

A Taxonomy of Point Cloud Search

Balthasar Teuscher¹, Paul Walther¹, Jiapan Wang², Martin Werner¹

¹ Technical University of Munich, Germany; TUM School of Engineering and Design, Department of Aerospace and Geodesy, Professorship of Big Geospatial Data Management - balthasar.teuscher@tum.de, paul.walther@tum.de, martin.werner@tum.de

² Technical University of Munich, Germany; TUM School of Engineering and Design, Department of Aerospace and Geodesy, Professorship of Remote Sensing Applications - jiapan.wang@tum.de

Keywords: Point Cloud, Search Taxonomy, Query Semantics.

Abstract

Point cloud analysis is rapidly evolving, targeting new applications and use cases with novel information retrieval needs that challenge existing solutions' scalability, robustness, and reusability to manage and process point cloud data. Analytical approaches to gain insights are increasingly based on machine learning and tend to turn away from data management solutions in favour of internalizing custom, dedicated workflow-specific query capabilities, satisfying their requirements. Unfortunately, these ad-hoc solutions often fail to scale well with large point cloud datasets generated through terrestrial, aerial, or mobile laser scanning. To address these limitations, we propose a point cloud search taxonomy and use it to identify fundamental requirements for a scalable, robust, and reusable data management system for state-of-the-art point cloud retrieval and data analytics. Our findings build a foundational analysis serving as a basis for the holistic development of point cloud data management solutions to overcome current bottlenecks.

1. Introduction

In recent years, there has been a surge in interest and attention to point clouds driven by an ever-increasing deployment of data acquisition techniques such as light detection and ranging or photogrammetry. A wide variety of sensors can be mounted on different platforms, such as satellites, low-altitude drones, vehicles, and even mobile devices, greatly enhancing the perception of their environment, respectively, three-dimensional space, by machines for assisted and autonomous navigation, depth perception, or 3D scanning (Toth and Józków, 2016). Point clouds have become popular data sources in applications such as urban mapping, traffic modeling, environmental monitoring, and autonomous driving by providing accurate 3D geometry and attribute information of the entities in the real world. The usefulness and benefits of these applications are profoundly intertwined with the extent to which information needs can be satisfied by point cloud search.

Point cloud search is fundamental in almost any analytical procedure involving point cloud data. Examples of searching include querying the nearest neighborhood of a given point (Rousopoulos et al., 1995) or extracting subsets based on defined spatiotemporal extents. However, more elaborate searches like farthest point sampling (Qi et al., 2017b) or ray queries (Chang et al., 2023) are often implemented ad-hoc as a necessity to satisfy the contextual needs of a given method. Unfortunately, transferring these approaches to larger datasets is rarely applicable. Further, while scalable data management systems exist (Lokugam Hewage et al., 2022), they are often not considered viable for the implied retrieval needs. The reasons for this are high integration and maintenance costs, and the lack of query capabilities, among others. To overcome these shortcomings, we would like to revisit the problem of point cloud search from the information retrieval perspective.

From the information retrieval perspective, point cloud search is the procedure of obtaining information system resources relevant to an information need from a collection of those re-




Figure 1. Basic structure of the information retrieval process, adapted from Buckland and Plaunt (1994) and Rijsbergen (1979).

sources (Manning, 2009). It entails an *information need* or a question that needs to be translated into a *query* that the information system is *matching* against (*indexed*) documents to answer with *retrieved documents* (Buckland and Plaunt, 1994). A *feedback* loop finally completes the iterative information retrieval process by refining the query formulation and adaptations to the information need as visualized in Figure 1.

This paper aims to investigate point cloud search from an information retrieval perspective, systematically abstract the steps of point cloud search, summarize the common search scenarios, and evaluate the corresponding strategies, to provide a playground for discussing point cloud search problems and how to overcome them. From examining existing and emerging retrieval needs, a more formalized point cloud search taxonomy is built that is generic over many possible point cloud transformations presented thereafter. Based on gained insights, recommendations are formulated towards scalable, unified point cloud data management systems for contemporary retrieval tasks.

Our main contributions are (1) an introduction to point cloud retrieval and the underlying theoretical problems, such as search, optimization and ranking, (2) reviewing and analyzing existing point cloud information retrieval needs, (3) proposing a new point cloud search taxonomy including a reflection on the manifold possible point cloud representations and transformations and their implications, and, (4) formulate recommendations for improving point cloud data retrieval and management systems.

2. Background

At the core of information retrieval is *matching*, where queries are resolved by evaluating predicates against the indexed documents. As such, point cloud search is essentially a search, optimization, or ranking problem.

