

Similarity joins and clustering for SPARQL

Sebastián Ferrada^{a,b,*}, Benjamin Bustos^b and Aidan Hogan^b

^a *Department of Computer and Information Science (IDA), Linköping University, Linköping, Sweden*
E-mail: sebastian.ferrada@liu.se

^b *Millenium Institute for Foundational Research on Data, Department of Computer Science, Universidad de Chile, Santiago, Chile*
E-mails: bebustos@dcc.uchile.cl, ahogan@dcc.uchile.cl

Editor: Marta Sabou, Technische Universität Wien (TUWien), Austria

Solicited reviews: Agnieszka Lawrynowicz, Poznan University of Technology, Poland; One anonymous reviewer

Abstract. The SPARQL standard provides operators to retrieve exact matches on data, such as graph patterns, filters and grouping. This work proposes and evaluates two new algebraic operators for SPARQL 1.1 that return similarity-based results instead of exact results. First, a similarity join operator is presented, which brings together similar mappings from two sets of solution mappings. Second, a clustering solution modifier is introduced, which instead of grouping solution mappings according to exact values, brings them together by using similarity criteria. For both cases, a variety of algorithms are proposed and analysed, and use-case queries that showcase the relevance and usefulness of the novel operators are presented. For similarity joins, experimental results are provided by comparing different physical operators over a set of real world queries, as well as comparing our implementation to the closest work found in the literature, DBSimJoin, a PostgreSQL extension that supports similarity joins. For clustering, synthetic queries are designed in order to measure the performance of the different algorithms implemented.

Keywords: Similarity joins, clustering, SPARQL

1. Introduction

Query languages such as SPARQL offer users the ability to match data using precise criteria. The results of evaluating the query contain *exact* matches that satisfy the specified criteria. This is useful, for example, to query for places with a population of more than one million (filters), for the names of famous musicians that fought in World War I (joins), or for the average life expectancy over all countries of each continent (grouping and aggregation). However, users may also require results that are based on imprecise criteria, such as similarity; for instance, obtaining countries with a population similar to Italy (similarity search), pairs of similar planetary systems from two different constellations (similarity join), or groups of similar chemical elements based on certain properties (clustering). The potential applications for efficient similarity-based queries in SPARQL are numerous, including: entity comparison and linking [35,40], multimedia retrieval [11,23], similarity graph management [8,12], pattern recognition [4], query relaxation [16], as well as domain-specific use-cases, such as protein similarity queries [2].

A similarity join $\bowtie_{\mathfrak{s}}$ operates over two sets of objects X, Y , bringing together all pairs $(x, y) \in X \times Y$ such that y is similar to x according to the given similarity criteria \mathfrak{s} (which may specify a distance measure, a threshold, the

* Corresponding author. E-mail: sebastian.ferrada@liu.se.

attributes to compare, etc.). Clustering operates over a single set of objects and partitions them into groups, called clusters, such that within a cluster, all members are similar to each other with respect to a given similarity criteria.

Similarity is usually measured through distance functions defined over a metric or vector space, where two objects are as similar as they are close in the given space [49]. Distances are usually defined over real, d -dimensional vector spaces (e.g., L_p distances); however, there are distances to compare strings (e.g., edit distance), sets (e.g., Jaccard distance), and many other types of objects. Two main types of similarity criteria are distinguished: range-based and nearest neighbours. Range-based similarity queries require a distance threshold r and assume that two objects $x \in X$, $y \in Y$ are similar if the distance among them is at most r . Nearest neighbour similarity queries require a natural k and state that x and y are similar if there are fewer than k other elements in Y with lower distance to x .

It must be noted that *some* similarity joins could be expressed in SPARQL 1.1; for example, a range-based similarity join that uses the Manhattan distance might be expressed using subqueries, numeric operations and filters. However, general similarity queries over metric spaces cannot be expressed using SPARQL, and those that can be expressed will typically be evaluated by computing all pair-wise distances and applying filters to each one, since SPARQL engines are not able to produce efficient planning and optimisations for such queries. Although several techniques for the efficient evaluation of similarity joins in metric spaces are available [6,17,33,48], the appropriate techniques depend on the similarity criteria to be used, where in order to efficiently compile vanilla SPARQL queries that express a similarity join into optimised physical operators that include these techniques, the engine would need to prove the equivalence of the query to a particular similarity join, which is not clear to be decidable.

Some works have proposed extending SPARQL with similarity features. However, these works have either focused on (1) unidimensional similarity measures that consider similarity with respect to one attribute at a time [20,44], or (2) domain-specific fixed-dimensional similarity measures, such as geographic distance [1,50]. Other approaches rather pre-compute and index similarity scores as part of the RDF graphs [8,16,30] or support metric distance measures external to a query engine [29,40].

In terms of clustering, Ławrynowicz [21] proposed an extension for SPARQL 1.0 that provides the solution modifier `CLUSTER BY`. This modifier receives several arguments, such as the clustering algorithm and its parameters; however their proposal was not implemented and it was not updated for SPARQL 1.1 so, for instance, it is unclear how the solution modifier works with the newly added modifiers such as `GROUP BY`. Thus, it is still unclear how the results of a SPARQL 1.1 query should be grouped with respect to similarity criteria, rather than exact criteria, and it is unclear what the performance cost of such an operator might be.

In this paper, two operators that enrich SPARQL with similarity-based queries are presented: similarity joins and clustering. Our work is, to our knowledge, the first to consider multidimensional similarity queries in the context of SPARQL, where the closest proposal to this one is DBSimJoin [43]: a PostgreSQL extension, which – though it lacks features arguably important for the RDF/SPARQL setting, namely nearest neighbours semantics – is considered as a baseline for experiments. Our work is also pioneer on the implementation and evaluation of clustering features on top of SPARQL with read-only capabilities (other works use SPARQL Update queries for some form of clustering [37]).

Motivating example To motivate combining similarity joins and clustering with SPARQL, we present some examples from the astronomy domain. Consider the RDF graph of Fig. 1, which describes, in Turtle syntax: the star TRAPPIST-1, including its type, constellation, mass (ronnagrames), radius (kilometres), temperature (Kelvin), and its orbiting planets; the exoplanet TRAPPIST-1B, including its type, mass (ronnagrames), radius (kilometres), temperature (Kelvin), orbital period (days), and semi-major axis of orbit (Astronomical Units); the Sun and the Earth, including analogous properties and units; and a class hierarchy of astronomic bodies. The example shown is based on Wikidata [45], using a fictional namespace and more readable identifiers for clarity.¹ Wikidata describes millions of astronomic bodies in this way. Combining similarity joins and clustering features with the broader SPARQL 1.1 query language enables the user to pose a wide range of queries over such data. Using similarity joins, the user can find, for example:

1. The planet orbiting TRAPPIST-1 most similar to Earth in terms of mass, radius, temperature, orbital period and semi-major axis.

¹We also directly standardise the units of measure. In practice, Wikidata allows for using a broad selection of different units, but the RDF export additionally provides reification properties that map these values to normalised values with the base S.I. unit of measure.

```

:TRAPPIST1 a :UltraCoolDwarfStar ; :constellation :Aquarius ; :mass 160000 ; radius 81396.9 ; :temperature 2550 ;
  :child :TRAPPIST1b , :TRAPPIST1c , :TRAPPIST1d , :TRAPPIST1e , :TRAPPIST1f , :TRAPPIST1g , :TRAPPIST1h .

:TRAPPIST1b a :Exoplanet ; :mass 8.244 ; :radius 7118 ; :temperature 400 ; :orbitalPeriod 1.51 ; semiMajorAxis 0.01154 .

:Sun a :GTypeStar ; :mass 1988550 ; :radius 696000 ; :temperature 15700000 ;
  :child :Mercury , :Venus , :Earth , :Mars , :Jupiter , :Saturn , :Uranus , :Neptune .

:Earth a :Planet ; :mass 5.972 ; :radius 6378 ; :temperature 288 ; :orbitalPeriod 1.00 ; semiMajorAxis 1.00 .

:Exoplanet rdfs:subClassOf :Planet .
:Planet rdfs:subClassOf :AstronomicBody .

:UltraCoolDwarfStar rdfs:subClassOf :DwarfStar .
:DwarfStar rdfs:subClassOf :Star .
:GTypeStar rdfs:subClassOf :MainSequenceStar .
:MainSequenceStar rdfs:subClassOf :Star .
:Star rdfs:subClassOf :AstronomicBody .

```

(a) RDF graph in Turtle syntax describing two stars, two planets, and an astronomic class hierarchy

```

SELECT ?planet ?dist WHERE {
  { :Earth :mass ?m1 ; :radius ?r1 ; :temperature ?t1 ; :orbitalPeriod ?o1 ; :semiMajorAxis ?s1 }
  SIMILARITY JOIN ON (?m1 ?r1 ?t1 ?o1 ?s1) (?m2 ?r2 ?t2 ?o2 ?s2) TOP 1 DISTANCE sim:mantattan AS ?dist
  { :TRAPPIST1 :child ?planet . ?planet :mass ?m2 ; :radius ?r2 ; :temperature ?t2 ; :orbitalPeriod ?o2 ; :semiMajorAxis ?s2 } }

```

(b) Query 1: Most similar planet to Earth orbiting the star TRAPPIST-1

```

SELECT ?s1 ?s2 ?dist WHERE {
  { SELECT ?s1 ?m1 ?r1 ?t1 (COUNT(DISTINCT ?c) AS ?c1)
    WHERE { ?s1 rdf:type/rdfs:subClassOf* :Star ; :constellation :Aquarius ;
      :mass ?m1 ; :radius ?r1 ; :temperature ?t1 ; :child ?c . }
    GROUP BY ?s1 ?m1 ?r1 ?t1 }
  SIMILARITY JOIN ON (?m1 ?r1 ?t1 ?c1) (?m2 ?r2 ?t2 ?c2) TOP 1 DISTANCE sim:mantattan AS ?dist
  { SELECT ?s2 ?m2 ?r2 ?t2 (COUNT(DISTINCT ?c) AS ?c2)
    WHERE { ?s2 rdf:type/rdfs:subClassOf* :Star ; :constellation :Orion ;
      :mass ?m2 ; :radius ?r2 ; :temperature ?t2 ; :child ?c . }
    GROUP BY ?s2 ?m2 ?r2 ?t2 } }

```

(c) Query 2: Most similar star in Orion for each star in Aquarius

```

SELECT ?star ?c WHERE {
  { SELECT ?s ?m ?r ?t (COUNT(DISTINCT ?p) AS ?n)
    WHERE { ?star rdf:type/rdfs:subClassOf* :Star ; :mass ?m ; :radius ?r ; :temperature ?t ; :child ?p . } }
  CLUSTER BY (?m ?r ?t ?n) WITH sim:kmeans(4) AS ?c }

```

(d) Query 3: Clusters of similar stars

```

SELECT ?planet ?dist ?con WHERE {
  { :Earth :mass ?m1 ; :radius ?r1 ; :temperature ?t1 }
  SIMILARITY JOIN ON (?m1 ?r1 ?t1) (?m2 ?r2 ?t2) TOP 5 DISTANCE sim:mantattan AS ?dist
  { { SELECT ?s ?c WHERE { ?s rdf:type/rdfs:subClassOf* :Star ; :mass ?m ; :radius ?r ; :temperature ?t . }
    CLUSTER BY (?m ?r ?t) WITH sim:kmeans(4) AS ?c }
  { SELECT * WHERE {
    { SELECT ?c WHERE { ?s rdf:type/rdfs:subClassOf* :Star ; :mass ?m ; :radius ?r ; :temperature ?t . }
      CLUSTER BY (?m ?r ?t) WITH sim:kmeans(4) AS ?c } FILTER(?s == :Sun) } }
    ?s :child ?planet . ?planet :mass ?m2 ; :radius ?r2 ; :temperature ?t2 . }
  ?planet (^:child|:parent)*:/constellation ?con }

```

(e) Query 4: Planets similar to Earth whose star is also in the same cluster as the Sun, returning also their constellation

Fig. 1. Examples of queries using similarity joins and clustering in the astronomy domain.

2. The most similar star in the constellation Orion for each star in the constellation Aquarius, based on mass, radius, temperature and number of planets.

Using clustering, the user can find, for example:

3. Clusters of stars based on mass, radius, temperature and number of planets.

Combining both similarity joins and clustering, the user can find:

4. The top 5 most similar planets to the Earth based on their mass, radius and temperature whose stars are in the same cluster as the Sun based on mass, radius and temperature, along with each planet's constellation.

Figure 1b–1e shows how these queries are expressed in the concrete syntax we will propose and describe in detail later in the paper. Of note is that we can use SPARQL's other features to apply similarity joins and clustering on computed attributes of an entity; for example, the second and third queries reference the number of planets, and although this numeric attribute is not explicitly given in the data, we can compute it in SPARQL using a count aggregation. We can likewise cluster planets based on properties of their stars (e.g., their temperature), based on attributes calculated from explicit values (e.g. the density calculated from mass and radius), etc. Furthermore, we can use SPARQL's features to retrieve additional details of results returned by similarity joins and clustering, such as the constellations of planets returned by the fourth query, retrieved via a property path.

Later we will present further examples of queries that are enabled by these features in the context of other domains, including chemistry, politics, demography, and multimedia.

Contributions This paper is an extension of our previous work, where we proposed, defined, implemented and evaluated novel types of similarity joins in SPARQL [9]. In this extended version, we discuss new use-case queries over different RDF data, and we expand the evaluation of the similarity join implementations by building a benchmark based on the query logs of Wikidata. Furthermore, we extend the proposal with a novel clustering solution modifier; we describe its syntax, semantics and implementation, and we evaluate its performance. Extended discussion and examples are also provided.

Paper outline In Section 2 we present definitions regarding RDF, the SPARQL algebra and syntax, and similarity-based operations, and in Section 3, related work. The similarity join operator, its properties, and implementation are introduced in Section 4. Section 5 is dedicated to the clustering solution modifier: its definition, a proposed syntax, the implemented clustering algorithms, and use cases. Evaluation of the similarity join operator and the clustering solution modifier is presented and discussed in Section 6. Finally, Section 7 presents our concluding remarks.

2. Preliminaries

In this section, the relevant definitions and notation that is used throughout this paper regarding RDF, SPARQL and Similarity Operations are presented.

2.1. RDF and SPARQL

RDF is the standard graph-based data model for the Semantic Web. *RDF terms* can be either IRIs (**I**), literal values (**L**) or blank nodes (**B**). An *RDF triple* is a tuple $(s, p, o) \in \mathbf{IB} \times \mathbf{I} \times \mathbf{IBL}$,² where s is referred to as the *subject*, p as the *predicate*, and o as the *object*. A finite set of RDF triples makes an *RDF graph*.

