Coarsening in Algebraic Multigrid using Gaussian Processes*

Hanno Gottschalk[†] Karsten Kahl[†]

April 24, 2020

Abstract

Multigrid methods have proven to be an invaluable tool to efficiently solve large sparse linear systems arising in the discretization of partial differential equations (PDEs). Algebraic multigrid methods and in particular adaptive algebraic multigrid approaches have shown that multigrid efficiency can be obtained without having to resort to properties of the PDE. Yet the required setup of these methods poses a not negligible overhead cost. Methods from machine learning have attracted attention to streamline processes based on statistical models being trained on the available data. Interpreting algebraically smooth error as an instance of a Gaussian process, we develop a new, data driven approach to construct adaptive algebraic multigrid methods. Based on Gaussian a priori distributions, Kriging interpolation minimizes the mean squared error of the a posteriori distribution, given the data on the coarse grid. Going one step further, we exploit the quantification of uncertainty in the Gaussian process model in order to construct efficient variable splittings. Using a semivariogram fit of a suitable covariance model we demonstrate that our approach yields efficient methods using a single algebraically smooth vector.

1 Introduction

The solution of large sparse linear systems of equations

$$Ax = b, \quad b \in \mathbb{R}^n, \quad A \in \mathbb{R}^{n \times n} \text{ symmetric positive definite}, \quad (1)$$

that arise in the discretization of partial differential equations, typically makes up the bulk of computations in modern scientific computing. It is thus of utmost importance to come up with efficient algorithms to solve these systems of equations. By exploiting a separation of scales multigrid methods can achieve optimal linear complexity for this task, but heavily rely on the availability of expert knowledge about the particular partial differential equation that the linear systems originates from as well as the employed discretization scheme; cf. [13]. Due to the fact that this information might not readily be available or that there is no known geometric multigrid construction, the concept of algebraic multigrid methods has been introduced in [4, 8, 29, 31, 35]. Efficiency in algebraic multigrid methods is achieved by pairing a simple iterative scheme, the *smoother*, with a variational coarse grid correction. Assuming a smoother

*This work was partially funded by Deutsche Forschungsgemeinschaft (DFG) Transregional Collaborative Research Centre 55 (SFB/TRR55)

[†]Bergische Universität Wuppertal, IMACM, Faculty of Mathematics and Natural Sciences, {hgottsch, kkahl}@uni-wuppertal.de

is defined by $M \approx A^{-1}$, the error propagator of a two-grid algebraic multigrid method with Galerkin coarse grid construction is given by

$$E_{2g} = (I - MA)(I - P(P^T AP)^{-1}P^T A)(I - MA). \quad (2)$$

Due to the variational construction of coarse grid correction the whole setup of an algebraic multigrid method can be reduced to the definition of the interpolation operator P . In particular, the dimension of the coarse space, n_c , the interpolation relations, i.e., the sparsity pattern of $P \in \mathbb{R}^{n \times n_c}$ and its entries need to be defined. Typically these tasks are split into two parts. Finding n_c and the sparsity pattern of P is often referred to as the *coarsening* problem, while determining the entries of P is known as the *interpolation* problem.

In the first algebraic multigrid methods [8, 29] operator based approaches have been suggested to solve both problems. In case A has M -matrix structure, e.g., as a particular discretization of an elliptic partial differential equation, it can be shown that these approaches lead to methods with fast convergence. However, these early approaches rely heavily on assumptions about the underlying problem and therefore cannot be extended significantly beyond the M -matrix case. A huge step in overcoming this limitation has been the introduction of adaptivity in algebraic multigrid methods [7, 14, 15]. Common in all adaptive approaches in algebraic multigrid methods is the idea to guide the definition of interpolation, posed in terms of the coarsening and interpolation problem either by using spectral information about A and/or the smoothing iteration. Due to the fact that explicit calculation of (partial) spectra is prohibitively expensive these methods rely on an iterative approximation process that makes use of the emerging multigrid hierarchy. In contrast to the interpolation problem, where several approaches showed promising performance [5, 7, 23, 27, 34, 36], the coarsening problem turned out to be harder to tackle.

Many of the approaches that have been tried to solve the coarsening problem in adaptive algebraic multigrid methods revolve around the idea of strength of connection, a concept introduced in classical algebraic multigrid. This includes approaches based on binary variable relations such as [5, 10, 11, 27] but approaches that take relations of more than two variables into account such as [20]. The detection of strongly connected pairs of variables is also found in aggregation-based approaches such as [10, 22, 25, 26]. In some sense the idea of compatible relaxation, introduced in [12], comes closest to general applicability, but is hard to integrate and mesh with typical solutions to the interpolation problems, i.e., the definition of the entries of P , in adaptive algebraic multigrid approaches.

In this paper we propose a new way of solving the coarsening problem which resembles the least squares interpolation approach of the bootstrap algebraic multigrid framework [7, 20]. By considering the test vectors of the bootstrap framework as instances of a Gaussian process [1, 3, 33] we are able to apply techniques from machine learning, especially the concept of conditional (a posteriori) distributions. To calibrate this statistical model to the data, we use parametric semivariogram models to fit the covariance structure to the data provided by algebraically smooth test vectors. This enables us to efficiently solve the coarsening and interpolation problem at the same time. The idea of Kriging interpolation is to view the values of test vectors at coarse grid variables as partial observations of a Gaussian process. Based on these observed values at the coarse grid variables, expected values and errors at the fine grid variables can be derived by computing conditional expectations and variances. This method originally stems from spatial interpolation in geostatistics, see e.g. [16] but has been widely used in various machine learning and engineering tasks in the past [18, 21]. Its main advantage lies in its statistical properties as best linear unbiased predictor, given

the data of the field on the ‘observed locations’ – the coarse grid – and the spatial correlation structure of the data. Coincidentally, the Kriging interpolator bears some resemblance to the least-squares formulation of interpolation introduced in [7] and more specifically the operator-based modification of it found in [24].

Inheriting from the spatial correlation structure of the Gaussian process, the conditional variance given the observations on certain points, is low at points with sufficient observations in the neighborhood and high elsewhere. This additional piece of information can now be used if one has to decide on where to make the next observation. In this paper we apply a greedy optimization procedure, picking the point of highest conditional variance given coarse grid results, to adaptively refine the coarse grid by subsequently adding fine grid points to the coarse grid until the conditional variance on all remaining fine grid points is small.

A statistical view on the adaptive setup in algebraic multigrid methods is not completely new. It has been used in [22] to motivate a definition of strength of connection based on a measure of correlation present in test vectors, but this construction lacks the framing of Gaussian processes. Other related work in the context of the integration of ideas from stochastics into multigrid method has been presented in [28]. While this work also contains a deep mathematical analysis of convergence, there are several points where our work takes a different route. First of all, our work is purely algebraic and only uses structures that can be derived from the matrix A . Hence we do not use any structures that stem from the spatial structures of the underlying PDE and thus we do not have techniques at our disposal that use spectral equivalence and other techniques based on harmonic analysis. Instead, we only use distance notions between ‘nodes’ that can intrinsically be derived from A . Also, we resort to data driven estimation of correlation structures instead of an analytic derivation of these structures from the matrix A . This empirical approach largely reduces the computational cost in the choice of priors. Furthermore, while the prior chosen in [28] is supported by error analysis in the A -norm, it has the disadvantage of producing a generalized random field with distributional paths [19], which is a disputable choice for the prior belief on the solution to a PDE. The price we pay is to use a more experimental and less deeply founded approach.

