

Quantifying cultural-intensity distributions in three dimensions: an anisotropic Gaussian spatial model with confidence weighting, demonstrated on the Neolithic cemetery–altar complex of Lingjiatan, China

Chenglei Meng

13227@nnu.edu.cn

Nanjing Normal University

Article

Keywords: spatial archaeology, kernel density estimation, anisotropy, confidence weighting, Lingjiatan, Neolithic China

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10894258/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

How is cultural activity distributed across an archaeological site, and can the directionality of that distribution be measured rather than smoothed away? We introduce the three-dimensional Gaussian spatial model (3DGSM), which represents each cultural feature unit as an anisotropic Gaussian kernel whose position, attribute, confidence weight and influence covariance are all archaeologically interpretable, and we apply it to 85 cultural feature units—70 tombs and 15 altar or sacrificial units—digitised from the published excavation plans of the Neolithic cemetery–altar complex at Lingjiatan, Anhui, China, with recorded intensities taken from the artefact counts documented in the five excavation campaigns (1987–2007) and the 2016 clearing of tomb 07M22. The fitted influence geometry is moderately anisotropic: influence extends c. 2.6 times farther along a principal axis (azimuth 146.8° , NNW–SSE) than perpendicular to it, following the long axis of the ridge on which the cemetery–altar complex sits, while the short vertical axis (0.6 m; horizontal-to-vertical decay ratio c. 29) quantifies the stratigraphic compactness of the deposits. Under a deliberately non-circular leave-one-out evaluation against held-out recorded intensities, the model attains an RMSE of 0.133, significantly outperforming ordinary kriging (0.140, $p = 0.020$) and reproducing the layout of the recorded distribution (zone-level agreement 75.3%; Spearman rank correlation 0.479); differences against isotropic KDE, IDW and the model's own ablations are small and non-significant, and we report them as such. A controlled simulation verifies that the estimator recovers known anisotropy whenever it is present (azimuth error below 5.4° for axis ratios $\rho \geq 2$, and below 0.5° at $\rho = 8$, $n = 200$) and that its predictive advantage over isotropic alternatives grows with true anisotropy and sample size—without inventing gains where none exist. The fitted geometry is stable under publication-scale digitization error, and a sensitivity analysis over an alternative composite intensity protocol changes the fitted azimuth by only 1.5° . Treating directionality and observational uncertainty as estimable quantities rather than smoothing choices, the 3DGSM offers an interpretable complement to the KDE-based workflows familiar to most archaeologists.

1. Introduction

A recurring question in spatial archaeology is deceptively simple: how is cultural activity distributed across a site, and what does the geometry of that distribution mean? For most of the discipline's quantitative history the question has been answered with density surfaces: findspots are smoothed into a continuous intensity, and the surface is read for hotspots. The smoothing geometry, however, is almost always chosen for convenience rather than estimated for meaning—a circular kernel applied in the horizontal plane, with elevation, where recorded, relegated to a secondary surface. Two things are lost in that choice. The first is directionality: cemeteries elongate along ridges, artefact distributions track paths and terraces, and ritual deposits respect bounded spaces, so that the way intensity decays with distance is itself evidence of how past communities organised space. The second is reliability: excavation records are uneven in quality, with disturbed contexts and partially excavated margins carrying systematically less certain information, yet conventional smoothing trusts every findspot equally. These are not abstract concerns: at the Neolithic cemetery–altar complex of Lingjiatan in Anhui Province, China,

whether the spatial organisation of the cemetery and altar was oriented to the river terrace on which it sits is a live archaeological question—and one that conventional density-surface workflows cannot answer quantitatively.

This paper asks whether both quantities can be estimated directly. We introduce the three-dimensional Gaussian spatial model (3DGSM), in which every cultural feature unit contributes an anisotropic Gaussian influence kernel whose four parameters—the centre position μ_i , the cultural attribute C_i , the confidence weight α_i and the influence covariance Σ_i —each carry an explicit archaeological interpretation. The model takes its representational cue from 3D Gaussian splatting in computer graphics (Kerbl et al. 2023), where a scene is modelled explicitly as a set of weighted anisotropic Gaussians, but it is estimated and evaluated with classical spatial-statistical machinery. The influence covariance is the key object: its eigenvalues quantify how far cultural influence extends in each direction, and its eigenvectors report the orientation of that spread, interpretable archaeologically as the orientation of movement corridors, terrace-guided spread or deliberately bounded ritual space. The confidence weight, by contrast, gives formal expression to excavation and taphonomic bias, allowing uncertain observations to contribute proportionally less to the estimated cultural landscape.

We develop and evaluate the 3DGSM with four aims: (i) to give each model parameter an explicit archaeological interpretation, so that estimation outputs translate directly into archaeological arguments; (ii) to quantify the anisotropy and the vertical structure of cultural-intensity distributions in three dimensions; (iii) to demonstrate the workflow on a real prehistoric case—the Neolithic cemetery–altar complex at Lingjiatan in Anhui Province, China—using positions digitised from published excavation plans and artefact counts documented in the excavation reports, together with a formal assessment of digitization error; and (iv) to benchmark predictive performance against isotropic 2D KDE, IDW and ordinary kriging, using a non-circular evaluation design with significance testing, supplemented by a controlled simulation study that verifies the recovery of known anisotropy parameters under favourable, exactly Gaussian conditions.

2. Background and related work

2.1 Quantitative spatial analysis in archaeology

Formal spatial statistics entered archaeology through the work of Hodder and Orton (1976), who adapted quadrat, nearest-neighbour and distance-based methods to the study of artefact scatters and settlement patterns. Ripley's K-function and its descendants supplied the inferential core of this tradition, testing aggregation and interaction against complete spatial randomness over a range of distances (Ripley 1977); Beardah and Baxter (1996) carried the companion technique of kernel density estimation explicitly into archaeology, where it has remained the standard route from discrete findspots to continuous surfaces ever since. At the intrasite scale, unconstrained clustering (Whallon 1984; Moclán et al. 2023) established the principle that activity areas should be recognised from the data rather than imposed from a grid, while the subsequent diffusion of GIS transformed landscape-scale analysis,

systematising predictive modelling of site locations, the study of movement across terrain and, more recently, the multivariate modelling of landscape affordances and connectivity (Wheatley and Gillings 2002; Conolly and Lake 2006; Kempf 2020). Across both scales, density estimation is the workhorse: following Baxter et al. (1997), KDE has been applied to artefact distributions, settlement intensities and chronological smoothing, and a substantial methodological literature now addresses bandwidth choice (Silverman 1986) and the interpretation of the resulting surfaces (VanPool and Leonard 2011). Localised approaches such as local indicators of spatial association (Anselin 1995) complement these global surfaces by detecting pockets of unusual concentration, but they share the same planar, direction-agnostic framing.

The present model sits squarely within this point-pattern mainstream, but asks a different question of it. The K-function family asks *where* patterning is significant, summarising second-order structure as a function of distance; the 3DGSM asks *what* its geometry is, summarising the same structure as an orientation, an axis ratio and a set of directional decay rates. Adaptive and anisotropic bandwidths have of course been explored within KDE precisely to cope with elongated clusters, but there the orientation enters as a smoothing parameter to be tuned; in the 3DGSM it is the estimand itself. The model thus continues, rather than replaces, the dialogue opened by Ripley (1977) and Beardah and Baxter (1996): from describing that a distribution is structured, to measuring how.

