Exploring Multivariate Geospatial Data Through Interactive Feature-Space Mapping and Visualization

Jürgen Döllner[1,2], Josafat-Mattias Burmeister[1]

[1] University of Potsdam, Digital Engineering Faculty, Potsdam, Germany [2] Hasso Plattner Institute, Potsdam, Germany
(doellner, burmeister)@uni-potsdam.de

Keywords: Geovisualization, Geospatial Data Analysis, Visual Analytics, Multidimensional Visualization, Interactive Visualization, Spatial Analysis.

Abstract

Geospatial objects are embedded in geographic space and simultaneously described by multiple thematic attributes. Conventional maps preserve spatial context but offer limited support for exploring thematic relations and multivariate similarity structures beyond geographic proximity. We present an interactive feature-space visualization approach that maps each geospatial object to a multidimensional feature vector and represents it as a particle in a two-dimensional reference plane. Particle positions are determined by thematic, analytical, and spatial magnets whose attraction strength depends on selected feature values or spatial relations. Starting from barycentric target positions, particles are iteratively arranged with collision-aware placement, revealing thematic neighborhoods, outliers, and latent groups independently of the original geographic layout. By configuring magnets interactively, analysts can formulate and test hypotheses about geospatial datasets. We demonstrate the approach using an urban tree inventory derived from mobile mapping LiDAR data and show how it supports the analysis of structural relations, isolated trees, similarity patterns, and terrain-dependent distributions. The results indicate that interactive feature-space visualization provides a useful complementary perspective for exploratory geovisualization and visual analytics.

1. Introduction

In geospatial datasets such as forest inventories, city models, or infrastructure databases, geospatial objects form the primary entities of analysis. They are characterized not only by their location but also by multiple thematic attributes, often derived from heterogeneous sources and attached to the objects. In the following, we refer to them as *multivariate geospatial objects*, i.e., objects with spatial coordinates and a multidimensional attribute vector. From a mathematical perspective, they can be interpreted as *feature vectors* in a multidimensional feature space.

Maps are fundamental for visualizing geospatial data (Kraak and Ormeling, 2020). A map is a 'symbolised representation of geographical reality, representing selected features or characteristics [...] and is designed for use when spatial relationships are of primary relevance' (ICA Executive Committee, 2004). In particular, thematic maps encode multivariate geospatial objects using visual variables such as color, size, or glyph shape. However, cartographic design and symbolization offer only a limited number of effective visual channels, which are quickly exhausted when several attributes must be considered simultaneously. As Bertin argues, once a cartographic representation must express geographic position together with multiple additional components, 'the totality of relationships cannot be perceived instantaneously' (Bertin, 1983, p. 407). Two objects may be spatially close but thematically dissimilar, while thematically similar objects may be distributed across distant regions. As a result, thematic relations, clusters, outliers, and other latent structures remain visually subordinated in map space.

In our approach, we transfer feature-space visualization to the geospatial domain by representing geospatial objects as particles derived from their feature vectors. We deliberately detach objects from their geographic arrangement and rearrange them according to thematic or analytical criteria. The approach builds on the *Dust & Magnet* metaphor (Yi et al., 2005) and *star-coordinate* principles (Kandogan, 2000). Each feature vector is represented as a particle in a reference plane. Particle positions result from the influence of interactively placed magnets that are associated with feature dimensions or with analytical criteria derived from feature vectors (Fig. 1).

Geospatial analysis often requires relating thematic patterns back to spatial structure. To this end, we extend the particle–magnet concept by introducing *spatial magnets* that derive their weights from spatial proximity or distance-dependent relations. They allow thematic and spatial criteria to be combined within the same interaction model, enabling analysts to explore not only which objects are thematically similar, but also how thematic similarity interacts with geographic proximity, regional concentration, or spatial separation.

To ensure readability and scalability for large datasets, we use collision detection and splat-based rendering. Magnet and particle configurations can be modified in real time, as 'visualization and interaction complement each other to support human analytical thinking' (Andrienko et al., 2020). This supports the identification of thematic neighborhoods and outliers as well as hypothesis-driven exploration through magnet placement and manipulation. Overall, our approach expands the methodological spectrum of exploratory geovisual analytics.

2. Related Work

Multivariate datasets represent each data item by multiple feature dimensions, often arising from heterogeneous sources (He et al., 2019), and are typically represented in tabular form. Example domains include geospatial information, financial indicators (Vaz de Melo Mendes and Pereira Câmara Leal, 2005), and socioeconomic statistics (Lengyel et al., 2022). Visualization techniques such as parallel coordinate plots (Heinrich

