ISSN: 2834-7706

Research Article

Journal of Mathematical Techniques and Computational Mathematics

Extending Kriging with Azimuthal Weighting and Material-Dependent Variance

Greg Passmore*

PassmoreLab, Austin, Texas, USA

***Corresponding Author**
Greg Passmore, PassmoreLab, Austin, Texas, USA.

Submitted: 2025, May 28; **Accepted:** 2025, Jul 04; **Published:** 2025, Sep 04

Citation: Passmore, G. (2025). Extending Kriging with Azimuthal Weighting and Material-Dependent Variance. *J Math Techniques Comput Math*, *4*(8), 01-12.

Abstract

This paper introduces a method called KRAM (Kriging with Regularization for Azimuth and Material), developed to address specific requirements in geospatial applications involving sparse, anisotropic, or heterogeneous elevation data. KRAM extends ordinary kriging by adding two targeted modifications: azimuthal regularization to reduce overreliance on collinear samples, and material-dependent variance modeling to capture localized surface roughness without requiring multiple covariance models. These enhancements improve interpolation accuracy in structured terrain, urban environments, and damage assessment scenarios, where geospatial data typically exhibit significant variability in both sampling geometry and surface materials. The method maintains the statistical foundation of ordinary kriging while improving structural accuracy, robustness to sampling geometry, and the interpretability of spatial uncertainty. KRAM is suited for use in terrain densification, surface reconstruction, and predictive modeling tasks where accuracy and auditability are operational requirements.

Keywords: Ordinary Kriging, Azimuthal Regularization, Material-dependent Spatial Variance, Anisotropic Sampling, Terrain Interpolation, Geostatistical Surface Modeling, Spatial Prediction Uncertainty

1. Introduction to KRAM

We developed KRAM to improve kriging performance in real-world terrain datasets that violate assumptions of isotropy and uniform variance. The method introduces a material-dependent variance model and an azimuthal penalty term to enhance structural fidelity and robustness. These datasets often violate key assumptions of traditional kriging, which presumes isotropic correlation and uniform residual variance.

To address this, KRAM adds two modifications to the ordinary kriging framework. The first is a material-dependent variance model, which adjusts the diagonal of the covariance matrix based on surface class. This captures local roughness and material behavior without requiring multiple variogram models or class segmentation. The second is an azimuthal regularization term that penalizes the overuse of collinear samples, improving angular diversity and reducing directional artifacts in structured terrain.

These extensions preserve the statistical properties of ordinary kriging, require minimal parameter tuning, and integrate directly into existing estimation workflows. KRAM was designed to improve spatial fidelity and uncertainty representation in high-consequence applications such as terrain densification, damage estimation, and structured surface modeling.

This work contributes:

- A material-dependent diagonal variance model for kriging
- An angular redundancy penalty for kriging weights
- Preservation of the unbiased, minimum-variance form of ordinary kriging

- Improved interpretability and fidelity in structured terrain scenarios

Limitations of Existing Densification Methods

Our lab processes large volumes of LiDAR, radar, photogrammetry, multispectral EO, thermal, and interferometric depth data. Depth fields derived from these sources often exhibit regions of sparsity that require densification for visualization and analysis. This paper addresses the use of kriging for depth field densification and introduces modifications to address several limitations we observe in practice. These include directional bias in sampling, failure to account for surface material properties, and undesirable geometric artifacts such as rounding or blob amplification in interpolated results. These problems are significant, since existing methods effectively blur and round geometry, leading to pattern recognition failures or misrepresentation of damage caused by natural disasters or war.

2. Method Build-Up

Before describing our method, we briefly review Shepard's method and ordinary kriging, which form the baseline for many geostatistical interpolators. We then present extensions intended to mitigate the persistent deficiencies common to all standard densification methods.

2.1. Shepard's Algorithm

The fastest and most straightforward method for depth field densification is inverse distance weighting, as originally defined by Shepard. In this approach, the interpolated value at an unsampled location is computed as a weighted average of nearby samples, where each weight is inversely proportional to the distance to that sample raised to a power parameter. When this parameter is large, the estimate is dominated by the nearest sample and the result approaches nearest-neighbor interpolation. When the parameter is small, the method approaches an unweighted average. This process requires no model fitting, has low computational cost, and is easy to implement. However, it does not account for spatial correlation structure, material variation, or geometric bias in the sample configuration. Despite these limitations, Shepard's method remains a widely used baseline due to its simplicity and speed [1]. Let's review the actual process.