SPARQL is the W3C standard query language for RDF [15]. The semantics of SPARQL operators are introduced by Pérez et al. [34], in terms of their evaluation over an RDF graph. The result of the evaluation of a SPARQL query is a set of *solution mappings*. Let \mathbf{V} denote a set of query variables disjoint with \mathbf{IBL} , a solution mapping is a partial function $\mu : \mathbf{V} \rightarrow \mathbf{IBL}$ defined for a set of variables, called its *domain*, $\text{dom}(\mu)$. Two solution mappings μ_1, μ_2 are

²Using that $\mathbf{IB} = \mathbf{I} \cup \mathbf{B}$.

compatible, noted as $\mu_1 \sim \mu_2$, if for all $v \in \text{dom}(\mu_1) \cap \text{dom}(\mu_2)$, $\mu_1(v) = \mu_2(v)$. Letting X and Y denote sets of solution mappings, the SPARQL algebra is defined as follows:

$$X \bowtie Y := \{\mu_1 \cup \mu_2 \mid \mu_1 \in X \wedge \mu_2 \in Y \wedge \mu_1 \sim \mu_2\} \quad (1)$$

$$X \cup Y := \{\mu \mid \mu \in X \vee \mu \in Y\} \quad (2)$$

$$\sigma_f(X) := \{\mu \in X \mid f(\mu) = \text{true}\} \quad (3)$$

$$X \triangleright Y := \{\mu \in X \mid \nexists \mu' \in Y : \mu \sim \mu'\} \quad (4)$$

$$X \bowtie Y := (X \bowtie Y) \cup (X \triangleright Y) \quad (5)$$

$$\pi_V(X) := \{\mu' \mid \exists \mu \in X : \mu' \subseteq \mu, \text{dom}(\mu') = V \cap \text{dom}(\mu)\} \quad (6)$$

$$\gamma_V(X) := \{(\mu, X') \mid \mu \in \pi_V(X), X' = \{\mu' \in X \mid \pi_V(\mu') = \mu\}\} \quad (7)$$

where f is a filter condition that, given a solution mapping, returns true, false, unbound or an error.

We denote the evaluation of a query Q over an RDF graph G as $Q(G)$. Before defining $Q(G)$, first let $t = (s, p, o)$ denote a triple pattern, such that $(s, p, o) \in \mathbf{VIL} \times \mathbf{VI} \times \mathbf{VIBL}$; then by $\text{vars}(t)$ we denote the set of variables appearing in t and by $\mu(t)$ we denote the image of t under a solution mapping μ . Finally, we can define $Q(G)$ recursively as follows:

$$t(G) := \{\mu \mid \mu(t) \in G, \text{dom}(\mu) = \text{vars}(t)\} \quad (8)$$

$$[Q_1 \text{ AND } Q_2](G) := Q_1(G) \bowtie Q_2(G) \quad (9)$$

$$[Q_1 \text{ UNION } Q_2](G) := Q_1(G) \cup Q_2(G) \quad (10)$$

$$[Q_1 \text{ OPTIONAL } Q_2](G) := Q_1(G) \bowtie Q_2(G) \quad (11)$$

$$\text{FILTER}_f(Q)(G) := \sigma_f(Q(G)) \quad (12)$$

$$\text{SELECT}_V(Q)(G) := \pi_V(Q(G)) \quad (13)$$

$$\text{GROUP}_V(Q)(G) := \gamma_V(Q(G)) \quad (14)$$

2.2. Similarity search

In this section, we introduce the concept of similarity search over a universe of objects, as well as different similarity search operations.

Definition 2.1 (Similarity Search). Let \mathbf{U} be the universe of objects, and $\mathbf{D} \subseteq \mathbf{U}$ a dataset. Given a query object $q \in \mathbf{U}$, similarity search obtains all elements $x \in \mathbf{D}$ such that x is similar to q .

Mathematically, similarity is modelled as a score. Let $s_{x,y} \in [0, 1]$ be the *similarity score* between objects x and y , where the higher the score, the more similar x and y are. Similarity scores are often measured in terms of distances within a given space; two objects being more similar if the distance among them is small. In this work, we consider a class of distance functions called metrics:

Definition 2.2 (Metric Distance). A distance function $\delta : \mathbf{U} \times \mathbf{U} \rightarrow \mathbb{R}$ is called a *metric* if the following properties hold for all $x, y, z \in \mathbf{U}$:

1. $\delta(x, y) \geq 0$ (positivity).
2. $\delta(x, y) = 0 \iff x = y$.
3. $\delta(x, y) = \delta(y, x)$ (symmetry).
4. $\delta(x, y) \leq \delta(x, z) + \delta(z, y)$ (triangle inequality).

Examples of metric distances are the Euclidean distance over real vectors, the edit distance over strings, and the Jaccard distance over sets.

Metric distances induce a topology over the set they are defined upon:

Definition 2.3 (Metric Space). A *metric space* is a tuple (\mathbf{U}, δ) where $\mathbf{U} \neq \emptyset$, and δ is a metric distance function defined over pairs of elements of \mathbf{U} .

There are two main types of similarity search queries in metric spaces: range and nearest neighbours queries.

Definition 2.4 (Range Similarity Search). Let $\mathbf{D} \subseteq \mathbf{U}$ be a dataset, and $q \in \mathbf{U}$ be a query object. The r -range similarity search, for $r \in \mathbb{R}$ is defined as:

$$Q_{\text{range}}(\mathbf{D}, q, r) = \{x | x \in \mathbf{D} \text{ and } \delta(x, q) \leq r\} \quad (15)$$

Definition 2.5 (Nearest Neighbours Similarity Search). Let $\mathbf{D} \subseteq \mathbf{U}$ be a dataset, and $q \in \mathbf{U}$ be a query object. The k -nearest neighbours similarity search, for $k \in \mathbb{N}$ is defined as:

$$Q_{\text{nn}}(\mathbf{D}, q, k) = \{x | x \in \mathbf{D} \text{ and } \kappa_{\delta}^{\mathbf{D}}(x, q) \leq k\} \quad (16)$$

where $\kappa_{\delta}^{\mathbf{D}}(x, q)$ is the rank of q in \mathbf{D} w.r.t. its distance to x (in ascending order). Note that two elements of \mathbf{D} at the same distance to x shall have the same rank.

If instead of a single query object q we consider a set of query objects, we call this operation a similarity join. A similarity join is an operation defined over two *comparable* sets. Two sets are deemed comparable if a (metric) distance function can be defined over their Cartesian product.

Definition 2.6 (Similarity Join). Let \mathbf{D}_1 and \mathbf{D}_2 be subsets of a universe of objects \mathbf{U} . Their *similarity join* is defined as follows:

$$Q_{\text{sim}}(\mathbf{D}_1, \mathbf{D}_2, \mathfrak{s}) = \{(x, y) | x \in \mathbf{D}_1, y \in \mathbf{D}_2 \text{ and } x \text{ is similar to } y \text{ according to sim under } \mathfrak{s}\} \quad (17)$$

where \mathfrak{s} refers to parameters for similarity criteria, which can be a range-based or k -nearest neighbours criterion, based on a (metric) distance δ , thus defining the following two operations:

$$Q_{\text{range}}(\mathbf{D}_1, \mathbf{D}_2, r) := \{(x, y) | x \in \mathbf{D}_1, y \in \mathbf{D}_2 \text{ and } \delta(x, y) \leq r\} \quad (18)$$

$$Q_{\text{nn}}(\mathbf{D}_1, \mathbf{D}_2, k) := \{(x, y) | x \in \mathbf{D}_1, y \in \mathbf{D}_2 \text{ and } \kappa_{\delta}^{\mathbf{D}_2}(x, y) \leq k\} \quad (19)$$

Note that similarity search – also known as *query-by-example* – is a special case of similarity join where $\mathbf{D}_2 = \{y\}$. We may also refer to the special case of *self similarity joins* where $\mathbf{D}_1 = \mathbf{D}_2$.

3. Related work

This section is divided in four parts: the first presents techniques to evaluate similarity joins; the second introduces works about the similarity operations available in query languages or database engines; the third is dedicated to clustering techniques; and the fourth to the inclusion of clustering in query languages.

3.1. Similarity joins

The brute force method for computing a similarity join between \mathbf{D}_1 and \mathbf{D}_2 is called a *nested loop*, which computes for each $x \in \mathbf{D}_1$ the distance to each $y \in \mathbf{D}_2$, returning pairs that satisfy the similarity condition, thus performing $|\mathbf{D}_1| \cdot |\mathbf{D}_2|$ distance computations (or $|\mathbf{D}_1| \cdot |\mathbf{D}_2|/2$ if it is a symmetric distance). For similarity joins over metric spaces, three strategies improve upon the brute force method: *indexing*, *space partitioning*, and/or *approximation*.

A common way to optimise similarity joins is to index the data using tree structures that divide the space in different ways (offline), then pruning distant pairs of objects from comparison (online). Several tree structures have been proposed as indexes for optimising similarity queries [14,31,41]. Among such approaches, we can find *vantage-point Trees* (*vp-Trees*) [48], which make recursive ball cuts of space centred on selected points, attempting to evenly distribute objects inside and outside the ball. *vp-Trees* have an average-case query-by-example search time of $O(n^\alpha)$ on n objects, where $0 \leq \alpha \leq 1$ depends on the distance distribution and dimensionality of the space, among other factors [27], thus having an upper bound of $O(n^{2\alpha})$ for a similarity join. Other tree indexes, such as the D-Index [6] and the List of Twin Clusters [33], propose to use clustering techniques over the data.

Other space partitioning algorithms are not used for indexing but rather for evaluating similarity joins online. The Quickjoin (QJ) algorithm [17] was designed to improve upon grid-based partition algorithms [3,5]; it divides the space into ball cuts using random data objects as pivots, splitting the data into the vectors inside and outside the ball, proceeding recursively until the groups are small enough to perform a nested loop. It keeps window partitions in the boundaries of the ball in case there are pairs needed for the result with vectors assigned to different partitions. QJ requires $O(n(1+w)^{\lceil \log n \rceil})$ distance computations, where w is the average fraction of elements within the window partitions. QJ was intended for range-based similarity joins and extending QJ to compute a k -nn similarity join appears far from trivial, since its simulation with a range-based join would force most of the data to fall within the window partitions, thus meaning that QJ will reach its quadratic worst case.

Another alternative is to apply approximations to evaluate the similarity joins, trading the precision of the results for more efficient computation. FLANN [26] is a library that provides several approximate k -nn algorithms based on randomised k - d -forests, k -means trees, locality-sensitive hashing, etc.; it automatically selects the best algorithm to index based, for example, on a target precision, which can be traded-off to improve execution time.

3.2. Similarity in databases

Though similarity joins do not form part of standard query languages, such as SQL or SPARQL, a number of systems have integrated variations of such joins within databases. In the context of SQL, DBSimJoin [43] implements a range-based similarity join operator for PostgreSQL. This implementation claims to handle any metric space, thus supporting various metric distances; it is based on the aforementioned index-free Quickjoin algorithm.

In the more specific context of the Semantic Web, a number of works have proposed online computation of similarity joins in the context of domain-specific measures. Zhai et. al [50] use OWL to describe the spatial information of a map of a Chinese city, enabling geospatial SPARQL queries that include the computation of distances between places. The Parliament SPARQL engine [1] implements an OGC standard called GeoSPARQL, which aside from various geometric operators, also includes geospatial distance. Works on link discovery may also consider specific forms of similarity measures [40], often string similarity measures over labels and descriptions [44].

Other approaches pre-materialise distance values that can then be incorporated into SPARQL queries. IMGpedia [8] pre-computes a k -nn self-similarity join offline over images and stores the results as part of the graph. Similarity measures have also been investigated for the purposes of SPARQL query relaxation, whereby, in cases where a precise query returns no or few results, relaxation finds queries returning similar results to those sought [16,30].

Galvin et al. [10] propose a multiway similarity join operator for RDF; however, the notion of similarity considered is based on semantic similarity that tries to match different terms referring to the same real-world entity. Closer to this work lies iSPARQL [20], which extends SPARQL with IMPRECISE clauses that can include similarity joins on individual attributes. A variety of distance measures are proposed for individual dimensions/attributes, along with aggregators for combining dimensions. However, in terms of evaluation, distances are computed in an attribute-at-a-time manner and input into an aggregator. For the multidimensional setting, a (brute-force) nested loop needs to be performed; the authors leave optimisations in the multidimensional setting for future work [20].

3.3. Clustering

Clustering is a process through which a set of objects can be automatically partitioned into groups – called *clusters* – such that the distance between objects of a cluster is minimised, and the distance between objects in different clusters is maximised. Clustering algorithms can be roughly divided into distance-based, density-based, grid-based,

distribution-based, and can also be hierarchical. A prime example of distance-based clustering algorithms is *k*-means [24], which first randomly selects *k* centroids and assigns each object to the cluster defined by the closest centroid; it then recomputes new centroids as the mean objects of each current cluster (which may not be a part of the set of objects), and reassigns the objects to the clusters defined by these new centroids until the centroids converge. The *k*-medoids algorithm is similar to *k*-means [19], but instead of centroids, it defines *medoids*: the element of a cluster such that its distance to every other element in the cluster is minimal. As for density-based algorithms, DBSCAN [7] groups together objects in dense regions (every object has several neighbours), separated by low density, sparse regions. DBSCAN also marks isolated objects as outliers, which are not included in any cluster. Grid-based methods, such as STING [46] and GridDBSCAN [25] divide the space into cells, aiming to quickly find dense areas to produce the clusters, and prune sparse areas as noise. However, these methods have been proven to not scale well with increasing dimensions [53]. Distribution-based methods assume that objects in the same cluster have a similar probability distribution: DBCLASD [47] assumes that the points inside a cluster are uniformly distributed. Finally, hierarchy-based approaches generate clusters within clusters forming a hierarchy that can be consumed at various levels. Hierarchical clustering methods include BIRCH [52], CURE [13], and Chameleon [18].

3.4. Clustering in databases

In databases, Ordonez and Omiecinski [32] provide an efficient disk-based *k*-means implementation for RDBMS using SQL INSERT and UPDATE queries. Also in SQL, Zhang and Huang [51] propose to extend the query language with a CLUSTER BY keyword, however, they do not provide an implementation. Li et al. [22] propose to extend SQL with fuzzy criteria for grouping and ordering, which they call clustering and ranking; they provide syntax and semantics, as well as a C++ implementation of the ranking and clustering algorithms based on so-called summary grids that provide a more time-efficient evaluation. Silva et al. [42] propose a similarity based GROUP BY that clusters the query results in three ways: 1) by defining in the query a maximum element separation or maximum cluster diameter, 2) by defining the centroids around which the data is clustered to, and 3) by defining a set of limit or threshold values. They provide a PostgreSQL implementation of these three clustering options.

Qi et al. [37] provide a *k*-means implementation for RDF data by using SPARQL Update queries. Ławrynowicz [21], proposed an extension to SPARQL 1.0, that introduces a CLUSTER BY solution modifier along with a clear syntax and semantics; however, they do not include an implementation nor evaluation, and it is not immediately clear how their proposal could be aligned with SPARQL 1.1 grouping and aggregates.

4. Similarity joins in SPARQL

In this section, the similarity join operator for SPARQL is presented. Along with the algebraic definition of the operator and its properties, we will present options for its implementation and optimisation, and further introduce use-case queries to illustrate its application. We first define the following general criteria that a novel similarity join operator for SPARQL 1.1 should satisfy:

CLOSURE: the new operator should be freely combinable with other available operators.

EXTENSIBILITY: the similarity join should allow the user to specify the required dimensions and metrics to be used in the query, and define new metrics as needed.

ROBUSTNESS: the new operator should make as few assumptions as possible about the input or the data, regarding completeness, comparability, etc.

USABILITY: the new operator should be easy for users to adopt, offering high expressiveness while avoiding unnecessary verbosity.