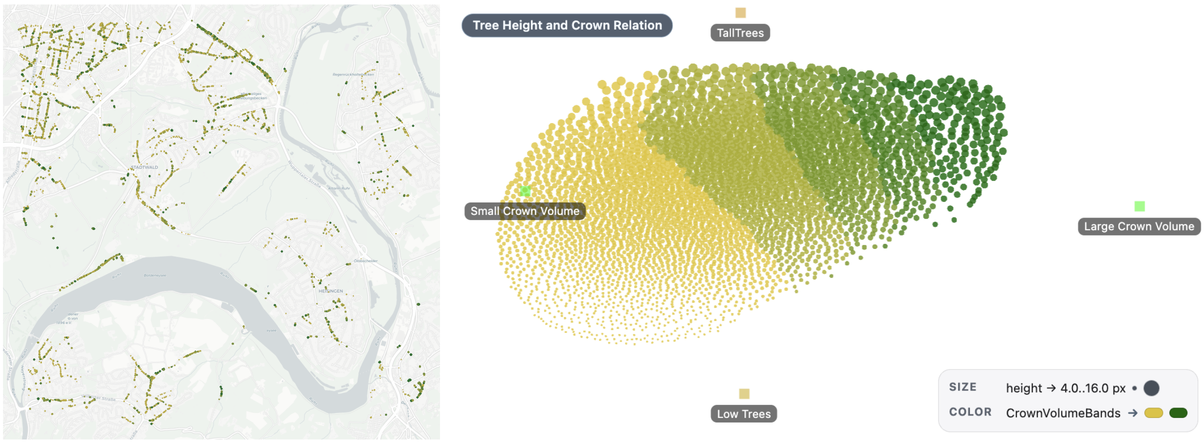

Figure 1. Dataset of 3563 urban trees shown both in geographic space and in a feature-space visualization. We use two single-feature magnet pairs: tree height is represented along the vertical axis (*Low Trees–Tall Trees*), and crown volume along the horizontal axis (*Small Crown Volume–Large Crown Volume*). Tree height is simultaneously encoded by particle size, and crown volume bands are represented by particle color (yellow: small; green: large). The particle cloud forms a broad diagonal gradient from small-crowned trees on the left to large-crowned trees on the right, while height increases from bottom to top. This pattern indicates a positive overall relation between tree height and crown volume, but not a strict linear dependency: medium and taller trees still exhibit substantial variation in crown volume. Extremely tall trees remain sparse, and the distribution is densest in the central region, suggesting that most trees belong to intermediate height and crown-volume ranges.

and Weiskopf, 2013), heatmaps (Trye et al., 2023), scatterplot matrices (Carr et al., 1987), and glyph-based methods (Ropinski et al., 2011) aim to reveal clusters, correlations, and outliers. With increasing dimensionality, these approaches face scalability and interpretability challenges, including factorial axis permutations and visual clutter in parallel coordinates.

Dimensionality reduction addresses this by projecting high-dimensional feature spaces into lower-dimensional embeddings (Espadoto et al., 2021). Linear techniques such as PCA (Hotelling, 1936; Ivosev et al., 2008) capture dominant variance directions, while nonlinear methods such as t-SNE (van der Maaten and Hinton, 2008) and UMAP (McInnes et al., 2020) preserve local or global similarity structures. However, these techniques visualize transformed projections rather than the original feature space, limiting direct interaction with individual feature dimensions and composite analytical criteria.

General Line Coordinates (GLCs) explicitly represent feature dimensions as axes and items as polylines (Heinrich and Weiskopf, 2013). Variants include conditional coordinates (Weidele, 2019) and non-parallel layouts (Antonini et al., 2023). Radar charts arrange axes radially to emphasize shape-based patterns (Saary, 2008). Recent work explores transitions between coordinate-based and scatterplot representations (Kiesel et al., 2023). While preserving explicit feature axes, these methods remain constrained in scalability and compositional expressiveness.

Interactive Spatialization of Feature Spaces. Interactive projection techniques allow users to steer mappings from high-dimensional data to visual space (Sacha et al., 2017). User-guided dimensionality reduction supports direct manipulation of projections, e.g., by adjusting cluster centroids or parameters (Fujiwara et al., 2022). Interactive scatterplot matrices facilitate navigation of pairwise feature relationships (Elmqvist et al., 2008). However, these approaches mainly encode linear combinations and offer limited support for higher-level analytical or statistical criteria within the spatial model.

Star coordinates represent feature dimensions as radial axes and spatialize multivariate objects as points using weighted difference vectors (Kandogan, 2000); repositioning axes reshapes the projection and reveals relationships between dimensions. In contrast, our method is not restricted to radial axes, but uses magnets as analytical operators for single-feature, inverse, composite, similarity-based, and spatially defined criteria.

Feature Space Exploration. The *Embedding Projector* supports interactive inspection of learned embeddings through projection, nearest-neighbor exploration, and metadata linking (Smilkov et al., 2016). *Latent Space Cartography* introduces multiscale navigation and structure-aware exploration of embeddings (Liu et al., 2019). These approaches align with explainable and human-in-the-loop machine learning, yet remain largely projection-centric and provide limited means for specifying compositional analytical criteria directly within feature space.

Visual Analytics and Analytical Feedback. Visual analytics integrates interactive visualization with analytical reasoning (Sacha et al., 2017). Analysts iteratively explore data, refine criteria, and steer analysis processes (Andrienko et al., 2020), particularly in high-dimensional and model-derived feature spaces where interpretability is essential.

Dust & Magnet Visualization. Dust & Magnet (D&M) visualization (Yi et al., 2005) spatializes multivariate data using a physics-inspired metaphor. By repositioning magnets, users interactively explore relationships as particles rearrange according to attraction strengths. It can be extended by a vector-based attraction model combining D&M with star-coordinate principles and 3D rendering (Vollmer and Döllner, 2020), but existing formulations remain primarily attribute-centric and do not address geospatial objects and spatial magnets.

Gap in research. In current geovisualization approaches, thematic attributes are rarely treated as first-class analytical entities. Instead, they are typically encoded as secondary properties of geographically fixed objects, which leaves thematic similarity structure and multivariate relations visually subordinate to map layout. This paper addresses the resulting gap by introducing a feature-space visualization in which thematic attributes directly structure the layout while analytical findings remain linked to geographic space.

3. Feature-Space Mapping and Visualization

The overall process consists of the following steps:

1. Map geospatial objects to feature vectors derived from selected thematic attributes.
2. Visualize feature vectors as particles in the reference plane.
3. Determine particle weights for active magnets.
4. Compute particle positions around barycentric target positions with collision avoidance.
5. Explore the resulting layout through interactive manipulation of magnet positions, strengths, and configurations.