Formula

The interpolated value $\hat{Z}(x)$ at a location x is:

$$\hat{Z}(x) = \frac{\sum_{i=1}^{n} w_i(x) \cdot Z(x_i)}{\sum_{i=1}^{n} w_i(x)}$$

where the weights are:

$$w_i(x) = \frac{1}{\|x - x_i\|^p}$$

- $\|x - x_i\|$ is the Euclidean distance between prediction point x and known point x_i
- p is the power parameter (typically $p = 2$).

When $p = 2$, the weight function corresponds to the inverse square law, which aligns with many physical phenomena such as gravitational and electromagnetic fields, where influence diminishes with the square of the distance. In interpolation, this choice gives closer points substantially more influence than distant ones without allowing a single nearest point to completely dominate, as would occur with $p >> 2$.

Values of $p < 2$ tend to over-smooth the interpolation, giving distant points excessive influence, while values $p > 2$ result in sharper transitions and increased sensitivity to the nearest points. The selection $p = 2$ is therefore a practical compromise that provides stable, localized interpolation with minimal tuning and without requiring spatial modeling or variogram fitting.

In our application, we are dealing with terrain and man-made structures, which do not follow the behavior assumed by idealized inverse-distance models. Features such as ridgelines, excavation cuts, and building edges introduce sharp discontinuities and directionally persistent gradients that violate the assumptions inherent in Shepard's method. The influence of an observed elevation is not isotropic and depends on factors including geometry, occlusion, and surface material. Using a fixed power parameter, such as $p = 2$, imposes a uniform decay rate that does not reflect these localized effects. The result is smearing across discontinuities, loss of structural fidelity, and distortion in areas with irregular sampling.

Ordinary kriging, as formalized by Matheron (1963), uses a statistical model of spatial covariance to assign weights that minimize prediction error under an unbiasedness constraint [2]. The fitted variogram or covariance function allows kriging to account for

correlation structure, point redundancy, and anisotropy in the data [3,4]. In addition to providing a deterministic estimate, it yields an associated prediction variance. These properties make it well-suited for structured elevation data and other applications where spatial behavior must be respected.

2.2. Ordinary Kriging

Ordinary kriging may be used to estimate the value of a random elevation $\boldsymbol{Z}$ at an unobserved location $s_0 = (x_0, y_0)$, based on observed values at n known locations $s_1, \ldots, s_n$.

The method assumes the mean of $\boldsymbol{Z}$ is unknown but constant within the local neighborhood of interpolation. It uses the second-order statistical structure (covariance or variogram) of the field to compute weights that minimize prediction error while ensuring unbiasedness.

2.2.1. Process

Given:

-Observation locations: $s_1, \ldots, s_n \in \mathbb{R}^2$
-Observed values: $Z = [Z_{s_1}, \ldots, Z_{s_n}]^T$
-Prediction location: $s_0 \notin \{s_1, \ldots, s_n\}$
-Covariance matrix: $\Sigma \in \mathbb{R}^{n \times n}$, where $\Sigma_{ij} = \mathrm{Cov}(Z_{s_i}, Z_{s_j})$
-Covariance vector: $c \in \mathbb{R}^n$, where $c_i = \mathrm{Cov}(Z_{s_i}, Z_{s_0})$

We seek a linear estimator of the form:

$$\hat{Z}_0 = \sum_{i=1}^{n} \lambda_i Z_{s_i} = \lambda^T Z$$

2.2.2. Constraint

To ensure unbiased behavior under an unknown constant mean, we require:

$$\sum_{i=1}^{n} \lambda_i = 1$$

2.2.3. Augmented Linear System

The weights $\lambda = [\lambda_1, \ldots, \lambda_n]^T$ are obtained by solving the following augmented linear system:

$$\begin{bmatrix} \Sigma & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} \lambda \\ \mu \end{bmatrix} = \begin{bmatrix} c \\ 1 \end{bmatrix}$$

Where:
- μ is a Lagrange multiplier enforcing the unbiasedness constraint
- $\mathbf{1}$ is a column vector of ones of length n
- The system has dimension $(n+1) \times (n+1)$

2.2.4. Kriging Predictor

Once the weights λ are computed, the estimated value at s_0 is:

$$\hat{Z}_0 = \lambda^T Z = \sum_{i=1}^{n} \lambda_i Z_{s_i}$$

2.2.5. Kriging Variance

The prediction error variance at s_0 is given by:

$$\tau^2(s_0) = \mathrm{Var}(Z_0 - \hat{Z}_0) = \sigma_0^2 - \lambda^T c + \mu$$

where $\sigma_0^2 = K(0)$ is the marginal spatial variance at the prediction location, equal to the process variance σ^2 in a stationary field.