With respect to CLOSURE, the similarity join will be defined in a manner analogous to other forms of joins that combine graph patterns in the WHERE clause of a SPARQL query; furthermore, the computed distance measure is allowed to be bound to a variable, facilitating its use beyond the similarity join (e.g., for solution modifiers such as ORDER BY, for aggregates such as MAX, MIN, AVG, for expressions within FILTER or BIND, etc). With respect to EXTENSIBILITY, rather than assume one metric, the type of distance metric used is explicit in the semantics and


```

SimilarityGraphPattern ::= 'SIMILARITY JOIN ON (' Var+ ') (' Var+ ')
                        ( 'TOP' INTEGER | 'WITHIN' DECIMAL ) 'DISTANCE' iri 'AS' Var
                        GroupGraphPattern
GraphPatternNotTriples ::= GroupOrUnionGraphPattern | [...] | SimilarityGraphPattern

```

Fig. 2. The extended production rules from SPARQL 1.1 to support similarity joins.

syntax, allowing other types of distance metrics to be used in future. Regarding ROBUSTNESS, the precedent of SPARQL's error-handling when dealing with incompatible types or unbound values is followed. Finally, regarding USABILITY, syntactic features for both range-based similarity joins and k -nn similarity joins are supported, noting that specifying particular distances for ranges can be unintuitive in abstract, high-dimensional metric spaces.

4.1. Syntax

To define the syntax for a similarity join, it is necessary to consider the syntax of other binary operators such as OPTIONAL or MINUS [15], as well as the required extra parameters: the variables representing the dimensions with respect to which the distance is to be computed, the distance function to be used, a fresh variable to bind the computed distances to, and the similarity parameter corresponding to a threshold distance (for range similarity joins) or to a threshold number of nearest neighbours (for k -nn similarity joins).

The syntax for the similarity join operator extends the SPARQL 1.1 EBNF grammar [15] by adding one production rule for `SimilarityGraphPattern`, and extending the rule for `GraphPatternNotTriples`. All other definitions are left unchanged. The exact changes to the grammar can be found in Fig. 2.

The keyword ON is used to define the variables in both graph patterns upon which the distance is computed; the keywords TOP and WITHIN denote a k -nn query and an r -range query respectively; the keyword DISTANCE specifies the IRI of the distance function to be used for the evaluation of the join, whose result will be bound to the variable indicated with AS, which is expected to be *fresh*, i.e., to not appear elsewhere in the `SimilarityGraphPattern`, similarly to BIND. The syntax may be extended in the future to provide further customisation, such as supporting different normalisation functions, or to define default parameters.

Depending on the metric, it could be possible, in principle, to express such queries as plain SPARQL 1.1 queries, by taking advantage of features such as variable binding, numeric expressions, sub-selects, etc. However, there are two key advantages of the dedicated syntax: (1) similarity join queries in vanilla syntax would be complex to express, particularly in the case of k -nn queries or metrics without the corresponding numeric operators in SPARQL; (2) optimising queries written in the vanilla syntax (beyond nested-loop performance) would be practically infeasible, requiring an engine that can prove equivalence between the distance metrics and semantics for which similarity join algorithms are optimised and the plethora of ways in which they can be expressed in vanilla syntax. This is why we propose to make similarity joins for multidimensional distances a first class feature, with dedicated syntax and physical operators offering sub-quadratic performance in average/typical cases.

4.2. Semantics

We now introduce some definitions to extend the SPARQL semantics defined in Section 2 to cover the similarity joins whose syntax is defined in Section 4.1. These definitions adapt the standard definitions of similarity joins as presented in Section 2.2 to make them compatible with the relational algebra underlying SPARQL, i.e., to perform an equi-join on shared variables, and a similarity join on other variables selected by the user.

We start with the definition of a *similarity join expression*, which indicates the criteria under which the similarity of two SPARQL solution mappings can be computed.

Definition 4.1 (Similarity Join Expression). A *similarity join expression* is a 4-tuple $\mathfrak{s} := (\mathcal{V}, \delta, v, \phi)$, where $\mathcal{V} \subseteq \mathbf{V} \times \mathbf{V}$ contains pairs of variables to be compared; δ is a distance metric that accepts a set of pairs of RDF terms and returns a value in $[0, \infty)$ or an error (interpreted as ∞) for incomparable inputs; $v \in \mathbf{V}$ is a fresh variable to which distances shall be bound; and $\phi \in \{\text{rg}_r, \text{nn}_k\}$ is a filter expression based on r -range or k -nn criteria.

In the following, we present several definitions that build up to the formal specification of our similarity join operators. We begin with a definition that, given two compatible solutions from either side of the similarity join μ_1 and μ_2 , and given a set of paired variables \mathcal{V} , produces a set of pairs of constants projected from the solutions over which distances can be computed.

Definition 4.2. Given two compatible solution mappings $\mu_1 \sim \mu_2$, we define the set of mapping pairs to be compared in the similarity join as:

$$\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1} := \{((\mu_1 \cup \mu_2)(x), (\mu_1 \cup \mu_2)(y)) \mid (x, y) \in \mathcal{V}\} \quad (20)$$

Note that per this definition, $\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1} = \llbracket \mathcal{V} \rrbracket_{\mu_1}^{\mu_2}$.

Next we define some useful notation for the result of a natural equi-join (applying a join with equality on shared variables, i.e., based on compatible solutions) that forms the basis of our similarity join operator.

We denote by $[\mu_1, \dots, \mu_n] \in X_1 \bowtie \dots \bowtie X_n$ the condition that $\mu_1 \in X_1, \dots, \mu_n \in X_n$, and $\mu_i \sim \mu_j$ (for $1 \leq i \leq n, 1 \leq j \leq n$). Thus $[\mu_1, \dots, \mu_n]$ denotes a solution mapping $\bigcup_{i=1}^n \mu_i$ for the natural equi-join $X_1 \bowtie \dots \bowtie X_n$.

Given a set of solutions, we now define notation for the union and intersection of the domains of the solutions (the sets of variables for which a solution is defined).

Definition 4.3. Given a set of solution mappings X we define the union and intersection domains as:

$$\text{udom}(X) := \bigcup_{\mu \in X} \text{dom}(\mu) \quad (21)$$

$$\text{idom}(X) := \bigcap_{\mu \in X} \text{dom}(\mu) \quad (22)$$

In other words, $\text{udom}(X)$ is the set of variables that are defined by some solution mapping $\mu \in X$, while $\text{idom}(X)$ is the set of variables that are defined for all solution mappings $\mu \in X$.

The next definition allows us to define a solution mapping that maps specific variables to specific SPARQL terms (similar to BIND). This will be used to bind the distance value to a specified (fresh) variable.

Definition 4.4. Letting $\{v_1, \dots, v_n\} \subseteq \mathbf{V}$ denote a set of variables, and $\{d_1, \dots, d_n\} \subseteq \mathbf{IBL}$ denote a set of RDF terms, the expression $\{v_1 \rightarrow d_1, \dots, v_n \rightarrow d_n\}$ denotes a mapping μ such that $\text{dom}(\mu) = \{v_1, \dots, v_n\}$ and $\mu(v_i) = d_i$ for all $1 \leq i \leq n$.

With all these definitions in hand, we introduce the similarity join operators in SPARQL as follows:

Definition 4.5 (Range Similarity Joins in SPARQL). Given two sets of solution mappings X and Y and a similarity join expression $\mathfrak{s} = (\mathcal{V}, \delta, v, \text{rg}_r)$ where r is the distance threshold to be used, the evaluation of an r -range similarity join is defined as:

$$X \bowtie_{(\mathcal{V}, \delta, v, \text{rg}_r)} Y := \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \mu_v(v) \leq r\} \quad (23)$$

when $v \notin \text{udom}(X) \cup \text{udom}(Y)$ or `error` otherwise.

Definition 4.6 (k -nn Similarity Joins in SPARQL). Given two sets of solution mappings X and Y and a similarity join expression $\mathfrak{s} = (\mathcal{V}, \delta, v, \text{nn}_k)$, where k is the number of nearest neighbours to be obtained, the evaluation of a k -nn similarity join is defined as:

$$X \bowtie_{(\mathcal{V}, \delta, v, \text{nn}_k)} Y := \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_\delta^Y(\mu_1, \mu_2) \leq k\} \quad (24)$$

when $v \notin \text{udom}(X) \cup \text{udom}(Y)$ or `error` otherwise, and where $\kappa_\delta^Y(\mu_1, \mu_2)$ is the rank of μ_2 in the sequence of solution mappings $\mu \in Y$ sorted by $\delta(\mu_1, \mu)$ in an ascending order. Note that two mappings at the same distance shall have the same rank.

Per Definition 4.6, more than k results can be returned for μ_1 in the rare case of ties in distance, which keeps the semantics of the k -nn similarity joins deterministic.

An error is returned when $v \in \text{udom}(X) \cup \text{udom}(Y)$ to emulate a similar behaviour to BIND in SPARQL.

Bag semantics are defined for similarity joins in the natural way, where the multiplicity of $\mu \in X \bowtie_{\mathfrak{s}} Y$ is defined to be the product of the multiplicities of the solutions $\mu_1 \in X$ and $\mu_2 \in Y$ that produce it.

Next, given the semantics of the similarity join operators, we present several algebraic properties for the similarity join and its relationship with other SPARQL operators. These properties and equivalences are useful for query planning and optimisation. For instance, having a commutative and associative operator can allow the Query Planner to choose the best possible order in which to evaluate a given sequence of these operators; similarly, having equivalences between operators under certain circumstances can allow for query rewriting and further optimisation. Particularly for the case of similarity joins, some rewriting rules can help to reduce the sets of intermediate solutions considerably before applying the similarity join, and thus reduce query execution time. The proofs for each of the claims along with some illustrative examples can be found in the Appendix:

1. \bowtie_{τ} is commutative and distributive over \cup .
2. \bowtie_{n} is not commutative nor is it distributive over \cup .
3. \bowtie_{n} is right-distributive over \cup .
4. $(X \bowtie_{\mathfrak{s}} Y) \bowtie_{\mathfrak{s}'} Z \not\equiv X \bowtie_{\mathfrak{s}} (Y \bowtie_{\mathfrak{s}'} Z)$ holds, i.e., $\bowtie_{\mathfrak{s}}$ is not associative.
5. Let $\mathfrak{s} = (\mathcal{V}, \delta, v, \phi)$. If each mapping in $X \bowtie Y$ binds all variables in \mathcal{V} and $v \notin \text{udom}(X) \cup \text{udom}(Y) \cup \text{udom}(Z)$, then the following hold:
 - (a) $(X \bowtie_{\mathfrak{s}} Y) \bowtie Z \equiv (X \bowtie Z) \bowtie_{\mathfrak{s}} Y$
 - (b) $(X \bowtie_{\mathfrak{s}} Y) \triangleright Z \equiv (X \triangleright Z) \bowtie_{\mathfrak{s}} Y$ if $\text{udom}(Z) \cap (\text{udom}(Y) - \text{idom}(X)) = \emptyset$
 - (c) $(X \bowtie_{\mathfrak{s}} Y) \bowtie_{\mathfrak{s}'} Z \equiv (X \bowtie_{\mathfrak{s}'} Z) \bowtie_{\mathfrak{s}} Y$ if $\text{udom}(Z) \cap (\text{udom}(Y) - \text{idom}(X)) = \emptyset$
 - (d) $\sigma_f(X \bowtie_{\mathfrak{s}} Y) \equiv \sigma_f(X) \bowtie_{\mathfrak{s}} Y$ if f is scoped to $\text{idom}(X)$.

A key finding here is that properties that hold for (equi) joins in SPARQL do not necessarily hold for similarity joins (commutativity, associativity, distributivity), and these equivalences show how and why this is the case. This means that common optimisations applied in SPARQL engines (such as join reordering) cannot be applied *as is* for similarity joins, and thus there are interesting open challenges on how to optimise these queries.

4.3. Implementation

The implementation of the similarity join operators extends ARQ – the SPARQL engine of Apache Jena³ – which indexes an RDF dataset and receives as input a (similarity) query in the syntax discussed in Section 4.1. The steps of the evaluation of an extended SPARQL query follow a standard flow, namely Parsing, Algebra Optimisation, Algebra Execution, and Result Iteration. The implementation is available at <https://github.com/scferrada/jena>.

The Parsing stage receives the query string written by a user, and outputs the algebraic representation of the similarity query. Parsing is implemented by extending the Parser of Jena using JavaCC,⁴ wherein the new keywords and syntax rules are defined, as introduced in Section 4.1. The result of the Parsing stage is an abstract syntax tree of algebraic operators, e.g., `knnsimjoin(leftjoin(...), triplepattern(...))`.

In terms of physical operators to evaluate the similarity joins, initially we considered applying offline indexing of literal values in the RDF graph in order to allow lookups of similar values. However, we discarded this option for the following two reasons:

- Indexing approaches typically consider a fixed metric space, with a fixed set of attributes specified for the distance measure (e.g., longitude and latitude). Conversely, in our scenario, the user can choose the attributes according to which similarity will be computed for a given query. In the setting of an RDF graph such as DBpedia or Wikidata, the user can choose from thousands of attributes. Indexing all possible combinations would lead to an exponential blow-up and would be prohibitively costly in terms of both time and space.

³<http://jena.apache.org>

⁴<https://javacc.github.io/javacc/>

- Aside from selecting explicit attributes in the data – and per the CLOSURE requirement – the user may often wish to compute attributes upon which distance can be measured. For example, when comparing countries, the user may wish to use the number of cities in each country (computed with COUNT), or the GDP per capita (dividing the GDP by the population). It would not be feasible to pre-index these computed attributes.

For this reason, we rather consider online physical operators that receive two sets of solution mappings and compute the similarity join. Some of these physical operators include indexing techniques, but these indexes are built online – for a specific query – over intermediate solutions rather than the RDF graph itself.

The similarity join physical operators apply a normalisation function to each of the dimensions of the space, in order to avoid differences in scale of these dimensions to affect disproportionately the similarity of two solution mappings. For example, a similarity join query computing distances w.r.t. the dimensions human development index, HDI, (a number in $[0, 1]$) and population (ranging from thousands to even millions), as defined, the HDI dimension would have a negligible effect on the results of the similarity join. However, as a future direction, we might provide special syntax to allow users to decide if the data should be automatically normalised or not.

Algebra Optimisation then applies static rewriting rules over the query, further turning logical operators (e.g., `knnsimjoin`) into physical operators. In terms of the physical operators, we implement three variants:

Nested Loops involves a brute-force pairwise comparison, incurring quadratic cost for similarity joins. It computes exact results for both range and k -nn queries and is intended as a performance baseline.

vp-Trees is an index-based approach based on applying ball-cuts in the metric space. It computes exact results for both range and k -nn queries. Though the worst-case for similarity joins is still quadratic, depending on the distance distribution, dimensionality, etc., vp-Trees can approach linear performance.

FLANN is an ensemble of approximate methods for k -nn queries only, where the best approach is chosen based on the query, data, target precision, etc.

Algebra Execution begins to evaluate the physical operators chosen in the previous step, with low-level triple/quad patterns and path expression operators feeding higher-level operators. Finally, Result Iteration streams the final results from the evaluation of the top-level physical operator. All physical similarity-join operators follow the same lazy evaluation strategy used for the existing join operators in Jena.

Besides the implementation of sub-quadratic average-case similarity join physical operators, there are other avenues for query optimisation, namely: query rewriting and result caching. As for rewriting rules, there is little opportunity for applying such rules because of a) similarity joins are a relatively expensive operation to compute, thus the best option is to reduce the size of the outputs of the sub-operands as early as possible and therefore, making typical rewriting rules impractical (such as delaying filters and optionals); and b) we see in Section 4.2 that the k -nn similarity join operator is not commutative nor associative, further preventing rewriting such as join reordering.