Recent years have seen a renewed emphasis on the formal treatment of uncertainty and reproducibility in archaeological analysis. Crema (2025) argues that measurement error, hierarchical structure and simulation-based generative inference should become standard components of quantitative workflows; community assessments now benchmark quantitative methods against simulated or experimentally controlled ground truth (Cortell-Nicolau et al. 2025; Liu et al. 2026); and reproducibility reviews have exposed how often published analyses cannot be re-run (Marwick 2025). The 3DGSM's confidence weight is a direct response to this agenda: excavation and taphonomic uncertainty enters the model as an explicit, per-unit parameter rather than as unmodelled noise. What the tradition still lacks, however, is a treatment of directional structure: directional variograms and allied tools have long existed in spatial statistics (Cressie 1993), but archaeological applications describe anisotropy after the fact rather than estimating it as an interpretable model parameter.

Two limitations of this tradition motivate the present work. The first is isotropy. Although anisotropic kernels are well established in spatial statistics (Cressie 1993), archaeological KDE is overwhelmingly implemented with a circular (or, at best, a separately scaled axis-aligned) kernel, and the orientation of the smoothing geometry is almost never itself an object of estimation and interpretation. The second is dimensional. Truly three-dimensional analysis of artefact proveniences remains rare relative to the volume of two-dimensional work, and where elevation is incorporated it is commonly treated as a secondary surface rather than as a third axis of the estimation space. The result is that stratigraphic separation—one of the most archaeologically meaningful forms of spatial structure—can be smoothed into apparent horizontal proximity.

2.2 From Gaussian splatting to Gaussian spatial models

The 3DGS technique of Kerbl et al. (2023) achieved real-time, high-fidelity novel-view synthesis by optimising an explicit representation of a scene as anisotropic 3D Gaussians, extending the implicit neural representation paradigm of Mildenhall et al. (2021); subsequent work has structured the representation for efficiency and fidelity, anchoring Gaussians to scaffolded point sets (Lu et al. 2024) and removing aliasing artefacts across scales (Yu et al. 2024). Its relevance here is conceptual rather than technical: 3DGS demonstrates that a complex, directionally structured spatial phenomenon can be recovered by optimising the positions, opacities and covariances of a set of weighted Gaussian primitives against observations. We transpose that representational insight into a statistical model of cultural-intensity distributions. The opacity parameter becomes an observation-reliability weight; the colour channels become a cultural-attribute vector; and the covariance matrix, freed from its rendering role, becomes a direct quantitative description of how cultural influence decays in three-dimensional space. No component of the splatting rasteriser is used, and the model is estimated with classical spatial-statistical machinery (weighted maximum likelihood and kernel regression) rather than photometric losses. The 3DGSM should therefore be understood as a member of the family of weighted Gaussian mixture and kernel methods long familiar in spatial statistics (Cressie 1993; Silverman 1986), parameterised in a way that makes archaeological interpretation immediate. The transposition of 3DGS concepts to cultural heritage has begun—spanning artefact visualisation (Khedr et al. 2025), architectural-heritage recording (Yu et al. 2025) and crossmedia representation (Mirocha 2025)—but in such work the Gaussian primitives remain rendering elements; the 3DGSM instead promotes the anisotropic Gaussian to the status of a direct estimand within a classical spatial-statistical framework.

3. Data and methods

3.1 The Lingjiatan cemetery–altar complex

Lingjiatan (31°27'39"N, 118°02'53"E), in Hanshan County, Anhui Province, is a late Neolithic settlement and ceremonial centre of c. 1.6 km² on a terraced tributary of the Yuxi River, dated to c. 3600–3300 BCE (Anhui Provincial Institute of Cultural Relics and Archaeology 2006, 2008; Liu and Chen 2012, pp 145–148; Wang et al. 2022). The coordinates refer to the cemetery–altar complex itself and are taken from the published excavation literature (Anhui Provincial Institute of Cultural Relics and Archaeology 2006). Its best-known component is a cemetery–altar complex at the highest point of a low ridge—the southern end of a c. 5 km ridge descending from the Taihu hills to the Yuxi River—where a central altar area, a surrounding cemetery of richly furnished tombs, and adjacent activity areas together yielded the jade assemblage for which the site is famous (China Cultural Relics News 2022). The complex was investigated in five excavation campaigns between 1987 and 2007: trial and formal excavations in 1987 uncovered 15 tombs (Anhui Provincial Institute of Cultural Relics and Archaeology 1989); the 1998 campaign exposed 29 further tombs and first recognised the stone-paved altar, a discovery listed among China's top ten archaeological finds of that year (Anhui Provincial Institute of Cultural Relics and Archaeology and Hanshan County Office for the Management of Cultural Relics 1999; Zhang 1999); the 2000 campaign added 21 tombs north-west of the altar (Hu et al. 2021); and the 2007 campaign

excavated four tombs, three ash-pits and a stone feature near the altar, including the richly furnished tomb 07M23 with its documented 330 grave goods and the c. 88 kg jade pig found above its fill (Anhui Provincial Institute of Cultural Relics and Archaeology 2008). Tomb 07M22, located in 2007 beneath a modern grave, was cleared in December 2016 after that grave was relocated (Lingjiatan Site Archaeological Team 2020). Because the complex occupies an elongated ridge segment, the directionality of the cultural distribution is itself an archaeological question: if the spatial organisation of the cemetery and altar respected the ridge, the orientation of the distribution should follow the ridge axis; if it reflected bounded ritual space alone, elongation should instead relate to the documented arrangement of the altar and activity areas. Renewed excavation since 2020 has substantially enlarged the documented ritual precinct south-east of the ridge (China Cultural Relics News 2022, 2026), and we return to these results as independent evidence in Section 5.1.

3.2 Data and digitization protocol

The evaluation data consist of the 85 cultural feature units—70 tombs and 15 altar or sacrificial units—documented in the published plans and reports of the cemetery–altar complex. The corpus assembles every unit reported by the five excavation campaigns and the 2016 clearing: the 44 tombs of the 1987 and 1998 campaigns, whose per-tomb artefact counts are catalogued in the excavation report (Anhui Provincial Institute of Cultural Relics and Archaeology 2006, Appendix Table 1); the 21 tombs of the 2000 campaign (Hu et al. 2021); the four tombs of the 2007 campaign (07M12, 07M19, 07M20 and 07M23; Anhui Provincial Institute of Cultural Relics and Archaeology 2008); tomb 07M22, cleared in 2016 (Lingjiatan Site Archaeological Team 2020); and 15 altar or sacrificial units—the altar itself, sacrificial pits J1–J3, four stone-piled circles, two burnt-clay deposits, a stone feature near the altar, three ash-pits of the 2007 campaign, and ash-pit 00H5 of the 2000 campaign. All documented units of the complex are included without selection: no unit meeting the feature definitions of the excavation reports was excluded, so the corpus carries no digitizer-side selection effect beyond the excavators' own recording decisions. Each unit is recorded by its position μ_i , a recorded cultural intensity o_i , a confidence weight $\alpha_i \in [0, 1]$ grading the reliability of the observation (Table 2), and an elevation z_i from the published contour and stratigraphic information (Fig. 1). Positions and intensities are referenced to a local metric grid with axes aligned to easting and northing.