Ordinary kriging provides the best linear unbiased estimate (BLUP) of an unknown value Z_0 by optimally weighting surrounding observations based on their spatial configuration and covariance structure. Unlike simple kriging, it does not assume a known mean, requiring only that the mean be constant within the local neighborhood. The method minimizes the expected squared prediction error while enforcing an unbiasedness constraint. The final estimate is expressed as a weighted sum of the observed values, where the weights are obtained by solving a linear system augmented with a constraint that their sum equals one. This formulation ensures that the prediction is both statistically consistent and sensitive to the spatial structure of the data.

3. Extension 1: Incorporating Terrain Material Types

We extend ordinary kriging to include the influence of terrain material types (e.g., rock, sand, clay, vegetation) and cultural constructions (e.g., buildings, runways, and military targets) by modeling the spatial variability and roughness characteristics associated with each class.

3.1. Material-Specific Variability

Different surface types exhibit distinct elevation variability profiles:

- Rocky terrain: high local variability, short-range correlation
- Sandy terrain: smoother surfaces, long-range correlation
- Vegetated surfaces: intermediate variability with moderate correlation range
- Built structures (e.g., buildings, runways): abrupt discontinuities, very short correlation range, directionally persistent edges
- Military targets: heterogeneous structure, often mixed material classes with both smooth and abrupt features

We incorporate this variability into kriging by modifying the diagonal elements of the covariance matrix to include class-specific residual variance. This adjustment reflects local roughness and structural behavior tied to material or construction type, improving prediction accuracy in heterogeneous environments.

3.2. Class-Based Variance Adjustment

If a single global covariance model is used, terrain-dependent variability can be included by assigning different local variance values to each material class. This is implemented by modifying the diagonal terms of Σ:

$$\Sigma_{ii} = \sigma^2 + \delta^2(m_i)$$

where $\delta^2(m_i)$ is the material-dependent residual variance assigned to class .

For example:

- Rock: $\delta^2 = 4.0$
- Sand: $\delta^2 = 0.5$
- Grassland: $\delta^2 = 1.2$

These adjustments allow the model to express that observations from rougher surfaces are more spatially variable, while smoother materials are more predictable.

In this modified method, the terrain material class must be specified at each observation point to enable the variance adjustment. This information determines the diagonal entry of the covariance matrix corresponding to each sample, allowing local roughness to be expressed without modifying the global correlation structure. The prediction variance $\tau^2(s_0)$ then reflects material-dependent variability directly, without the need for additional structural changes to the kriging system. Importantly, the constraint that the sum of the weights equals one remains unchanged, preserving the unbiasedness condition inherent to ordinary kriging.

3.3. Details of Material-Adjusted Covariance

Observations are located at positions (x_i, y_i), and the associated values z_i represent terrain height. The goal is to predict Z at an unmeasured location (x_0, y_0) using ordinary kriging with a covariance model that now also accounts for terrain material–dependent variability.

As above, we take $(x_1, y_1, z_1, m_1), \ldots, (x_n, y_n, z_n, m_n)$ where m_i is the terrain material class at point and a prediction point: $(x_0, y_0) \rightarrow \hat{z}_0$, with known or assumed material m_0.

3.3.1. Step 1: Construct Vectors and Covariance Structures

Define:

$$\vec{s}_i = (x_i, y_i)$$
$$\vec{s}_0 = (x_0, y_0)$$
$$d_{ij} = \|\vec{s}_i - \vec{s}_j\|$$
$$d_i = \|\vec{s}_i - \vec{s}_0\|$$

Covariance function isotropic characteristics:

$$K(d) = \sigma^2 \cdot \exp\left(-\frac{d^2}{2\ell^2}\right)$$

where ℓ is the spatial range parameter.
Next, we need to build the covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$:

- Off-diagonal entries:

$$\Sigma_{ij} = K(d_{ij}) \text{ for } i \neq j$$

-Diagonal entries (material-adjusted residual variance):
$\Sigma_{ii} = K(0) + \delta^2(m_i)$ where $\delta^2(m_i)$ is the residual variance for material class (m_i)

Covariance vector $c \in \mathbb{R}^n$:

$$c_i = K(d_i)$$

3.3.2. Step 2: Solve for Weights with Unbiasedness Constraint

To enforce $\sum \lambda_i = 1$, as with ordinary kriging, solve the augmented linear system:

$$\begin{bmatrix} \Sigma & \mathbf{1} \\ \mathbf{1}^T & 0 \end{bmatrix} \begin{bmatrix} \lambda \\ \mu \end{bmatrix} = \begin{bmatrix} c \\ 1 \end{bmatrix}$$

where:
- λ are the kriging weights
- μ is a Lagrange multiplier
- $\mathbf{1}$ is an n-vector of ones