4.4. Use-case queries

Back in Section 1, we presented some motivating example queries that can be written using our proposed extension to SPARQL over a small artificial dataset. Now, in this section, we present three use-case queries demonstrating different capabilities of the extension, which are based on real-world data from Wikidata [45] and IMGpedia [8]. All prefixes used in the queries of this paper are listed in Table 1.

Similar chemicals Chemical elements contain several numeric/ordinal properties, such as boiling and melting point, mass, electronegativity and so on. In Fig. 3 we request for metals similar to non-metals in terms of those properties, along with sample results of the similar metal/non-metal query. Figure 3 also presents a selection of results of the query, where we see that Selenium is most similar to Aluminium, and Hydrogen to Sodium.

Similar laureates: Figure 4, presents a more complex similarity query over the Nobel Prize Linked Dataset⁵ to find the 10 countries that have the most similar number of Nobel Prize laureates in literature, chemistry, physics and medicine than Canada. Figure 4 also presents the query results.

⁵<https://data.nobelprize.org/>

Table 1
IRI prefixes used in the paper

Prefix	IRI
wdt:	<code><http://www.wikidata.org/prop/direct/></code>
wd:	<code><http://www.wikidata.org/entity/></code>
sim:	<code><http://sj.dcc.uchile.cl/sim#></code>
imo:	<code><http://imgpedia.dcc.uchile.cl/ontology#></code>
tdb:	<code><https://triplifydb.com/tripty/iris/def/measure/></code>
stp:	<code><http://stars.dcc.uchile.cl/prop/></code>
dbo:	<code><http://dbpedia.org/ontology/></code>
nobel:	<code><http://data.nobelprize.org/terms/></code>
country:	<code><http://data.nobelprize.org/resource/country/></code>

```

SELECT ?element1 ?element2 ?dist WHERE {
  {?element1 wdt:P31 wd:Q11344; #chemical elements
   wdt:P2102 ?boil1; wdt:P2101 ?melt1;
   wdt:P1108 ?eneg1;
   wdt:P2054 ?dens1; wdt:P2067 ?mass1.
   (wdt:P279+) wd:Q11426. # metals
  }
  SIMILARITY JOIN ON (?boil1 ?melt1 ?eneg1 ?dens1 ?mass1)
                    (?boil2 ?melt2 ?eneg2 ?dens2 ?mass2)
  TOP 1 DISTANCE sim:manhattan AS ?dist
  {?element2 wdt:P31 wd:Q11344;
   wdt:P2102 ?boil2; wdt:P2101 ?melt2;
   wdt:P1108 ?eneg2;
   wdt:P2054 ?dens2; wdt:P2067 ?mass2.
  }
  MINUS {?element2 (wdt:P279+) wd:Q11426.}} # non-metals

```

Element 1	Element 2	Distance
wd:Q876 [Se]	wd:Q663 [Al]	0.4644
wd:Q556 [H]	wd:Q658 [Na]	0.3107

Fig. 3. Query for non-metal chemical elements similar to metallic elements in terms of atomic properties.

Similar images: Figure 5 presents a similarity query over IMGpedia [8]: a multimedia Linked Dataset. The query retrieves images of the Entel Tower in Chile, and computes a 10-nn similarity join for images of towers based on a precomputed Histogram of Oriented Gradients (HOG) descriptor, which extracts the distribution of edge directions of an image. Note that this query uses a different distance function IRI than the previous two, since variables `?hog1` and `?hog2` are bound to string serializations of vectors, and thus the function `sim:manhattanvec` deals with the parsing of such strings as well. The figure includes a sample of the results for one of the images of the Entel Tower (in index 0), showing for each similar image, the related Wikidata entity, its label, and the distance to the source image.

5. Clustering solution modifier

In this section, we analyse, implement and evaluate a `CLUSTER BY` solution modifier for SPARQL queries – inspired by the proposal of Ławrynowicz [21] – based on the aforementioned notions of similarity and distance. First, we discuss how this extension of SPARQL was initially conceived, how it can be formalised syntactically and semantically and how it interacts with the rest of the solution modifiers. We then discuss implementation details for the solution modifier, and present a prototype built on top of Apache Jena. Finally, use-case queries are presented.

As it has been argued before in this paper, users sometimes require non-exact matches over the data and instead wish for similar matches. This is also the case when using the solution modifier `GROUP BY` in SPARQL, which

```

SELECT ?country ?literature ?chemistry ?physics ?medicine ?d WHERE {
  {SELECT DISTINCT ?caLiterature ?caChemistry ?caPhysics ?caMedicine WHERE {
    {SELECT (COUNT(DISTINCT ?caLiteratureLaureate) AS ?caLiterature) WHERE {
      ?caLiteratureLaureate nobel:nobelPrize ?caLiteraturePrize; dbo:birthPlace country:Canada .
      ?caLiteraturePrize nobel:category nobel:Literature .}}
    {SELECT (COUNT(DISTINCT ?caChemistryLaureate) AS ?caChemistry) WHERE {
      ?caChemistryLaureate nobel:nobelPrize ?caChemistryPrize; dbo:birthPlace country:Canada .
      ?caChemistryPrize nobel:category nobel:Chemistry .}}
    {SELECT (COUNT(DISTINCT ?caPhysicsLaureate) AS ?caPhysics) WHERE {
      ?caPhysicsLaureate nobel:nobelPrize ?caPhysicsPrize; dbo:birthPlace country:Canada .
      ?caPhysicsPrize nobel:category nobel:Physics .}}
    {SELECT (COUNT(DISTINCT ?caMedicineLaureate) AS ?caMedicine) WHERE {
      ?caMedicineLaureate nobel:nobelPrize ?caMedicinePrize; dbo:birthPlace country:Canada .
      ?caMedicinePrize nobel:category nobel:Physiology_or_Medicine .}}}}
  SIMILARITY JOIN ON (?caLiterature ?caChemistry ?caPhysics ?caMedicine) (?literature ?chemistry ?physics ?medicine)
  TOP 10 DISTANCE sim:manhattan AS ?d
  {SELECT DISTINCT ?country ?literature ?chemistry ?physics ?medicine WHERE {
    {SELECT ?country (COUNT(DISTINCT ?literatureLaureate) AS ?literature) WHERE {
      ?country rdf:type dbo:Country .
      OPTIONAL {
        ?literatureLaureate nobel:nobelPrize ?literaturePrize; dbo:birthPlace ?country .
        ?literaturePrize nobel:category nobel:Literature .}} GROUP BY ?country}}
    {SELECT ?country (COUNT(DISTINCT ?chemistryLaureate) AS ?chemistry) WHERE {
      ?country rdf:type dbo:Country .
      OPTIONAL {
        ?chemistryLaureate nobel:nobelPrize ?chemistryPrize; dbo:birthPlace ?country .
        ?chemistryPrize nobel:category nobel:Chemistry .}} GROUP BY ?country}}
    {SELECT ?country (COUNT(DISTINCT ?physicsLaureate) AS ?physics) WHERE {
      ?country rdf:type dbo:Country .
      OPTIONAL {
        ?physicsLaureate nobel:nobelPrize ?physicsPrize; dbo:birthPlace ?country .
        ?physicsPrize nobel:category nobel:Physics .}} GROUP BY ?country}}
    {SELECT ?country (COUNT(DISTINCT ?medicineLaureate) AS ?medicine) WHERE {
      ?country rdf:type dbo:Country .
      OPTIONAL {
        ?medicineLaureate nobel:nobelPrize ?medicinePrize; dbo:birthPlace ?country .
        ?medicinePrize nobel:category nobel:Physiology_or_Medicine .}} GROUP BY ?country}}}}
} ORDER BY DESC(?d)

```

?country	?literature	?chemistry	?physics	?medicine	?d
country:Canada	2	4	6	4	0.0000
country:Prussia	2	5	4	2	0.0721
country:Austria	2	5	3	6	0.0866
country:China	2	1	5	2	0.0933
country:Austria-Hungary	2	6	2	3	0.1063
country:India	2	1	3	2	0.1223
country:Switzerland	1	3	6	6	0.1265
country:Norway	2	2	1	2	0.1334
country:South_Africa	2	1	0	3	0.1531
country:Ukraine	2	1	0	2	0.1658

Fig. 4. Query requesting countries that have a number of Literature, Chemistry, Physics and Medicine Nobel Prizes similar to those of Canada.

makes a partition of the resulting mappings of a query based on certain variable values that need to match exactly in order to assign a binding into a group. In SPARQL 1.1 [15] the algebra of `GROUP BY` is defined as described in Section 2.

As presented, the `GROUP BY` modifier partitions the bindings into equivalence classes (the X' multisets) based on the different key values (the values of μ'). The similarity-based alternative to `GROUP BY` is to perform clustering over the bindings, so that solution mappings that are similar (instead of equal) are assigned to the same cluster.







5.1. Syntax

In order to support clustering in SPARQL, we must first define its syntax and semantics in a way that a user can express the clustering variables, and the algorithms and their required parameters. Ławrynowicz [21] proposes the following extension for SPARQL 1.0 [36], where the cluster identifier is bound AS a fresh variable (cf. Fig. 6).

```

SELECT ?img2 ?tower ?dist WHERE {
  { ?img1 imo:associatedWith wd:1421270 .
    ?vector1 a imo:HOG ;
      imo:describes ?img1 ;
      imo:value ?hog1 .}
  SIMILARITY JOIN ON (?hog1) (?hog2) TOP 10 DISTANCE sim:manhattanvec AS ?dist
  { ?tower wd:P31/wdt:P279* wd:Q12518 .
    ?img2 imo:associatedWith ?tower .
    ?vector2 a imo:HOG ;
      imo:describes ?img2 ;
      imo:value ?hog2 .}}

```

Index	0	1	2	3	4	5
?img2						
?tower	wd:Q1421270	wd:Q635	wd:Q7793202	wd:Q7653429	wd:Q7011062	wd:Q4943563
Label	Torre Entel	Tower of Fourvière	Thomas Point Light	Swan Bells	New Presque Isle Light	Boon Island Light
?dist	0	22.282	24.224	24.439	25.052	25.911






Index	6	7	8	9	10
?img2					
?tower	wd:Q1413217	wd:Q7208057	wd:Q5250767	wd:Q6391315	wd:Q7092627
Label	Žižkov TV Tower	Point Judith Light	Deer Island Light	Kenosha North Pier Light	One Fathom Bank Lighthouse
?dist	25.984	26.397	26.850	26.940	27.016

Fig. 5. Query for the 10 images of towers most similar to each image of the Entel Tower in terms of the HOG visual descriptor.

```

SolutionModifier ::= OrderClause? LimitOffsetClauses? ClusterByClause?
ClusterByClause  ::= 'CLUSTER BY' Var+ 'AS' Var UsingClause?
UsingClause      ::= 'USING' IRIRef ('('Params')')?
Params           ::= TriplesBlock

```

Fig. 6. Syntactic extension to the SPARQL 1.0 grammar for clustering [21].

The first thing that can be noticed from the definition is that, since the extension was conceived for SPARQL 1.0, it lacks the definition of the `GROUP BY` solution modifier. The definition of a clustering solution modifier for SPARQL 1.1 needs to take into account the existence of `GROUP BY`, and how it shall interact with `CLUSTER BY`.

Since clustering is the similarity-based version of grouping (which is based on exact matching), one alternative is to make the modifiers `CLUSTER BY` and `GROUP BY` be mutually exclusive in a query. However, the two modifiers can be made compatible by defining the clustering modifier as binding the cluster identifier of each mapping to a fresh variable, potentially allowing users to later group the results by the identifier value; in other words, `CLUSTER BY` extends the solution mappings with an extra variable bound to a cluster identifier, rather than directly grouping them. This decision is made since any new operation included into a query language should be as non-disruptive as possible, and to preclude the usage of `GROUP BY` when clustering is indeed disruptive.

```

SolutionModifier ::= ClusterByClause? GroupClause? HavingClause? OrderClause? LimitOffsetClauses?
ClusterByClause  ::= 'CLUSTER BY' Var+ WithClause ClusterAlias
WithClause       ::= 'WITH' iriOrFunction
ClusterAlias     ::= 'AS' Var

```

Fig. 7. Syntactic extension to the SPARQL 1.1 grammar for clustering.

```

SELECT ?lifex ?growth ?code ?c WHERE {
  ?c1 wdt:P31 wd:Q6256 ; #countries
      wdt:P2250 ?lifex ;
      wdt:P2219 ?growth;
      wdt:P297 ?code .
}
CLUSTER BY ?lifex ?growth WITH sim:kmeans(4) AS ?c

```

Fig. 8. Example SPARQL query with the CLUSTER BY solution modifier.

Since the CLUSTER BY solution modifier is not mutually exclusive with GROUP BY, CLUSTER BY does not interact with the HAVING clause and the different aggregators, which can rather be computed *after* the clustering takes place. Conversely, in order to apply clustering with respect to values computed through aggregation functions, the aggregation must be computed in a subquery. Thus, following the SPARQL spec, the order of the solution modifiers including CLUSTER BY is: clustering, grouping, having, ordering, projection, distinct, reduced, offset, and limit.

Ławrynowicz [21] makes use of the SPARQL Query Results Clustering Ontology (SQRCO) to define the clustering algorithm and its parameters as a set of triples in the query. However, at the time of writing, the ontology is not found on the Web and is not registered in prefix.cc. We instead take a syntax-based approach, using special keywords for the clustering configuration, similarly to what is done with the similarity join operator.

Our proposed extension of the SPARQL 1.1 syntax that includes CLUSTER BY is presented in Fig. 7. The extension adds a new production rule, `ClusterByClause`, as an alternative in the `SolutionModifier` production rule. We make use of the Extension Function framework defined for SPARQL 1.1 to allow users to define the desired clustering algorithm and its parameters by using an `iriOrFunction` rule after the keyword `WITH`, thus allowing the function to be called without any parameters as well, which permits to define default values for each supported algorithm. Finally, we provide syntax to define the fresh variable on which the cluster identifier shall be bound `AS`. We permit only variables to be used as clustering attributes, but expressions could be supported in a future version. Nevertheless, expressions could be indirectly used for clustering by using the `BIND` function.

An example of a SPARQL query that uses this extension can be appreciated in Fig. 8, where the countries available in Wikidata are clustered based on their life expectancy and economic growth. The query specifies that the clustering shall use k -means returning 4 clusters. Notice that we also obtain the 2-letter ISO code for each country (property `wdt:P297`) and we are able to project that variable even if it is not part of the CLUSTER BY variables.

5.2. Semantics

Given the SPARQL semantics presented in Section 2, as well as the discussed syntax for CLUSTER BY, we now propose the semantics for the CLUSTER BY solution modifier, per the following definition.

Definition 5.1. Given a multiset of solution mappings X , the clustering algorithm \mathcal{C} , the variables with respect to which the clustering is computed $\mathcal{V} \subseteq \text{udom}(X)$ and a fresh variable $c \in \mathbf{V}$ that stores the cluster identifier, the evaluation of the clustering solution modifier is defined as follows:

$$\chi_{\mathcal{V},c,c}(X) := \{ \mathcal{C}(\pi_{\mathcal{V}}(X), \pi_{\mathcal{V}}(\mu), c) \mid \mu \in X \} \quad (25)$$

where \mathcal{C} is a clustering function that, given a multiset of solution mappings X , a solution mapping μ and a variable c , returns a mapping $\mu \cup c/i$ where i is an RDF term used as the identifier of the cluster selected for μ by \mathcal{C} . Concrete values for the identifiers are left to each specific algorithm, where a natural solution is to enumerate the clusters with integers. Note that the original solution mapping μ is being extended, not the projected version $\pi_{\mathcal{V}}(\mu)$.