Positions were digitised from the published general plan and campaign plans with an estimated plan-reading precision of ± 0.5 m. The recorded intensity o_i is, in the primary protocol, the documented artefact count of the unit: 59 units carry counts stated in the reports (the 44 catalogued tombs; five individually reported tombs of 2000, with 9 to 34 artefacts each; 07M23 with 330, 07M22 with 60, 07M12 with 1, and sacrificial pit J1 with 4; and six excavated units whose reports document no portable finds, counted as a lower bound of zero). A further 18 units carry transparent, reproducible estimates derived from the reports' own totals: the 16 tombs of 2000 not individually reported are assigned the residual season mean of 6.5 artefacts (the documented season total minus the five individually reported assemblages, distributed over the remaining tombs), and 07M19 and 07M20 are assigned 34.5 each (the documented campaign total of c. 400 artefacts minus those of 07M23 and 07M12, split evenly).

Structural units without portable artefacts—the altar, the stone-piled circles, the burnt-clay deposits and the stone feature—carry zero in the primary protocol: a documented lower bound, not a claim that these constructions were insignificant. Because that choice could be contested, Section 4.6 repeats the entire analysis under a composite protocol in which structural units instead carry area-based intensity proxies, and reports which conclusions change and which do not. Elevations z_i combine the documented opening depths of the tombs below the recorded surface with the published topography of the ridge (the cemetery slopes from c. 26 m above sea level in the north to c. 20 m in the south; the altar surface lies at c. 20 m). Confidence weights α_i follow the rubric of Table 2, applied to the reports' disturbance and recording statements: tomb 87M1, excavated intact with its three jade figures in place, grades 0.85; tomb 98M16, disturbed by a modern grave, grades 0.5; the tomb numbers, count totals, depths and disturbance statements underlying every entry are given unit-by-unit in the deposited corpus (Data availability).

Because the corpus is digitised from published plans rather than total-station records, the effect of plan-reading error on the fitted model must be assessed directly. We do this in Section 4.2 with a perturbation experiment: each coordinate is displaced by Gaussian noise of $\sigma = 0.5$ m (and the vertical coordinate by $\sigma = 0.2$ m, isolating the sensitivity of the vertical axis), the model is refitted, and the distribution of the fitted parameters over 100 replicates is reported. An honest appraisal of digitization error belongs to the methods, not to a caveat: if the fitted anisotropy were an artefact of ± 0.5 m reading error, the model would have nothing to say about the past.

Table 1
The four 3DGSM parameters and their archaeological interpretation

Parameter	Symbol	Archaeological interpretation
Position	$\mu_i \in \mathbb{R}^3$	Provenience of the cultural feature unit in three dimensions, including stratigraphic elevation
Attribute	C_i	Intensity of significant cultural material: artefact density, assemblage richness, structural investment
Confidence	$\alpha_i \in [0, 1]$	Reliability of the observation against excavation and taphonomic bias (Table 2)
Influence covariance	Σ_i	Anisotropic spatial influence of the unit: orientation, axis ratio and directional decay rates of cultural intensity

3.3 The three-dimensional Gaussian spatial model

Each cultural feature unit i is modelled as an anisotropic Gaussian influence kernel (Eq. 1), and the cultural landscape at an arbitrary point p is the superposition of all unit influences (Eq. 2).

$$I_i(p) = \alpha_i \cdot C_i \cdot \exp \left[-1/2 (p - \mu_i)^T \Sigma_i^{-1} (p - \mu_i) \right] \tag{1}$$

$$S(p) = \sum_{i=1}^N I_i(p) \tag{2}$$

The covariance Σ_i is the shared influence geometry; in the present paper a single global Σ is estimated for the complex, so that its eigenstructure summarises the directional organisation of the whole distribution. Fitting proceeds in two steps. First, the global influence geometry is estimated by α -weighted maximum likelihood: the weighted mean $\hat{\mu}$ and weighted covariance Σ of the unit positions, with weights $w_i = o_i \alpha_i$. The square roots of the eigenvalues of Σ are the axis lengths of the influence ellipsoid; the eigenvectors give their orientations. Second, prediction of the cultural landscape at unobserved locations uses the anisotropic Gaussian-kernel regression form of the model, in which the contribution of unit i to point p is weighted by the Gaussian kernel with bandwidth matrix $H = h^2 \Sigma$ and $h = n^{(-1/7)}$ in three dimensions. Formally, the predicted recorded intensity at a target unit of confidence $\alpha(p)$ is

$$\hat{s}(p) = \alpha(p) \cdot \frac{\sum_j K_H(p - \mu_j) (o_j / \alpha_j)}{\sum_j K_H(p - \mu_j)} \tag{3}$$

the kernel estimate of the reliability-corrected intensity field, rescaled by the target unit's confidence—the kernel-regression analogue of Eq. (1). The confidence weight enters both steps, so that disturbed or uncertain observations shape neither the fitted geometry nor the predicted surface at full strength.

Table 2
Confidence grading rubric for α_i

α range	Context grade	Archaeological meaning
0.9–1.0	Undisturbed, fully recorded	Sealed contexts documented at full plan and section control
0.7–0.9	Minor disturbance	Contexts with limited later intrusion but secure provenience
0.5–0.7	Partially excavated or recorded	Units at trench margins or with incomplete attribute recording
0.3–0.5	Disturbed	Contexts disturbed by later activity; position or attribution uncertain

3.4 Evaluation design

Evaluation is leave-one-out cross-validation (LOO-CV) against the held-out recorded intensities o_i : for each unit i the model is refitted on the remaining $n - 1$ units, the predicted intensity at μ_i is computed, and the error is measured against the recorded o_i —not against the model's own fitted surface, which would make the evaluation circular. All errors are computed on the max-normalised intensities ($o_i / \max o$), so that RMSE values are commensurable across methods; the primary protocol uses the documented artefact counts (Section 3.2), and the composite-proxy protocol is reported separately as a sensitivity analysis (Section 4.6), never pooled with the primary results. The 3DGSM is compared with four

conventional baselines under identical folds: (i) Ordinary kriging on the three-dimensional coordinates, with an exponential semivariogram fitted to the empirical semivariances of the recorded intensities; (ii) Isotropic 2D KDE, planar kernel regression with a circular Gaussian kernel and bandwidth $h = n^{(-1/6)}$; (iii) IDW (power 2) on the horizontal coordinates—the conventional 2.5D treatment of elevation; and (iv) Isotropic 3D Gaussian ablation, identical to the 3DGSM except that Σ is replaced by its isotropic equivalent (mean variance times identity), isolating the contribution of anisotropy. Two further ablations target specific components: (v) a No-confidence ablation setting all $\alpha_i = 1$ while retaining the full covariance, isolating the contribution of confidence weighting (Section 4.5); and (vi) an Anisotropic 2D ablation, identical to the 3DGSM but with the vertical coordinate discarded and a full 2×2 covariance retained, isolating the contribution of the third dimension (Section 4.3). Error distributions are compared by Wilcoxon signed-rank tests on paired absolute errors; 95% confidence intervals are percentile bootstrap intervals over 2000 resamples of the squared errors. As robustness checks, Section 4.3 also reports the sensitivity of the results to a $\pm 20\%$ change of the bandwidth factor and a directional-variogram cross-examination of the fitted anisotropy. Predictions are further compared with the recorded intensities at the level of functional zones (high- versus low-intensity areas, split at the median), reported as a zone-level agreement rate and a Spearman rank correlation. All computations were performed in NumPy/SciPy; the analysis code reproducing every number reported below is deposited in a public repository (see Data availability).