2.1 The Search Problem

A *search problem* is a fundamental computational problem of theoretical computer science. It entails looking for the set $L(R)$ formed by the values of x such that the predicate, respectively, the binary relation $R(x, y)$ is true for a given y (Eq. 1).

$$L(R) = \{x \mid \exists y R(x, y)\} \quad (1)$$

For example, x is a data point in the dataset, and y is a query or reference point. The equation 1 defines the subset of data that matches the query condition (i.e., when $R(x, y)$ is true).

Search algorithms describe implementable solutions to the theoretical search problem. Since search problems are expected to always return an explicit result, meaning the predicate invariably evaluates to true or false, this can be denoted as an *exact predicate search*. There are two basic types of search algorithms for exact predicate search: linear and binary. *Linear search* approaches perform a sequential scan on a set of records or documents. *Binary search* approaches repeatedly divide the search space in the middle of the data structure by performing a comparison. The latter requires the records to be represented in an ordered set. In the context of point cloud data retrieval, basic examples are evaluating whether a point exists in the set or a range query to extract a spatial subset.

2.2 Optimization Problem

Finding the best solution from all feasible solutions is denoted as an optimization problem. Depending on whether the variables are continuous or discrete, it is a continuous optimization or a discrete, respectively, combinatorial optimization. The latter can be expressed as a decision problem, while the former is a minimization problem or an equivalent maximization problem with a reversed sign. Finding a nonexact solution to an optimization problem is denoted as an *approximation*. Common optimization problems in point cloud retrieval are pose estimation for localization, registration, or surface intersection.

2.3 Ranking Search

The *ranking problem* and *constraint satisfaction problem* are other similarly fuzzy approaches like approximation. Ranking is widely explored in information retrieval. It addresses the *ordering* of matched results based on criteria corresponding to the importance and measure of goodness to answer the given information need. A classic example of ranking in point

clouds is nearest neighbor queries, where the points are ordered by increasing distance to a query point, often accompanied by a parameter k to constrain the size of the result set. Addressing the large result set problem occurring from queries with low selectivity is a key aspect of ranking.

3. Related Work

Taniar and Rahayu (2013) proposed a taxonomy for nearest neighbor queries in spatial databases. Their taxonomy for queries on stationary objects comprises four perspectives: *Space*, *Result*, *Query-Point*, and *Relationship*. The *Space* perspective emphasizes the importance of algorithms to respect data semantics in measures by comparing Euclidean distance to network distance when querying against network data like roads. The *Result* perspective encompasses nearest neighbor searches that do not return a predefined number of results; occurring, for example, from using relative measures, adding constraints, or combinations with range searches. The *Query-Point* perspective includes queries with multiple points or ranges as input, like all-kNN, Group Nearest Neighbour, Group Nearest Group, or Range-kNN. Finally, the *Relationship* perspective highlights different cases of the traditionally uni-directional relation between point and neighbor, like bi-directional or reversed.

Subsequently, a taxonomy for range queries on stationary objects was proposed that categorizes into finding *objects of interest*, *forming regions*, and *determining centroids* (Taniar and Rahayu, 2015). Queries in the *objects of interest* category are about finding objects in a region and consist of three basic elements: the region, objects of interest, and the query point. The category is divided into six types: traditional region queries, approximate region queries, constrained region queries, clustered objects region queries, outer/inner fence object queries, and inverse range queries. The *forming regions* category is about forming regions based on a set of objects. It distinguishes further between creating polygons, zones of influence, optimum regions, safe regions, and reverse region queries. Finally, the *determining centroids* category covers finding the centroid of polygons and clustered objects.

For moving objects, a taxonomy exists that elaborates on the *location*, *motion*, *object*, *temporal*, and *patterns* perspectives of moving object queries (Alamri et al., 2014). The *location* perspective includes common spatial queries such as nearest neighbors (KNNs), range queries, and others. The *motion* perspective covers direction, velocity, distance, and displacement queries. The *object* perspective includes type and form status queries. The *temporal* perspective includes trajectory, timestamped, inside, disjoint, meet, equal, contain, overlap, and period queries. In the *patterns* perspective, the moving objects use undefined or predefined movement patterns, including numerous spatial and temporal movement patterns.

All the mentioned existing taxonomies and underlying search engines generally take planar spatial representations as granted. Some can easily be extended to a third dimension, such as promoting a bounding box range query to a bounding volume range query. However, the peculiarities of point cloud search are neglected. Research in the field of 3D information retrieval addresses many challenges of 3D search engines and possibilities to retrieve, respectively search for shapes (Funkhouser et al., 2003; Tangelder and Velkamp, 2008; Lian et al., 2013; Grabner et al., 2018; Williams et al., 2022). Unfortunately these efforts

are mainly geared towards enclosed small scale models, as used in computer-aided design or online shopping platforms, which makes it of little use for the open-ended nature of large scale terrestrial laser scans.