In section 2 we give a short introduction into the construction of algebraic multigrid methods and highlight how adaptivity can be used in order to capture the nature and underlying structure of the problem at hand, especially in terms finding suitable coarsenings. Then we give an overview on Gaussian processes in section 3, where we motivate the connection to the adaptive setup process and discuss interpolation in the context of these processes. This leads us directly to the formulation of Kriging interpolation in section 3.2, which we discuss and compare to the least squares interpolation approach in appendix A. Finally, using additional heuristics we present our new adaptive coarsening approach in section 4 before closing with numerical tests in section 5 demonstrating the potential of our approach and some final remarks in section 6.

2 Adaptive algebraic multigrid methods

The efficiency of multigrid methods lies in the complementarity of the smoothing iteration and the coarse grid correction. Algebraic multigrid methods construct complementarity without relying on knowledge of the underlying problem or the employed discretization strategy. Assuming that A is symmetric positive definite, it is common

to eigenvalues close to 1, i.e., components that are slow to converge—also known as algebraically smooth error components [29].

The adaptive construction of algebraic multigrid methods can thus be interpreted as the generation of a low dimensional, sparse representation of the space of algebraically smooth error. In the context of the bootstrap algebraic multigrid framework and its extensions [7, 20, 24] this is facilitated by the use of a set of algebraically smooth test vectors $\{v^{(1)}, \dots, v^{(K)}\}$ which are obtained by smoothing initially random test vectors with entries that stem from a normal distribution. By drawing on a connection to Gaussian processes we are now going to modify the construction of interpolation in this setting using ideas that originate in statistical geophysics.

3 Gaussian Processes

In this section, we introduce Gaussian processes [1] and the Kriging predictor [2, 30] to interpolate from coarse to fine grid points. We discuss various approaches for modeling the covariance of the underlying Gaussian process, including semivariogram estimation based on test vectors.

3.1 Gaussian Stochastic Processes and Kriging

A stochastic process with the index set \mathcal{I} is a collection of real valued random variables $X = \{X_i \mid i \in \mathcal{I}\}$ on a common probability space $(\Omega, \Sigma, \mathbb{P})$, where $\mathbb{E}[\cdot]$ denotes the expected value with respect to \mathbb{P} . Here Ω is the event set, Σ the sigma field and \mathbb{P} the probability measure. X is a Gaussian process, if the distribution of any finite collection $X_{\mathcal{H}} = (X_{i_1}, \dots, X_{i_q})^T$ for finite subsets $\mathcal{H} = \{i_1, \dots, i_q\} \subseteq \mathcal{I}$ is a multivariate Gaussian distribution $N(\mu_{\mathcal{H}}, C_{\mathcal{H}})$ with probability density

$$f_{\mathcal{H}}(x_{\mathcal{H}}) = \frac{1}{\sqrt{2\pi^q} |C_{\mathcal{H}}|^{\frac{1}{2}}} e^{-\frac{1}{2}(x_{\mathcal{H}} - \mu_{\mathcal{H}})^T C_{\mathcal{H}}^{-1} (x_{\mathcal{H}} - \mu_{\mathcal{H}})}, \quad x_{\mathcal{H}} \in \mathbb{R}^m, \quad (4)$$

where $|\cdot|$ denotes the determinant, $\mu_{\mathcal{H}} \in \mathbb{R}^q$ is the expected value and $C_{\mathcal{H}} \in \mathbb{R}^{q \times q}$ is a positive definite covariance matrix. Consistency of the finite dimensional distributions in the sense of Kolmogorov [32] implies that there exist functions

$$\mu : \mathcal{I} \rightarrow \mathbb{R} \text{ and } C : \mathcal{I} \times \mathcal{I} \rightarrow \mathbb{R}$$

such that $\mu_{\mathcal{H}} = (\mu_{i_1}, \dots, \mu_{i_q})^T$ and $C_{\mathcal{H}} = (C_{i,j})_{i,j \in \mathcal{H}}$.

In the given context, we interpret \mathcal{I} as the computational domain, and consider $\mathcal{V} \subseteq \mathcal{I}$ as a finite discretization of \mathcal{I} . X_i represents the epistemic uncertainty about the solution of (1) at the grid point $i \in \mathcal{V}$.

Let us suppose we gathered partial information on $X_{\mathcal{V}}$, e.g., by solving (1) on a coarse subset of variables $\mathcal{C} \subseteq \mathcal{V}$ and we would like to infer about the solution on the whole grid \mathcal{V} . By construction, any prediction $\hat{X}_i, i \in \mathcal{V}$ can only depend on information from $X_{\mathcal{C}}$ and therefore has to be measurable with respect to the sigma field $\sigma_{\mathcal{C}} \subseteq \Sigma$ associated with $X_{\mathcal{C}}$. We thus wish to find an optimal solution to the problem of making a prediction based on the values of $X_{\mathcal{C}}$ that minimizes the expected squared error, also called mean square error (*MSE*), for $X_i, i \in \mathcal{V}$

$$\mathbb{E} \left[\left(X_i - \hat{X}_i \right)^2 \right] \rightarrow \min \quad \text{s.t.} \quad \hat{X}_i \text{ is } \sigma_{\mathcal{C}}\text{-measurable.} \quad (5)$$

Let $L^2(\Omega, \Sigma, \mathbb{P})$ be the space of square integrable random variables, then $L^2(\Omega, \sigma_{\mathcal{C}}, \mathbb{P}) \subseteq L^2(\Omega, \Sigma, \mathbb{P})$ of $\sigma_{\mathcal{C}}$ -measurable L^2 functions is a closed subspace. Thus, the problem

where $L(w, \lambda) = w^T C_C w - \lambda(w^T \mathbb{1}_C - 1)$. It is now easily seen that the solution w_C of the equations (11) is given by

$$w = \frac{C_C^{-1} \mathbb{1}_C}{\mathbb{1}_C^T C_C^{-1} \mathbb{1}_C} \implies \hat{\mu} = \frac{\mathbb{1}_C^T C_C^{-1} X_C}{\mathbb{1}_C^T C_C^{-1} \mathbb{1}_C} \quad (12)$$

Inserting $\mu_C = \frac{\mathbb{1}_C^T C_C^{-1} x_C}{\mathbb{1}_C^T C_C^{-1} \mathbb{1}_C} \mathbb{1}_C$ into (7a) results in a prediction that depends linearly on the observed data. Thus this predictor corresponds to the construction of an interpolation matrix $P = P_{\mathcal{F}, \mathcal{C}}$ in (3).

From both alternatives described here, we follow the second, estimating the value of μ from the coarse grid data x_C and not setting it to zero. The reason is that, while for test vectors $\mu = 0$ is a natural choice (see subsection 3.4), this is not necessarily the case for the problem, the multigrid solver is applied to. A more thorough comparison of both approaches goes beyond the scope of this initial study. Note, that so far, we have only considered the calculation of the entries of $P_{\mathcal{F}, \mathcal{C}}$ given a subset \mathcal{C} and neglecting the sparsity requirement of P for now. The choice of \mathcal{C} and the construction of localized, i.e., sparse, interpolation will be discussed next.

3.3 Local Kriging

The computational cost of (7a) and (7b) in many cases is prohibitive, due to the fact that C_C in general is not sparse. However, we can localize (7a) and (7b) in the following sense. Assuming that there exists a (pseudo) distance $d_{\mathcal{V}}(i, j)$ on \mathcal{V} and that the correlation

$$\varrho(i, j) = \frac{C_{i,j}}{\sqrt{C_{i,i} C_{j,j}}}$$

decreases sufficiently fast when $d_{\mathcal{V}}(i, j)$ grows, we can neglect the effect of observations in far away points in \mathcal{C} on the prediction \hat{X}_i of X_i . In many cases, it is therefore sufficient to choose a subset \mathcal{C}_i of \mathcal{C} containing the q_{\max} points in \mathcal{C} that are closest to $i \in \mathcal{F}$ for the calculation of a suitable predictor. The number of neighbours q_{\max} can be chosen independently of the size of \mathcal{C} and \mathcal{V} and is referred to as the *caliber* of interpolation.