4. Results

4.1 Fitted influence geometry

The α -weighted maximum-likelihood estimate of the global influence covariance is moderately anisotropic. The square roots of its eigenvalues—the axis lengths of the influence ellipsoid—are 18.5 m, 7.1 m and 0.6 m: influence extends c. 2.6 times farther along the principal axis than perpendicular to it in the horizontal plane, and the horizontal-to-vertical decay ratio is c. 29, so the model reports a distribution that is elongated and stratigraphically compact. The principal axis runs at azimuth 146.8° , NNW–SSE, following the long axis of the ridge segment on which the cemetery–altar complex sits (Section 3.1). Two properties of the fitted orientation deserve emphasis. First, it is not the cemetery–to–altar bearing: the α -weighted centroids of the tomb group and the sacrificial group are 20.7 m apart at bearing 82.3° (east–west), c. 64° away from the fitted axis—the principal axis therefore expresses the internal NNW–SSE organisation of the cemetery distribution itself, including the extension towards the 2000-campaign tombs north-west of the altar, rather than the contrast between the two functional groups. Second, the axis is estimated from the distribution as a whole and was not constrained by any documented orientation; that it falls in the NNW–SSE quadrant of the ridge, rather than along the east–west group separation, is an outcome of the data, not an assumption of the model (Fig. 2).

The fitted vertical axis is short—0.6 m against 18.5 m and 7.1 m in the plane—so the model reads the complex as a set of thin, elongated stratigraphic lenses rather than as a vertically diffused scatter. This is precisely the structure that planar KDE cannot represent: projected onto the plane, stratigraphically

separated units appear superimposed, and their separation is smoothed away. The compact vertical span accords with the excavators' reading of a focused period of use of the cemetery–altar complex within the late Neolithic occupation.

4.2 Stability under digitization error

The perturbation experiment of Section 3.2 was run with 100 replicates. Under coordinate displacement of $\sigma = 0.5$ m in the horizontal plane and $\sigma = 0.2$ m vertically, the fitted principal azimuth varies by only $\pm 0.3^\circ$ (SD around the unperturbed 146.8°), the axis lengths are essentially unchanged (relative SD of the second horizontal axis 1.3% and of the vertical axis 5.6% of their means), and the LOO-CV RMSE moves only from 0.1331 to 0.1333 ± 0.0024 . The fitted geometry is therefore stable at the plan-reading precision of the published sources: it reflects the documented configuration, not digitizer noise. The vertical coordinate, the dimension most exposed to contour reading, is perturbed together with the horizontal ones in this experiment; the stratigraphic compactness reported above (horizontal-to-vertical decay ratio c. 29) is a robust feature of the data, not of the reading precision (Table 3).

Table 3
Fitted parameters under coordinate perturbation (100 replicates, $\sigma = 0.5$ m horizontal + 0.2 m vertical; mean $\pm$ SD)

Quantity	Unperturbed	Perturbed
Principal azimuth	146.8°	$146.8^\circ \pm 0.3^\circ$
Axis lengths (m)	18.5 / 7.1 / 0.6	relative SD — / 1.3% / 5.6%
Horizontal-to-vertical decay ratio	28.9	stable to within the axis SDs above
LOO-CV RMSE	0.1331	0.1333 ± 0.0024

4.3 Comparative performance under non-circular evaluation

Table 4 reports LOO-CV performance against the held-out recorded intensities, and we report it as it stands. The 3DGSM attains an RMSE of 0.1331. Its advantage over ordinary kriging (0.1403) is systematic and significant (Wilcoxon $p = 0.020$), for reasons analysed in Section 5.4. Against the remaining smoothers the picture is one of parity: isotropic 2D KDE (0.1310, $p = 0.598$), IDW (0.1341, $p = 0.604$) and the anisotropic 2D ablation (0.1316, $p = 0.897$) are statistically indistinguishable from the 3DGSM, and the isotropic 3D ablation (0.1293, $p = 0.434$) is nominally—but not significantly—more accurate. On this corpus, with its moderate fitted anisotropy (2.6) and nugget-dominated recorded counts (below), the anisotropic covariance therefore adds interpretability rather than demonstrable predictive edge; Section 4.4 uses controlled simulation to show where that edge does and does not exist. The no-confidence ablation is discussed in Section 4.5. Zone-level agreement between predicted and recorded high- versus low-intensity areas is 75.3%, and the Spearman rank correlation between predicted and recorded intensities is 0.479: the model reproduces the layout of the recorded distribution far better than its pointwise RMSE alone suggests.

Table 4
LOO-CV performance against held-out recorded intensities (primary counts protocol, $n = 85$)

Method	RMSE	p vs 3DGSM	Bootstrap 95% CI of mean squared-error difference
3DGSM (this study)	0.1331	—	—
Ordinary kriging (3D)	0.1403	0.020	[− 0.0028, 0.0078]
Isotropic 2D KDE	0.1310	0.598	[− 0.0034, 0.0026]
IDW (2.5D)	0.1341	0.604	[− 0.0026, 0.0032]
Isotropic 3D Gaussian (ablation)	0.1293	0.434	[− 0.0031, 0.0006]
Anisotropic 2D (ablation)	0.1316	0.897	[− 0.0022, 0.0013]
No-confidence (ablation)	0.1359	0.754	[− 0.0007, 0.0030]

As a cross-examination of the fitted anisotropy we examined the empirical semivariogram of the recorded intensities, globally and by direction. The fitted partial sill collapses to zero: at $n = 85$ the recorded counts are effectively nugget-dominated, so variographic anisotropy cannot be estimated at all, and the short-lag (< 30 m) semivariance is if anything higher along the fitted axis than perpendicular to it (0.023 versus 0.018). We read this honestly: the anisotropy the model reports is a property of the α -weighted spatial arrangement of the units—which is elongated NNW–SSE—not of a directional covariance of the counts themselves. The two statements are compatible: a distribution can be geometrically elongated while its intensity values are spatially uncorrelated. The fitted orientation is corroborated by its perturbation stability (SD 0.3°; Section 4.2) and by the simulation-based recovery analysis of Section 4.4. Finally, the results are not tuned to the bandwidth: a $\pm 20\%$ change of the bandwidth factor h moves the 3DGSM's RMSE only between 0.1320 and 0.1345, and the fitted orientation and axis ratio are unaffected because they derive from the covariance estimate, not from the bandwidth.