3.1 Feature Vectors

Let $\mathcal{G} = \{g_1, \ldots, g_n\}$ denote the set of geospatial objects in a dataset. Each geospatial object $g \in \mathcal{G}$ is represented by a d-dimensional feature vector

$$\mathbf{x}(g) = (x_1(g), \ldots, x_d(g)) \in \mathbb{R}^d, \tag{1}$$

whose components correspond to selected thematic attributes, such as structural, semantic, temporal, geometric, or contextual properties. The feature vector thus determines the position of a geospatial object in thematic space rather than in geographic space. Because the magnet-based layout requires comparable weights across heterogeneous feature dimensions, feature values are normalized to the interval $[0,1]$. For each feature dimension $k \in \{1, \ldots, d\}$, let

$$x_k^{\min} = \min_{g \in \mathcal{G}} x_k(g), \qquad x_k^{\max} = \max_{g \in \mathcal{G}} x_k(g). \tag{2}$$

The normalized feature value can then be defined, in a min–max normalization, as

$$\tilde{x}_k(g) = \begin{cases} \dfrac{x_k(g) - x_k^{\min}}{x_k^{\max} - x_k^{\min}}, & \text{if } x_k^{\max} \neq x_k^{\min}, \\ 0, & \text{if } x_k^{\max} = x_k^{\min}. \end{cases} \tag{3}$$

We denote the normalized feature vector by

$$\tilde{\mathbf{x}}(g) = (\tilde{x}_1(g), \ldots, \tilde{x}_d(g)). \tag{4}$$

3.2 Feature Space

By normalizing each feature dimension to the interval $[0,1]$, the feature space can be defined as the **d-dimensional unit hypercube** $\mathcal{F} = [0,1]^d$. Each object is mapped to a point

$$\tilde{\mathbf{x}}(g) \in \mathcal{F}, \text{ with } \tilde{\mathbf{x}}(g) = (\tilde{x}_1(g), \ldots, \tilde{x}_d(g)), \tag{5}$$

where each coordinate expresses the relative position in the corresponding original value domain. The unit hypercube provides a compact and uniform representation for the subsequent layout process. Magnets can then operate on individual feature dimensions or on combinations thereof, independent of the original data type or measurement scale of the underlying attributes.

3.3 Magnets

A magnet represents an analytical criterion that influences the placement of particles in the reference plane. Let $\mathcal{M} = \{M_1, \ldots, M_m\}$ be the set of active magnets, and let $\mathbf{s}_j \in \mathbb{R}^2$ denote the position of magnet M_j in the reference plane. Each magnet defines a weight function

$$w_j : \mathcal{G} \to [0,1] \tag{6}$$

that determines the unsigned strength with which a geospatial object responds to the magnet. In the simplest case, a magnet is associated with a single feature dimension k, yielding

$$w_j(g) = \tilde{x}_k(g). \tag{7}$$

More generally, magnets may represent analytical combinations of several feature dimensions or spatial criteria, e.g., conjunctive, disjunctive, or proximity-based relations. In this way, magnets define the current interpretation of the feature space and can be adapted interactively to different exploratory aspects.

3.4 Selecting Thematic Attributes

The choice of thematic attributes determines which analytical structures become visible in feature space. *Core thematic attributes* describe essential properties of geospatial objects and are directly relevant to the analysis task. Examples include tree height, crown volume, stem diameter, inclination, and site conditions in a forest inventory, or functional type, condition, and construction period in an infrastructure dataset. The overall aim is a compact and semantically meaningful feature space in which each dimension adds distinct analytical value.

Feature dimensions should represent primary thematic descriptors rather than derived attributes, which often restate existing information and introduce redundancy. Such relations are often better explored through analytical magnets, for example by applying minimum, maximum, or sink magnets. Feature dimensions should contribute complementary information and should therefore not be tightly correlated. Mapping two highly correlated attributes as separate dimensions overemphasizes one thematic aspect and reduces the discriminative power of the representation. Conjunctive, disjunctive, comparative, or variability-based relations can be modeled more explicitly through analytical magnets.

3.5 Transforming Attributes into Feature Dimensions

Thematic attributes must be transformed into feature dimensions in the d-dimensional unit cube. Since geospatial datasets often contain heterogeneous attribute types, this transformation depends on the measurement level of the attribute.

Numerical Attributes. If an attribute is numerical, its values can be mapped directly to a normalized feature dimension as defined in Eq. (3). In practice, direct min–max normalization may be overly sensitive to extreme outliers. To obtain more robust feature dimensions, the normalization range can be restricted to inner percentiles.

Let p_k^- and p_k^+ denote lower and upper percentile bounds, for example the 2.5th and 97.5th percentiles. Then values below p_k^- are clamped to 0, values above p_k^+ are clamped to 1, and intermediate values are linearly mapped to $[0,1]$:

$$\tilde{x}_k(g) = \begin{cases} 0, & x_k(g) \leq p_k^-, \\ \dfrac{x_k(g) - p_k^-}{p_k^+ - p_k^-}, & p_k^- < x_k(g) < p_k^+, \\ 1, & x_k(g) \geq p_k^+. \end{cases} \tag{8}$$

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-89-2026 | © Author(s) 2026. CC BY 4.0 License.

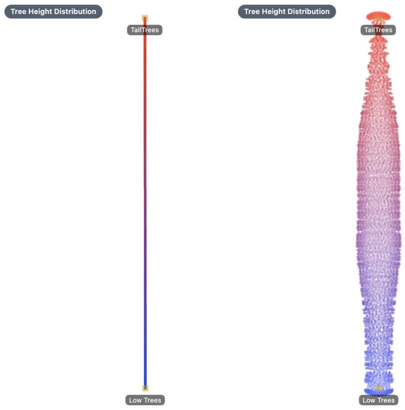

Figure 2. Opposing single-feature magnets, spanning an axis along which trees are sorted by height. Left: Barycentric placement, leading to overplotting. Right: Iterative placement with collision avoidance.

This reduces the dominance of extreme values and yields more stable attraction behavior in the layout.

Ordinal Attributes. Ordinal attributes possess an inherent order but not necessarily a metric interpretation. Such attributes can be mapped to ranks and then normalized to $[0,1]$. This preserves order relations while making the attribute compatible with the feature-space model.

