# Practical Construction of $\boldsymbol{k}$-Nearest Neighbor Graphs in Metric Spaces 

Rodrigo Paredes ${ }^{1}$, Edgar Chávez ${ }^{2}$, Karina Figueroa ${ }^{1}$, and Gonzalo Navarro ${ }^{1}$<br>${ }^{1}$ Center for Web Research, Dept. of Computer Science, University of Chile.<br>${ }^{2}$ Univ. Michoacana, Morelia, Mich. Mexico.


#### Abstract

Let $\mathbb{U}$ be a set of elements and $d$ a distance function defined among them. Let $N N_{k}(u)$ be the $k$ elements in $\mathbb{U}-\{u\}$ which have the smallest distance towards $u$. The $k$-nearest neighbor graph ( $k \mathrm{NNG}$ ) is a weighted directed graph $G(\mathbb{U}, E)$ such that $E=\left\{(u, v), v \in N N_{k}(u)\right\}$. We focus on the metric space context, so $d$ is a metric. Several $k$ NNG construction algorithms are known, but they are not suitable to general metric spaces. We present a general methodology to construct $k$ NNGs that exploits several features of metric spaces, requiring empirically around $O\left(n^{1.27}\right)$ distance computations for low and medium dimensional spaces, and $O\left(n^{1.90}\right)$ for high dimensional ones.


Keywords: Graph Algorithms, Metric Spaces, Nearest Neighbors.

## 1 Introduction

Let $\mathbb{U}$ be a set of elements and $d$ a distance function defined among them. Let $N N_{k}(u)$ be the $k$ elements in $\mathbb{U}-\{u\}$ which have the smallest distance towards $u$ according to the function $d$. The $k$-nearest neighbor graph ( $k \mathrm{NNG}$ ) is a weighted directed graph $G(\mathbb{U}, E)$ connecting each element to its $k$ nearest neighbors, thus $E=\left\{(u, v), v \in N N_{k}(u)\right\}$. Building the $k N N G$ is a direct generalization of the well known all-nearest-neighbor (ANN) problem, so ANN corresponds to the 1 NNG construction problem. $k$ NNGs are central in many applications, for instance cluster and outlier detection $[14,4]$, VLSI design, spin glass and other physical process simulations [6], pattern recognition [12], and query or document recommendation systems [3].

There are several algorithms to construct $k$ NNGs $[1,5,6,9-11,13,15,18,19$, 21]. Nevertheless, most of them assume that nodes are points in $\mathbb{R}^{D}$ and that $d$ is the Euclidean or some $L_{p}$ Minkowski distance. However, this is not the case in several applications where $k$ NNGs are required. An example is collaborative filters for Web searching such as query or document recommendation systems, where $k$ NNGs are used to find clusters of similar queries, to later improve the quality of the results shown to the final user by exploiting cluster properties [3].

To handle this problem one must resort to a more general model called metric spaces. A metric space is a pair $(\mathbb{X}, d)$, where $\mathbb{X}$ is the universe of objects and $d$ is a distance function among them that satisfies the triangle inequality.

Another appealing problem in metric spaces is similarity searching [8]. Given a finite metric database $\mathbb{U} \subseteq \mathbb{X}$, the goal is to build an index for $\mathbb{U}$ such that
later, given a query object $q \in \mathbb{X}$, one can find the elements of $\mathbb{U}$ close to $q$ with as few distance computations as possible. See [8] for a comprehensive survey.

We have already demonstrated the searching capabilities of the $k$ NNG in general metric spaces [20], where we give $k$ NNG-based search algorithms with practical applicability in low-memory scenarios or in metric spaces of medium or high dimensionality. Therefore, in this paper we focus on a methodology to construct $k$ NNGS in general metric spaces. Based on our methodology, we propose two algorithms requiring empirically around $O\left(n^{1.27}\right)$ distance computations for low and medium dimensionality spaces, and $O\left(n^{1.90}\right)$ for high dimensionality ones. Note that a naive construction requires $O\left(n^{2}\right)$ distance evaluations.

### 1.1 A summary of metric space searching

Given the universe of objects $\mathbb{X}$, a metric space is a pair $(\mathbb{X}, d)$, where $d: \mathbb{X} \times \mathbb{X} \rightarrow$ $\mathbb{R}^{+}$is any distance function in $\mathbb{X}$ that is symmetric and satisfies the triangle inequality. Some examples are $\left(\mathbb{R}^{D}, L_{p}\right)$; the space of strings under the edit distance; or the space of documents under the cosine distance. Note that objects do not necessarily have coordinates.

The metric database is a finite set $\mathbb{U} \subseteq \mathbb{X}, n=|\mathbb{U}|$. A similarity query is an object $q \in \mathbb{X}$, and can be of two basic types: $(a)$ Range query $(q, r)$ : retrieve all objects $u \in \mathbb{U}$ such that $d(u, q) \leq r ; \quad$ and $(b) k$-Nearest neighbor query $N N_{k}(q)$ : retrieve the $k$ objects in $\mathbb{U}$ closest to $q$ according to the distance $d$. A $N N_{k}(q)$ algorithm is called range-optimal [16] if it uses the same number of distance evaluations as the equivalent range query whose radius retrieves exactly $k$ objects. We call this radius covering radius.