4.4 Controlled simulation: recovery of known anisotropy

Because the case study offers a single realised distribution, we verified the estimator in a controlled simulation, joining the growing practice of benchmarking archaeological quantitative methods against simulated ground truth (Cortell-Nicolau et al. 2025). The simulation has two parts, answering two different questions. Part A asks whether the α -weighted covariance estimator recovers a known geometry: point sets of $n \in \{50, 100, 200\}$ units were sampled from an anisotropic Gaussian of known in-plane axis ratio $\rho \in \{1, 2, 4, 8\}$ (azimuth 45°, compact vertical axis), weighted as in the case study, and the estimator was refitted; 20 replicates per configuration. Part B asks when that geometry translates into predictive skill: Gaussian random fields were generated with the same anisotropic covariance structures, intensities were drawn as $o = \exp(0.8 \cdot g)$ of the field values, and the 3DGSM was compared with its isotropic 3D ablation and isotropic 2D KDE under the same LOO-CV design as the case study

(Table 6). Two results matter (Table 5, Table 6, Fig. 5). First, the estimator recovers the true geometry whenever it exists: the recovered axis ratio tracks p within Monte Carlo error at all sample sizes (Table 5), and for $p \geq 2$ the mean azimuth error is at most 5.4° even at $n = 50$, falling to 0.5° at $p = 8$, $n = 200$; at $p = 1$ the principal axis is undefined and the apparent error of $40\text{--}52^\circ$ reflects uniform orientation noise, as expected. Second, the predictive advantage of the full-covariance model appears exactly where anisotropy exists and nowhere else: at $p = 1$ the three methods are on par (differences of a few percent either way, including a small cost at $n = 50$, where estimating an undefined axis adds variance); at $p \geq 2$ the 3DGSM is the most accurate method in every configuration, its advantage growing with anisotropy and sample size (RMSE 0.670 versus 0.810 and 0.858 at $p = 8$, $n = 200$), and the advantage is significant ($p < 0.001$) for all $p \geq 2$ at $n \geq 100$ and for $p \geq 4$ already at $n = 50$. The case-study corpus, with fitted ratio 2.6 at $n = 85$, sits near the lower edge of the regime in which the simulation finds clear predictive gains—which is precisely the parity outcome observed in Section 4.3. The simulation is a best case for the model in the sense that the truth is exactly Gaussian; its purpose is not to demonstrate superiority in general, but to verify that the estimator recovers the parameters it claims to estimate, and to locate the sample-size and anisotropy thresholds below which no advantage should be claimed.

Table 5
Simulation part A: recovery of a known influence geometry (20 replicates per configuration; recovered in-plane axis ratio and principal-axis azimuth error, mean $\pm$ SD)

True ratio p	$n = 50$	$n = 100$	$n = 200$
1 (axis undefined)	1.23 ± 0.15 ; 40.5°	1.15 ± 0.10 ; 52.3°	1.10 ± 0.05 ; 43.3°
2	2.00 ± 0.29 ; 5.4°	2.05 ± 0.21 ; 3.4°	2.05 ± 0.15 ; 2.4°
4	3.97 ± 0.58 ; 2.1°	4.08 ± 0.41 ; 1.4°	4.10 ± 0.30 ; 1.0°
8	7.93 ± 1.16 ; 1.0°	8.16 ± 0.82 ; 0.6°	8.20 ± 0.59 ; 0.5°

Table 6
Simulation part B: LOO-CV RMSE under controlled anisotropy (20 replicates per configuration; Wilcoxon p against the 3DGSM in parentheses)

True ratio ρ	n	3DGSM	iso-3D ablation	iso-2D KDE
1	50	1.147	1.119 (—)	1.129 (—)
1	100	1.184	1.230 (—)	1.265 (—)
1	200	0.999	1.032 (—)	1.067 (—)
2	50	1.011	1.049 (0.0 64)	1.064 (0.011)
2	100	1.119	1.182 (< 0 .001)	1.224 (< 0.001)
2	200	0.833	0.920 (< 0 .001)	0.958 (< 0.001)
4	50	1.001	1.062 (0.0 15)	1.086 (0.004)
4	100	0.889	1.013 (< 0 .001)	1.060 (< 0.001)
4	200	0.769	0.892 (< 0 .001)	0.937 (< 0.001)
8	50	0.804	0.902 (0.0 04)	0.928 (< 0.001)
8	100	0.890	1.017 (< 0 .001)	1.067 (< 0.001)
8	200	0.670	0.810 (< 0 .001)	0.858 (< 0.001)

4.5 The contribution of confidence weighting

The confidence weight is the model's formal answer to excavation and taphonomic bias, and its contribution was assessed in two ways, whose results we report without embellishment. First, the no-confidence ablation of Section 3.4: removing the confidence weights (setting all $\alpha_i = 1$ while retaining the full covariance) changes the LOO-CV RMSE from 0.1331 to 0.1359—an improvement of the full model that does not reach significance ($p = 0.754$). The effect is small because the confidence field of this corpus is essentially unstructured: high- and low-confidence units are intermingled across the complex, so the reliability correction acts as a per-unit discount with little spatial leverage, and most low- α units carry low recorded counts whose correction barely moves the max-normalised errors. Second, a perturbation experiment in which every α_i is independently displaced by ± 0.1 across 100 replicates shifts the RMSE only to 0.1336 ± 0.0016 and leaves the fitted azimuth at $146.8^\circ \pm 0.3^\circ$: whatever the confidence weighting contributes is robust to substantial errors in the grading itself. We therefore do not claim a predictive gain from confidence weighting on this corpus. What the experiment establishes is that (i) grading the observations costs nothing in accuracy, (ii) the resulting reliability surface is an interpretable output in its own right—it tells the reader where the estimated field can be trusted—and (iii) the machinery is in place for corpora in which low-confidence units do cluster spatially, where the same weighting would act as a regional mask rather than a per-unit discount.

4.6 Sensitivity to the intensity protocol

The primary protocol records intensities as documented artefact counts, with structural units set to their documented lower bound of zero (Section 3.2). Because the choice of intensity variable is the main modelling decision a user of the 3DGSM will face, the entire analysis was repeated under a composite protocol in which structural units instead carry area-based proxies (the altar c. 1200 m², the stone feature c. 20 m², the smaller constructions 2–5 m²) alongside the documented counts. Three findings matter. First, the fitted geometry is robust to the protocol: the composite fit gives axis lengths of 15.1, 7.3 and 0.5 m, an in-plane ratio of 2.06 and—critically—a principal azimuth of 148.3°, within 1.5° of the primary estimate; the NNW–SSE organisation of the complex is not an artefact of the intensity definition. Second, the predictive ranking is not robust, and we say so plainly: under the composite protocol the 3DGSM (RMSE 0.1188) is nominally the least accurate of the methods tested (ordinary kriging 0.1134, $p = 0.072$; all other comparisons $p \geq 0.12$), because the altar's dominant area proxy forms a single extreme spike that the anisotropic kernel smears across its elongated neighbourhood. Third, this divergence is itself informative: it shows that the model's behaviour is sensitive to whether the intensity variable is commensurable across units—counts of portable artefacts and construction areas are not the same quantity—and we regard the counts protocol as the defensible primary choice on exactly those grounds. Protocols should be chosen for their archaeological meaning and stated in advance; the geometry the model recovers is stable across the defensible choices, but accuracy claims must always be tied to a stated protocol.