Categorical Attributes. Nominal or categorical attributes must first be transformed into a numerical representation. A direct integer encoding of categories is generally problematic because it introduces artificial order and distance relations. Therefore, categorical attributes should only be mapped to a single feature dimension if a meaningful ordering exists or if the application admits a domain-specific scalar encoding. Otherwise, they are better handled through category-specific filters, similarity relations, or binary indicator features. If a categorical attribute is mapped numerically, the transformation must be accompanied by an inverse mapping that allows the original semantic value, such as a species name or land-use class, to be recovered from the encoded representation. More generally, each transformed feature dimension should maintain a traceable link to the original thematic attribute domain so that analytical findings in feature space remain interpretable in domain terms.

3.6 Inverse Mapping and Semantic Interpretability

Normalization is necessary for the proposed magnet-based layout because heterogeneous attributes must be expressed on a common scale. However, feature-space exploration remains analytically useful only if normalized feature values can be related back to the original data semantics. For interpretation, each feature transformation is accompanied by an inverse or reconstruction function

$$\psi_k : [0,1] \to \mathcal{D}_k, \tag{9}$$

where $\mathcal{D}_k$ denotes the original value domain of attribute k. For numerical attributes, ψ_k maps normalized values back to the original scale, possibly within the clamped percentile range. For ordinal or categorical attributes, ψ_k maps the normalized representation to the corresponding rank or category. ψ_k can be used to expose the original values in the user interface, for example through axis labels or tooltips.

4. Particle Mapping

Particles are positioned according to the magnet-induced structure of the feature space such that spatial proximity in the reference plane reflects analytical similarity with respect to the currently active exploration criteria. Particle size and color may encode additional thematic cues, but the primary analytical signal is conveyed by the spatial arrangement itself. As magnets are moved or reconfigured, particles reorganize accordingly, making changes in analytical criteria visible as spatial motion. Particle mapping therefore provides both a static arrangement of geospatial objects in feature space and a dynamic process through which thematic structure, group formation, and outliers become perceptible.

4.1 Barycentric Target Position

For a given magnet configuration, each geospatial object has a target position in the reference plane that summarizes the influences of all active magnets. We compute this position as the weighted barycenter of the magnets:

$$\mathbf{q}^*(g) = \frac{\sum_{j=1}^{m} w_j(g)\,\mathbf{s}_j}{\sum_{j=1}^{m} w_j(g) + \varepsilon}, \tag{10}$$

where $\varepsilon > 0$ is a small constant to avoid division by zero. As a convex combination, the barycentric target remains within the convex hull of the attracting magnets. This target position provides a compact geometric summary of the relation of a geospatial object to the active magnets: objects with similar feature vectors tend to receive nearby barycentric positions, whereas objects that respond differently to the magnet configuration are separated in the plane.

4.2 Iterative Particle Placement

Placing all particles exactly at their barycentric targets would often lead to excessive overlap in dense regions (Fig. 2). To support readable layouts, particle positions are therefore computed by an iterative placement process that combines attraction toward the target position with local collision avoidance.

Let $\mathbf{q}_t(g) \in \mathbb{R}^2$ denote the reference-plane position of object g at iteration t. Particles are iteratively moved according to

$$\mathbf{q}_{t+1}(g) = \mathbf{q}_t(g) + \eta\left(\kappa\big(\mathbf{q}^*(g) - \mathbf{q}_t(g)\big) + \lambda\,\mathbf{R}(g;\mathbf{q}_t)\right), \tag{11}$$

where $\eta > 0$ is a step size, $\kappa > 0$ controls attraction toward the target position, and $\lambda \geq 0$ scales the contribution of collision avoidance. The repulsion term $\mathbf{R}(g;\mathbf{q}_t)$ models local separation from neighboring particles and reduces overlap in dense regions. Iteration continues until the layout stabilizes, a maximum number of iterations is reached, or the time budget is exhausted. The resulting position is therefore close to the barycentric target induced by the magnets, but can deviate locally to ensure visual separability. This yields compact layouts in which clusters, gradients, and outliers remain visible even for dense datasets.

5. Magnet Types

Different magnet types support different analytical tasks in the interactive exploration of multivariate geospatial datasets.

https://doi.org/10.5194/isprs-annals-XII-4-W1-2026-89-2026 | © Author(s) 2026. CC BY 4.0 License.

Single-Feature Magnet. A single-feature magnet is associated with one feature dimension $k \in \{1,\ldots,d\}$ and uses the corresponding normalized feature value directly as the weight:

$$w_k(g) = \tilde{x}_k(g). \tag{12}$$

A single-feature magnet therefore assigns attraction in proportion to the strength of one selected thematic attribute. Single-feature magnets are useful for inspecting the distribution of individual attributes and for composing layouts from several thematic criteria. They remain limited to a one-dimensional analytical perspective.

Inverse Magnet. An inverse magnet inverts the weight of another magnet:

$$w_k^{-1}(g) = 1 - w_k(g). \tag{13}$$

It is particularly useful for forming opposing magnet pairs, for example to span a thematic axis.

Minimum Magnet. A minimum magnet combines two feature dimensions in a conjunctive manner. Let $i, j \in \{1,\ldots,d\}$ denote two selected dimensions. The weight is defined as

$$w_{\min}(g) = \min\big(\tilde{x}_i(g), \tilde{x}_j(g)\big). \tag{14}$$

This magnet yields a strong attraction only if both feature values are simultaneously high. The smaller of the two values acts as the limiting factor, so particles with imbalanced values receive a lower weight. Minimum magnets are therefore suitable for expressing AND-like analytical relations.

Maximum Magnet. A maximum magnet combines two feature dimensions in a disjunctive manner:

$$w_{\max}(g) = \max\big(\tilde{x}_i(g), \tilde{x}_j(g)\big). \tag{15}$$

A particle is strongly attracted if at least one of the two feature values is high. Dominance in one dimension is sufficient, and imbalance is not penalized. Maximum magnets are therefore useful for expressing OR-like analytical relations and for revealing objects that stand out in at least one thematic aspect.