4. Information Need

A basic premise of database management systems for point clouds is to offer the required data retrieval needs of users and applications formulated as explicit queries. These information needs are as manifold as the branches where point clouds are used: Fields such as architecture (Zhao et al., 2023), construction (Mirzaei et al., 2022), agriculture (Song et al., 2025), heritage preservation and archeology (Yang et al., 2023), and geomorphology (Dumic and da Silva Cruz, 2025) all rely on point cloud data to model the real world. From a high-level perspective, one can loosely distinguish between backward- and forward-oriented use cases.

4.1 Backward-Oriented Information Need

In *backward-oriented* use cases, virtual representations are typically static; observations are a snapshot in time, and changes are captured through discrete, recurring surveying. Such point clouds can be large, covering a wide area with varying spatio-temporal resolution. For example, point clouds in real-world environmental observations may range from whole districts to even larger areas, with a centimeter-level spatial resolution and annually, monthly, or even hourly temporal frequency, resulting in tens of thousands of point clouds (Vos et al., 2022). The retrieval needs concerning such point clouds are characterized by prevalent analytical approaches like object and change detection, classification, or semantic segmentation over large spatiotemporal extents.

4.2 Forward-Oriented Information Need

In comparison, *forward-oriented* applications, such as simultaneous localization and mapping for autonomous navigation based on point cloud data from LiDAR sensors or computer graphics and game design, require (close to) real-time feedback on their information need. Environment changes are comparatively fast and form the basis for deriving actionable insights. However, such point clouds have a relatively short temporal lifespan and predominantly target the environment in the close vicinity of the surveying entity (car, viewer in computer games). Consequently, they are relatively small but require low latency and high query throughput.

4.3 Nature of the Information Need

Meanwhile, the nature of needed information is relatively similar in many cases: *aggregated* or *summarized* data from the point clouds (Lu et al., 2021), retrieval of a relevant *subset* of the whole point cloud (Asad and Savva, 2023), *visualize* information stored in the point cloud (Remondino, 2003; Dumic et al., 2020), *detect* and *localize objects* (Kadam et al., 2022), or *feature extraction* and *classification* (Wu et al., 2019; Qi et al., 2017b). Depending on the application, the *exactness* of the needed information can vary considerably. Point cloud data in the construction environment may require high-precision results, while geomorphology observations in post-event evaluations may allow more error tolerance. Some applications may even display dynamic exactness requirements, like visualizations where the desired accuracy negatively correlates with the distance to the viewpoint.

To summarize, the information needed from point clouds mainly differs in five dimensions, aligning with big data characteristics:

- *Volume* - size of the queried point cloud data and result set.
- *Velocity* - retrieval speed, latency, and throughput.
- *Variety* - data modalities, representations, and transformations.
- *Veracity* - exactness requirements.
- *Variability* - robustness of the acquired information.

If all of these dimensions are adequately addressed for a specific application, it results in the highest achievable *value* of the point cloud search in this particular use case.

So far, no known point cloud retrieval system conforms to the various information needs of all the applications mentioned above. Additionally, there is no comprehensive overview of contemporary information needs for point cloud search problems in database management systems. Consequently, in the following, we strive to structure different aspects of point cloud search into a taxonomy viable for a wide variety of (even future) information needs, to offer a perspective to induce and inform discussions on requirements for such systems.

5. Point Cloud Search

A point cloud search engine is a computational system that takes a point cloud and provides a query language for retrieving results. Applying the Input-Process-Output (IPO) model (Goel, 2010) to this definition, the input is a query, the process entails the query transformation and execution, and the output is the representation of the query results of the ingested, transformed, and indexed point cloud. In this work, we formalize a point cloud P as

$$P = (c_i, a_i), \quad (2)$$

where c_i is the spatial location and a_i are additional attributes. It needs to be noted that the spatial dimensions, though required to qualify as a point cloud, are considered attributes as well, from which follows

$$P \subseteq \prod_{i=1}^m D_i, \quad (3)$$

where D represents the data space, equivalent to the search space. An advantage of these definitions is that they do not assume any concrete spatial encoding and are therefore suitable to cover various transformations like voxelization, rasterization, meshing, and arbitrary skeletons. For example, considering range queries in point clouds to generate patches for machine learning approaches, the architecture of the specific algorithm defines whether a sparse voxelization (Engelcke et al., 2017) or unordered point sets (Qi et al., 2017a) is required. This is adding a level of complexity to the information needed. Therefore, we will look at different point cloud transformations in Subsection 5.5.

5.1 A Taxonomy

To build our point cloud search taxonomy, we rely on a faceted approach on the first level that incorporates the overall retrieval structure and the Input-Processing-Output model view on point cloud search engines. The input is denoted as *query modality*, the processing relates to the *query semantics* emphasizing user control over execution, and the output is dubbed *result modality* as depicted in Figure 2.