3.3.3. Step 3: Compute Predicted Value

Once the kriging weights λ_i have been determined by solving the augmented system, the predicted elevation $\hat{z}_0$ at the location (x_0, y_0) is computed as a weighted sum of the observed values:

$$\hat{z}_0 = \sum_{i=1}^{n} \lambda_i z_i$$

Each weight λ_i depends on the spatial arrangement of the observations and the assigned material classes through the material-adjusted covariance matrix $\mathbf{\Sigma}$ and the covariance vector c. Writing the prediction more explicitly, we substitute the definition of $c_i = K(d_i)$ and define $\mathbf{\Sigma}$ component-wise:

$$\hat{z}_0 = \sum_{i=1}^{n} z_i \left[\sum_{j=1}^{n} (\Sigma^*)^{-1}_{ij} K\left(\|\vec{s}_j - \vec{s}_0\|\right) - \left(\frac{\mathbf{1}^T (\Sigma^*)^{-1} c - 1}{\mathbf{1}^T (\Sigma^*)^{-1} \mathbf{1}} \right) (\Sigma^*)^{-1}_{ij} \right]$$

and the covariance vector is:

$$c_j = K(|\vec{s}_j - \vec{s}_0|)$$

This method ensures that the weights reflect both the spatial configuration and material properties of the observed points. The Lagrange multiplier term, which enforces the constraint $\Sigma \lambda_i = 1$, adjusts the weights to maintain unbiasedness regardless of material distribution or sample geometry. The final prediction $\hat{z}_0$ incorporates anisotropy, material class, and spatial redundancy without changing the fundamental form of ordinary kriging.

3.4. Benefits of Material Consideration

Incorporating material-based variance into ordinary kriging introduces local surface variability while preserving the core structure of the method. Implementation is straightforward and involves a class-specific adjustment to the diagonal of the covariance matrix. This modification allows the model to account for material-dependent surface behavior, such as the roughness of rock or the smoothness of sand, without requiring multiple covariance models or altering the global correlation structure. In standard kriging, residual variance is uniform and surface roughness is treated as constant. By contrast, this material-based approach introduces heteroscedasticity aligned with material properties, enabling terrain-adaptive variance while maintaining a single global covariance model. This results in improved spatial realism, more meaningful uncertainty quantification, and increased interpretability without increasing underlying algorithmic complexity.

4. Extension 2: Add Angular Regularization

If sampling geometry is biased, define pairwise angles:

$$\theta_{ij} = \angle(\vec{s}_i - \vec{s}_0, \vec{s}_j - \vec{s}_0)$$

and penalize alignment:

$$A(\lambda) = \sum_{i<j} \lambda_i \lambda_j \cdot \exp\left(-\frac{(\theta_{ij} - \pi/2)^2}{2\sigma_\theta^2} \right)$$

Augmented objective (if minimizing directly):

$$\Psi(\lambda) = \lambda^T \Sigma \lambda - 2c^T \lambda + \rho \cdot A(\lambda)$$

In this process, input coordinates (x, y) are specified in meters and used directly to compute Euclidean distances between observation points and the prediction location. The terrain material at each sample point influences the interpolation by modifying the diagonal of the covariance matrix $\mathbf{\Sigma}$, allowing residual variance to reflect local surface variability. Kriging weights are computed by solving an

augmented linear system that enforces the unbiasedness constraint, ensuring that the sum of the weights equals one. The interpolated elevation $\hat{z}_0$ is then obtained as a weighted sum of the observed elevation values z_i. When directional bias is present in the sampling geometry, an optional angular penalty can be applied to reduce the influence of collinear points and promote more isotropic sampling support. This extension improves the reliability of the estimate in anisotropic or under-sampled regions without altering the core structure of ordinary kriging.