Similarity Magnet. For a given reference object g^*, similarity can be derived from the distance between normalized feature vectors in the d-dimensional unit hypercube:

$$w_{\text{similarity}}(g) = 1 - \frac{\|\tilde{\mathbf{x}}(g^*) - \tilde{\mathbf{x}}(g)\|}{d^{\max} + \varepsilon} \tag{16}$$

where

$$d^{\max}(g^*) = \max_{q \in \mathcal{G}} \|\tilde{\mathbf{x}}(g^*) - \tilde{\mathbf{x}}(q)\|. \tag{17}$$

This magnet attracts objects whose feature vectors are close to the selected reference object in feature space.

Vector-Length Magnet. For a selected subset S of feature dimensions, we interpret the corresponding components as a vector in the induced $|S|$-dimensional subspace and use its length as a weight:

$$w_{\text{length}}(g) = \frac{\|\tilde{\mathbf{x}}_S(g)\|}{d_S^{\max} + \varepsilon},\ \tilde{\mathbf{x}}_S(g) = (\tilde{x}_k(g))_{k \in S} \tag{18}$$

It allows us to express joint prominence or combined strength of feature dimensions. Unlike the minimum magnet, it rewards combined magnitude rather than strict conjunction.

Sink Magnet. A sink magnet redirects particles considered irrelevant into a grid-like side layout, where they remain visible as contextual feedback. Let

$$\chi_{\text{sink}}(g) = \begin{cases} 1, & \tilde{x}_i(g) = 0, \\ 0, & \tilde{x}_i(g) > 0, \end{cases} \tag{19}$$

be the sink predicate for feature dimension i. If $\chi_{\text{sink}}(g) = 1$, particle g is excluded from the main layout and assigned to a grid position in a designated sink region. Otherwise, it remains in the regular barycentric layout. This allows the analysis to focus on relevant particles while preserving filtered-out objects as visible contextual feedback. Color and size encodings remain active for particles placed in the sink.

Spatial Magnets. Spatial magnets derive their weights from spatial relations between geospatial objects and thus integrate spatial analysis into feature-space exploration. They incorporate geographic context into a formally non-spatial layout model, enabling thematic and spatial reasoning to be combined without re-establishing a fixed cartographic arrangement.

If a spatial property is fixed in the dataset, such as nearest-neighbor distance, local density, or a precomputed isolation measure, it can be treated as an ordinary thematic feature. Spatial magnets, in contrast, integrate spatial analysis into the interaction process. They derive weights from spatial relations to interactively selected, moved, or edited reference objects, so that these relations become first-class criteria for the layout computation and adapt to the magnet configuration.

Point-Distance Spatial Magnet. Let $\Phi = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}]$ denote the 2D spatial extent of the dataset. Each $g \in \mathcal{G}$ is associated with a spatial position

$$\mathbf{p}(g) = (x(g), y(g)) \in \Phi. \tag{20}$$

Let $\mathbf{c} \in \Phi$ denote a point of interest, for example the center of the study area or the position of a selected geospatial object. We define the dataset-wide maximum distance

$$d_{\max} = \max_{g \in \mathcal{G}} \|\mathbf{p}(g) - \mathbf{c}\|. \tag{21}$$

The spatial weight of g with respect to $\mathbf{c}$ is then given by

$$w_{\text{center}}(g) = 1 - \frac{\|\mathbf{p}(g) - \mathbf{c}\|}{d_{\max} + \varepsilon}. \tag{22}$$

Objects close to $\mathbf{c}$ receive weights near 1, while more distant objects receive smaller weights. This magnet can be used, for example, to emphasize trees near a selected sampling location, road junction, water source, or plot center. By changing $\mathbf{c}$ interactively, analysts can immediately observe how proximity to different locations affects the particle distribution.

Magnet Selection. The magnet types support complementary modes of exploration. Single-feature and inverse magnets provide simple, interpretable axes for inspecting individual attributes. Minimum, maximum, and vector-length magnets allow testing hypotheses about feature interactions, ranging from strict conjunctions to disjunctions and joint magnitude across multiple dimensions. However, because these magnets aggregate several attributes, individual feature contributions may require additional inspection. Similarity magnets enable example-based exploration with respect to a reference object, while sink magnets support filtering and reduce visual clutter. Spatial magnets reintroduce geographic relations into feature-space layouts and can be interpreted together with a linked map.

6. Case Study: Urban Tree Inventory

Over the past decades, LiDAR has become an established technique for collecting detailed surface representations of objects such as trees, buildings, and infrastructure. Using deep learning-based processing, such data can be transformed into semantically rich geospatial datasets (Bello et al., 2020; Gao et al., 2026). For example, LiDAR is widely applied to monitor urban and forest trees, as it enables comprehensive structural measurements in a non-destructive manner (Holvoet et al., 2025; Maeda et al., 2025). In our case study, we consider a tree inventory dataset for an urban area derived from automatically segmented mobile mapping LiDAR data. The dataset raises hypotheses about correlations, outliers, and structural subgroups that require simultaneous reasoning over several attributes.

6.1 Dataset and Feature Representation

Each tree is represented by its centered ground position and the following thematic attributes:

- *Height*: maximum height of tree points above the terrain
- *Under-branch height*: height up to visible stem points
- *Crown volume*: volume of all crown voxels
- *Crown width*: average width and depth of the axis-aligned crown bounding box
- *Stem diameter*: diameter at breast height (dbh)
- *Stem inclination*: magnitude of deviation of the first principal component of stem points from the vertical axis
- *Nearest-neighbor distance*: distance to the nearest tree neighbor
- *Elevation and roughness*: terrain descriptors for a tree.