We note that (7a) and (7b) remain valid when replacing \mathcal{C} with \mathcal{C}_i and \mathcal{F} with i . Also, the estimates eqs. (10) and (11) can be localized accordingly. The complexity of computing the predictions on \mathcal{F} therefore acquires the optimal linear growth in the size of this set, provided the construction \mathcal{C}_i is either negligible for practical purposes or can be carried out with optimal complexity as well. Due to the applied localization this can be guaranteed by employing efficient graph based techniques.

3.4 Non-Parametric Covariance Estimation

The selection of the expected value $\mu_{\mathcal{V}}$ and the covariance $C_{\mathcal{V}}$ has to essentially capture the correlation structure of the problem at hand. Assuming that a set of test vectors $V = [v^{(1)} \mid \dots \mid v^{(K)}]$ is given on \mathcal{V} , we define the average value and the empirical covariance matrix of these test vectors by

$$\hat{\mu}_{\mathcal{V}} = \frac{1}{K} V \mathbb{1}_K, \quad (13a)$$

$$\hat{C} = \frac{1}{K} \left(V - \frac{1}{K} V \mathbb{1}_K \mathbb{1}_K^T \right) \left(V - \frac{1}{K} V \mathbb{1}_K \mathbb{1}_K^T \right)^T, \quad (13b)$$

where $\mathbb{1}_K \in \mathbb{R}^K$ is the column vector with value 1 in every entry. By construction the rank of the empirical covariance matrix is bounded from above by K , the number of test vectors, so that it is in general not possible to simply replace the theoretical covariance C in (7a) and (7b) by the empirical counterparts. As long as $K < |\mathcal{C}|$, i.e., the number of test vectors is smaller than the number of coarse grid variables, \widehat{C}_C^{-1} does not exist. There are two ways to fix this problem. First, we can regularize the calculation by replacing $\widehat{C}_C \rightarrow \widehat{C}_C + \varepsilon I$ where $\varepsilon > 0$ stands for a substitution $X \rightarrow X + N$ where $N = \{N_i\}, i \in \mathcal{I}$ is spatially uncorrelated Gaussian white noise with zero mean and variance ε . Second, in the context of local Kriging described in subsection 3.3 it is only required that \widehat{C}_{C_i} is positive definite for all $i \in \mathcal{F}$. As we can choose $|C_{C_i}| = q_{\max} \leq K$, this guarantees the non-singularity of \widehat{C}_{C_i} for almost all sets of randomly generated testvectors V so that in this case an ε -regularisation is not required.

Note that in many cases, the test vectors $v^{(j)}$ will be constructed by the application of one or a few smoothing, e.g., Gauss-Seidel, steps applied to a vector with noise data on \mathcal{V} , which is statistically centred around 0. In such cases, $\widehat{\mu}_{\mathcal{V}} = 0$ by theoretical considerations and we can replace the statistical estimation in (13a). Likewise, in this case we could as well simplify (13b) by omission of the terms $\frac{1}{K} V \mathbb{1}_K \mathbb{1}_K^T$.

There exist several choices for the distance function in local Kriging. The first option is to use the coordinate distance $d^c(i, j)$, if an embedding of \mathcal{V} in \mathbb{R}^d is known and the underlying continuum problem is isotropic. This will not always be the case, especially in the context of algebraic multigrid. Instead we use the graph distance $d^A(i, j)$, which measures the shortest path in the undirected graph associated with the system matrix A of eq. (1) over \mathcal{V} assuming that the length of an edge is defined as the inverse of the corresponding matrix entry, i.e., edge $\{i, j\}$ has length $\frac{1}{|A_{i,j}|}$.

In the context of local Kriging, it is not necessary to assemble the full matrix \widehat{C} , but only $n_{\mathcal{F}} = |\mathcal{F}|$ submatrices \widehat{C}_{C_i} of size $q_{\max} \times q_{\max}$ and $n_{\mathcal{F}}$ matrices \widehat{C}_{i, C_i} of size $1 \times q_{\max}$. Each entry of these matrices requires flops proportional to K , which gives linear complexity in $n_{\mathcal{F}}$, provided the search for C_i is either negligible in terms of compute time on the relevant problem sizes or is implemented with optimal complexity in the sense that the search for C_i has bounded complexity.

3.5 Parametric Semivariogram Estimation

The disadvantage in the procedure described above lies in the fact that a relatively large number K of test vectors $v^{(j)}$ is required in order to obtain a reasonable estimate \widehat{C} for the underlying covariance structure C . Given that each test vector is calculated with a computational cost proportional to n , the generation of up to $K \approx 100$ test vectors can be a significant computational burden.

Also, the large number n of estimates of matrices \widehat{C}_{C_i} increases the probability that there is at least one $i \in \mathcal{F}$ for which the estimate for \widehat{C}_{C_i} is poor, making it more difficult to obtain guarantees for the estimation of \widehat{X}_i for all $i \in \mathcal{F}$.

To reduce K down to numbers in the range 1 – 10 and to stabilise the estimation of single elements in C_{C_i} , we follow methods from geostatistics that allow, under appropriate assumptions, a more efficient estimate.

To this purpose, we assume that the underlying continuous problem has some kind of translation invariance. This can either be caused by a strict invariance of the underlying operator which is discretized by A (neglecting the effect of boundary conditions) or an invariance in some statistical sense, where the local inhomogeneity is statistically the same around all points in the computational domain. In both cases

it is legitimate, to work with translation invariant models for the a priori distribution of the Gaussian process.

A stochastic process $X = \{X_i\}, i \in \mathcal{I}$ with the special choice $\mathcal{I} = \mathbb{R}^d$ is stationary, if the process $X_h = \{X_{i+h} \mid i \in \mathcal{I}\}$ has the same finite dimensional distributions as X for all $h \in \mathbb{R}^d$. By (4) for Gaussian processes this amounts to a constant μ_i independent of i and $C_{i+h,j+h} = C_{i,j}$ for all $i, j, h \in \mathbb{R}^d$. Consequently, the covariance function $C_{i,j} = C_{0,i-j} =: C(i-j)$ defines a function $C : \mathbb{R}^d \rightarrow \mathbb{R}$. Furthermore, the process is called isotropic, if $C(h) = C(\Lambda h)$ for any rotation matrix $\Lambda \in SO(d)$ and $h \in \mathbb{R}^d$. This can be justified if the underlying operator in the continuum (approximately) shares this property. In this case, we can model (with a slight abuse of notation) $C(h) = C(|h|)$ where only a function $C : \mathbb{R}_+ \rightarrow \mathbb{R}$ has to be estimated from the test vectors V . This is a standard task in geostatistics [2, 30], which we briefly review.

In geostatistics it is customary to fit semivariograms instead of the covariance function. Both are connected via

$$\begin{aligned} \gamma(|h|) &= \frac{1}{2} \mathbb{E} [(X(0) - X(h))^2] \\ &= \frac{1}{2} \mathbb{E} [(X(0) - \mu)^2 + 2(X(0) - \mu)(X(h) - \mu) + (X(h) - \mu)^2] \\ &= C(0) - C(|h|) \end{aligned} \quad (14)$$

and contain the same information, as the asymptotic sill value $C(0)$ can be obtained by letting $h \rightarrow \infty$ and thereby $C(|h|) \rightarrow 0$ as correlations decline at large distances. Note, that in (14) we made use of the stationarity assumption.