5.3. Implementation

Analogous to the similarity join operators, we implement the `CLUSTER BY` solution modifier on top of Apache Jena's ARQ, where the code is publicly available in the following repository: <https://github.com/scferrada/jena>. The implementation supports three clustering algorithms: k -means, k -medoids, and DBSCAN. These algorithms are currently implemented in-memory. All three algorithms use the Manhattan distance as the dissimilarity measure. All three algorithms enumerate the clusters in no particular order using natural numbers starting from 1, and use said numbers as the identifiers of the clusters. For k -medoids we implemented the FastPAMI algorithm proposed by Schubert and Rousseeuw [38] which is $O(k)$ faster than the naive algorithm and uses $O(k)$ extra memory.

Each algorithm requires different parameters, but the implementation offers default values for them, if the user does not wish to specify them. The parameters of each algorithm and the chosen default values are:

- k -means [24]: k , the number of clusters to be returned, default 3; m , the number of iterations to terminate the algorithm, default 10.
- k -medoids [19]: k , the number of clusters to be returned, default 3.
- DBSCAN [7]: ϵ , the distance to search for neighbours of an object, default 0; p , the minimum number of neighbours that an object needs to be considered part of a cluster, default 1.

DBSCAN is different from k -means and k -medoids in that it does not require the number of clusters to be specified, but also may mark some objects as outliers, not being in any cluster. To comply with the semantics of `CLUSTER BY`, it assigns all solution mappings marked as outliers to a special cluster (with identifier -1). Notice that the default parameters for DBSCAN result in all objects marked as outliers. The implementation of DBSCAN uses vp-Trees in order to improve performance, since for each element a range-based similarity search must be conducted.

Other clustering algorithms were not chosen for various reasons: both grid-based and distribution-based methods do not scale well when dimension increases [53]. Hierarchical clustering methods are left for future versions, because given the syntax and semantics of the clustering solution modifier, it would be hard to add multiple variables bound to different levels of clustering. An initial option to support these methods would be to force the user to choose a clustering level to retrieve, but we argue that in most cases, the user might not know which level of clustering is most significant for their purposes.

5.4. Use-cases

Expanding on the motivating example queries of Section 1 (see Fig. 1d–1e), we now showcase three queries that use the clustering solution modifier here proposed in different domains. First, in Fig. 9 we display the result of the query presented in Fig. 8, that clusters the countries in Wikidata with respect to their life expectancy and economic growth rate. Some results were omitted to avoid clutter in the figure. Each point has a different colour and marker depending on their assigned cluster (whose identifier is bound to variable $?c$). It can be seen in this case that, since there are differences in the magnitudes of the values of life expectancy and economic growth rate, the largest values tend to weigh more in the distance value and thus, in this example, we see a stronger division along the vertical axis. To prevent that any given dimension has a disproportionate influence over the distance function, it is recommended to normalise each dimension before clustering. Additionally, we note that, unlike `GROUP BY`, we can keep and project any variables that are not part of the `CLUSTER BY` clause. In Fig. 9 we also present the 2-letter ISO code of each country (which is directly obtained from the query) annotating each point in the graph, so that the reader can know which countries belong to each cluster more easily.

Then, the Iris Dataset in RDF⁶ is used to run a clustering query on petal length and width. In Fig. 10 the extended SPARQL query that divides the objects into three clusters using the k -medoids algorithm is presented. Figure 11 shows the results of the clustering, where the green cluster is iris setosa, the blue cluster is mostly iris virginica, and the red cluster is mostly iris versicolor. Notice that the iris types can also be retrieved with the query.

⁶Retrieved from <https://triplifydb.com/RosalinedeHaan/iris>. Credit to Rosaline de Haan.




Fig. 9. Clusters of countries by growth and life expectancy.

```

SELECT ?iris ?type ?plength ?pwidth ?c WHERE {
  ?iris a ?type;
  tdb:petalLength ?plength ;
  tdb:petalWidth ?pwidth . }
CLUSTER BY ?plength ?pwidth WITH sim:kmedoids(3) AS ?c

```

Fig. 10. Extended SPARQL query to obtain three clusters of iris.

Next, we present a Hertzsprung–Russell (HR) diagram, that is used in astronomy to classify stars into main sequence, giants, white dwarfs, etc. HR diagrams have the normalised B–V Colour (a measure of the temperature or “brightness” of the star) on the X-axis, and M_V (the absolute magnitude) on the Y-axis. To achieve this, we generate an RDF version of a real-world catalogue of stars.⁷ In Fig. 12 a sample of the generated data can be found. The full RDF data is available at <https://github.com/scferrada/starsrdfdata>. Then, we cluster the stars by their absolute magnitude (M_V) and their B–V colour using DBSCAN, since the main sequence cluster is better obtained by density-based methods. We also filter in the query the stars that have a high parallax uncertainty (a measure of the chance that the distance of the star, and thus its magnitude, is incorrect), because these stars may have an unreliable magnitude measure and just add noise to the diagram. Since magnitude and colour have very different scales, we also apply min–max normalisation in two BIND clauses, and perform the clustering over the normalised data. The SPARQL query that represents this process can be found in Fig. 13. The plot of the result of the query can be seen in Fig. 14, where the central purple cluster represents stars in the main sequence, the top right blue and sky blue clusters represent the giant and super giant stars, and the bottom-left red cluster contains the white dwarfs.

Finally, we showcase how the CLUSTER BY solution modifier can interact with GROUP BY in a SPARQL query. To do this, we modify the query from Fig. 8, which obtains the life expectancy and growth rate of the countries in Wikidata, to now GROUP BY cluster identifier and compute the average life expectancy and the minimum growth rate of each cluster. This query is presented in Fig. 15, and the results of its evaluation over Wikidata can be seen

⁷The CSV data is available at <http://burro.astr.cwru.edu/Academics/Astr221/HW/HW5/yaletrigplx.dat>.




Fig. 11. Clusters obtained with query 10.

<http://stars.dcc.uchile.cl/star/2.00>	<http://stars.dcc.uchile.cl/prop/color>	1.04 .
<http://stars.dcc.uchile.cl/star/2.00>	<http://stars.dcc.uchile.cl/prop/absMagnitude>	0.261669 .
<http://stars.dcc.uchile.cl/star/2.00>	<http://stars.dcc.uchile.cl/prop/parallaxUncertainty>	149.0 .

Fig. 12. Extract of the RDF data for stars.

```

SELECT ?temp ?lum ?c WHERE {
  ?star stp:color ?temp ;
        stp:absMagnitude ?lum ;
        stp:parallaxUncertainty ?error .
  FILTER (?error < 75)
  BIND((?lum+1.4444)/(14.7121 + 1.4444) AS ?lum2)
  BIND((?temp+0.17)/(1.88+0.17) AS ?temp2)}
CLUSTER BY ?temp2 ?lum2 WITH sim:dbscan(0.05, 10) AS ?c

```

Fig. 13. SPARQL query for clustering star data.

in Table 2 where cluster 1 corresponds to the dark blue cluster in Fig. 9, which includes Japan and Italy, and has the highest life expectancy and a minimum growth rate of 0.9%; cluster 2 is the pink cluster, which contains New Zealand and Chile; cluster 3 is the sky blue cluster, which includes Brazil and Bulgaria, and presenting the lowest growth rate at -18% ; and cluster 4 is the red cluster, presenting the lowest life expectancy average.

6. Experimental results

Experimental results for both the similarity join operator and the clustering solution modifier are presented in this section. Unless stated otherwise, experiments were run on a machine with Ubuntu 20, an 8-core Intel® Core™ i7-6700 CPU @3.4 GHz, 16 GB of RAM and a 931 GB HDD. We provide the experiment code, queries, etc., online.⁸

6.1. Similarity join evaluation

We compare the performance of different physical operators for similarity joins in the context of increasingly complex SPARQL queries, as well as a comparison with the baseline system DBSimJoin. We conduct experiments with respect to two benchmarks: the first is a benchmark we propose for Wikidata, while the second is an existing benchmark used by DBSimJoin based on visual descriptors of images.

⁸<https://github.com/scferrada/SimilaritySPARQL>




Fig. 14. Visual representation of the star clusters obtained with query from Fig. 13. Outliers not presented.

```
SELECT AVG(?lifex) MIN(?growth) ?c WHERE {
?c1 wdt:P31 wd:Q6256 ; #countries
    wdt:P2250 ?lifex ;
    wdt:P2219 ?growth .
}
CLUSTER BY ?lifex ?growth WITH sim:kmeans(4) AS ?c
GROUP BY ?c
```

Fig. 15. Example SPARQL query with the CLUSTER BY solution modifier and GROUP BY.

Table 2
Results of query of Fig. 15

?c	AVG(?lifex)	MIN(?growth)
1	83.399	-5.42
2	78.696	-3.20
3	73.121	-18.00
4	61.909	-13.80

6.1.1. Wikidata: k -nn self-similarity queries

Extending the performance experimentation presented in our previous work [9], we now present similar experiments using real SPARQL queries. We use a truthful version of the Wikidata RDF dump.⁹ We start with SPARQL queries from the Wikidata query logs, and apply several filters: first, at least one numeric attribute must be projected in the SELECT clause; second, these numeric attributes must not be obtained in an OPTIONAL clause; third, they produce at least 8 solution mappings; and fourth, they are executed in less than 3 hours using Jena. We obtain 2,726 SPARQL queries. For each of these base queries, we wrote a k -nn self-similarity join version for our experiments.

We executed self-similarity joins with $k \in \{2, 4, 8\}$, and compare 3 different techniques: Nested Loops (NL), vp-Trees (VPT), and FLANN. We run the queries sequentially, to avoid caching. In Fig. 16 we present our findings, where we provide the execution time of the base query as a baseline for comparison. VPT perform the best in 2,841 out of 8,178 scenarios (34.7%), NL in 2,770 (33.9%), and FLANN in 2,567 (31.4%). NL tend to perform better when the number of solution mappings is relatively smaller, which makes sense, since building indices has an overhead cost. The mean set size of the cases where NL performs the best is 178 elements. FLANN tends to perform

⁹We use the dump from February 2020.

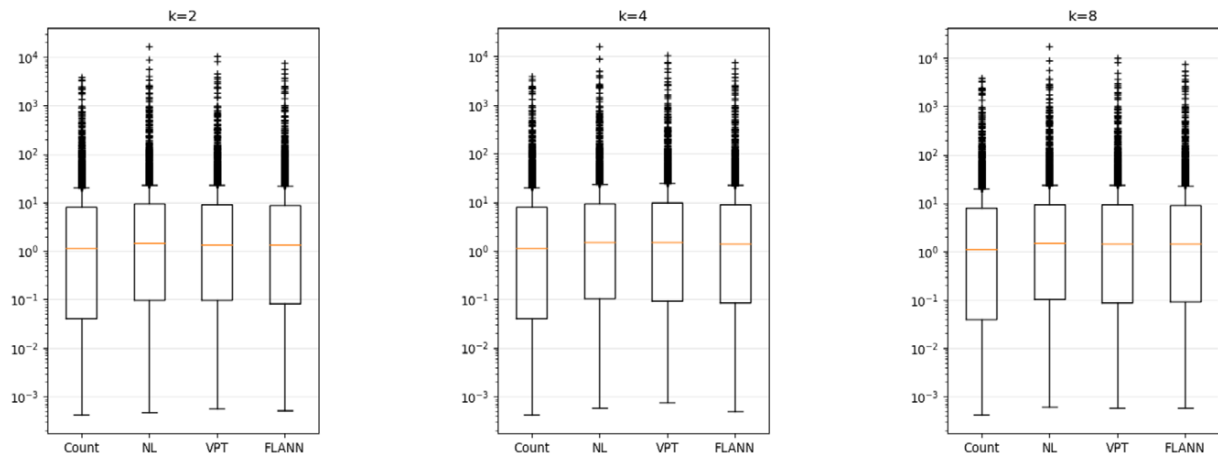


Fig. 16. Execution time for self-similarity join queries, for different values of k .

better, when the set is comparably larger (mean size of 2,155 elements). As VPT perform better in the majority of the times and provides an exact result it should be preferred for larger result sets, and NL for smaller result sets.

6.1.2. Corel colour moments: Range similarity queries

The system proposed in this work is compared with the closest found in literature, DBSimJoin [43], a PostgreSQL extension that supports range-based similarity joins in metric spaces implementing the Quickjoin (QJ) algorithm. As DBSimJoin only supports range queries, it is to be compared with the vp-Tree implementation in Jena (Jena-vp). The DBSimJoin system was originally tested with synthetic datasets and with the Corel Colour Moments (CCM) dataset, which consists of 68,040 records of images, each described with 9 dimensions representing the mean, standard deviation and skewness for the colours (red, green and blue) of the pixels of the image. For CCM, the DBSimJoin paper only reports the effect of the number of QJ pivots on the execution time when $r = 1\%$ of the maximum possible distance in the space [43]. Here, the comparison between the systems takes into account more general performance metrics, using CCM. The CCM dataset was converted to an RDF graph using a direct mapping.

To find suitable distances, we first compute a 1-nn self-similarity join, where we take the maximum 1-nn distance in the result: 5.9; this distance ensures that each object has at least one neighbour in a range-based query with $r = 5.9$. Therefore, we compare the runtime of both systems with increasing values of r , using $r = 5.9$ as an upper bound.

DBSimJoin is implemented and distributed an extension for PostgreSQL 8. When installed in the current hardware, it crashed with any similarity join query, because of the use of outdated I/O directives, making it impossible to re-run the experiments. Thus, we present the results previously reported [9] where we compare them with Jena-vp using the same dataset, the same queries, and the same hardware reported in [9]. Here, we provide extended discussion over these results, which are presented in Table 3. The execution time of DBSimJoin grows rapidly with r because of the quick growth of the size and number of window partitions in the QJ algorithm. DBSimJoin crashes with $r \geq 0.4$ so results between 0.3 and 0.4 are included to illustrate the growth in runtime closer to the failing point. The runtime of Jena-vp increases slowly with r , where more tree branches need to be visited as the result-set size increases up to almost 400 million records. The implementation of Jena-vp does not crash with massive results because it fetches the bindings lazily: it obtains all pairs within distance r for a single object and returns them one by one, only computing the next batch when it runs out of pairs; on the other hand, QJ starts piling up window partitions that replicate the data at a high rate, thus causing DBSimJoin to crash. In Fig. 17 the execution times of DBSimJoin and vp-Jena are presented visually, along with a polynomial trend of order 2 that predicts the execution times for when DBSimJoin crashes. Notice that both axes are in logarithmic scale.

Considering range-based similarity joins, these results indicate that Jena-vp outperforms DBSimJoin in terms of efficiency (our implementation is faster) and scalability (our implementation can handle larger amounts of results).