5. Discussion

5.1 What the model says about Lingjiatan

The fitted geometry makes three specific statements about the Lingjiatan complex. First, the principal axis runs NNW–SSE (azimuth 146.8°), following the long axis of the ridge segment on which the cemetery–altar complex sits, and expressing the internal organisation of the cemetery distribution rather than the east–west contrast between the cemetery and the altar areas (Section 4.1): the spatial organisation of the complex respected the ridge, and the model quantifies that organisation rather than asserting it. Second, the in-plane axis ratio of c. 2.6 shows that cultural influence extends about two and a half times farther along the complex than across it—the quantitative signature of an elongated, bounded ritual precinct, but a moderately rather than extremely elongated one; we report the moderate value because it is what the documented units support, and the sensitivity analysis of Section 4.6 shows it is not a product of the intensity protocol. Third, the short vertical axis (0.6 m; horizontal-to-vertical decay ratio c. 29) reports deposits that are stratigraphically compact: the complex was used intensively over a shallow depth of deposit, consistent with the excavators' reading of a focused period of use.

The reading of the complex as a deliberately planned ritual core has since received independent support. Excavations in 2020–2022 south-east of the ridge exposed a large sacrificial pit (4.1 × 3.5 m, 0.32 m deep) containing more than 260 artefacts, including the largest stone *yue* battle-axe yet found at the site

and a dragon-headed jade object; an extensive burnt-clay deposit (a restored building area of c. 3400 m², and c. 5500 m² together with its platform); and an outer moat c. 45 m wide at the top, 25 m at the bottom and 2 m deep, whose spoil likely furnished the platform's fill. The 2023 season confirmed that these sacrificial features sit on a large, purpose-built platform, designed and piled as a single construction (China Cultural Relics News 2022, 2026; China Daily 2023). A ritual precinct of this scale and deliberation is exactly the kind of spatially bounded, anisotropic phenomenon the model is built to measure, and it independently justifies the treatment of recording heterogeneity as structured rather than random noise: the more monumental and integrated the construction, the more the quality of its recording varies with the history of excavation rather than with the past itself. We are explicit that the correspondence between the fitted geometry and the new exposures is in kind—an elongated, bounded, vertically engineered ritual precinct—and not in the exact axis direction, which the newly exposed area, lying outside the digitised corpus, cannot corroborate.

These statements are made with the model's own uncertainty attached. The fitted geometry is estimated under ± 0.5 m digitization error and shifts by $\pm 0.3^\circ$ in azimuth and barely at all in RMSE under that error (Section 4.2), and it is stable to within 1.5° across the two intensity protocols (Section 4.6). The 3DGSM does not replace the excavators' interpretations; it converts them into estimable quantities whose uncertainty can be reported alongside them.

5.2 The single-case limitation

The empirical demonstration rests on one site, and we are explicit about what that does and does not establish. The case study demonstrates the workflow—digitization, confidence grading, fitting, perturbation assessment, non-circular evaluation—and shows that the fitted parameters are archaeologically legible for a complex whose spatial organisation is independently understood. It cannot by itself establish that the same legibility holds for sites of different structure. The simulation study of Section 4.4 is the complement: it verifies that the estimator recovers known anisotropy across a controlled range of anisotropy ratios and sample sizes, and it locates the threshold ($n \geq 100$ advisable for predictive claims; at $n = 50$ orientation is already recovered to within 5.4° for $p \geq 2$) below which the geometry is not reliably identifiable. Applying the model to further complexes—particularly to sites whose organisation is *not* already understood, where the model's output would be genuinely predictive—is the necessary next step. Positioned against recent spatial-analytical work in this journal and allied venues—point-process models of settlement location (Zhang et al. 2025), least-cost-path reconstructions of hunter-gatherer mobility (Rafart-Vidal et al. 2025), machine-learning predictive modelling for survey design (Wang et al. 2025) and quantitative readings of spatial organisation in Palaeolithic cave art (Intxaurbe et al. 2024) and spatial analysis of original excavation data (Sánchez-Romero and Benito-Calvo 2026)—the 3DGSM asks a complementary question: not where activity concentrates, but in which direction, and at what rate, cultural intensity decays.

5.3 The confidence weight in practice

The α -grading rubric of Table 2 formalises judgements excavators already make informally. In the present corpus its predictive contribution is small and non-significant (Section 4.5), and we present it as what it is: an honesty mechanism and an interpretable output rather than a guaranteed accuracy gain. Three practical points follow. First, the grading need not be fine: the ± 0.1 perturbation experiment shows the results are robust to substantial grading error, so a four-level rubric applied from the excavation report is sufficient, and mis-grading individual units cannot drive the conclusions. Second, the α -field can be mapped to show where the model's output is trustworthy. In the present case the confidence field is essentially unstructured—high- and low-confidence units are intermingled across the complex—so confidence weighting acts as a per-unit reliability discount rather than a regional mask; where low-confidence units do cluster, for instance on partially excavated margins, the same machinery would instead flag whole regions whose output should be read with care. Third, including confidence information costs nothing: the full model is never significantly worse than its no-confidence ablation in any experiment we ran, so the interpretive benefit of the reliability surface comes without a measurable accuracy penalty.

5.4 An honest comparison with kriging

The comparison with ordinary kriging requires honesty about what drives it. The recorded intensities are nugget-dominated to an extreme degree: the fitted partial sill of the exponential semivariogram collapses to zero, so per-unit recording heterogeneity, not spatial correlation, dominates the raw counts (Section 4.3). Ordinary kriging built on such a semivariogram has almost no covariance structure to exploit, and its leave-one-out predictions collapse towards the grand mean (Fig. 3), discarding the layout of the distribution; the 3DGSM instead smooths along an influence geometry estimated from the weighted arrangement of the units, and retains a graded response (RMSE 0.1403 versus 0.1331, $p = 0.020$). We therefore read the kriging comparison as evidence that a model whose geometry is estimated from the point arrangement can outperform a model whose geometry must be estimated from the noisy intensities themselves—not as evidence against kriging, which remains the appropriate default wherever the intensity field carries genuine spatial covariance, as it does in the simulated fields of Section 4.4.

It is equally honest to note where the model's advantage comes from in general and where it does not exist: the anisotropic covariance adapts the smoothing geometry to the data, and the confidence weighting discounts unreliable observations. On nearly isotropic data neither mechanism can help, and the model should be expected to perform on a par with conventional methods. The simulation study shows exactly this behaviour at $p = 1$, where the three methods are within a few percent of one another (Section 4.4), and the case study shows it directly: against isotropic KDE, IDW and the isotropic 3D ablation, the 3DGSM's differences are small and non-significant (Section 4.3). The 3DGSM's contribution is not a universal accuracy gain; it is that directionality, vertical structure and observational reliability become estimable, interpretable quantities, with the accuracy gain appearing where structure exists to be measured.