In the next step we generate the empirical semivariogram based on the test vectors V on \mathcal{V} and the coordinate distance function $d^c(i, j)$. To this purpose, we collect all pairs of values $(d^c(i, j), (v_i^{(\ell)} - v_j^{(\ell)})^2), i, j \in \mathcal{V}$ and $\ell = 1, \dots, K$ in the so-called variogram cloud. Discretising the range of all values $d^c(i, j)$ into bins of width Δ , we obtain the empirical semivariogram as

$$\hat{\gamma}(h) = \frac{1}{K \#\{(i, j) : |h| - \frac{\Delta}{2} \leq d^c(i, j) < |h| + \frac{\Delta}{2}\}} \sum_{\substack{\{(i, j) : |h| - \frac{\Delta}{2} \leq d^c(i, j) < |h| + \frac{\Delta}{2}\} \\ \ell=1, \dots, K}} (v_i^{(\ell)} - v_j^{(\ell)})^2. \quad (15)$$

In the following, parametric semivariogram functions γ_θ are fitted to $\hat{\gamma}(|h|)$, mostly using weighted least squares [2]. There are many known families of semivariograms, but in this work we use the exponential family

$$\gamma_\theta(|h|) = \sigma^2 \left(1 - e^{-\left(\frac{|h|}{\eta}\right)}\right), \quad \theta = (\sigma^2, \eta) \in \mathbb{R}_+^2 \quad (16)$$

and the spherical family defined by

$$\gamma_\theta(|h|) = \begin{cases} \sigma^2 \left(\frac{3|h|}{2a} - \frac{1}{2} \left(\frac{h}{\eta}\right)^3\right) & \text{for } |h| < \eta \\ \sigma^2 & \text{for } |h| \geq \eta \end{cases}, \quad \theta = (\sigma^2, \eta) \in \mathbb{R}_+^2. \quad (17)$$

Empirical semivariograms and fitted semivariograms of exponential and spherical type can be seen in Figures 2 and 3 below. For further models we refer to [2, 30]. From the fitted semivariogram one then computes the covariance function $C_\theta(|h|)$ which can be used to compute $C_{\mathcal{H}}$.

Provided that the correlation length η (also called range in the geostatistical literature) is much smaller than the size of the underlying domain, it is often enough to work with just a few or even just one test vector, $K = 1$: If spatial correlations quickly decrease, the random field effectively contains many resamplings of its statistics in just one realization, i.e., in one test vector over a sufficiently extended grid \mathcal{V} .

3.6 Parametric Semivariograms for Inhomogeneity and Anisotropy

In most situations where homogeneity and isotropy of X cannot be expected, we replace the coordinate metric $d^c(i, j)$ with the graph metric $d^A(i, j)$ introduced in subsection 3.3. This is very much in the spirit of algebraic multigrid approaches, where the knowledge of coordinate lists of the variables cannot be guaranteed and the use of the graph distance dates back at least to the definition of strength-of-connection in classical AMG [29].

This pragmatic approach however comes with a conceptual problem. It is not clear if the weighted graph obtained from \mathcal{V} and A with distance function $d^A(i, j)$ can be isometrically embedded to some space $\mathbb{R}^{d'}$. In fact, this is possible if and only if the condition

$$w^T [d^A(i, j)^2]_{i, j \in \mathcal{V}} w \leq 0, \text{ for all } w \in \mathbb{R}^n \text{ s.t. } \mathbb{1}_n^T w = 0 \quad (18)$$

holds, see [17]. If this is true, then the positivity of the covariance matrix C follows from the fact that for many families of parametric semivariograms $C_\theta(|h|)$ defines a positive definite function in *any* dimension d and in particular in the dimension d' of the isometric embedding, which can differ from the dimension d of the underlying continuum problem.

If the condition (18) is violated, formula (7a) still defines an interpolator, as is easily checked, but the Kriging variance (7b) needs no longer to be non-negative and the probabilistic interpretation of the Gaussian process is gone. In fact, we observe this in the examples presented below, despite observing high correlation between coordinate distance and graph distance.

In this situation we still can carry out the algebraic manipulations from both formulae and we use the Gaussian process in the probabilistic sense simply as a source of inspiration.

Note however that local models of Gaussian processes on $\mathcal{C}_i \cup \{i\}$, $i \in \mathcal{F}$ with metric $d^A(i, j)$ may well exist as the embedding problem for such smaller graphs is much alleviated. In fact, in the numerical examples given below we do not observe any non-positive covariance metrics $\widehat{C}_{\mathcal{C}_i \cup \{i\}}$ for moderate size of q_{\max} . We therefore suggest that the calculation of local Kriging predictors and local Kriging variances can still be used and still give reasonable results, even though the (global) probabilistic interpretation has to be used with caution and strictly speaking the localized version of (18) should be checked, at least if the method proposed in the following section does not show the expected performance. If this condition is violated, it seems to be reasonable to lower q_{\max} at least locally, to obtain smaller local graphs which are more easily embedded.

4 Adaptive Coarsening using Gaussian Processes

The fact that algebraically smooth error can be interpreted as instances of a spatial Gaussian process allows us to use the methodology of Gaussian processes and the Kriging interpolation developed in section 3 to solve the coarsening problem of algebraic multigrid methods. The calculation of a splitting of the variables \mathcal{V} into a set of coarse variables \mathcal{C} and the remaining fine variables $\mathcal{F} = \mathcal{V} \setminus \mathcal{C}$ as well as the computation of interpolation suited to this setting is carried out in three consecutive steps.

First, we need to determine the covariance structure C of algebraically smooth error when viewed as instances of a spatial Gaussian process. In order to do so,

we start with a number, K , of test vectors $v^{(1)}, \dots, v^{(K)}$, each with entries that are normally distributed with mean zero and variance 1. These initially random vectors are then subject to a number, ν of smoothing iterations with right-hand-side zero, where we employ the smoother that is going to be used in the algebraic multigrid method eq. (2) as well. These test vectors are then fed into the calculation of the covariance structure. In our tests we compare localized non-parametric approaches (cf. section 3.4), which use a sizeable number of test vectors, with parametric models (cf. section 3.5) that make use of only a small number of test vectors.

Once the covariance structure is determined we can use Kriging interpolation to tackle the coarsening problem. The set of coarse variables \mathcal{C} has to be chosen in such a way that interpolation of information from these variables to the remaining variables \mathcal{F} is as accurate as possible for algebraically smooth error. In accordance with the interpretation of Kriging interpolation as the interpolation that minimizes the MSE under the assumption that algebraically smooth error can be interpreted as instances of a Gaussian field, we can use the variance of the Kriging estimator eq. (7b) in order to define the coarse variable set. Starting with $\mathcal{C} = \emptyset$ and using the fact that any variable that ends up in \mathcal{C} during the coarsening has zero variance after interpolation we proceed to add those variables to \mathcal{C} with largest variance. In case there is a tie, we choose the first occurrence, but other selection strategies are possible as well. Exploiting the fact that the correlation distance can be used to limit the reach of the Kriging interpolation, we can actually add multiple variables to \mathcal{C} at the same time if they are spaced so far apart that interpolation between them is not considered due to the localization of the Kriging interpolation. After variables are added to \mathcal{C} , we update the Kriging interpolation of all affected variables. By using one of the pseudo-distances we first determine for each $i \in \mathcal{F}$ the set of interpolatory variables $\mathcal{C}_i \subseteq \mathcal{C}$, again respecting the limitation of reach due to a finite correlation distance. Based on these sets, the Kriging estimator and the corresponding variances or MSE are calculated/updated. This process of adding variables to \mathcal{C} based on the uncertainty with which we can predict the value at the respecting variable and updating Kriging interpolation is repeated until a prescribed tolerance on either the size of the coarse variable set \mathcal{C} , typically relative to the total number of variables \mathcal{V} , or the largest remaining uncertainty of the Kriging estimator is reached. The resulting process is roughly summarized in algorithm 4.1.