In cases where the spatial distribution of samples is directionally biased, such as when points lie predominantly along a single axis, kriging can overweight collinear sources and underrepresent directions that are sparsely sampled. This results in anisotropic smoothing, loss of angular fidelity, and amplification of alignment artifacts in the interpolated field. To mitigate this, we introduce an angular diversity penalty based on the pairwise directions of samples relative to the prediction point.
Define the unit vectors from the prediction location $\vec{s}_0$ to each sample $\vec{s}_i$ as $\vec{v}_i = \vec{s}_i - \vec{s}_0$. For each unordered sample pair (*i, j*), compute the angle between vectors $\vec{v}_i$ and $\vec{v}_j$:

$$\theta_{ij} = \angle(\vec{v}_i, \vec{v}_j)$$

To encourage angular coverage, we penalize configurations where sample pairs are nearly aligned. The penalty function is constructed as a Gaussian centered at $\pi / 2$, which promotes orthogonal sampling directions:

$$A(\lambda) = \sum_{i<j} \lambda_i \lambda_j \cdot \exp\left(-\frac{(\theta_{ij} - \pi/2)^2}{2\sigma_\theta^2}\right)$$

Here, σ_θ controls the angular selectivity and ρ is a scalar that sets the strength of the angular regularization. This term is added to the ordinary kriging objective, yielding the augmented functional:

$$\Psi(\lambda) = \lambda^T \Sigma \lambda - 2c^T \lambda + \rho \cdot A(\lambda)$$

This form remains quadratic in λ with a smooth additive correction that introduces no additional constraints or unknowns. It favors solutions where influence is distributed across diverse angular sectors while still minimizing prediction variance.

Combining this angular regularization with the material-aware kriging formulation, the final expanded prediction for the interpolated value $\hat{z}_0$ becomes:

$$\hat{z}_0 = \sum_{i=1}^{n} z_i \left[\sum_{j=1}^{n} (\Sigma^*)_{ij}^{-1} K\left(\|\vec{s}_j - \vec{s}_0\|\right) - \left(\frac{\mathbf{1}^T (\Sigma^*)^{-1} c - 1}{\mathbf{1}^T (\Sigma^*)^{-1} \mathbf{1}} \right) \sum_{j=1}^{n} (\Sigma^*)_{ij}^{-1} \right] - \rho \cdot A(\lambda)$$

In this expression, Σ* encodes both spatial proximity and material-specific variance, *c* contains the spatial covariance terms between the prediction location and samples, and $A(\lambda)$ penalizes angular redundancy in the influence map. This combined form improves robustness in heterogeneous sampling geometries and anisotropic surfaces while preserving the unbiasedness and covariance-informed structure of ordinary kriging.

4.1. Formulation

For clarity and finality, we can fully expand this with kriging weights and the angular penalty as:

$$\hat{Z}_0 = \sum_{i=1}^{n} z_i \left[\sum_{j=1}^{n} \left(\Sigma^{*-1}\right)_{ij} K(\|\vec{s}_j - \vec{s}_0\|) - \left(\frac{\mathbf{1}^T \Sigma^{*-1} c - 1}{\mathbf{1}^T \Sigma^{*-1} \mathbf{1}} \right) \left(\Sigma^{*-1}\right)_{ij} \right]$$
$$-\rho \sum_{i<j} \left[\sum_{k=1}^{n} \left(\Sigma^{*-1}\right)_{ik} K(\|\vec{s}_k - \vec{s}_0\|) - \left(\frac{\mathbf{1}^T \Sigma^{*-1} c - 1}{\mathbf{1}^T \Sigma^{*-1} \mathbf{1}} \right) \left(\Sigma^{*-1}\right)_{ik} \right]$$
$$\left[\sum_{\ell=1}^{n} \left(\Sigma^{*-1}\right)_{j\ell} K(\|\vec{s}_\ell - \vec{s}_0\|) - \left(\frac{\mathbf{1}^T \Sigma^{*-1} c - 1}{\mathbf{1}^T \Sigma^{*-1} \mathbf{1}} \right) \left(\Sigma^{*-1}\right)_{j\ell} \right] \cdot \exp\left(-\frac{(\theta_{ij} - \pi/2)^2}{2\sigma_\theta^2} \right)$$

Where:
- $\vec{s}_i = (x_i, y_i)$: coordinate of sample i
- z_i : observed elevation at $\vec{s}_i$
- m_i : terrain material class at $\vec{s}_i$
- $K(d) = \sigma^2 \cdot \exp\left(-\frac{d^2}{2\ell^2}\right)$: isotropic spatial covariance kernel
- $\delta^2(m_i)$: residual variance for material class m_i
- $\Sigma^*_{ij} = K(\|\vec{s}_i - \vec{s}_j\|)$ for $i \neq j, \Sigma^*_{ii} = K(0) + \delta^2(m_i)$

$$\Sigma^*_{ij} = \begin{cases} K(\|\vec{s}_i - \vec{s}_j\|), & i \neq j \\ K(0) + \delta^2(m_i), & i = j \end{cases}$$