5.5 Applicability and limitations

Three limitations bound the present claims. First, the case-study corpus is digitised from published plans at ± 0.5 m precision, and although its intensities are documented artefact counts for 59 units, 18 units carry estimates derived transparently from the reports' own totals, and structural units enter the primary protocol at their documented lower bound of zero (Section 3.2); the sensitivity analysis of Section 4.6 shows the fitted geometry is stable across protocols but the predictive ranking is not, so the accuracy figures should be read as belonging to the stated counts protocol. Application to total-station provenience data, where positions need no plan digitization and counts need no estimation, is the natural next step and will exercise the model on data whose recording bias is entirely the excavators' own. Second, a single global influence geometry is assumed; complexes with multiple episodes or modes of organisation may require mixture or locally fitted geometries, at the cost of additional parameters. Third, the model inherits the assumptions of its kernel form: Gaussian influence is a smooth, unimodal idealisation, and sharply bounded or linear features (walls, rows) are approximated rather than represented exactly. None of these limitations affects the paper's central claim, which is methodological: that anisotropy, vertical structure and confidence can be estimated as explicit, interpretable parameters within a classical spatial-statistical workflow.

6. Conclusions

The 3DGSM treats an archaeological site as a set of weighted anisotropic Gaussian influences and estimates, rather than assumes, the geometry of the cultural-intensity distribution. Applied to 85 documented units of the Lingjiatan cemetery–altar complex—positions digitised from the published plans, intensities from the artefact counts of the excavation reports—the fitted geometry is moderately anisotropic (in-plane axis ratio c. 2.6, azimuth 146.8° NNW–SSE, following the ridge the complex occupies) and stratigraphically compact (vertical axis 0.6 m; horizontal-to-vertical decay ratio c. 29). Under a deliberately non-circular evaluation the model significantly outperforms ordinary kriging ($p = 0.020$), whose variogram sees only nugget in the recorded counts, and reproduces the layout of the recorded distribution (zone-level agreement 75.3%; Spearman 0.479); against isotropic KDE, IDW and its own ablations its accuracy differences are small and non-significant, and we report them as such. Confidence weighting costs nothing and yields an interpretable reliability surface, robust to substantial grading error; the fitted geometry is stable under publication-scale digitization error (azimuth $\pm 0.3^\circ$) and across intensity protocols ($\pm 1.5^\circ$); and a two-part controlled simulation verifies that the estimator recovers known anisotropy whenever it is present (azimuth error below 5.4° for $p \geq 2$, below 0.5° at $p = 8$, $n = 200$) and that predictive advantages appear exactly where anisotropy exists—without inventing gains where none exist. Treating directionality and observational uncertainty as estimable quantities, the 3DGSM offers an interpretable complement to the KDE-based workflows familiar to most archaeologists, and a direct route from point-pattern description to the measurement of spatial organisation.

Declarations

Funding. [Funder name and grant number to be completed by the authors.]

Conflicts of interest. The authors declare no conflicts of interest.

Author contributions. [CRediT roles to be completed by the authors, e.g. Conceptualization: ...; Methodology: ...; Software: ...; Formal analysis: ...; Writing – original draft: ...; Writing – review and editing:]

Data availability. The digitised 85-unit evaluation corpus (with per-unit documentation of counts, depths, disturbance grades and their sources) and the analysis code reproducing every reported number are deposited in a Zenodo repository (<https://doi.org/10.5281/zenodo.21740527>). The excavation data underlying the published plans belong to the Anhui Provincial Institute of Cultural Relics and Archaeology.

References

1. Anhui Provincial Institute of Cultural Relics and Archaeology (1989) Excavation report of the Neolithic cemetery at Lingjiatan, Hanshan, Anhui. *Wenwu (Cultural Relics)* 1989(4):1–9 (in Chinese)
2. Anhui Provincial Institute of Cultural Relics and Archaeology (2006) Lingjiatan: excavation report of the Neolithic cemetery and altar site. Cultural Relics Publishing House, Beijing. (in Chinese)
3. Anhui Provincial Institute of Cultural Relics and Archaeology (2008) New discoveries in the fifth excavation at the Lingjiatan site in Hanshan County, Anhui. *Kaogu (Archaeology)* 2008(3):7–17 (in Chinese)
4. Anhui Provincial Institute of Cultural Relics and Archaeology, Hanshan County Office for the Management of Cultural Relics (1999) Report on the third excavation at the Lingjiatan site in Hanshan County, Anhui. *Kaogu (Archaeology)* 1999(11):1–12 (in Chinese)
5. Anselin L (1995) Local indicators of spatial association—LISA. *Geographical Anal* 27(2):93–115
6. Baxter MJ, Beardah CC, Wright RVS (1997) Some archaeological applications of kernel density estimates. *J Archaeol Sci* 24(4):347–354
7. Beardah CC, Baxter MJ (1996) The archaeological use of kernel density estimates. *Internet Archaeol* 1. <https://doi.org/10.11141/ia.1.1>
8. China Cultural Relics News (2022) New archaeological discoveries at the Lingjiatan site, Anhui. China Cult Relics News, 8 December 2022 (in Chinese)
9. China Cultural Relics News (2026) Major advances in the 2020–2023 excavations at Lingjiatan: the large ritual platform and outer moat. China Cult Relics News, 19 May 2026 (in Chinese)
10. China Daily (2023) New light shed on beginnings of Chinese civilization. China Dly, 6 January 2023
11. Conolly J, Lake M (2006) Geographical information systems in archaeology. Cambridge University Press, Cambridge
12. Cortell-Nicolau A, Carrignon S, Rodríguez-Palomo I, Hromada D, Kahlenberg R, Mes A et al (2025) Assessing quantitative methods in archaeology via simulated datasets: the Archaeoriddle

- challenge. Concept, project and motivations. *J Archaeol Sci* 177:106179.
<https://doi.org/10.1016/j.jas.2025.106179>
13. Crema ER (2025) Statistical modelling in archaeology: some recent trends and future perspectives. *J Archaeol Sci* 180:106295. <https://doi.org/10.1016/j.jas.2025.106295>
 14. Cressie NAC (1993) *Statistics for spatial data*, revised edn. Wiley, New York
 15. Hodder I, Orton C (1976) *Spatial analysis in archaeology*. Cambridge University Press, Cambridge
 16. Hu P, Yao J, Zhang X, Wu W (2021) Report on the fourth excavation at the Lingjiatan site in Hanshan County, Anhui. *Southeast Cult* 2021(5):35–47 (in Chinese)
 17. Intxaurbe I, Garate D, Arriolabengoa M (2024) Drawing in the depths: spatial organization patterns related to Magdalenian cave art. *Archaeol Anthropol Sci* 16:104. <https://doi.org/10.1007/s12520-024-02007-3>
 18. Kempf M (2020) Modeling multivariate landscape affordances and functional ecosystem connectivity in landscape archeology. *Archaeol Anthropol Sci* 12:159.
<https://doi.org/10.1007/s12520-020-01127-w>
 19. Kerbl B, Kopanas G, Leimkühler T, Drettakis G (2023) 3D Gaussian splatting for real-time radiance field rendering. *ACM Trans Graphics* 42(4):139:1–139
 20. Khedr M, Metawie M, Eldin AS, Hamdy A, El-Shihy A, Azab S, Abdelaty A, Marzouk M (2025) 3D visualization of damaged statues using Gaussian splatting and web interface integration. *npj Herit Sci* 13:587. <https://doi.org/10.1038/s40494-025-02063-5>
 21. Liu C, Valletta F, Baena Preysler J, Falcucci A (2026) Reconstructing the sequentiality of adjacent flake removal scars on lithic 3D models: a curvature-based computational approach. *Archaeol Anthropol Sci* 18:150. <https://doi.org/10.1007/s12520-026-02510-9>
 22. Lingjiatan Site Archaeological Team (2020) The clearing of a Neolithic tomb at the Lingjiatan site in Hanshan County, Anhui. *Kaogu (Archaeology)* 2020(11):16–28 (in Chinese)
 23. Liu L, Chen X (2012) *The archaeology of China: from the late Paleolithic to the early Bronze Age*. Cambridge University Press, Cambridge
 24. Lu T, Yu M, Xu L, Xiangli Y, Wang L, Lin D, Dai B (2024) Scaffold-GS: structured 3D Gaussians for view-adaptive rendering. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp 20654–20664
 25. Marwick B (2025) Is archaeology a science? Insights and imperatives from 10,000 articles and a year of reproducibility reviews. *J Archaeol Sci* 180:106281.
<https://doi.org/10.1016/j.jas.2025.106281>
 26. Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R (2021) NeRF: representing scenes as neural radiance fields for view synthesis. *Commun ACM* 65(1):99–106
 27. Mirocha Ł (2025) Beyond digital twins: 3D Gaussian splatting, game engines and crossmedia cultural heritage representations. In: *Proceedings of SIGGRAPH 2025 (Conference Talks)*.
<https://doi.org/10.1145/3721239.3734094>