Table 3
Execution times in seconds for range similarity joins over the CCM dataset

Distance	# of Results	DBSimJoin (s)	Jena-vp (s)
0.01	68,462	47.22	6.92
0.1	68,498	84.00	7.45
0.3	72,920	775.96	9.63
0.39	92,444	1,341.86	11.17
0.4	121,826	–	11.30
1.0	4,233,806	–	35.01
5.9	395,543,225	–	1,867.86

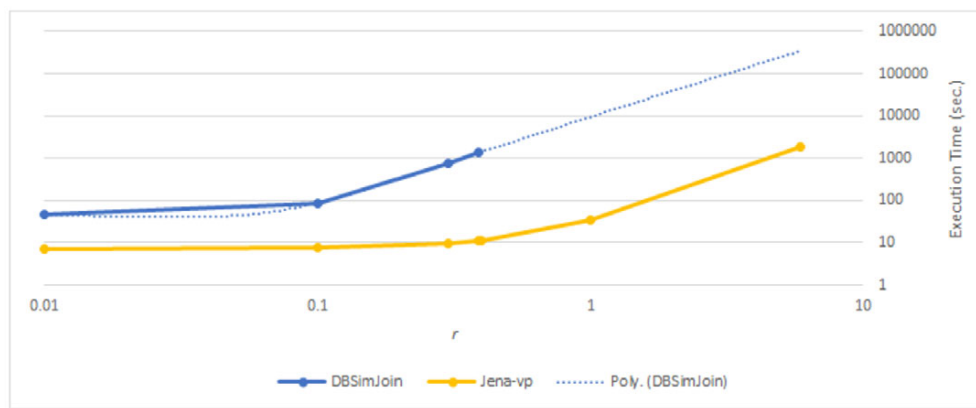


Fig. 17. Execution time of range-based self-similarity joins, including polynomial trend line for DBSimJoin times beyond $r = 0.4$.

6.2. Clustering experimental results

To gain insight on the performance of the clustering solution modifier, we measure the overhead produced by the clustering algorithms. To achieve this, we prepared 1,000 synthetic SPARQL queries that draw values from the numeric properties of Wikidata. The queries are composed of only one Basic Graph Pattern (BGP), and we apply clustering over its solutions. To produce these queries, we begin with some data exploration. Specifically, from a Wikidata dump, we compute the *ordinal/numeric characteristic sets* [28] by: (1) filtering triples with properties that do not take numeric datatype values, or that take non-ordinal numeric datatype values (e.g., Universal Identifiers); (2) extracting the *characteristic sets* of the graph [28] along with their cardinalities, where, specifically, for each subject in the resulting graph, we extract the set of properties for which it is defined, and for each such set, we compute for how many subjects it is defined. Finally, we generate `CLUSTER BY` queries from 1,000 ordinal/numeric characteristic sets containing more than 3 properties that were defined for at least 500 subjects.

Then, we execute these queries with different parameters, in order to compare the execution time of the three implemented algorithms under various realistic configurations. These parameters are:

- For *k*-means: $k \in \{4, 8, 12, 16\}$, and $m \in \{10, 20, 30\}$.
- For DBSCAN: $\varepsilon \in \{0.01, 0.1, 0.5, 1\}$, and $p \in \{1, 10, 100\}$.
- For *k*-medoids: $k \in \{4, 8, 12, 16\}$.

We use as a baseline the execution time of the query that `COUNTs` the results of the underlying BGP of each benchmark query, so as to measure the overhead produced by the `CLUSTER BY` operators.

In Fig. 18, we present the results for the queries using the *k*-means algorithm. We use same-coloured lines for runs with the same value of *k*, and different dashing for different values of *m*. The baseline is shown with the light green line. We observe that the execution time increases as the value of *k* and *n* increases, and that the value of *m* does not

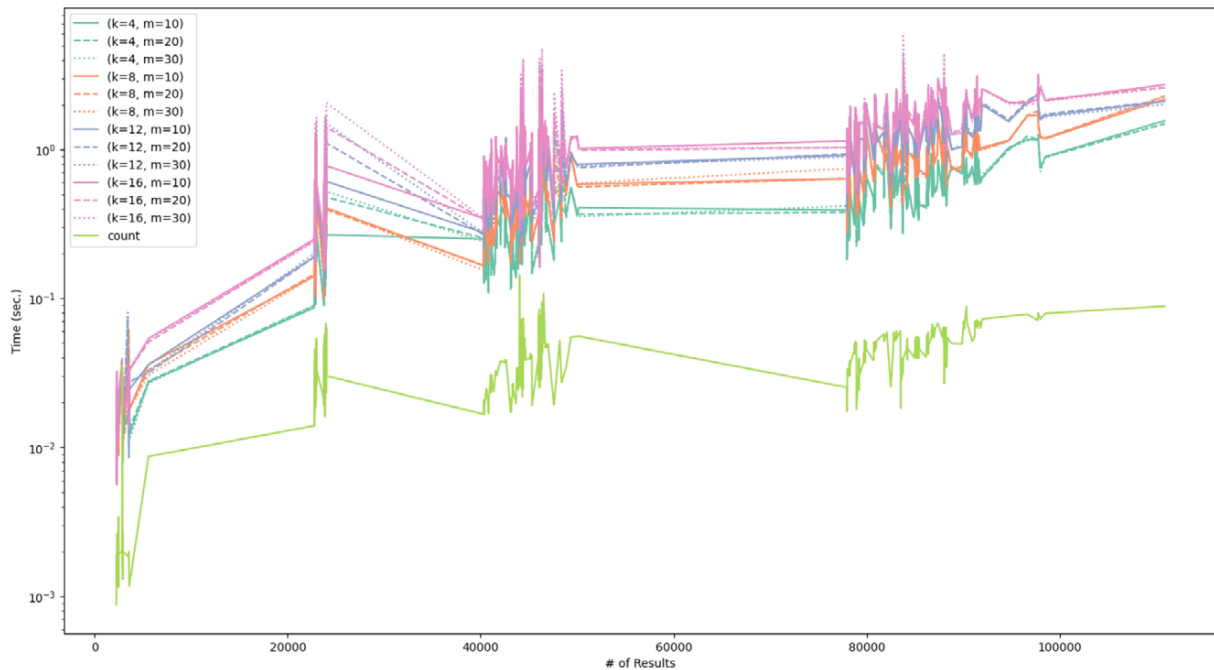


Fig. 18. Execution time of the CLUSTER BY queries using k -means clustering.

significantly influence the execution time. This is expected, since k -means has an upper bound time complexity of $O(knm)$ distance computations. However, k -means can terminate before m iterations if the clusters have converged.

The results for the queries using DBSCAN are harder to interpret, as the execution time also depends on the distance distribution of the space defined by the solution mappings obtained from evaluating each BGP. This is because DBSCAN's cost depends on the cost of performing ε -range searches, and on the number of found neighbours. This can be noticed in Fig. 19, which presents a plot where queries with the same value for ε are shown in the same colour, and queries with the same value for p are shown with the same line dashing. The baseline is shown with the light green line. In Fig. 19, we observe unpredictable changes in the magnitude of the execution time, even for queries with a similar n . This can also be explained by the fact that we use fixed values of ε for this evaluation instead of using distance values depending on the metric space derived from the query evaluation, as it is recommended [39]. Notice that normalisation of the underlying metric space would not guarantee more consistent results, as the cost of DBSCAN depends on the distance distribution of the space. Nevertheless, queries with a larger ε and a smaller p tend to take more time. For these results, we marked queries taking more than 3 hours as timeouts, and thus not reported in the graph. This leads to most queries with $\varepsilon = 1$ to not be included.

k -medoids is the costliest clustering algorithm of the three, as it minimises intra-cluster distance while maximising inter-cluster distance. Since the implementation requires to compute all pairwise distances first, $O(n^2)$ memory is required, with an extra memory of $O(k)$ used by the FASTPAM algorithm [38]. This memory requirement resulted in queries with n larger than 5,584 to fail with out of memory errors. In Fig. 20, we present the execution times for those queries that finished successfully. It can be noticed that k -medoids takes much longer to execute than the underlying query, and around four orders of magnitude more time than k -means. For $k = 16$, queries with $n > 2908$ take longer than 12 hours and were manually cancelled and therefore not reported.

Given these experimental results, we can argue that the preferred way to cluster query results would be to use k -means, which shall have a similar execution time as the underlying query. Moreover, the k -medoids clustering algorithm should be avoided given its intensive memory and time usage. For the case of DBSCAN, it is necessary to be familiar with the vector space that is being generated, and carefully choose the parameters, for which we refer to the recommendations of Schubert et al. [39], of using $p = 2 \cdot d$, where d is the dimensionality of the space, and ε as the “elbow” value in a graph of the distance, of each object to its $2d - 1$ -nearest neighbour.




Fig. 19. Execution time of the CLUSTER BY queries using DBSCAN clustering.




Fig. 20. Execution time of the CLUSTER BY queries using k -medoids clustering.

7. Concluding remarks

In this paper we have presented two extensions to SPARQL that incorporate two new similarity-based operators: similarity join and clustering. We define each operator's syntax and semantics as part of the SPARQL grammar and the SPARQL algebra, and show how these new operators can be combined with the rest. In the case of the similarity join, we prove algebraic properties and rewriting equivalences. For the case of the clustering solution modifier, we discuss how it interacts with the other solution modifiers, such as GROUP BY. We motivate the importance of these extensions by defining use-case queries ranging several domains (images, stellar data, encyclopaedic data).

These operators have been implemented on top of Apache Jena, and the implementation is publicly available. We provide experimental results to evaluate the performance of the implementation, where we show that our similarity join outperforms the closest system to ours, DBSimJoin. We also offer comparisons among different methods and

configurations to evaluate both the similarity join and the clustering. We determine that the better way to compute the similarity join is the vp-tree when there is a large number of solution mappings, and the nested loop when there are fewer solution mappings. For these experiments, we use 2,726 real SPARQL queries obtained from the logs of Wikidata. We also show that, when clustering, users may prefer to use the k -means algorithm, as it is faster and consumes less memory than k -medoids. We conclude by observing that DBSCAN is rather unpredictable in terms of execution time, and that it can produce useful and fast results if the parameters are carefully chosen.

As a future direction for this work, we plan to implement the physical operators using algorithms for secondary memory. This would allow for more efficient query processing and improved scalability, which has been an issue in particular for k -medoids clustering. Furthermore, we could design tailored algorithms that can take advantage of the database setting and the features of SPARQL, to further optimise the query processing.

Acknowledgements

This work was partly funded by ANID – Millennium Science Initiative Program – Code ICN17_002, by the Swedish Research Council (Vetenskapsrådet, project reg. no. 2019-05655), and the CENIIT program at Linköping University (project no. 17.05). Hogan was supported by Fondecyt Grant 1221926

Appendix. Proofs of algebraic properties

This section presents proofs for the algebraic properties of the similarity join operator introduced in Section 4.2.

Example A.1. To provide examples of the algebraic properties, and illustrate the ideas behind the proofs, we consider the following sets of solution mappings \bar{A} , \bar{B} , \bar{C} , \bar{D} and \bar{E} given as follows:

$$\bar{A} = \{\{?x \rightarrow 2, ?y \rightarrow 3\}, \{?x \rightarrow 4, ?y \rightarrow 4\}\}$$

$$\bar{B} = \{\{?z \rightarrow 1, ?w \rightarrow 1\}\}$$

$$\bar{C} = \{\{?z \rightarrow 3, ?w \rightarrow 2\}\}$$

$$\bar{D} = \{?u \rightarrow 3\}$$

$$\bar{E} = \{?y \rightarrow 3, ?u \rightarrow 3\}$$

Let l_1 denote Manhattan distance. For ranges, we take the similarity criteria $\bowtie_{\bar{\tau}} = (\{(?x, ?z), (?y, ?w)\}, l_1, ?vr, \text{rg}_3)$, whereby we want to find pairs of bindings for $(?x, ?y)$ and $(?z, ?w)$ that are within Manhattan distance of three (note that in the similarity criteria definition, the variables are transposed to pairs of variables that will be compared). For nearest neighbours, we assume the similarity criteria $\bowtie_{\bar{\tau}} = (\{(?y, ?w)\}, l_1, ?vk, \text{nn}_1)$, whereby we want to find, for each binding of $?y$, its nearest neighbour bound to $?w$ according to Manhattan distance.

Proposition A.1. \bowtie_{τ} is commutative and distributive over \cup .

Proof. Assume $\tau = (\mathcal{V}, \delta, v, \text{rg}_r)$. For any pair of sets of mappings X and Y :

$$\begin{aligned} X \bowtie_{\tau} Y &:= \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Y \bowtie \{ \{v \rightarrow \delta([\mathcal{V}]_{\mu_2}^{\mu_1}) \} \mid \mu_v(v) \leq r \} \\ &\equiv \{[\mu_2, \mu_1, \mu_v] \in Y \bowtie X \bowtie \{ \{v \rightarrow \delta([\mathcal{V}]_{\mu_1}^{\mu_2}) \} \mid \mu_v(v) \leq r \} \\ &\equiv Y \bowtie_{\tau} X \end{aligned}$$

Which proves the commutativity. For left-distributivity over \cup :

$$\begin{aligned}
X \bowtie_{\tau} (Y \cup Z) &:= \{[\mu_1, \mu_2, \mu_v] \in X \bowtie (Y \cup Z) \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq r\} \\
&\equiv \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Y \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq r\} \\
&\quad \cup \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Z \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq r\} \\
&\equiv (X \bowtie_{\tau} Y) \cup (X \bowtie_{\tau} Z)
\end{aligned}$$

Commutativity and left-distributivity over \cup imply distributivity over \cup . \square

Example A.2. Referring back to Example A.1, with respect to commutativity, note that:

$$\begin{aligned}
\bar{A} \bowtie_{\bar{\tau}} \bar{B} &= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{v \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq 3\} \\
&= \{?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{z} \rightarrow 1, ?\bar{w} \rightarrow 1, ?\bar{v} \rightarrow 3\} \\
&= \{?\bar{z} \rightarrow 1, ?\bar{w} \rightarrow 1, ?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{v} \rightarrow 3\} \\
&= \{[\mu_2, \mu_1, \mu_v] \in \bar{B} \bowtie \bar{A} \bowtie \{v \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_1}^{\mu_2})\} \mid \mu_v(v) \leq 3\} \\
&= \bar{B} \bowtie_{\bar{\tau}} \bar{A}
\end{aligned}$$

Since the natural join operator is commutative, the distance measure symmetric, and the solution mappings unordered (recall that $[\mu_1, \mu_2, \mu_v]$ corresponds to the solution mapping $\mu_1 \cup \mu_2 \cup \mu_v$), $X \bowtie_{\tau} Y$ and $Y \bowtie_{\tau} X$ are equivalent.

With respect to left-distributivity, for example, note that:

$$\begin{aligned}
\bar{A} \bowtie_{\bar{\tau}} (\bar{B} \cup \bar{C}) &:= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie (\bar{B} \cup \bar{C}) \bowtie \{v \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq 3\} \\
&= \{?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{z} \rightarrow 1, ?\bar{w} \rightarrow 1, ?\bar{v} \rightarrow 3\}, \\
&\quad \{?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{z} \rightarrow 3, ?\bar{w} \rightarrow 2, ?\bar{v} \rightarrow 2\}, \\
&\quad \{?\bar{x} \rightarrow 4, ?\bar{y} \rightarrow 4, ?\bar{z} \rightarrow 3, ?\bar{w} \rightarrow 2, ?\bar{v} \rightarrow 3\} \\
&= \{?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{z} \rightarrow 1, ?\bar{w} \rightarrow 1, ?\bar{v} \rightarrow 3\} \\
&\quad \cup \{?\bar{x} \rightarrow 2, ?\bar{y} \rightarrow 3, ?\bar{z} \rightarrow 3, ?\bar{w} \rightarrow 2, ?\bar{v} \rightarrow 2\}, \\
&\quad \{?\bar{x} \rightarrow 4, ?\bar{y} \rightarrow 4, ?\bar{z} \rightarrow 3, ?\bar{w} \rightarrow 2, ?\bar{v} \rightarrow 3\} \\
&= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{v \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq 3\} \\
&\quad \cup \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{C} \bowtie \{v \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \mu_v(v) \leq 3\} \\
&\equiv (\bar{A} \bowtie_{\bar{\tau}} \bar{B}) \cup (\bar{A} \bowtie_{\bar{\tau}} \bar{C})
\end{aligned}$$

Taking the similarity join of the union $\bar{A} \bowtie_{\bar{\tau}} (\bar{B} \cup \bar{C})$ or the union of the similarity join $(\bar{A} \cup \bar{C}) \bowtie_{\bar{\tau}} \bar{B}$ yields the same result in this case, and per the proof, in all such cases. Right-distributivity then follows the same lines.