Figure 2. Point Cloud Search Taxonomy.

The query modality encompasses how a query is expressed and encoded, distinguished between *Sub-Cloud*, *Mesh*, *Image*, or *Concepts*. The query semantics entail how the query is interpreted and matched against the point clouds, either *Invariant*, *Approximate*, or *Probabilistic*. The result modality covers different kinds of results, namely *Pose*, *Frequency*, *Subset*, or *Summaries*. In the following, the different taxonomy facets are examined in more detail.

5.2 Query Modality

Queries for point cloud retrieval are often formulated in terms of spatial concepts such as geometric entities and topological relations. For the query modality, we concentrate on the geometric aspect of queries, specifically their encoding.

Sub-Cloud A sub-cloud defines a query encoded similarly to the queried documents. The term relates to the mathematical concept of a subspace in a vector space. It indicates that the query input should be representable or mapped to the search space defined by the documents. Therefore, it is not limited only to spatial dimensions (i.e., x , y , z) but also includes attributes (e.g., time, features) adhering to our definition of a point cloud in Equations 2 and 3. Exemplary information needs are to know whether the point cloud contains specific points, search points with a particular label or attribute value, find similar parts in the point cloud, spatio-temporal slicing, or find points along a trajectory.

Mesh Similar to unordered point clouds, meshes can faithfully represent manifolds such as surfaces and volumes. Their benefit is accurately defining split planes, including potentially inside and outside semantics. The inputs of many queries, such as arbitrary range queries, bounding box queries, view frustum queries, or even spatio-temporal slicing, are representable as a mesh and therefore considered in this modality.

Image Querying based on images is a prominent use case when working with point clouds. Examples are image registration in photogrammetry and image localization through depth

maps. The primary information retrieval need is the pose estimation and alignment of images’ discrete two-dimensional projections to the unordered three-dimensional point clouds. This often means finding out the extrinsic camera matrix in the reference system of the queried point cloud.

Concepts Concepts entail queries formulated based on an abstract notion, such as a “car” or a “church”. Herby, the “car” or “church” is not encoded as a sub-cloud, mesh, or image but as a textual description or term. Unlike the other modalities, where a transformation is enough to match against the search space, concepts require a translation into the search space before matching and retrieving results. However, if the point cloud is semantically enriched with an attribute for a car label, such a query would be demoted to the sub-cloud modality.

5.3 Query Semantics

The query semantics facet can be considered an extension of the query modality. It captures essential aspects for scalability and performance, such as sampling and approximation, crucial for efficient information retrieval in large datasets.

Invariant Queries with invariant semantics are expected to return the complete and exact result set given the predicates. This is often implied because one expects to get all matching documents when issuing a particular query. For point clouds, this means they require, e.g., the retrieval of every and all points matching the query. A significant drawback of invariant query semantics is their subpar scalability when the dataset size increases and predicates are not selective enough.

Approximate Queries with approximate semantics emphasize not expecting an exact result. Most prominent are approximate nearest neighbors queries, where the results are in close proximity to the query point but not necessarily the closest ones. The benefits of approximate queries are fuzzy searching methods and early stopping of the matching process, which can help to increase retrieval performance by sacrificing accuracy.

Probabilistic Queries that include probabilistic semantics. A straightforward example is to reduce the result set to meet a specific budget by random sampling for visualization purposes. Another is common to machine learning methods, where probabilistic samples generate multiple inputs from the same perspective.

5.4 Result Modality

The result modality distinguishes queries based on the result produced and returned.

Pose A pose consists of a location and orientation. Queries with the intent to localize a given input return a pose. This is commonly done for image and point cloud registration and localization tasks.

Frequency Queries returning frequencies comprise the number of matching observations instead of the matched documents or data points. Given the name, the results can be aggregated continuously or discretely across a given domain, like time.

Subset A subset implies that the result set is part of the original data set. This is often the case when queries contain predicates that translate into a binary search and, hence, resemble filtering.

Summaries A summary is derived from the original dataset through descriptive statistics or other functions. Summary results can relate to the search, respective data space (e.g., centroid, mean), or be unrelated (e.g., count, size). The commonality is that the returned information is condensed and representative of a subset of the dataset.

5.5 Representations

Querying and matching against raw point clouds does not necessarily serve the information needs or the presented query modalities, query semantics, and result modalities sufficiently. Hence, retrieval systems often operate on derived high-level descriptions (e.g., a skeleton) to find matching results. The multitude of transformations possible to generate a searchable representation from raw point cloud data represents a principal challenge in point cloud search and in building the taxonomy. Several transformations are presented in the following, including their implications for searching.




Figure 3. Illustration of various point cloud representations: a) raw points, b) mesh, c) grid rounded, d) voxel, e) raster, f) skeleton.