Algorithm 4.1: Coarsening based on Kriging interpolation

Data:

- 1 Initialize $\mathcal{C} = \emptyset, \mathcal{F} = \mathcal{V}$
- 2 **while** $|\mathcal{C}| < n_c$ **do**
- 3 Choose $i \in \mathcal{F}$ with largest variance μ_i
- 4 Add i to \mathcal{C} , remove i from \mathcal{F}
- 5 **for** $j \in \mathcal{F}$ with $i \in \mathcal{C}_j$
- 6 Update Kriging interpolation with new \mathcal{C}_j set
- 7 Compute corresponding updated variances μ_j (set $\mu_i = 0$)

5 Numerical Case Studies

In order to gauge the efficiency of the new coarsening scheme based on Gaussian fields and Kriging interpolation we consider the general diffusion problem

$$-\left(c_1 \frac{\partial^2}{\partial x^2} + c_2 \frac{\partial^2}{\partial y^2} + c_3 \left(\frac{\partial}{\partial x} \frac{\partial}{\partial y} + \frac{\partial}{\partial y} \frac{\partial}{\partial x}\right)\right) u = f, \quad (19)$$

for constant and anisotropic coefficients by choosing c_1, c_2 and c_3 accordingly. We further choose the computational domain as the unit square $(0,1)^2$, employing a finite difference discretization on a regular mesh and the unit circle $U(1)$, where the discretization is defined by linear finite elements on a triangularization generated in MATLAB. We considered 4 parameter combinations in our tests collected in table 1.

name	s-iso	s-aniso	c-iso	c-aniso
domain	square	square	circle	circle
n	2025	2025	2521	2521
$\begin{bmatrix} c_1 & c_3 \\ c_3 & c_2 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 \\ 0 & 10^{-2} \end{bmatrix}$	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 1 & 0 \\ 0 & 10^{-2} \end{bmatrix}$

Table 1: Parameter choices of the considered test cases.

The overall aim of these case studies is to get a first impression of the performance of the new approach of constructing coarse variable sets and interpolation. Thus we compare methods with respect to the underlying covariance model, i.e., the empirical covariance function vs. semivariogram models (spherical and exponential). In addition, we test the robustness of the approach with respect to the number of test vectors used to construct the covariance model. In all tests we employ localization in the calculation of Kriging interpolation, which is an inevitable technique to guarantee overall linear complexity and well-posedness of the Kriging interpolation as discussed in section 3.3 especially for the empirical construction.

The test vectors are generated by applying just one iteration of a colored Gauss-Seidel iteration to white noise vectors and we use the same iterative method for the smoother of our resulting two-grid method as well. That is, in the reported results we run a $V(1,1)$ -cycle two-grid method with a direct solve for the coarse grid system of equations.

Pseudo-distances The use of pseudo-distances is an important aspect when it comes to the independence from problem specific knowledge as is required by a true algebraic multigrid approach. Thus we first compare the correlation of the graph pseudo-distance and the true geometric distances based on the coordinates of the unknowns. In this we found a correlation of the distances for the cases c-iso of 97.68% and for s-iso of 99.65%, which leads us to believe that using the algebraic graph distance as a pseudo distance in our Gaussian field analysis is viable.

Covariance models In the following we fit the exponential and the spherical semivariogram models to the empirical semivariogram generated from 1, 10 and 100 test vectors, see equations (17), (16) and (15). This is done for the isotropic and the anisotropic case, both for the circle (Figure 2) and the square grid (Figure 3). The fits and models are performed using the R library `gstat`. The fits expose reasonable quality and not much variation caused by the number of test vectors used.

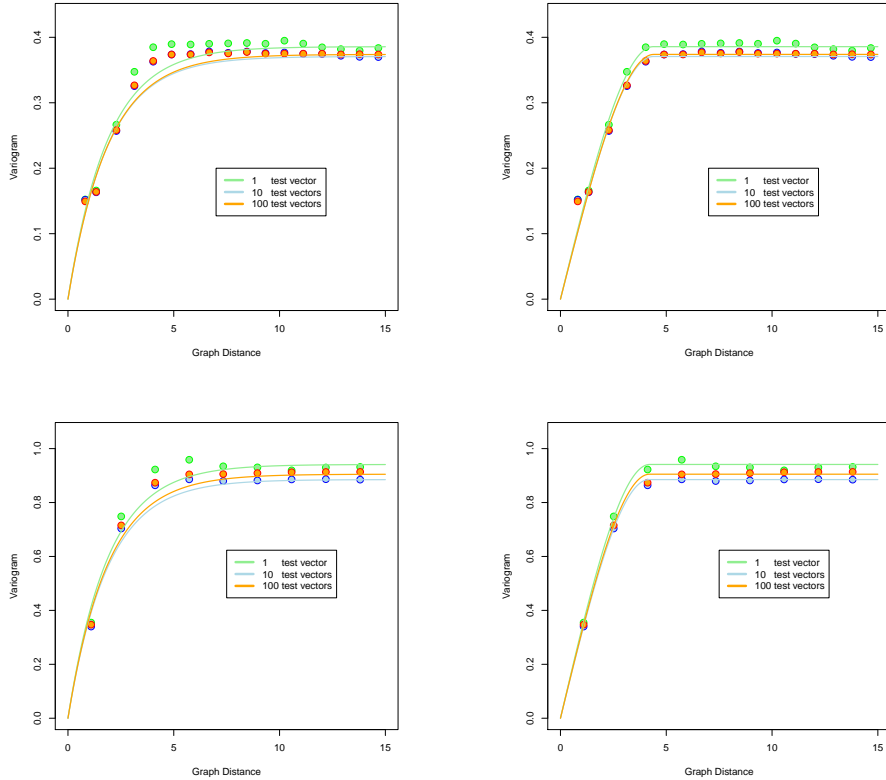


Figure 2: Fitted semivariogram models for the unstructured mesh on the circle: The top row displays the isotropic and the bottom row the anisotropic case. On the left the exponential semivariogram model is used and on the right the spherical.

Coarsening and two-grid results Finally, after analyzing the components of the metrics and fits underlying the Gaussian process framing of the coarse grid construction we are ready to apply our coarsening algorithm and present two-grid results using the Kriging interpolation. In this we use the test cases described in table 1 combined with

- an empirical construction of the covariance structure, which we term *emp-K*,
- a spherical covariance model based on a semivariogram fit eq. (16), termed *sph-K*,
- an exponential covariance model based on a semivariogram fit eq. (17), termed *exp-K*.

In this naming convention K represents the number of test vectors used in the construction of the model or semivariogram, respectively.