Proposition A.2. \bowtie_n is not commutative nor is it distributive over \cup .

Proof. As counter examples for commutativity and distributivity, note that there exist sets of mappings X, Y, Z with $|X| = n, |Y| = |Z| = 2n, n \geq k > 0$, and such that $\delta(x, t_1) \neq \delta(x, t_2)$ for all $x \in X$ and $t_1, t_2 \in Y \cup Z, t_1 \neq t_2$, we have that:

- Commutativity: $|X \bowtie_n Y| = nk$ and $|Y \bowtie_n X| = 2nk$.
- Distributivity: $|X \bowtie_n (Y \cup Z)| = nk$ and $|(X \bowtie_n Y) \cup (X \bowtie_n Z)| = 2nk$.

Hence the result holds: \bowtie_n cannot be commutative nor distributive over \cup since the numbers of solutions differ. \square

Example A.3. Referring back to Example A.1, $\bar{A} \bowtie_{\bar{n}} \bar{B}$ finds the nearest neighbour in \bar{B} for each solution mapping in \bar{A} . Since \bar{A} has two solutions, $\bar{A} \bowtie_{\bar{n}} \bar{B}$ will return two solutions, one for each solution in \bar{A} . On the other hand, $\bar{B} \bowtie_{\bar{n}} \bar{A}$ finds the nearest neighbour in \bar{A} for each solution mapping in \bar{B} . Since \bar{B} has one solution, $\bar{B} \bowtie_{\bar{n}} \bar{A}$ will return only one solution. Hence $\bar{A} \bowtie_{\bar{n}} \bar{B} \neq \bar{B} \bowtie_{\bar{n}} \bar{A}$ and \bowtie_n cannot be commutative. Turning to distributivity, for each solution in \bar{A} , the join $\bar{A} \bowtie_{\bar{n}} (\bar{B} \cup \bar{C})$ returns the nearest neighbour from \bar{B} or \bar{C} , again yielding $|\bar{A}| = 2$ solutions. Conversely, for each solution in \bar{A} , the union $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \cup (\bar{A} \bowtie_{\bar{n}} \bar{C})$ gives a nearest neighbour from \bar{B} and a nearest neighbour \bar{C} , yielding $2|\bar{A}| = 4$ solutions. Hence $\bar{A} \bowtie_{\bar{n}} (\bar{B} \cup \bar{C}) \neq (\bar{A} \bowtie_{\bar{n}} \bar{B}) \cup (\bar{A} \bowtie_{\bar{n}} \bar{C})$, and thus \bowtie_n cannot be distributive over \cup .

Proposition A.3. \bowtie_n is right-distributive over \cup .

Proof. Assume $n = (\mathcal{V}, \delta, v, nn_k)$. We see that:

$$\begin{aligned} (X \cup Y) \bowtie_n Z &:= \{[\mu_1, \mu_2, \mu_v] \in (X \cup Y) \bowtie Z \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{\delta}^Z(\mu_1, \mu_2) \leq k\} \\ &\equiv \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Z \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{\delta}^Z(\mu_1, \mu_2) \leq k\} \\ &\quad \cup \{[\mu_1, \mu_2, \mu_v] \in Y \bowtie Z \bowtie \{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{\delta}^Z(\mu_1, \mu_2) \leq k\} \\ &\equiv (X \bowtie_n Z) \cup (Y \bowtie_n Z) \end{aligned} \quad \square$$

Example A.4. Referring back to Example A.1, we will look at the case of $(\bar{B} \cup \bar{C}) \bowtie_{\bar{n}} \bar{A}$. We see that:

$$\begin{aligned} (\bar{B} \cup \bar{C}) \bowtie_{\bar{n}} \bar{A} &= \{[\mu_1, \mu_2, \mu_v] \in (\bar{B} \cup \bar{C}) \bowtie \bar{A} \bowtie \{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{l_1}^{\bar{A}}(\mu_1, \mu_2) \leq 1\} \\ &= \{?z \rightarrow 1, ?w \rightarrow 1, ?x \rightarrow 2, ?y \rightarrow 3\}, \{?z \rightarrow 3, ?w \rightarrow 2, ?x \rightarrow 2, ?y \rightarrow 3\} \\ &= \{?z \rightarrow 1, ?w \rightarrow 1, ?x \rightarrow 2, ?y \rightarrow 3\} \cup \{?z \rightarrow 3, ?w \rightarrow 2, ?x \rightarrow 2, ?y \rightarrow 3\} \\ &= \{[\mu_1, \mu_2, \mu_v] \in \bar{B} \bowtie \bar{A} \bowtie \{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{l_1}^{\bar{A}}(\mu_1, \mu_2) \leq 1\} \\ &\quad \cup \{[\mu_1, \mu_2, \mu_v] \in \bar{C} \bowtie \bar{A} \bowtie \{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\} \mid \kappa_{l_1}^{\bar{A}}(\mu_1, \mu_2) \leq 1\} \\ &= (\bar{B} \bowtie_{\bar{n}} \bar{A}) \cup (\bar{C} \bowtie_{\bar{n}} \bar{A}) \end{aligned}$$

Hence we can apply the union of the left-hand side of a k -nn similarity join first (per $(X \cup Y) \bowtie_n Z$), or we can apply the k -nn similarity join to each operand of the union individually and later apply a union (per $(X \bowtie_n Z) \cup (Y \bowtie_n Z)$).

Proposition A.4. $(X \bowtie_{\mathfrak{s}} Y) \bowtie_{\mathfrak{s}'} Z \neq X \bowtie_{\mathfrak{s}} (Y \bowtie_{\mathfrak{s}'} Z)$ holds, i.e., $\bowtie_{\mathfrak{s}}$ is not associative.

Proof. As a counter example, consider that \mathfrak{s} and \mathfrak{s}' bind distance variables v and v' respectively such that $v' \in \text{udom}(X)$, $v' \notin \text{udom}(Y) \cup \text{udom}(Z)$ and $v \notin \text{udom}(X) \cup \text{udom}(Y) \cup \text{udom}(Z)$. Now $(X \bowtie_{\mathfrak{s}} Y) \bowtie_{\mathfrak{s}'} Z$ returns an error as the left operand of $\bowtie_{\mathfrak{s}'}$ assigns v but $X \bowtie_{\mathfrak{s}} (Y \bowtie_{\mathfrak{s}'} Z)$ will not. \square

Example A.5. Taking $\bowtie_{\bar{n}} = (\{(?y, ?w)\}, l_1, ?vk, nn_1)$ from Example A.1, consider the similarity criteria $\bowtie_{\bar{n}'} = (\{(?vk, ?u)\}, l_1, ?vk', nn_1)$, as well as the queries $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie_{\bar{n}'} \bar{D}$ and $\bar{A} \bowtie_{\bar{n}'} (\bar{B} \bowtie_{\bar{n}'} \bar{D})$. Note that $\bowtie_{\bar{n}'}$ computes a similarity join that depends on the distance between solutions in \bar{A} and \bar{B} as bound to $?vk$, and hence while the query $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie_{\bar{n}'} \bar{D}$ computes these distances first and binds $?vk$ before joining with \bar{D} , this is not the case for $\bar{A} \bowtie_{\bar{n}'} (\bar{B} \bowtie_{\bar{n}'} \bar{D})$, which computes $\bar{B} \bowtie_{\bar{n}'} \bar{D}$ first, wherein $?vk$ will be unbound. Hence the results are not the same.

Proposition A.5. Let $\mathfrak{s} = (\mathcal{V}, \delta, v, \phi)$. If each mapping in $X \bowtie Y$ binds all variables in \mathcal{V} and $v \notin \text{udom}(X) \cup \text{udom}(Y) \cup \text{udom}(Z)$, then the following hold:

1. $(X \bowtie_{\mathfrak{s}} Y) \bowtie Z \equiv (X \bowtie Z) \bowtie_{\mathfrak{s}} Y$
2. $(X \bowtie_{\mathfrak{s}} Y) \triangleright Z \equiv (X \triangleright Z) \bowtie_{\mathfrak{s}} Y$ if $\text{udom}(Z) \cap (\text{udom}(Y) - \text{idom}(X)) = \emptyset$
3. $(X \bowtie_{\mathfrak{s}} Y) \bowtie Z \equiv (X \bowtie Z) \bowtie_{\mathfrak{s}} Y$ if $\text{udom}(Z) \cap (\text{udom}(Y) - \text{idom}(X)) = \emptyset$
4. $\sigma_f(X \bowtie_{\mathfrak{s}} Y) \equiv \sigma_f(X) \bowtie_{\mathfrak{s}} Y$ if f is scoped to $\text{idom}(X)$.

Proof. We prove each claim in the following:

1. The third step here is possible as ϕ does not rely on Z (per the assumptions).

$$\begin{aligned} (X \bowtie_{\mathfrak{s}} Y) \bowtie Z &:= \{[\mu_1, \mu_2, \mu_v] \in X \bowtie Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \phi(\mu_v)\} \bowtie Z \\ &\equiv \{[\mu_1, \mu'_1, \mu_2, \mu_v] \in X \bowtie Z \bowtie Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \phi(\mu_v)\} \\ &\equiv \{[\mu_1, \mu_2, \mu_v] \in (X \bowtie Z) \bowtie Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \phi(\mu_v)\} \\ &\equiv (X \bowtie Z) \bowtie_{\mathfrak{s}} Y \end{aligned}$$

2. For a mapping $\mu = [\mu_1, \mu_2]$ such that $\text{udom}(Z) \cap (\text{dom}(\mu_2) - \text{dom}(\mu_1)) = \emptyset$, there does not exist $\mu' \in Z$ such that $\mu \sim \mu'$ if and only if there does not exist $\mu' \in Z$ such that $\mu_1 \sim \mu'$. Taking $\mu_1 \in X$ and $\mu_2 = [\mu'_2, \mu''_2] \in Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\}$ from $X \bowtie_{\mathfrak{s}} Y$, the result then holds per the given assumptions.
3. The second step here uses the previous two results. The third step uses the right-distributivity of \bowtie_r and \bowtie_n over \cup proven in previous propositions.

$$\begin{aligned} (X \bowtie_{\mathfrak{s}} Y) \bowtie Z &:= ((X \bowtie_{\mathfrak{s}} Y) \bowtie Z) \cup ((X \bowtie_{\mathfrak{s}} Y) \triangleright Z) \\ &\equiv ((X \bowtie Z) \bowtie_{\mathfrak{s}} Y) \cup ((X \triangleright Z) \bowtie_{\mathfrak{s}} Y) \\ &\equiv ((X \bowtie Z) \cup (X \triangleright Z)) \bowtie_{\mathfrak{s}} Y \\ &\equiv (X \bowtie Z) \bowtie_{\mathfrak{s}} Y \end{aligned}$$

4. For a mapping $\mu = [\mu_1, \mu_2]$ and filter f scoped to $\text{dom}(\mu_1)$, $f(\mu)$ is true if and only if $f(\mu_1)$ is true. Taking $\mu_1 \in X$ and $\mu_2 = [\mu'_2, \mu''_2] \in Y \bowtie \{\{v \rightarrow \delta(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\}$ from $X \bowtie_{\mathfrak{s}} Y$, the result holds per the given assumptions. \square

Example A.6. We will look at examples for each claim:

1. Take $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie \bar{E}$, which gives:

$$\begin{aligned} (\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie \bar{E} &= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \bowtie \bar{E} \\ &= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}, \\ &\quad \{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 3\}\} \\ &\quad \bowtie \{?y \rightarrow 3, ?u \rightarrow 3\} \\ &= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2, ?u \rightarrow 3\}\} \\ &= \{[\mu_1, \mu_2, \mu_v] \in (\bar{A} \bowtie \bar{E}) \bowtie \bar{B} \\ &\quad \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \\ &= (\bar{A} \bowtie \bar{E}) \bowtie_{\bar{n}} \bar{B} \end{aligned}$$

This equality would not hold if $\bowtie_{\bar{n}}$ depended on a variable bound by the right operand of the natural join (\bar{E}), but this case is forbidden by the condition that “each mapping in $X \bowtie Y$ binds all variables in \mathcal{V} ” in Proposition A.5. It would also not hold if this operand (\bar{E}) mentioned the variable to which the distance is bound by $\bowtie_{\bar{n}}$ since $(\bar{A} \bowtie \bar{E}) \bowtie_{\bar{n}} \bar{B}$ would throw an error, while $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie \bar{E}$ would not; this case is covered by the proposition’s condition that $v \notin \text{udom}(X) \cup \text{udom}(Y) \cup \text{udom}(Z)$.

2. Recalling that \triangleright denotes an antijoin,¹⁰ take the query $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \triangleright \bar{E}$, which gives:

$$\begin{aligned}
(\bar{A} \bowtie_{\bar{n}} \bar{B}) \triangleright \bar{E} &= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \triangleright \bar{E} \\
&= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}, \\
&\quad \{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 3\}\} \\
&\quad \triangleright \{?y \rightarrow 3, ?u \rightarrow 3\} \\
&= \{\{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}\} \\
&= \{[\mu_1, \mu_2, \mu_v] \in (\bar{A} \triangleright \bar{E}) \bowtie \bar{B} \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \\
&= (\bar{A} \triangleright \bar{E}) \bowtie_{\bar{n}} \bar{B}
\end{aligned}$$

Aside from the conditions previously discussed for the first claim, such equalities apply for antijoins so long as the right operand of the antijoin (in this case \bar{E}) does not “depend” on variables bound by the right operand of the similarity join (in this case \bar{B}), i.e., so long as all variables that can be bound by \bar{E} are either always bound by \bar{A} or never bound by \bar{B} . This relates to the condition that queries are *well-designed* [34].

3. Recalling that \bowtie denotes a left-join (or OPTIONAL in SPARQL), take the query $(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie \bar{E}$, which gives:

$$\begin{aligned}
(\bar{A} \bowtie_{\bar{n}} \bar{B}) \bowtie \bar{E} &= \{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \bowtie \bar{E} \\
&= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}, \\
&\quad \{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 3\}\} \\
&\quad \bowtie \{?y \rightarrow 3, ?u \rightarrow 3\} \\
&= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2, ?u \rightarrow 3\}, \\
&\quad \{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}\} \\
&= \{[\mu_1, \mu_2, \mu_v] \in (\bar{A} \bowtie \bar{E}) \bowtie \bar{B} \\
&\quad \bowtie \{\{?vk \rightarrow l_1(\llbracket \mathcal{V} \rrbracket_{\mu_2}^{\mu_1})\}\} \mid \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\} \\
&= (\bar{A} \bowtie \bar{E}) \bowtie_{\bar{n}} \bar{B}
\end{aligned}$$

The same conditions as applied before in the second case apply again here.