Meshing and Sampling Meshing and sampling are two interrelated methods. Sampling is merely the process that occurs when scanning a real-world object to create a point cloud. Sampling spatial objects, like polygons and other geometries, offers the interesting trait of producing a consistent and generic representation for any object as a point cloud. Meshing, on the other hand, is the reverse process of reconstructing the surface representation by a point cloud as a mesh. One advantage of meshes is their continuous, well-defined surface representation, which allows faithful interpolation and surface intersection, including normals.

Grid Rounding and Voxelization Grid rounding moves the coordinates of the points to the nearest center or corner of a cell in a predefined grid. This offers the possibility of transforming the coordinates from continuous to discrete variables and thus moving them from floating point to integer types. One can further aggregate the points corresponding to a cell into a so-called voxel. Depending on how the aggregation is done, this process of voxelization can reduce the complexity of the data by reducing it to a single representation of the cell. However, it comes at the cost of information loss.

Projections (Rasterization) Projections of point clouds are commonly used to reduce dimensionality and discretize the information. Further, this is used to map raw point clouds in array data structures. Prominent applications are, for example, rasterized digital elevation and surface models.

Aggregation (Skeleton / Graph) Aggregation is a common principle to derive a higher-level skeleton for point clouds or

parts thereof. Aggregation often involves semantic enrichment of points through clustering, classification, or semantic segmentation to group points into more meaningful entities. These objects or segmented parts are generally composed of individual observations and are more suitable for formulating queries. This is similar to words being composed of letters in a text search. For certain retrieval needs, it can be beneficial to put such a higher semantic layer at the core of the retrieval system, as documents, respectively, records in a database, instead of single observations.

6. Discussion

Our developed taxonomy consists of the three facets *Query Modality*, *Query Semantics*, and *Result Modality*. To evaluate its applicability to a wide variety of point cloud queries, we identify typical examples and situate them in the taxonomy for validation. Based on this integration and the reflection in the context of the existing system capabilities, we derive recommendations for improving existing and future point cloud retrieval systems.

6.1 Examples

Based on various literature, we identified about twenty typical point cloud queries, which are presented in Table 1, including the applied taxonomy and prominent application context. They can loosely be categorized as neighbourhood, range, sampling, visualization, registration, aggregation, or semantic queries, and discussed accordingly hereafter.

Neighbourhood Queries Neighbourhood queries, such as radius and kNN queries, are the most specific for point cloud analysis. They often either return points (Subset) or derived neighborhood descriptors thereof (Summary) in the vicinity of a query point (Sub-Cloud). Apart from kNN with approximate semantics, they are generally invariant. A special case are vector queries in higher dimensions, which are not widely used for point cloud analysis primarily focusing on spatial dimensions. However, we predict this will become more important for semantically enriched point cloud analysis with large language models (LLM).

Range Queries Range queries are the most elemental in database management systems as they integrate well with typical indexing facilities. Likewise, axis-aligned bounding box predicates based on a spatial extent are fundamental to spatial operations for extracting subsets. More complex queries, such as slicing and polygons (often generated by buffering geometries), support more elaborate predicates for spatio-temporal and topological analysis. Even though topological relationships with polygons are not directly matchable through range queries, they are often used in a filter and refine strategy with their bounding box. The query modality is mesh, given that all the geometries involved are representable with meshes, and the result modality is subset. Current systems predominantly only support invariant semantics.

Sampling Queries Sampling queries can have any query modality since this refers to our taxonomy’s semantic facet. Sampling can be discrete, random, or pseudo-random, like poison disk sampling. Especially for large point clouds, sampling becomes paramount; nevertheless, this is rarely considered in existing solutions apart from the explicit data organization for visualization in formats. Further, modern machine learning based approaches heavily use downsampling and scaling; these models