In table 2 we collect results of the resulting two-grid methods for the isotropic test cases s-iso and c-iso. The quantities we report are the asymptotic convergence rates, ρ of the two-grid methods, which are indicative of the overall compatibility

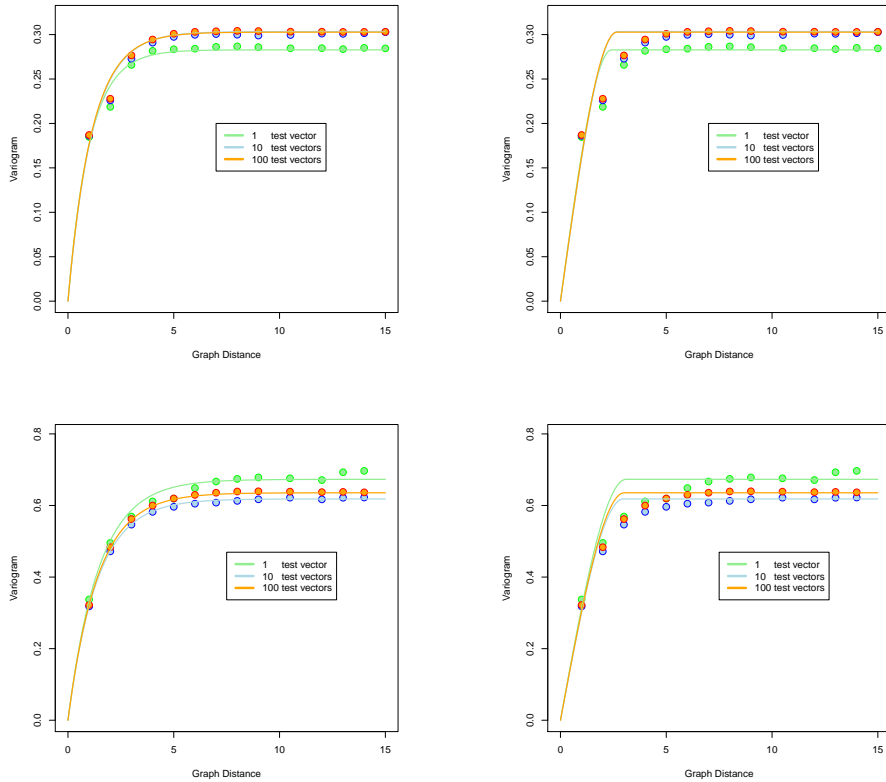


Figure 3: Fitted semivariogram models for the structured mesh on the square: The top row displays the isotropic and the bottom row the anisotropic case. On the left the exponential semivariogram model is used and on the right the spherical.

of the smoother and coarse-grid correction, and iteration numbers, k , of the conjugate gradients method preconditioned with the two-grid method that are required to reduce the initial residual by a factor of 10^8 . This latter quantity gives insight into the flaws of the two-grid construction. Oftentimes, as also explored in [7, 6, 20], the two-grid method might show bad asymptotic convergence rates even though the preconditioned conjugate gradients iteration converges rapidly. This typically corresponds to the presence of a few outliers in the spectrum of the preconditioned matrix and indicates that the two-grid construction provides better complementarity of coarse grid construction and smoothing than the asymptotic convergence rate suggests. Taking into account that all results of table 2 are generated with test vectors that are smoothed by only a single iteration of colored Gauss-Seidel and that a $V(1,1)$ -cycle is employed the results are surprisingly good; cf. [7]. In part this can be explained by the implicit preservation of the constant vector in the Kriging interpolation, a modification that has been shown to be particularly effective in improving the performance of bootstrap AMG in [7, 5, 24]. The most interesting observation is the fact that the model based constructions of Kriging interpolation are competitive when using only a single test vector. Due to the fact that the descriptive power of

s-iso	emp-10	emp-100	sph-1	sph-10	sph-100	exp-1	exp-10	exp-100
ρ	.387	.302	.256	.251	.253	.224	.225	.222
k	10	9	9	9	10	9	8	9

c-iso	emp-10	emp-100	sph-1	sph-10	sph-100	exp-1	exp-10	exp-100
ρ	.563	.319	.275	.27	.273	.314	.294	.303
k	14	10	10	10	10	11	10	10

Table 2: Asymptotic convergence rates ρ of the two-grid V(1,1) cycle and number of iterations k of the conjugate gradients method preconditioned with the two-grid method to reduce the initial residual by a factor of 10^8 for both isotropic test cases. All approaches generate a coarse variable set with $n_c = \frac{n}{4}$ variables, use a localization of distance 4 and a caliber of 4, i.e., $|C_i| \leq 4$ for all $i \in \mathcal{F}$.

the semivariogram approach should get better the more points are available for sampling, i.e., the finer the discretization, this finding should scale well with respect to the problem size. In fig. 4 we collected some sections of the resulting coarsenings of the methods considered in table 2.

The coarsenings do not show any particularly interesting features. As only sections in the bulk are shown and the coarsening ratio of the depicted section is smaller than the preset $\frac{1}{4}$, a mild agglomeration of coarse grid variables at the boundary takes place. It remains to be seen if this poses a problem when recursing on the construction in a multigrid fashion.

The more interesting test cases arise when anisotropy is present in the model, especially for discretization on unstructured grids, where a canonical coarsening that follows the anisotropy is not available. Testing automatic coarsening approaches on the simple grid-aligned case in order to gauge their robustness and ability to reproduce the canonical coarsenings has been common in the past [5, 20]. Analogous to the results for the isotropic test cases we report in table 3 asymptotic convergence rates and iterations counts of the preconditioned conjugate gradients method. The grid

s-aniso	emp-10	emp-100	sph-1	sph-10	sph-100	exp-1	exp-10	exp-100
ρ	.463	.305	.06	.06	.154	.224	.225	.222
k	9	8	6	6	6	9	8	9

c-aniso	emp-10	emp-100	sph-1	sph-10	sph-100	exp-1	exp-10	exp-100
ρ	.648	.533	.69	.684	.681	.704	.71	.685
k	19	15	22	21	21	22	24	22

Table 3: Asymptotic convergence rates ρ of the two-grid V(1,1) cycle and number of iterations k of the conjugate gradients method preconditioned with the two-grid method to reduce the initial residual by a factor of 10^8 . All approaches generate a coarse variable set with $n_c = \frac{n}{2}$ variables, use a localization of distance 4 in graph distance and a caliber of 2 and 3 for all $i \in \mathcal{F}$ for the test cases formulated on the square and circle, respectively.

aligned anisotropy in the square test case s-aniso does not pose any difficulty for any of the model based approaches and again a single test vector is sufficient to obtain a good enough statistics for a suitable fit of the model as already suggested by figs. 2 and 3. The results suggest that the spherical model is better suited for this problem than the exponential model and we see a clear advantage of the model based approaches over the empirical approach, showing extremely fast convergence. Similar to the isotropic




Figure 4: Illustration of the variable splittings obtained by emp-10 (top), sph-1 (mid) and exp-1 (bottom) for the isotropic test cases on the square (left) and circle (right). In order to avoid cluttering of the illustrations we just show a section of the bulk of the domains (i.e., $[.35, .4] \times [.65, .6]$ and $[-.3, -.2] \times [.3, .2]$ for square and circle, resp.).

case we see a notable improvement when increasing the number of the test vectors for the empirical model from 10 to 100, which is quite frankly an unfeasible number of test vectors, but serves the illustrative purpose quite well. As can be seen in fig. 5 the coarsenings obtained by the different approaches for the square test case show that the anisotropy has been clearly detected and the coarsening constructed accordingly. When it comes to the results for the anisotropic problem on the circular domain and unstructured grid the results are comparable to results reported, e.g., in [5] for non-grid aligned anisotropies. Whilst all approaches yield good preconditioners for the conjugate gradients method with comparable iteration numbers, the asymptotic convergence rates are considerably worse at around $.7$ compared to the grid-aligned case for the model based approaches. Interestingly the empirical construction is able to cope with the unstructured grid better than the model based approaches. Taking into account that the model based approaches implicitly assume shift invariance of the problem, which might be violated more strongly in this test case compared to the circular isotropic problem, this does not come as a large surprise. The coarsenings depicted in fig. 5 demonstrate the capability of the approach to detect the anisotropy and construct coarsenings that are suitable to construct an efficient two-grid method.