- $c_j = K(\|\vec{s}_j - \vec{s}_0\|)$: covariance vector between sample j and prediction point $\vec{s}_0$
- $\theta_{ij} = \angle(\vec{s}_i - \vec{s}_0,\ \vec{s}_j - \vec{s}_0)$: angular separation between samples relative to $\vec{s}_0$
- ρ : angular penalty weight
- σ_θ : angular spread parameter for directional regularization
- 1 : n-vector of ones

This formulation for $\hat{Z}_0$ encodes:
- Spatial weighting via $K(d)$
- Local material roughness via $\delta^2(m_i)$
- Constraint enforcement via the Lagrange multiplier term
- Directional sampling regularization via angular penalty coupling all pairs (i, j)

4.2. Benefits of Angular Sensitivity

Directional sensitivity improves the interpolation of structured surfaces by penalizing the overreliance on collinear or clustered input points. In terrain and constructed environments, elevation changes often follow dominant gradients; along ridge-lines, edges of runways, or the faces of buildings. When samples are aligned along such features, conventional kriging may assign excessive weight to these directions, leading to elongated artifacts and loss of transverse detail. Incorporating directional sensitivity addresses this by favoring angular diversity in the spatial support of the prediction.

This is implemented by reducing the combined influence of pairs of samples that contribute similar directional information relative to the prediction location. The result is a weighting pattern that spreads influence more uniformly across available angular sectors, reducing directional bias. This is particularly useful in cases where sampling geometry is non-uniform or dictated by platform constraints, such as nadir-only sensors or structured flight lines. Directional regularization increases robustness in sparse, anisotropic fields and preserves fine-scale orthogonal structure often lost in purely distance-based methods.

By explicitly accounting for angular redundancy, the method improves the local representativeness of the estimate while preserving the statistical consistency of kriging. It maintains unbiasedness and allows the existing covariance structure to dominate where geometry is already diverse. This ensures minimal disruption to the formulation while improving fidelity in directionally structured domains.

5. Prior Work

Prior methods addressing spatial heterogeneity include Locally Varying Anisotropy (LVA) for directional effects and Multi-Unit Variograms (MUVs) for material variation. LVA introduces spatial anisotropy by deforming the distance metric used in the covariance function, typically through a Mahalanobis transform derived from locally estimated orientation and scale tensors. While this allows modeling of directional continuity in layered or elongated features, it requires dense and reliable estimates of local anisotropy, which are difficult to obtain in sparse or irregular datasets.

MUVs address material heterogeneity by partitioning the domain into discrete facies or units, each interpolated independently using indicator or proportion kriging. Variogram models are fitted separately for each unit, and estimates are combined by weighted averaging. This process introduces structural discontinuities and restricts estimation to same-class data, which limits fidelity near transitions and in mixed-material regions.

KRAM avoids these limitations by retaining the standard kriging framework and introducing heterogeneity through two additive modifications. Material-dependent variance is modeled through diagonal scaling of the covariance matrix based on class membership, allowing all observations to contribute while modulating influence by residual uncertainty. Directional imbalance is addressed by adding a regularization term to the objective function that penalizes angular redundancy in the sampling configuration. These changes preserve the statistical consistency of ordinary kriging and require no space warping, variogram splitting, or class segmentation.

5.1. Locally Varying Anisotropy (LVA)

LVA addresses directional variation in spatial correlation by replacing the scalar Euclidean distance $\boldsymbol{d}$ in the covariance function $\boldsymbol{K(d)}$ with a locally adapted metric that reflects anisotropy at each interpolation point. This is typically implemented using a Mahalanobis-style transformation that warps space via a rotation and scaling defined by a local anisotropy tensor. The result is a directionally dependent distance measure that enables kriging to model elongated features such as stratified orebodies, flow-aligned sediments, or structural layering [5].

While theoretically well-suited for anisotropic fields, LVA requires accurate estimation of the anisotropy direction and aspect ratio at each prediction site. This typically involves local variogram fitting or gradient estimation, which is data-intensive and unstable in sparse or irregular sampling conditions. The deformation of space also complicates variogram modeling, increases sensitivity to noise, and introduces implementation overhead due to the need for a full tensor field across the domain.

This approach achieves similar goals without modifying the geometry of the space. Rather than altering $\boldsymbol{K(d)}$, we retain the isotropic form of the covariance function and introduce a penalty term into the kriging objective that discourages angular redundancy in the spatial support. This avoids the need for local tensor estimation while maintaining compatibility with existing kriging workflows and preserving convexity and interpretability. The result is improved directional balance with fewer parameters and greater robustness under sparse or structured sampling.