Query	Query Modality	Query Semantic	Result Modality	Use Cases / Prominent Application
k-Nearest Neighbor (kNN)	Sub-Cloud	I	Subset / Summary	Neighbourhood descriptors and features
Approximate kNN	Sub-Cloud	A	Subset / Summary	Neighbourhood descriptors and features
Radius	Sub-Cloud	I, (A, P)	Subset / Summary	Neighbourhood descriptors and features
Vector	Sub-Cloud	P, (I, A)	Subset / Summary	Similarity search
Range / Box	Mesh	I, (A, P)	Subset	Spatio-Temporal filtering
Slicing	(Mesh)	I, (A, P)	Subset	Spatio-Temporal subsetting
Polygon / Cone / Solid	Mesh	I, (A, P)	Subset	Manifold intersection, Topological analysis
Discrete Sampling	Any	I, (A)	Subset	Environmental Data Retrieval
Random Sampling	Any	P	Subset	continuous Level of Detail, Thinning, Sampling
Poisson Disk Sampling	Any	P	Subset	Disk-/blue noise sampling, Downsampling
Furthest Points Sampling (FPS)	Any	I	Subset	Downsampling, Representative points
View Frustum	Mesh	I, (P)	Subset	Visibility analysis (Field of View), Visualization
Ray	Sub-Cloud	I, (P)	Pose / Frequency	Rendering, Surface intersection, Visibility analysis
Trajectory	Sub-Cloud	I, (P)	Subset	Movement analysis, Collision detection
Co-registration	Sub-Cloud	Any	Pose	Iterative closest point (ICP) / Alignment
Point set registration	Sub-Cloud	Any	Pose / Frequency	Part search, Alignment
Image registration	Image	A, (P)	Pose	Camera position and orientation / Photogrammetry
Aggregate	Any	Any	Summary	Point density, Centroid, Representative points/values
Window	Any	I	Summary	Spatio-temporal change detection
Join	Any	I	Subset	Elevation, Data fusion, Topological analysis
Attribute	Sub-Cloud	I, (A, P)	Subset	Filter by class or label, Semantic search
Text	Concept	P, (I, A)	Any	Natural language, Semantic search

Table 1. Example queries and their situation in the query taxonomy. Query semantics are invariant (I), approximate (A), and probabilistic (P).

inevitably lead to creating probabilistic patches. The sampling strategy and quality have been shown to substantially influence the performance and quality. However, more research is needed to investigate sampling strategies based on semantic importance.

Visualization Queries Visualization queries, such as ray tracing queries or view frustum, are similar to other queries, such as trajectory or solid queries. Even though visualization requires relatively simple queries, it heavily relies on augmenting points with an importance semantic, such as continuous or discrete Level of Detail (LoD). The reliance on custom formats that incorporate this in the data organization and representation, and the failure of database management systems to easily provide such features, indicates a general problem of such solutions to cope with adding additional attributes and executing predicates on them efficiently.

Registration Queries Registration queries, such as image registration or point set registration, intend to localize a given query input by returning a pose. While iterative query semantics are possible (e.g., when having an exact subset), approximate or probabilistic semantics are the norm. They are common in 3D reconstruction workflows and present higher-level queries that are often solved iteratively while maintaining intermedi-

ate results generated and updated by many simpler range and neighbourhood queries. Even though such complex retrieval workflows are not expected to be directly integrated by retrieval systems, they should at least be able to serve as a backend for doing so. Similarly, this is also desirable for machine learning workflows such as object detection, classification, and semantic segmentation.

Aggregation Queries Aggregation queries, such as window functions, generally entail some subquery run over the windows or partitions and can therefore have any query modality, semantic and result modality, of which summarization is most common for the latter. Use cases are extracting representative values/points (Summary), such as change or centroid, or partitioning into tiles (Subset). Interestingly, such queries are predestined to apply transformations such as voxelization, rasterization, or generating skeletons.

Semantic Queries Semantic queries, such as attribute or text queries, instead target non-spatial information of points. Still, attribute queries are considered to belong to the sub-cloud modality, as they make up the search space given our inclusive definition of point clouds (see Equation 2 and 3). Text queries, conversely, embody concepts of information in point clouds, from well-defined objects such as cars and trees up to arbitrary ques-

tions expressed in natural language. Necessary semantic enrichment of points to facilitate such queries, such as classification and semantic segmentation, is a vividly explored topic. However, their integration into retrieval and data management systems based on and for semantically enriched point clouds is still in its infancy. Some current systems even struggle to support simple attribute queries, like, for example, key-value stores, which only expose the key transparently, whereas the rest of the information, even spatial dimensions, are in a binary blob.

6.2 Recommendations

Investigating point cloud search from a retrieval perspective revealed a high variety of information needs and possible representations. A key issue of the taxonomy surfaced from the openness of what dimensions are part of the search space. From an inclusive viewpoint, one could argue that every existing attribute of any point in a point cloud constitutes a dimension. The implication would be to assign a mere attribute query to the sub-cloud query modality. However, existing retrieval systems seldom treat all attributes equally.

Spatial dimensions often get treated preferentially while others are neglected. This is usually deeply projected into the data model, hindering queries from seamlessly supporting semantic enrichment and probabilistic query semantics. Further, point cloud data is frequently integrated into existing systems as an extension, poorly integrated with their capabilities, diminishing their usefulness. For example, using space-filling curves because only binary tree indices are supported is a viable workaround, but it will ultimately become inefficient as the system fails to understand this concept and act upon it.