Figure 5: Illustration of the variable splittings obtained by emp-10 (top), sph-1 (mid) and exp-1 (bottom) for the anisotropic test cases on the square (left) and circle (right). In order to avoid cluttering of the illustrations we just show a section of the bulk of the domains (i.e., $[-.35, .4] \times [.65, .6]$ and $[-.3, -.2] \times [.3, .2]$ for square and circle, resp.).

Even though the resulting two-grid methods for the model based approaches yield worse asymptotic convergence rates for the non grid-aligned anisotropy in c-aniso this cannot be traced to a defect in the coarsening structure. The respective sections of the coarsenings show clearly that the approach detected the anisotropy, which leads us to believe that the poor performance is more due to a poor choice of interpolatory set or interpolation weights or both rather than a poor coarsening structure.

6 Conclusion and Outlook

In this paper we presented a new approach to adaptive algebraic multigrid construction using ideas from geostatistics and statistical learning theory. Based on the resemblance of algebraically smooth error to instances of spatial Gaussian fields we develop an empirical and semivariogram based approach to recover the covariance structure of the unknown, underlying Gaussian process. Once the covariance structure is known, efficient interpolation can be formulated by Kriging interpolation. Further exploiting the interpretation of the local interpolation error as the variance of the representation, we are able to formulate a coarsening approach that is seamlessly integrated into the

determination of interpolation. Finally, by using graph distance and assuming shift invariance of the operator with respect to this pseudo-distance we are able to obtain good statistics for the semivariogram fit using only a single test vector. This is due to the fact that the semivariogram collects information at all variables to formulate a covariance function that depends solely on the distance between variables. Combined with the observation that the correlation distance of the underlying Gaussian process is very small, a single vector provides enough information about the short range correlation of values. In addition, the short correlation distance allows us to strictly localize all calculations which preserves the linear complexity of the whole process.

One apparent aim of future development is the integration of the Gaussian process approach into a multigrid setup and potentially a bootstrap type setup that is able to generate additional information about the underlying process on coarse scales as well. In line with [28] we plan to investigate the connection between the partial differential operator, its discretization and the smoothing scheme with the resulting covariance structure of the Gaussian fields. To some extent this development can be seen in line with the investigation of optimal interpolation in algebraic multigrid methods in [9], where an explicit influence of the smoother on the optimal construction of interpolation has been shown. Insight into this might allow us to translate the demonstrated potential for efficient adaptive algebraic multigrid constructions using a minimal amount of test vectors to more complex problems.

A Least squares and Kriging interpolation.

In this appendix we present a direct comparison between the least squares [7] and the Kriging interpolation for the case where $\mu = 0$. As explained in Subsection 3.2, an alternative to estimating the stationary mean μ from the data, one can also assume $\mu = 0$ if the test vectors are sampled from a centred distribution. In this case, equations (13a) and (13b) simplify and we obtain

$$C_{ij} = \frac{1}{K} \langle V_i, V_j \rangle_2 = \frac{1}{K} (V \cdot V^T)_{ij}.$$

Correspondingly define the correlation matrix by

$$X_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}} = \frac{\langle V_i, V_j \rangle_2}{\|V_i\|_2 \|V_j\|_2}.$$

Now consider the following definition of algebraic coupling strength μ_{ij} between variables i and j , based on the notion of least squares interpolation (cf. [5, 7, 22]),

$$\sigma_{ij}^2 = \frac{\sum_k \left(v_i^{(k)} - p_{ij} v_j^{(k)} \right)^2}{\sum_k \left(v_i^{(k)} \right)^2}.$$

using the short-hand notation this can be written as

$$\sigma_{ij}^2 = \frac{\|V_i - p_{ij} V_j\|_2^2}{\|V_i\|_2^2}.$$

Then the minimizing p_{ij} is given by

$$p_{ij}^\# = \frac{\langle V_i, V_j \rangle_2}{\|V_j\|_2^2} = C_{ij} \cdot C_{jj}^{-1}$$

and the corresponding minimal value by

$$(\sigma_{ij}^\#)^2 = 1 - \frac{|C_{ij}|^2}{C_{ii}C_{jj}} = 1 - X_{ij}^2.$$

Thus the notion of algebraic distances yields a distance with an interpretation that is similar to our approach, i.e.,

$$\hat{\sigma}_{ij}^2 = \frac{1}{1 + K|C_{ij}|},$$

yields large distances for small correlations and small distances for large correlations.

Based on these findings it might be worthwhile to take a closer look at the multiple interpolation variables case. Again starting with the least-squares interpolation setting and comparing it to the Kriging approach.

In general the least squares interpolation for variable i from variables in $C_i = \{j_1, \dots, j_\ell\}$ is defined as the minimizer of

$$\sigma_{i,C_i}^2 = \sum_k \left(v_i^{(k)} - \sum_{j \in C_i} p_{ij} v_j^{(k)} \right)^2 = \sum_k \left(v_i^{(k)} - p_{i,C_i} V_{C_i} \right)^2.$$

Again using the short-hand notation with $V_{C_i} \in \mathbb{R}^{\ell \times K}$ this can be written as

$$\sigma_{i,C_i}^2 = \|V_i - p_{i,C_i} V_{C_i}\|_2^2.$$

With $C_{C_i,C_i} := V_{C_i} \cdot V_{C_i}^T$ we obtain the minimizer

$$p_{i,C_i}^\# = C_{i,C_i} C_{C_i,C_i}^{-1}$$

Correspondingly, the minimal value of μ_{i,C_i} is then given by

$$\sigma_{i,C_i}^\# = C_{i,i} - C_{i,C_i} C_{C_i,C_i}^{-1} C_{C_i,i},$$

the Schur complement of the correlation matrix.

Interestingly, this last quantity is, up to appropriate scaline, equal to the estimate of the variance in our Kriging formulation. The interpolation weights of the Kriging formulation can be obtained from $p_{i,C_i}^\#$ by enforcing $\sum_{j \in C_i} p_{i,j} = 1$. That is, they are given by

$$\begin{aligned} p_{i,C_i}^{\text{krig}} &= C_{i,C_i} C_{C_i,C_i}^{-1} + \frac{1 - C_{i,C_i} C_{C_i,C_i}^{-1} \mathbb{1}^T}{\mathbb{1} C_{C_i,C_i}^{-1} \mathbb{1}^T} \mathbb{1} C_{C_i,C_i}^{-1} \\ &= p_{i,C_i}^\# + \frac{1 - C_{i,C_i} C_{C_i,C_i}^{-1} \mathbb{1}^T}{\mathbb{1} C_{C_i,C_i}^{-1} \mathbb{1}^T} \mathbb{1} C_{C_i,C_i}^{-1}, \end{aligned}$$

where the latter part simply ensures $p_{i,C_i}^{\text{krig}} \mathbb{1}^T = 1$.