5.1.1. Preservation of Isotropic Covariance Form

LVA modifies the spatial distance metric directly by applying locally rotated and scaled transformations to the coordinate space. This requires defining a local anisotropy tensor at each point, which can be difficult to estimate robustly in sparse or irregularly sampled domains. We avoid this by preserving the original scalar distance function $\boldsymbol{K(d)}$ and instead adding a regularization term that operates on the kriging weights. This allows the existing covariance structure to remain unchanged, simplifying both implementation and interpretation.

5.1.2. Regularization Instead of Structural Redefinition

LVA changes the geometry of the model itself by warping space, which can cause sensitivity to local estimation errors and complicate variogram fitting. We introduce angular sensitivity through an additional penalty term in the objective function: $\Psi(\lambda) = \lambda^T \Sigma \lambda - 2c^T \lambda + \rho A(\lambda)$, where $A(\lambda)$ penalizes angular redundancy. This formulation augments the standard kriging optimization without changing the underlying system structure. It is mathematically stable, preserves convexity, and can be tuned via a single penalty parameter ρ.

5.1.3. Minimal Parameter Overhead

LVA requires estimating a local anisotropy tensor, including direction and magnitude, at each interpolation point. This typically involves directional variogram fitting or local gradient analysis, both of which are sensitive to noise and difficult to perform reliably in sparse or irregular fields. Moreover, anisotropy estimation is highly geospecific and context-dependent, often relying on prior geologic or structural knowledge. KRAM avoids this complexity by introducing only two global parameters: a regularization weight ρ and an

angular spread parameter σ_θ. This reduces tuning effort, avoids overfitting, and improves robustness under general conditions.

5.1.4. Compatibility with Existing Models

Because KRAM does not modify the distance metric or require directionally dependent variogram models, it remains compatible with standard kriging workflows. It can be implemented as a post-hoc weighting adjustment, avoiding the need for additional preprocessing or structural parameter estimation.

5.1.5. Improved Control over Angular Redundancy

Finally, we directly target angular redundancy through a controlled functional form. This allows for fine-tuning of directional influence without requiring the transformation of coordinate space or introduction of potentially unstable local metrics. It offers a more transparent and interpretable way to enforce directional diversity in the interpolation support.

5.2. Multi-Unit Variograms (MUVs)

MUVs are commonly used in the mining industry to address material heterogeneity by treating each geological unit (or facies) as a separate interpolation domain [6]. Each unit, such as oxide, sulfide, or mixed ore, is interpolated independently using class-specific kriging or indicator methods. Grade and volumetric proportion are estimated separately per unit and then combined into a final block-level prediction.

This partitioned strategy requires a separate variogram model for each class and restricts estimation to within-class samples, excluding potentially informative data from other classes. The domain must be segmented prior to interpolation, and results from each class must be recombined post-hoc, increasing both modeling complexity and operational overhead. Transitional zones between material types are especially problematic, often requiring manual interpretation or introducing artifacts at unit boundaries.

KRAM eliminates the need for segmentation by embedding material influence directly into the kriging system. Rather than isolating classes, we adjust the diagonal of the covariance matrix to reflect class-specific residual variance. This allows all samples to contribute simultaneously, with material differences encoded as varying levels of confidence rather than strict separation. The global covariance structure remains intact, ensuring statistical consistency while enabling cross-class interpolation in mixed or uncertain regions. The result is a unified, stable, and more easily implemented model for terrain with heterogeneous surface properties.

5.2.1. Unified Framework Instead of Class Partitioning

MUV approaches interpolate each material class independently and recombine results at the block level. KRAM retains a single kriging system by embedding material-dependent variability directly into the covariance matrix. All observations contribute simultaneously, with class effects represented through diagonal variance adjustments.

5.2.2. Elimination of Multiple Variogram Models

MUVs require separate variogram fitting for each material type, increasing modeling effort and introducing discontinuities at class interfaces. Our approach uses a single global covariance model and incorporates material heterogeneity through scalar variance terms, eliminating the need for model switching.

5.2.3. Cross-Class Interpolation Support

In MUV workflows, estimation is restricted to observations of the same material class. This limits predictive quality in regions near facies boundaries or containing mixed materials. We allow all samples to contribute to the prediction, with influence modulated by class-dependent variance, improving accuracy in transitional zones.

5.2.4. Simplified Implementation

MUV systems require class segmentation, multiple kriging passes, and post-processing to combine results. Our approach modifies only the diagonal of the covariance matrix based on class labels, preserving the standard kriging system and reducing implementation complexity.

5.2.5. Preservation of Kriging Optimality

Because KRAM perturbs only the diagonal of the covariance matrix, it preserves the unbiasedness and minimum variance properties of ordinary kriging. The statistical structure remains intact while extending applicability to geologically heterogeneous domains.