Supporting a wide variety of queries is necessary to empower point cloud database management systems to back up modern analytical retrieval needs. While the multitude of possible transformations of point clouds presents an obstacle, it also opens up an interesting view of a foundational multimodal data representation intertwined with respective indices for a dedicated system to support a wide range of use cases. Towards building reliable and robust systems that manage to address the presented and future needs, the following recommendations are made:

- Even though spatial attributes are prominent, allow equal treatment of all attributes for query support.
- Facilitate dynamic schema evolution to support semantic enrichment by adding attributes transparently and readily queryable.
- Support heterogeneous data representations and seamless transformations between them.
- Assume probabilistic and approximate query semantics as the default to incorporate scalability.
- Consider a multimodal data model tightly integrated with indexing capabilities.

A step in this direction has been explored in a retrieval system for point cloud data from vehicles integrating semantic information of points combined with sampling and approximate query execution (Li et al., 2025). Another system based on the lake-house architecture, incorporating ideas and recommendations presented in this paper, has been evaluated by exemplifying

probabilistic range queries for scalable visualization (Teuscher and Werner, 2025). Even for existing systems, these recommendations are interesting to consider for incremental improvements.

7. Conclusions

This work investigated point cloud search from an information retrieval perspective. Founded on emerging and novel information needs, a query taxonomy encompassing query modalities, query semantics, and result modalities was introduced that is generic over the manifold possible representations of point clouds. The gained insights allowed us to derive recommendations for the evolution of point cloud data management systems to serve the retrieval needs of contemporary analytical methods. Key features surfaced from this are the support for probabilistic query semantics, the possibility for extending points with additional attributes, their equal treatment compared to the spatial dimensions for query and indexing, and support for variable point cloud representations and seamless transformation between them. Following up on these, we envision manifesting a point cloud database management system for scalable point cloud analytics where the peculiarities of point clouds are an integral part instead of an afterthought towards semantic point cloud data retrieval with text queries expressed in natural language.

Acknowledgements

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 507196470 and 460036893.

References

- Alamri, S., Taniar, D., Safar, M., 2014. A taxonomy for moving object queries in spatial databases. *Future Generation Computer Systems*, 37, 232–242. <https://linkinghub.elsevier.com/retrieve/pii/S0167739X14000326>.
- Asad, R., Savva, M., 2023. 3dssr: 3d subscene retrieval. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition workshops*, IEEE, Piscataway, NJ, 2708–2716.
- Buckland, M., Plaunt, C., 1994. On the construction of selection systems. *Library Hi Tech*, 12(4), 15–28.
- Chang, J.-H. R., Chen, W.-Y., Ranjan, A., Yi, K. M., Tuzel, O., 2023. Pointersect: Neural Rendering with Cloud-Ray Intersection. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Vancouver, BC, Canada, 8359–8369.
- Dumic, E., Battisti, F., Carli, M., da Silva Cruz, L. A., 2020. Point cloud visualization methods: a study on subjective preferences. A. Marques, B. Hunyadi (eds), *28th European Signal Processing Conference (EUSIPCO 2020)*, IEEE, Piscataway, NJ, 595–599.
- Dumic, E., da Silva Cruz, L. A., 2025. Three-Dimensional Point Cloud Applications, Datasets, and Compression Methodologies for Remote Sensing: A Meta-Survey. *Sensors*, 25(6), 1660.