Let $\mathcal{T} = \{t_1, \ldots, t_n\}$ denote the set of trees. Each tree $t \in \mathcal{T}$ is associated with a spatial position $\mathbf{p}(t) = (x(t), y(t))$ and a normalized feature vector

$$\tilde{\mathbf{x}}(t) = \big(\tilde{x}_{\text{height}}(t), \tilde{x}_{\text{ubh}}(t), \tilde{x}_{\text{cvol}}(t), \tilde{x}_{\text{cwidth}}(t), \tilde{x}_{\text{dbh}}(t),\\ \tilde{x}_{\text{incl}}(t), \tilde{x}_{\text{nndist}}(t), \tilde{x}_{\text{elevation}}(t), \tilde{x}_{\text{roughness}}(t)\big) \in [0,1]^9.$$

6.2 Crown Volume – Stem Diameter Analysis

According to allometric theory, trees with larger stem diameters are expected to have larger crowns, but this relation may not be uniform across the whole dataset. Different species, site conditions, pruning, or space constraints may lead to deviations (Franceschi et al., 2022). In Fig. 3, we visualize the conjunction of both attributes using a minimum magnet

$$w_{\min}(t) = \min\big(\tilde{x}_{\text{dbh}}(t), \tilde{x}_{\text{cvol}}(t)\big).$$

The *Big Stem & Large Crown* magnet attracts trees that simultaneously have large stem diameter and crown volume; *Big Stem*, *Small Stem* and *Large Crown Volume* further separate the tree population. Particle color is linked to tree height and provides an additional plausibility check by exposing structurally implausible combinations of stem diameter and crown volume.

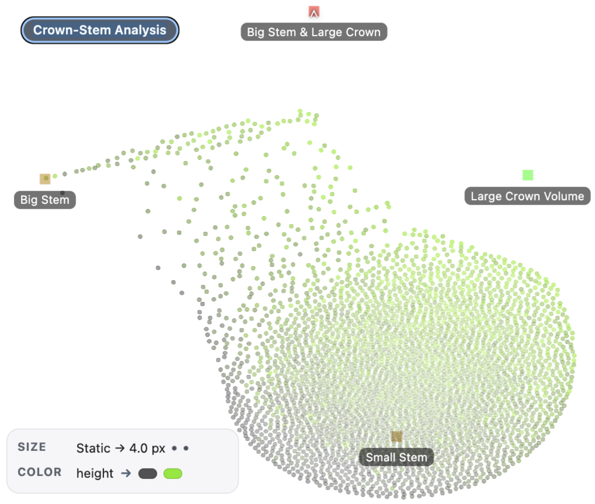

Figure 3. A small subset of trees is strongly attracted to *Big Stem & Large Crown*; the majority remains in the complementary region, indicating that a large stem diameter is not a sufficient condition for large crown volume and that the conjunction is satisfied by only a few trees.

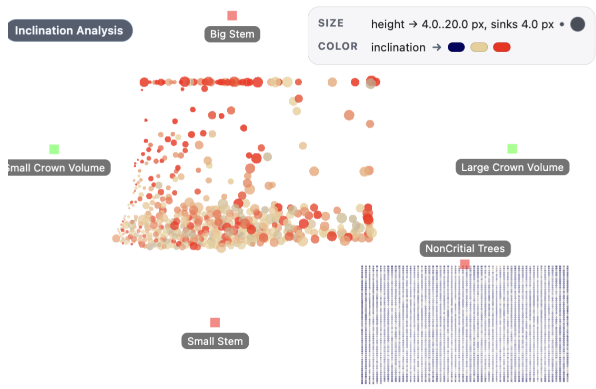

Figure 4. Pronounced inclination occurs only in a minority of trees (ca. 10 %) and is spread across different crown volumes. Critically inclined trees tend to have either small or large stem diameters.

6.3 Stem Inclination Analysis

A high degree of stem inclination may be caused by environmental stress, growth conditions, human intervention, wind exposure, or limited space. To test whether high inclination is related to other structural attributes such as crown volume or stem diameter, we first filter out all trees with inclination in a range considered to be non-critical. For this purpose, we use a sink magnet *NonCritical Trees* (Fig. 4); a threshold transfer function sets inclination values $< \tau$ to zero.

$$w_{sink}(t) = threshold\big(\tilde{x}_{\text{incl}}(t), \tau\big).$$

Particle color encodes inclination, and particle size encodes tree height to reinforce visual discrimination.

6.4 Isolated Tree Analysis

In urban space, trees can occur in dense rows or groups, whereas others stand isolated. Isolated trees may develop broader crowns, larger crown volume, or different branching behavior due to reduced competition and more available space. We apply a single-feature magnet and its inverse (Fig. 5) $w_{isolated}(t) = \tilde{x}_{\text{nndist}}(t)$ and $w_{nonisolated}(t) = 1 - \tilde{x}_{\text{nndist}}(t)$. The crown size of a tree is computed by a length magnet for vectors composed of the inverted under-branch height, crown volume, and crown width:

$$w_{\mathrm{csize}}(t) = \frac{\left\| \left(1-\tilde{x}_{\mathrm{ubh}}(t),\ \tilde{x}_{\mathrm{cvol}}(t),\ \tilde{x}_{\mathrm{cwidth}}(t)\right) \right\|}{d^{\max}+\varepsilon}.$$

The horizontal axis spans the range from *Non-Isolated Trees* to *Isolated Trees*. Vertically, the crown size magnets arrange trees according to structural size. Particle color encodes tree height.

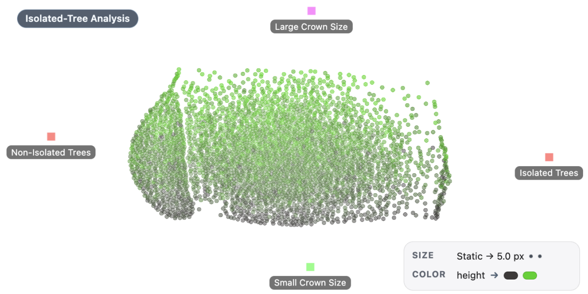

Figure 5. A compact group of non-isolated trees is revealed on the left, representing denser tree zones, whereas only a small subset of strongly isolated trees appears on the right. These are not exclusively large; smaller trees still form the majority. Overall, most trees occupy the central range of the distribution, indicating moderate isolation and size.