28. Moclán A, Cobo-Sánchez L, Domínguez-Rodrigo M, Méndez-Quintas E, Rubio-Jara S, Panera J, Pérez-González A, Santonja M (2023) Spatial analysis of an Early Middle Palaeolithic kill/butchering site: the case of the Cuesta de la Bajada (Teruel, Spain). *Archaeol Anthropol Sci* 15:91. <https://doi.org/10.1007/s12520-023-01792-7>
29. Rafart-Vidal E, Rubio-Campillo X, Mangado X, González-Olivares C, Sánchez de la Torre M (2025) Spatial analysis of long-term mobility dynamics among hunter-gatherer groups in mountain areas: the case of Montlleó (Eastern Pyrenees). *Archaeol Anthropol Sci* 17:245. <https://doi.org/10.1007/s12520-025-02351-y>
30. Ripley BD (1977) Modelling spatial patterns. *J Roy Stat Soc: Ser B (Methodol)* 39(2):172–212
31. Sánchez-Romero L, Benito-Calvo A (2026) Inferring the role of hominids at Torralba: spatial analysis of original excavation data. *Archaeol Anthropol Sci* 18:157. <https://doi.org/10.1007/s12520-026-02517-2>
32. Silverman BW (1986) Density estimation for statistics and data analysis. Chapman and Hall, London
33. VanPool TL, Leonard RD (2011) Quantitative analysis in archaeology. Wiley-Blackwell, Chichester
34. Wang CX, Zhang X, Zhang Y, Wu Y, Huang C, Fan A (2022) Luminescence dating of heated quartz extracted from burnt clay and pottery excavated from the Lingjiatan archaeological site, China. *Front Earth Sci* 10:933342. <https://doi.org/10.3389/feart.2022.933342>
35. Wang J, Soroush M, Maziar S, Khazraee E (2025) Predictive modeling for targeted archaeological survey of Arsacid period sites in the Iranian borderland region of the Araxes River valley. *J Archaeol Method Theory* 33(1):3. <https://doi.org/10.1007/s10816-025-09733-5>
36. Wheatley D, Gillings M (2002) Spatial technology and archaeology: the archaeological applications of GIS. Taylor & Francis, London
37. Whallon R (1984) Unconstrained clustering for the analysis of spatial distributions in archaeology. In: Hietala HJ (ed) *Intrasite spatial analysis in archaeology*. Cambridge University Press, Cambridge, pp 242–253
38. Yu Y, Verbree E, van Oosterom P, Pottgiesser U (2025) 3D Gaussian splatting for modern architectural heritage: integrating UAV-based data acquisition and advanced photorealistic 3D techniques. *AGILE: GIScience Ser* 6:51. <https://doi.org/10.5194/agile-giss-6-51-2025>
39. Yu Z, Chen A, Huang B, Sattler T, Geiger A (2024) Mip-Splatting: alias-free 3D Gaussian splatting. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp 19447–19456
40. Zhang J (1999) Main results of the third excavation at the Lingjiatan site in Hanshan. *Southeast Cult* 1999(5):14–17 (in Chinese)
41. Zhang X, Zhang Y, Tan L, Crema ER, Tian Y, Wang Z (2025) A comparative analysis of stone- and earth-wall settlement locations of the Lower Xiajiadian Culture in Aohan Banner, China. *Archaeol Anthropol Sci* 17:18. <https://doi.org/10.1007/s12520-024-02125-y>

Figures




Figure 1

The Lingjiatan evaluation corpus ($n = 85$). (a) Plan view in the local metric grid: circles are tombs ($n = 70$), triangles altar or sacrificial units ($n = 15$); point size is proportional to the confidence weight α and colour encodes elevation z (m above sea level). The compact northern group is the 2000-campaign trench area north-west of the altar, on the highest part of the ridge; the main southern mass combines the 1987/1998 cemetery, the altar area and the 2007 units. (b) East–west vertical section; colour encodes the recorded intensity (max-normalised artefact count). The deposits of the main mass occupy a compact vertical band of c. 0.8 m (19.2–20.0 m ASL), while the northern group sits c. 4–5 m higher on the ridge (23.7–24.8 m ASL).




Figure 2

Fitted influence geometry. Circles: tombs; triangles: altar or sacrificial units; colour: recorded intensity (max-normalised artefact count); point size proportional to $\alpha \hat{\mu}$. The ellipse is the 1σ horizontal section of the influence ellipsoid (18.5×7.1 m; principal axis at azimuth 146.8°); the dashed line is the principal eigenvector; the blue cross is the α -o-weighted mean $\hat{\mu}$. For comparison, the dotted line joins the α -o-weighted centroids of the tomb and sacrificial groups (grey squares; bearing 82.3°): the fitted axis expresses the internal organisation of the distribution, not the inter-group separation.




Figure 3

LOO-CV predictions against recorded intensities for the 3DGSM (left) and ordinary kriging (right); the dashed line is the identity. Kriging, fitted to an effectively nugget-only semivariogram, collapses towards the grand mean and loses the layout of the distribution; the 3DGSM preserves a graded response. Both methods underestimate the single extreme unit (tomb 07M23, recorded $o = 1.0$), which dominates the upper tail of the error distribution of every method in Table 4.




Figure 4

LOO-CV RMSE of the 3DGSM, the conventional baselines and the ablations (primary counts protocol, $n = 85$), with Wilcoxon signed-rank p-values against the 3DGSM.




Figure 5

Simulation study. (a) Estimator recovery: recovered in-plane axis ratio against the true ratio ρ (log-log; dashed line: identity). The estimator is unbiased for $\rho \geq 2$ at all sample sizes; at $\rho = 1$ the recovered ratio stays near 1 and the azimuth is undefined, as it should be. (b) Predictive skill at $n = 200$: the 3DGSM matches the isotropic alternatives at $\rho = 1$ and outperforms them by a margin that grows with true anisotropy.