Plugging the Kriging interpolation weights back into the least squares formulation yields a distance measure

$$\sigma_{i,C_i}^{\text{krig}} = \sigma_{i,C_i}^\# + \frac{(1 - C_{i,C_i} C_{C_i,C_i}^{-1} C_{C_i,i})^2}{\mathbb{1} C_{C_i,C_i}^{-1} \mathbb{1}^T}.$$

Similar constructions, explicitly preserving the constant vector, yet for reasons connected to the underlying PDE can be found in the operator based bootstrap AMG approach [24], where least squares interpolation is mixed with classical AMG constructions and inherits the preservation of constants from this approach. The connection between the correlation and covariance structure of test vectors has been used to define aggregation based interpolations in [22], yet again a consistent definition of the whole coarsening process in terms of Gaussian processes is lacking.

References

- [1] R. J. Adler. *The geometry of random fields*. SIAM, 2010.
- [2] R. S. Bivand, E. J. Pebesma, V. Gomez-Rubio, and E. J. Pebesma. *Applied spatial data analysis with R*. Vol. 747248717. Springer, 2008.
- [3] V. I. Bogachev. *Gaussian measures*. 62. American Mathematical Soc., 1998.
- [4] A. Brandt. “Algebraic Multigrid Theory: The Symmetric Case”. *Appl. Math. Comput.* 19.1–4 (1986), pp. 23–56.
- [5] A. Brandt, J. Brannick, K. Kahl, and I. Livshits. “Algebraic Distance for Anisotropic Diffusion Problems: Multilevel Results”. *Electron. Trans. Numer. Anal.* 44 (2015), pp. 472–496.
- [6] A. Brandt, J. Brannick, K. Kahl, and I. Livshits. “Bootstrap Algebraic Multigrid: Status Report, Open Problems, and Outlook”. *Numer. Math. Theory Me.* 8.1 (Feb. 2015), pp. 112–135.
- [7] A. Brandt, J. Brannick, K. Kahl, and I. Livshits. “Bootstrap AMG”. *SIAM J. Sci. Comput.* 33.2 (2011), pp. 612–632.
- [8] A. Brandt, S. F. McCormick, and J. W. Ruge. “Algebraic Multigrid (AMG) for Sparse Matrix Equations”. In: *Sparsity and its Applications*. Ed. by D. J. Evans. Cambridge: Cambridge University Press, 1985, pp. 257–284.
- [9] J. Brannick, F. Cao, K. Kahl, R. D. Falgout, and X. Hu. “Optimal Interpolation and Compatible Relaxation in Classical Algebraic Multigrid”. *SIAM J. Sci. Comput.* 40.3 (2018), A1473–A1493.
- [10] J. Brannick, Y. Chen, J. Kraus, and L. T. Zikatanov. “An Algebraic Multigrid Method Based on Matching in Graphs”. In: *Domain Decomposition Methods in Science and Engineering XX*. Ed. by R. Bank, M. Holst, O. Widlund, and J. Xu. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 143–150.
- [11] J. Brannick, Y. Chen, and L. T. Zikatanov. “An algebraic multilevel method for anisotropic elliptic equations based on subgraph matching.” *Numer. Linear Algebra Appl.* 19.2 (2012), pp. 279–295.
- [12] J. Brannick and R. D. Falgout. “Compatible relaxation and coarsening in algebraic multigrid”. *SIAM J. Sci. Comput.* 32.3 (2010), pp. 1393–1416.
- [13] S. C. Brenner and L. R. Scott. *The Mathematical Theory of Finite Element Methods*. Vol. 15. Texts in Applied Mathematics. Springer, 2008.
- [14] M. Brezina, R. D. Falgout, S. P. MacLachlan, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge. “Adaptive Algebraic Multigrid.” *SIAM J. Sci. Comput.* 27.4 (2006), pp. 1261–1286.
- [15] M. Brezina, R. D. Falgout, S. P. MacLachlan, T. A. Manteuffel, S. F. McCormick, and J. W. Ruge. “Adaptive Smoothed Aggregation (α SA).” *SIAM J. Sci. Comput.* 25.6 (2004), pp. 1896–1920.
- [16] G. Christakos. *Random field models in earth sciences*. Courier Corporation, 2012.
- [17] M. M. Deza and M. Laurent. *Geometry of cuts and metrics*. Vol. 15. Springer, 2009.
- [18] A. Forrester, A. Sobester, and A. Keane. *Engineering design via surrogate modelling: a practical guide*. John Wiley & Sons, 2008.

- [19] I. Gelfand and H. Vilenkin. *Generalized functions, part 4*. Academic Press New York, 1964.
- [20] K. Kahl and M. Rottmann. “Least Angle Regression Coarsening in Bootstrap Algebraic Multigrid”. *SIAM J. Sci. Comput.* 40.6 (2018), A3928–A3954.
- [21] J. P. Kleijnen. “Kriging metamodeling in simulation: A review”. *European J. Oper. Res.* 192.3 (2009), pp. 707–716.
- [22] O. E. Livne and A. Brandt. “Lean algebraic multigrid (LAMG): Fast graph Laplacian linear solver”. *SIAM J. Sci. Comput.* 34.4 (2012), B499–B522.
- [23] J. Mandel, M. Brezina, and P. Vaněk. “Energy Optimization of Algebraic Multigrid Bases”. *Computing* 62.3 (June 1999), pp. 205–228.
- [24] T. A. Manteuffel, S. F. McCormick, M. Park, and J. W. Ruge. “Operator-based interpolation for bootstrap algebraic multigrid.” *Numer. Linear Algebra Appl.* 17.2-3 (2010), pp. 519–537.
- [25] A. Napov and Y. Notay. “An Efficient Multigrid Method For Graph Laplacian Systems”. *Electron. Trans. Numer. Anal.* 45 (2016), pp. 201–218.
- [26] Y. Notay. “An aggregation-based algebraic multigrid method”. *Electron. Trans. Numer. Anal.* 37 (2010), pp. 123–146.
- [27] L. N. Olson, J. B. Schroder, and R. S. Tuminaro. “A General Interpolation Strategy for Algebraic Multigrid Using Energy Minimization”. *SIAM J. Sci. Comput.* 33.2 (2011), pp. 966–991.
- [28] H. Owhadi. “Multigrid with rough coefficients and multiresolution operator decomposition from hierarchical information games”. *SIAM Rev.* 59.1 (2017), pp. 99–149.
- [29] J. W. Ruge and K. Stüben. “Algebraic Multigrid”. In: *Multigrid Methods*. Ed. by S. F. McCormick. Frontiers in Applied Mathematics. Society for Industrial and Applied Mathematics, Jan. 1987. Chap. 4, pp. 73–130.
- [30] M. Sherman. *Spatial statistics and spatio-temporal data: covariance functions and directional properties*. John Wiley & Sons, 2011.
- [31] K. Stüben. “Algebraic Multigrid (AMG): Experiences and Comparisons”. *Appl. Math. Comput.* 13.3-4 (Jan. 1983), pp. 419–451.
- [32] T. Tao. *An introduction to measure theory*. Vol. 126. American Mathematical Society Providence, RI, 2011.
- [33] E. Vanmarcke. *Random fields: analysis and synthesis*. World Scientific, 2010.
- [34] W. Wan, T. Chan, and B. Smith. “An Energy-minimizing Interpolation for Robust Multigrid Methods”. *SIAM J. Sci. Comput.* 21.4 (1999), pp. 1632–1649.
- [35] J. Xu and L. Zikatanov. “Algebraic multigrid methods”. *Acta Numerica* 26 (2017), pp. 591–721.
- [36] J. Xu and L. T. Zikatanov. “On an energy minimizing basis for algebraic multigrid methods”. *Comput. Vis. Sci.* 7.3 (Oct. 2004), pp. 121–127.