6.5 Analysis of Similarity and Vicinity

We select a reference tree t^*, for example, one with a very large crown size, and arrange the particles according to their spatial distance to t^* by a spatial-distance magnet and by a similarity magnet, which computes similarity over the feature dimensions $\{\tilde{x}_{\mathrm{height}}(t), \tilde{x}_{\mathrm{cvol}}(t), \tilde{x}_{\mathrm{cwidth}}(t), \tilde{x}_{\mathrm{ubh}}(t)\}$. Together with their inverse magnets, this yields a layout in which five stacked horizontal bands encode increasing spatial distance to t^*, while the horizontal direction shows tree similarity (Fig. 6). Particle color encodes similarity (red: very similar; yellow: moderately similar; gray: dissimilar), and particle size encodes tree height.

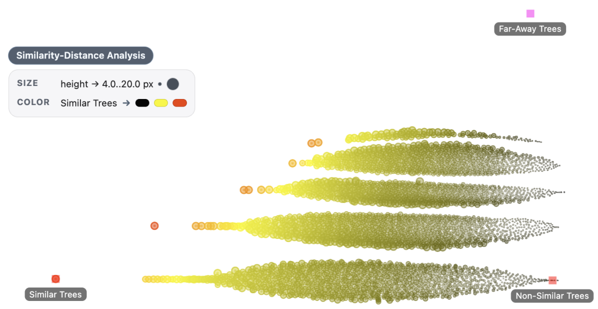

Figure 6. Only a few trees are highly similar to t^*, and these occur across several distance classes rather than within a single local neighborhood. The distribution therefore suggests that structural similarity to t^* is rare and only weakly coupled to geographic proximity.

6.6 Analysis of Terrain Factors

We analyze the relation between terrain aspect (cardinal direction: N-E-S-W) and terrain roughness (degree of planarity) with respect to tree occurrence (Fig. 7). The vertical axis spans terrain roughness from *Planar Terrain* to *Rough Terrain*, and the horizontal axis groups terrain aspect into cardinal-direction sectors. Staircase transfer functions discretize magnet weights into bins, so that each particle cluster represents one terrain-factor combination. Particle color encodes crown-volume bands.

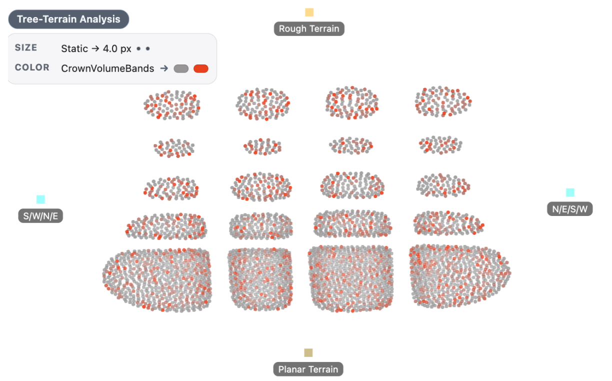

Figure 7. The majority of trees occur on planar or slightly rough terrain, with far fewer located in rough terrain. In contrast, the distribution across terrain-aspect classes is comparatively even, indicating that terrain orientation has only a weak influence in the study area. Particle colors do not reveal a strong directional concentration of large-crowned trees.

7. Conclusions

This paper presented an interactive feature-space visualization approach for multivariate geospatial objects. By representing geospatial objects as particles derived from feature vectors and structuring them through thematic, analytical, and spatial magnets, the approach complements map-based views with a non-cartographic layout in which thematic relations, outliers, and latent groups become directly explorable while remaining linkable to geographic space. The urban tree inventory case study showed that the method supports hypothesis-driven analysis of structural relations such as height–crown size dependencies, inclination patterns, isolation effects, similarity neighborhoods, and terrain-related distributions. Beyond revealing thematic structures that are difficult to perceive in maps, the approach also proved useful for plausibility checking and qualitative inspection of derived multivariate geospatial data.

In an ongoing project, we are applying this method to AI-based logistical planning of maintenance work for technical facilities: Magnets represent technical staff, and particles represent maintenance orders. Criteria include spatial distance (e.g., travel costs), facility complexity (e.g., materials), and order priority (e.g., scheduled maintenance, emergency maintenance).

Overall, the paper positions feature-space visualization as a complementary geovisual analytics technique for multivariate geospatial datasets. Future work will focus on extending the repertoire of analytical and spatial magnets, integrating interactive feature construction more tightly into iterative geospatial analysis workflows, and investigating possible extensions to 3D visualization. Such extensions could use either layered visualizations with stacked 2D particle layouts or true 3D layouts in which particles are arranged directly in 3D space.

Acknowledgements

We thank the Department for Geoinformation, Surveying and Cadastre of the City of Essen for providing mobile mapping LiDAR data. This work was partially funded by the Federal Ministry of Research, Technology and Space, Germany, through grant 033L305 ('TreeDigitalTwins') and by the German Ministry of Economics under the grants ZIM SPUR-OPT 16KN113522 and ZIM Insta.KI 16KN117820 & 16KN11782. We would also like to thank Jose Carvajal, Markus Portune, Rico Richter, and Norbert Riedel for their collaboration.

References

Andrienko, N., Andrienko, G., Fuchs, G., Slingsby, A., Turkay, C., Wrobel, S., 2020. *Principles of Interactive Visualisation*. Springer, Cham, 51–88.

Antonini, A. S., Luque, L., Ganuza, M. L., Castro, S. M., 2023. Toward a taxonomy for 2D non-paired General Line Coordinates: a comprehensive survey. *Int. J. Data Sci. Anal.*, 15(2), 133–158.

Bello, S. A., Yu, S., Wang, C., Adam, J. M., Li, J., 2020. Review: Deep Learning on 3D Point Clouds. *Remote Sens.*, 12(11), 1729.

Bertin, J., 1983. *Semiology of Graphics*. University of Wisconsin Press, Madison, WI, USA. (trans. Berg, W.).

Carr, D. B., Littlefield, R. J., Nicholson, W. L., Littlefield, J. S., 1987. Scatterplot Matrix Techniques for Large N. *J. Am. Stat. Assoc.*, 82(398), 424–436.