5.3. Comparison to Deep Learning REM Interpolation

Deep learning has been applied to the generation of Radio Environment Maps (REMs), where signal strength is predicted from spatial

sensor data [7]. These models learn spatial structure from labeled examples, often achieving high accuracy in dense and well-labeled datasets. However, they do not explicitly represent angular variation or material heterogeneity. Anisotropic effects may be learned implicitly during training, but the model lacks explicit control over directional influence or material variance.

No formal mechanism exists to encode surface type. Material differences are only learned if labeled examples are provided, and even then, interpretation is limited to patterns embedded in internal network weights. Once trained, the model becomes a fixed function with no parametric access to covariance, variance, or weighting structure. This makes error diagnosis and adjustment difficult, particularly in operational contexts where transparency and traceability are required.

In contrast, KRAM provides explicit, interpretable control over material variance and angular redundancy. KRAM retains the structure of ordinary kriging while introducing only two additional global parameters. Each component, the covariance model, material scaling, and angular penalty, remains accessible and auditable. This is important in applications where model behavior must be explained, modified, or validated. Deep learning models may perform well in narrow domains with sufficient training data, but lack the interpretability and control required for traceable geostatistical inference.

5.3.1. No Explicit Material or Angular Encoding

REM generation does not explicitly model angular heterogeneity or surface type. Any anisotropic behavior is learned implicitly through training, and potential material effects are only captured if such classes are present and well-labeled in the dataset. There is no mechanism in the network to encode physical material properties or directional bias directly.

5.3.2. Opaque Internal Representations

These models rely on internal states that lack direct physical interpretation. Once trained, prediction outputs result from abstract learned features and cannot be traced back to meaningful quantities like covariance structure or material variance. This opacity prevents auditing, limits model explainability, and restricts adaptability in operational contexts where interpretability is required.

5.3.3. Dependence on Training Coverage

Performance depends heavily on the diversity and density of the training set. If certain surface types or directional configurations are underrepresented, the model may fail to generalize. There is no guarantee of graceful degradation outside the training distribution, and error behavior is difficult to characterize or bound.

5.3.4. No Parametric Control

Unlike kriging, which allows user-defined variogram parameters, directional penalties, or material-based variance terms, deep learning methods offer limited parametric adjustment after training. Model behavior is effectively fixed, requiring retraining to adapt to new sampling geometries or material regimes.

6. KRAM Advantages

KRAM preserves full control over the interpolation behavior. Material variance, angular weighting, and spatial correlation are all explicitly defined and interpretable. The kriging system remains transparent, auditable, and tunable, offering consistent and explainable results across varying terrain and sampling configurations. This makes it well suited to applications requiring traceable reasoning, performance diagnostics, or hybrid modeling.

7. Conclusion

Ordinary kriging can be improved by addressing two major sources of bias: surface heterogeneity and directional imbalance in sampling. Material-dependent variance adjusts the kriging formulation to reflect local surface behavior without introducing separate models or structural complexity. Angular regularization complements this by discouraging overreliance on collinear inputs, improving coverage around the prediction site. These extensions maintain the statistical properties of kriging while producing estimates that more accurately reflect terrain and structural form. Combined, they offer a practical and robust enhancement for elevation interpolation in sparse, structured, or mixed-material datasets.

References

1. Shepard, D. (1968, January). A two-dimensional interpolation function for irregularly-spaced data. In *Proceedings of the 1968 23rd ACM national conference* (pp. 517-524).
2. Matheron, G. (1963). Principles of geostatistics. *Economic geology, 58*(8), 1246-1266.
3. Cressie, N. (2015). *Statistics for spatial data*. John Wiley & Sons.
4. Chiles, J. P., & Delfiner, P. (2012). *Geostatistics: modeling spatial uncertainty*. John Wiley & Sons.

5. Boisvert, J. B., Manchuk, J. G., & Deutsch, C. V. (2009). Kriging in the presence of locally varying anisotropy using non-Euclidean distances. *Mathematical Geosciences, 41*, 585-601.
6. Séguret, S. (2013). Mathematical Geosciences.
7. Burmeister, F., Palaios, A., Geuer, P., Krause, A., Jacob, R., Schulz, P., & Fettweis, G. (2023, June). Deep Learning-Supported Kriging for Interpolation of High-Resolution Indoor REMs. In *2023 Joint European Conference on Networks and Communications & 6G Summit* (*EuCNC/6G Summit*) (pp. 54-59). IEEE.