- Engelcke, M., Rao, D., Wang, D. Z., Tong, C. H., Posner, I., 2017. Vote3Deep: Fast object detection in 3D point clouds using efficient convolutional neural networks. *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, Singapore, Singapore, 1355–1361.
- Funkhouser, T., Min, P., Kazhdan, M., Chen, J., Halderman, A., Dobkin, D., Jacobs, D., 2003. A search engine for 3D models. *ACM Transactions on Graphics*, 22(1), 83–105. <https://dl.acm.org/doi/10.1145/588272.588279>.
- Goel, A., 2010. *Computer fundamentals*. Pearson Education India.
- Grabner, A., Roth, P. M., Lepetit, V., 2018. 3D Pose Estimation and 3D Model Retrieval for Objects in the Wild. *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Salt Lake City, UT, USA, 3022–3031.
- Kadam, P., Zhou, Q., Liu, S., Kuo, C.-C. J., 2022. Pcrp: Unsupervised point cloud object retrieval and pose estimation. G. Valenzise, T. Maugey (eds), *2022 IEEE International Conference on Image Processing*, IEEE, Piscataway, NJ, 1596–1600.
- Li, J., Yuan, H., Cong, G., Kiah, H. M., Zhang, S., 2025. MAST: Towards Efficient Analytical Query Processing on Point Cloud Data. *Proceedings of the ACM on Management of Data*, 3(1), 1–27. <https://dl.acm.org/doi/10.1145/3709702>.
- Lian, Z., Godil, A., Bustos, B., Daoudi, M., Hermans, J., Kawamura, S., Kurita, Y., Lavoué, G., Van Nguyen, H., Ohbuchi, R., Ohkita, Y., Ohishi, Y., Porikli, F., Reuter, M., Sipiran, I., Smeets, D., Suetens, P., Tabia, H., Vandermeulen, D., 2013. A comparison of methods for non-rigid 3D shape retrieval. *Pattern Recognition*, 46(1), 449–461. <https://linkinghub.elsevier.com/retrieve/pii/S0031320312003305>.
- Lokugam Hewage, C. N., Laefer, D. F., Vo, A.-V., LeKhac, N.-A., Bertolotto, M., 2022. Scalability and Performance of LiDAR Point Cloud Data Management Systems: A State-of-the-Art Review. *Remote Sensing*, 14(20), 5277. <https://www.mdpi.com/2072-4292/14/20/5277>.
- Lu, T., Wang, L., Wu, G., 2021. Cga-net: Category guided aggregation for point cloud semantic segmentation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Piscataway, NJ, 11688–11697.
- Manning, C. D., 2009. *Introduction to Information Retrieval*.
- Mirzaei, K., Arashpour, M., Asadi, E., Masoumi, H., Bai, Y., Behnood, A., 2022. 3D point cloud data processing with machine learning for construction and infrastructure applications: A comprehensive review. *Advanced Engineering Informatics*, 51, 101501.
- Qi, C. R., Hao, S., Mo, K., Guibas, L. J., 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, IEEE, Honolulu, HI, 77–85.
- Qi, C. R., Yi, L., Su, H., Guibas, L. J., 2017b. Pointnet++: deep hierarchical feature learning on point sets in a metric space. *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, Curran Associates Inc., Red Hook, NY, USA, 5105–5114.
- Remondino, F., 2003. From point cloud to surface: the modeling and visualization problem. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 34.
- Rijsbergen, C. J. V., 1979. *Information Retrieval*. 2nd edn, Butterworth-Heinemann, USA.
- Roussopoulos, N., Kelley, S., Vincent, F., 1995. Nearest neighbor queries. *Proceedings of the 1995 ACM SIGMOD international conference on Management of data - SIGMOD '95*, ACM Press, San Jose, California, United States, 71–79.
- Song, H., Wen, W., Wu, S., Guo, X., 2025. Comprehensive review on 3D point cloud segmentation in plants. *Artificial Intelligence in Agriculture*, 15(2), 296–315.
- Tangelder, J. W. H., Veltkamp, R. C., 2008. A survey of content based 3D shape retrieval methods. *Multimedia Tools and Applications*, 39(3), 441–471. <http://link.springer.com/10.1007/s11042-007-0181-0>.
- Taniar, D., Rahayu, W., 2013. A taxonomy for nearest neighbour queries in spatial databases. *Journal of Computer and System Sciences*, 79(7), 1017–1039. <https://linkinghub.elsevier.com/retrieve/pii/S0022000013000275>.
- Taniar, D., Rahayu, W., 2015. A taxonomy for region queries in spatial databases. *Journal of Computer and System Sciences*, 81(8), 1508–1531. <https://linkinghub.elsevier.com/retrieve/pii/S0022000014001895>.
- Teuscher, B., Werner, M., 2025. Point Cloud Data Management for Analytics in a Lakehouse. *AGILE: GIScience Series*, 6, 1–7. <https://agile-giss.copernicus.org/articles/6/47/2025/>.
- Toth, C., Jóźków, G., 2016. Remote Sensing Platforms and Sensors: A Survey. *ISPRS Journal of Photogrammetry and Remote Sensing*, 115, 22–36.
- Vos, S., Anders, K., Kuschnerus, M., Lindenbergh, R., Höfle, B., Aarninkhof, S., De Vries, S., 2022. A high-resolution 4D terrestrial laser scan dataset of the Kijkduin beach-dune system, The Netherlands. *Scientific Data*, 9(1), 191. <https://www.nature.com/articles/s41597-022-01291-9>.
- Williams, J., Scott, S., Wedig, S., Hindanov, T., Roedig, C., 2022. Data Management Challenges for Internet-scale 3D Search Engines. arXiv:2209.03913 [cs].
- Wu, W., Qi, Z., Fuxin, L., 2019. Pointconv: Deep convolutional networks on 3d point clouds. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, IEEE, Piscataway, NJ, 9613–9622.
- Yang, S., Hou, M., Li, S., 2023. Three-Dimensional Point Cloud Semantic Segmentation for Cultural Heritage: A Comprehensive Review. *Remote Sensing*, 15(3), 548.
- Zhao, J., Hua, X., Yang, J., Yin, L., Liu, Z., Wang, X., 2023. A review of point cloud segmentation of architectural cultural heritage. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, Copernicus GmbH, 247–254.