Elmqvist, N., Dragicevic, P., Fekete, J.-D., 2008. Rolling the Dice: Multidimensional Visual Exploration using Scatterplot Matrix Navigation. *IEEE Trans. Vis. Comput. Graph.*, 14(6), 1141–1148.

Espadoto, M., Martins, R. M., Kerren, A., Hirata, N. S. T., Telea, A. C., 2021. Toward a Quantitative Survey of Dimension Reduction Techniques. *IEEE Trans. Vis. Comput. Graph.*, 27(3), 2153–2173.

Franceschi, E., Moser-Reischl, A., Rahman, M. A., Pauleit, S., Pretzsch, H., Rötzer, T., 2022. Crown Shapes of Urban Trees–Their Dependences on Tree Species, Tree Age and Local Environment, and Effects on Ecosystem Services. *Forests*, 13(5), 748.

Fujiwara, T., Wei, X., Zhao, J., Ma, K.-L., 2022. Interactive Dimensionality Reduction for Comparative Analysis. *IEEE Trans. Vis. Comput. Graph.*, 28(1), 758–768.

Gao, Y., Xia, S., Wang, P., Xi, X., Nie, S., Wang, C., 2026. LiDAR remote sensing meets weak supervision: Concepts, methods, and perspectives. *ISPRS J. Photogramm. Remote Sens.*, 235, 72–104.

He, X., Tao, Y., Wang, Q., Lin, H., 2019. Multivariate spatial data visualization: A survey. *J. Vis.*, 22(5), 897–912.

Heinrich, J., Weiskopf, D., 2013. State of the Art of Parallel Coordinates. *Eurographics – State of the Art Reports*, The Eurographics Association, 95–116.

Holvoet, J., Eichhorn, M. P., Giannetti, F., Kükenbrink, D., Liang, X., Mokroš, M., Novotný, J., Pitkänen, T. P., Puliti, S., Skudnik, M., Stereńczak, K., Terryn, L., Vega, C., Torresan, C., 2025. Terrestrial and mobile laser scanning for national forest inventories: From theory to implementation. *Remote Sens. Environ.*, 329(114947).

Hotelling, H., 1936. Simplified Calculation of Principal Components. *Psychometrika*, 1(1), 27–35.

ICA Executive Committee, 2004. A Strategic Plan for the International Cartographic Association 2003–2011. *Map*, 42(2), 61–73.

Ivosev, G., Burton, L., Bonner, R., 2008. Dimensionality Reduction and Visualization in Principal Component Analysis. *Anal. Chem.*, 80(13), 4933–4944.

Kandogan, E., 2000. Star Coordinates: A Multi-dimensional Visualization Technique with Uniform Treatment of Dimensions. *IEEE VIS*, IEEE.

Kiesel, D., Riehmann, P., Froehlich, B., 2023. Smooth Transitions Between Parallel Coordinates and Scatter Plots via Polycurve Star Plots. *Comput. Graph. Forum*, 42(6), e14923.

Kraak, M.-J., Ormeling, F., 2020. *Cartography: Visualization of Geospatial Data*. CRC Press, Boca Raton, FL, USA.

Lengyel, J., Roux, S., Alvanides, S., 2022. Multivariate analysis of socioeconomic profiles in the Ruhr area, Germany. *J. Maps*, 18(3), 576–584.

Liu, Y., Jun, E., Li, Q., Heer, J., 2019. Latent Space Cartography: Visual Analysis of Vector Space Embeddings. *Comput. Graph. Forum*, 38(3), 67–78.

Maeda, E. E., Brede, B., Calders, K., Disney, M., Herold, M., Lines, E. R., Nunes, M. H., Raumonen, P., Rautiainen, M., Saarinen, N., Starck, I., Su, Y., Terschanski, J., Terryn, L., 2025. Expanding forest research with terrestrial LiDAR technology. *Nat. Commun.*, 16(8853).

McInnes, L., Healy, J., Melville, J., 2020. UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv preprint*, stat.ML, 1802.03426.

Ropinski, T., Oeltze, S., Preim, B., 2011. Survey of glyph-based visualization techniques for spatial multivariate medical data. *Comput. Graph.*, 35(2), 392–401.

Saary, M. J., 2008. Radar plots: A useful way for presenting multivariate health care data. *J. Clin. Epidemiol.*, 61(4), 311–317.

Sacha, D., Zhang, L., Sedlmair, M., Lee, J. A., Peltonen, J., Weiskopf, D., North, S. C., Keim, D. A., 2017. Visual Interaction with Dimensionality Reduction: A Structured Literature Analysis. *IEEE Trans. Vis. Comput. Graph.*, 23(1), 241–250.

Smilkov, D., Thorat, N., Nicholson, C., Reif, E., Viégas, F. B., Wattenberg, M., 2016. Embedding Projector: Interactive Visualization and Interpretation of Embeddings. *arXiv preprint*, stat.ML, 1611.05469.

Trye, D., Apperley, M., Bainbridge, D., 2023. Extending the Heatmap Matrix: Pairwise Analysis of Multivariate Categorical Data. *IEEE VIS*, IEEE, 29–36.

van der Maaten, L., Hinton, G., 2008. Visualizing data using t-SNE. *J. Mach. Learn. Res.*, 9(86), 2579–2605.

Vaz de Melo Mendes, B., Pereira Câmara Leal, R., 2005. Robust multivariate modeling in finance. *Int. J. Manag. Finance*, 1(2), 95–106.

Vollmer, J. O., Döllner, J., 2020. 2.5D Dust & Magnet Visualization for Large Multivariate Data. *VINCI*, ACM.

Weidele, D. K. I., 2019. Conditional Parallel Coordinates. *IEEE VIS*, IEEE, 221–225.

Yi, J. S., Melton, R., Stasko, J., Jacko, J. A., 2005. Dust & Magnet: Multivariate Information Visualization Using a Magnet Metaphor. *Inf. Vis.*, 4(4), 239–256.