4. Let $\bar{f} = “?y \neq 4”$ be a selection condition (FILTER in SPARQL) that returns true for a solution μ if $?y \notin \text{dom}(\mu)$ or $\mu(?y) \neq 4$, and false otherwise: i.e., the filter removes solutions where $?y$ is

¹⁰So long as X and Y always return some variable in common, then an antijoin $X \triangleright Y$ is analogous to SPARQL’s MINUS. However, for $\mu_x \in X$ that shares no variable in common with any $\mu_y \in Y$, the MINUS operator will return μ_x while an antijoin will not (since $\mu_x \sim \mu_y$ for any $\mu_y \in Y$).

bound to 4. Now consider the query $\sigma_{\bar{f}}(\bar{A} \bowtie_{\bar{n}} \bar{B})$, as follows:

$$\begin{aligned}
\sigma_{\bar{f}}(\bar{A} \bowtie_{\bar{n}} \bar{B}) &= \sigma_{\bar{f}}(\{[\mu_1, \mu_2, \mu_v] \in \bar{A} \bowtie \bar{B} \bowtie \{\{v \rightarrow l_1([\mathcal{V}]_{\mu_2}^{\mu_1})\}\} | \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\}\}) \\
&= \sigma_{\bar{f}}(\{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}, \\
&\quad \{?x \rightarrow 4, ?y \rightarrow 4, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 3\}\}) \\
&= \{\{?x \rightarrow 2, ?y \rightarrow 3, ?z \rightarrow 1, ?w \rightarrow 1, ?vk \rightarrow 2\}\} \\
&= \{[\mu_1, \mu_2, \mu_v] \in \sigma_{\bar{f}}(\bar{A}) \bowtie \bar{B} \bowtie \{\{v \rightarrow l_1([\mathcal{V}]_{\mu_2}^{\mu_1})\}\} | \kappa_{l_1}^{\bar{B}}(\mu_1, \mu_2) \leq 1\}\} \\
&= \sigma_{\bar{f}}(\bar{A}) \bowtie_{\bar{n}} \bar{B}
\end{aligned}$$

Added to the base conditions of claim 1, such equalities only hold so long as all the variables that \bar{f} mentions are always bound in the left operand of the similarity join (in this case \bar{A}).

References

- [1] R. Battle and D. Kolas, Enabling the geospatial Semantic Web with Parliament and GeoSPARQL, *Semantic Web* 3(4) (2012), 355–370. doi:10.3233/SW-2012-0065.
- [2] F. Belleau, M.-A. Nolin, N. Tourigny, P. Rigault and J. Morissette, Bio2RDF: Towards a mashup to build bioinformatics knowledge systems, *Journal of Biomedical Informatics* 41(5) (2008), 706–716, <https://www.sciencedirect.com/science/article/pii/S1532046408000415>. doi:10.1016/j.jbi.2008.03.004.
- [3] C. Böhm, B. Braunmüller, F. Krebs and H.-P. Kriegel, *Epsilon Grid Order: An Algorithm for the Similarity Join on Massive High-Dimensional Data*, Association for Computing Machinery, New York, NY, USA, 2001, pp. 379–388, ISSN 0163-5808. doi:10.1145/376284.375714.
- [4] C. Böhm and F. Krebs, Supporting KDD applications by the k -nearest neighbor join, in: *Database and Expert Systems Applications*, V. Mařík, W. Retschitzegger and O. Štěpánková, eds, Springer Berlin Heidelberg, Berlin, Heidelberg, 2003, pp. 504–516. ISBN 978-3-540-45227-0. doi:10.1007/978-3-540-45227-0_50.
- [5] J.-P. Dittrich and B. Seeger, GESS: A scalable similarity-join algorithm for mining large data sets in high dimensional spaces, in: *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '01*, Association for Computing Machinery, New York, NY, USA, 2001, pp. 47–56. ISBN 158113391X. doi:10.1145/502512.502524.
- [6] V. Dohnal, C. Gennaro, P. Savino and P. Zezula, D-index: Distance searching index for metric data sets, *Multimedia Tools and Applications* 21(1) (2003), 9–33. doi:10.1023/A:1025026030880.
- [7] M. Ester, H.-P. Kriegel, J. Sander, X. Xu et al., A density-based algorithm for discovering clusters in large spatial databases with noise, in: *Proceedings of the 2nd ACM International Conference on Knowledge Discovery and Data Mining (KDD)*, Vol. 96, 1996, pp. 226–231.
- [8] S. Ferrada, B. Bustos and A. Hogan, IMGpedia: A linked dataset with content-based analysis of Wikimedia images, in: *The Semantic Web – ISWC 2017*, Springer International Publishing, Cham, 2017, pp. 84–93. ISBN 978-3-319-68204-4. doi:10.1007/978-3-319-68204-4_8.
- [9] S. Ferrada, B. Bustos and A. Hogan, Extending SPARQL with similarity joins, in: *The Semantic Web – ISWC 2020*, Springer International Publishing, Cham, 2020, pp. 201–217. ISBN 978-3-030-62419-4. doi:10.1007/978-3-030-62419-4_12.
- [10] M. Galkin, M. Vidal and S. Auer, Towards a multi-way similarity join operator, in: *New Trends in Databases and Information Systems*, Springer International Publishing, Cham, 2017, pp. 267–274. ISBN 978-3-319-67162-8. doi:10.1007/978-3-319-67162-8_26.
- [11] G. Giacinto, A nearest-neighbor approach to relevance feedback in content based image retrieval, in: *Proceedings of the 6th ACM International Conference on Image and Video Retrieval, CIVR '07*, Association for Computing Machinery, New York, NY, USA, 2007, pp. 456–463. ISBN 9781595937339. doi:10.1145/1282280.1282347.
- [12] R. Guerraoui, A.-M. Kermarrec, O. Ruas and F. Taïani, Fingerprinting big data: The case of KNN graph construction, in: *2019 IEEE 35th International Conference on Data Engineering (ICDE)*, 2019, pp. 1738–1741. doi:10.1109/ICDE.2019.00186.
- [13] S. Guha, R. Rastogi and K. Shim, CURE: An efficient clustering algorithm for large databases, *SIGMOD Rec.* 27(2) (1998), 73–84. doi:10.1145/276305.276312.
- [14] A. Guttman, *R-Trees: A Dynamic Index Structure for Spatial Searching*, Vol. 14, ACM, 1984.
- [15] S. Harris, A. Seaborne and E. Prud'hommeaux, SPARQL 1.1 Query Language, 2013, <https://www.w3.org/TR/sparql11-query/>.
- [16] A. Hogan, M. Mellotte, G. Powell and D. Stampouli, Towards fuzzy query-relaxation for RDF, in: *The Semantic Web: Research and Applications*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2012, pp. 687–702. doi:10.1007/978-3-642-30284-8_53.
- [17] E.H. Jacox and H. Samet, Metric space similarity joins, *ACM Trans. Database Syst.* 33(2) (2008). doi:10.1145/1366102.1366104.
- [18] G. Karypis, E.-H. Han and V. Kumar, Chameleon: Hierarchical clustering using dynamic modeling, *Computer* 32(8) (1999), 68–75. doi:10.1109/2.781637.

- [19] L. Kaufman and P.J. Rousseeuw, *Finding Groups in Data: An Introduction to Cluster Analysis*, John Wiley & Sons, 2009.
- [20] C. Kiefer, A. Bernstein and M. Stocker, The fundamentals of iSPARQL: A virtual triple approach for similarity-based Semantic Web tasks, in: *The Semantic Web*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2007, pp. 295–309. doi:10.1007/978-3-540-76298-0_22.
- [21] A. Ławrynowicz, Query results clustering by extending SPARQL with CLUSTER BY, in: *On the Move to Meaningful Internet Systems: OTM 2009 Workshops*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 826–835. ISBN 978-3-642-05290-3. doi:10.1007/978-3-642-05290-3_101.
- [22] C. Li, M. Wang, L. Lim, H. Wang and K.C.-C. Chang, Supporting ranking and clustering as generalized order-by and group-by, in: *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data, SIGMOD '07*, Association for Computing Machinery, New York, NY, USA, 2007, pp. 127–138. ISBN 9781595936868. doi:10.1145/1247480.1247496.
- [23] H. Li, X. Zhang and S. Wang, Reduce pruning cost to accelerate multimedia kNN search over MBRs based index structures, in: *2011 Third International Conference on Multimedia Information Networking and Security*, 2011, pp. 55–59. doi:10.1109/MINES.2011.85.
- [24] J. MacQueen et al., Some methods for classification and analysis of multivariate observations, in: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, Vol. 1, Oakland, CA, USA, 1967, pp. 281–297.
- [25] S. Mahran and K. Mahar, Using grid for accelerating density-based clustering, in: *2008 8th IEEE International Conference on Computer and Information Technology*, 2008, pp. 35–40. doi:10.1109/CIT.2008.4594646.
- [26] M. Muja and D.G. Lowe, Fast approximate nearest neighbors with automatic algorithm configuration, in: *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISSAPP)*, INSTICC Press, 2009, pp. 331–340.
- [27] G. Navarro, Analyzing metric space indexes: What for? in: *2009 Second International Workshop on Similarity Search and Applications*, 2009, pp. 3–10. doi:10.1109/SISAP.2009.17.
- [28] T. Neumann and G. Moerkotte, Characteristic sets: Accurate cardinality estimation for RDF queries with multiple joins, in: *2011 IEEE 27th International Conference on Data Engineering*, 2011, pp. 984–994. doi:10.1109/ICDE.2011.5767868.
- [29] A.N. Ngomo and S. Auer, LIMES – a time-efficient approach for large-scale link discovery on the web of data, in: *International Joint Conference on Artificial Intelligence (IJCAI)*, 2011, pp. 2312–2317.
- [30] R. Oldakowski and C. Bizer, SemMF: A framework for calculating semantic similarity of objects represented as RDF graphs, poster at ISWC. 2005.
- [31] B.C. Ooi, Spatial kd-tree: A data structure for geographic database, in: *Datenbanksysteme in Büro, Technik und Wissenschaft*, Springer, 1987, pp. 247–258. doi:10.1007/978-3-642-72617-0_17.
- [32] C. Ordonez and E. Omiecinski, Efficient disk-based k -means clustering for relational databases, *IEEE Transactions on Knowledge and Data Engineering* **16**(8) (2004), 909–921. doi:10.1109/TKDE.2004.25.
- [33] R. Paredes and N. Reyes, Solving similarity joins and range queries in metric spaces with the list of twin clusters, *Journal of Discrete Algorithms* **7**(1) (2009), 18–35. doi:10.1016/j.jda.2008.09.012.
- [34] J. Pérez, M. Arenas and C. Gutiérrez, Semantics and complexity of SPARQL, *ACM TODS* **34**(3) (2009), 16:1–16:45. doi:10.1145/1567274.1567278.
- [35] A. Petrova, E. Sherkhonov, B. Cuenca Grau and I. Horrocks, Entity comparison in RDF graphs, in: *The Semantic Web – ISWC 2017*, Springer International Publishing, Cham, 2017, pp. 526–541. ISBN 978-3-319-68288-4. doi:10.1007/978-3-319-68288-4_31.
- [36] E. Prud'hommeaux, SPARQL query language for RDF, 2008, <http://www.w3.org/TR/rdf-sparql-query/>.
- [37] L. Qi, H.T. Lin and V. Honavar, Clustering remote RDF data using SPARQL update queries, in: *2013 IEEE 29th International Conference on Data Engineering Workshops (ICDEW)*, 2013, pp. 236–242. doi:10.1109/ICDEW.2013.6547456.
- [38] E. Schubert and P.J. Rousseeuw, Fast and eager k -medoids clustering: $O(k)$ runtime improvement of the PAM, CLARA, and CLARANS algorithms, *Information Systems* **101** (2021), 101804. doi:10.1016/j.is.2021.101804.
- [39] E. Schubert, J. Sander, M. Ester, H.P. Kriegel and X. Xu, DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN, *ACM Trans. Database Syst.* **42**(3) (2017). doi:10.1145/3068335.
- [40] M.A. Sherif and A.N. Ngomo, A systematic survey of point set distance measures for link discovery, *Semantic Web* **9**(5) (2018), 589–604. doi:10.3233/SW-170285.
- [41] K. Shim, R. Srikant and R. Agrawal, *High-Dimensional Similarity Joins*, 2002, pp. 156–171. doi:10.1109/69.979979.
- [42] Y.N. Silva, W.G. Aref and M.H. Ali, Similarity group-by, in: *2009 IEEE 25th International Conference on Data Engineering*, 2009, pp. 904–915. doi:10.1109/ICDE.2009.113.
- [43] Y.N. Silva, S.S. Pearson and J.A. Cheney, Database similarity join for metric spaces, in: *Similarity Search and Applications*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 266–279. doi:10.1007/978-3-642-41062-8_27.
- [44] J. Volz, C. Bizer, M. Gaedke and G. Kobilarov, Discovering and maintaining links on the web of data, in: *The Semantic Web – ISWC 2009*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2009, pp. 650–665. ISBN 978-3-642-04930-9. doi:10.1007/978-3-642-04930-9_41.
- [45] D. Vrandečić and M. Krötzsch, Wikidata: A free collaborative knowledgebase, *Communications of the ACM* **57** (2014), 78–85. doi:10.1145/2629489.
- [46] W. Wang, J. Yang, R. Muntz et al., STING: A statistical information grid approach to spatial data mining, in: *Very Large Data Bases Conference (VLDB)*, Vol. 97, 1997, pp. 186–195.
- [47] X. Xu, M. Ester, H.-P. Kriegel and J. Sander, A distribution-based clustering algorithm for mining in large spatial databases, in: *Proceedings 14th International Conference on Data Engineering*, 1998, pp. 324–331. doi:10.1109/ICDE.1998.655795.
- [48] P.N. Yianilos, Data structures and algorithms for nearest neighbor search in general metric spaces, in: *Symposium on Discrete Algorithms (SODA)*, Vol. 93, 1993, pp. 311–321.
- [49] P. Zezula, G. Amato, V. Dohnal and M. Batko, *Similarity Search: The Metric Space Approach*, Vol. 32, Springer Science & Business Media, 2006.

- [50] X. Zhai, L. Huang and Z. Xiao, Geo-spatial query based on extended SPARQL, in: *2010 18th International Conference on Geoinformatics*, 2010, pp. 1–4. doi:[10.1109/GEOINFORMATICS.2010.5567605](https://doi.org/10.1109/GEOINFORMATICS.2010.5567605).
- [51] C. Zhang and Y. Huang, Cluster by: A new SQL extension for spatial data aggregation, in: *Proceedings of the 15th Annual ACM International Symposium on Advances in Geographic Information Systems, GIS '07*, Association for Computing Machinery, New York, NY, USA, 2007. ISBN 9781595939142. doi:[10.1145/1341012.1341077](https://doi.org/10.1145/1341012.1341077).
- [52] T. Zhang, R. Ramakrishnan and M. Livny, BIRCH: An efficient data clustering method for very large databases, *SIGMOD Rec.* **25**(2) (1996), 103–114. doi:[10.1145/235968.233324](https://doi.org/10.1145/235968.233324).
- [53] A. Zimek, E. Schubert and H.-P. Kriegel, A survey on unsupervised outlier detection in high-dimensional numerical data, *Science Journal* **5**(5) (2012), 363–387. doi:[10.1002/sam.11161](https://doi.org/10.1002/sam.11161).