An index is a data structure built over $\mathbb{U}$, by selecting some precomputed cells from the whole $\mathbb{U} \times \mathbb{U}$ distance matrix, that permits solving the above queries without comparing $q$ against every element in $\mathbb{U}$. There are two kinds of indices: pivot based and compact partition based. Search algorithms use the index and some distance evaluations to discard (using the triangle inequality) as many objects as they can to produce a small candidate set that could be relevant to $q$. Later, they exhaustively check the candidate set computing distances from $q$ towards each candidate to obtain the query result.

As the distance is considered expensive to compute, it is customary to use the number of distance evaluations as the complexity measure both for index construction and object retrieving. For instance, each computation of the cosine distance in the document metric space takes 1.4 msecs in our machine (Pentium IV of 2 GHz ). This is really costly even compared with the operations introduced by the graph, such as shortest path computation using Dijkstra's algorithm.

Many authors agree that the proximity query cost worsens quickly as the dimension of the space grows, which is called the curse of dimensionality. However, there is not and accepted criterion to measure the dimension. In $\mathbb{R}^{D}$ with points chosen randomly, the dimension is simply $D$. In other metric spaces (strings, documents, $\mathbb{R}^{D}$ where points are not uniformly distributed, and so on), the dimension can be empirically defined in terms of the distance histogram [8]. In general terms, the dimension grows as the histogram concentrates.

### 1.2 Related work on $k$ NNG construction

The naive approach to construct $k$ NNGs uses $\frac{n(n-1)}{2}=O\left(n^{2}\right)$ distance computations and $O(k n)$ memory. For each $u \in \mathbb{U}$ we compute the distance towards all the others, and select the $k$ lowest-distance objects. However, there are several alternatives to speed up the procedure. The proximity properties of the Voronoi diagram [2] or its dual, the Delaunay triangulation, allow solving the problem more efficiently. The ANN problem can be optimally solved in $O(n \log n)$ time in the plane [13] and in $\mathbb{R}^{D}$ for any fixed $D[9,21]$, but the constant depends exponentially on $D$. In $\mathbb{R}^{D}, k$ NNGs can be built in $O(n k \log n)$ time [21] and even in $O(k n+n \log n)$ time [5, 6, 11]. Approximation algorithms for the problem have also been proposed [1]. Nevertheless, none of these alternatives, except the naive one, is suitable for metric spaces, as they use coordinate information that is not necessarily available in general metric spaces.

Clarkson states the first generalization of ann to general metric spaces [10], where the problem is solved using randomization in $O\left(n \log ^{2} n \log ^{2} \Gamma(\mathbb{U})\right)$ expected time, where $\Gamma(\mathbb{U})$ is the distance ratio between the farthest and closest pairs of points in $\mathbb{U}$. The author argues that in practice $\Gamma(\mathbb{U})=n^{O(1)}$, in which case the approach is $O\left(n \log ^{4} n\right)$ time. However, the analysis needs a sphere packing bound in the metric space. Otherwise the cost must be multiplied by "sphere volumes", that is, extra factors that are similar to those that, in $\mathbb{R}^{D}$, are exponential on $D$. Moreover, the algorithm needs $\Omega\left(n^{2}\right)$ space for high dimensions, which is too much for practical applications.

In [15], another technique for general metric spaces is given. It solves $n$ range queries of decreasing radius by using a pivot-based index. As it is well known, the performance of pivot-based algorithms worsens quickly as the dimension of the space grows, limiting the applicability of this technique. Our pivot based algorithm (Section 2.4) can be seen as an improvement over this technique.

Recently, Karger and Ruhl present the metric skip list [18], an index that uses $O(n \log n)$ space and can be constructed with $O(n \log n)$ distance computations. The index answers $N N_{1}(q)$ queries using $O(\log n)$ distance evaluations with high probability. Krauthgamer and Lee introduce navigating nets [19], another index that can be constructed also with $O(n \log n)$ distance computations, yet using $O(n)$ space, and which gives an $(1+\epsilon)$-approximation algorithm to solve $N N_{1}(q)$ queries in time $O(\log n)+(1 / \epsilon)^{O(1)}$. Both of them could serve to solve the ANNG problem with $O(n \log n)$ distance computations but not to build $k$ NNGs. In addition, the hidden constants are exponential on the intrinsic dimension, which makes these approaches useful only in low dimensional metric spaces.

## 2 Our methodology

We are interested in practical $k$ NNG construction algorithms for general metric spaces. This problem is equivalent to solve $n N N_{k}(u)$ queries for all $u \in \mathbb{U}$. Thus, a straightforward solution has two stages: the first is to build some known metric index $\mathcal{I}$ [8], and the second is to use $\mathcal{I}$ to solve the $n$ queries. However, this basic scheme can be improved if we take into account these observations:

- We are solving queries for all the elements in $\mathbb{U}$, not for general objects in $\mathbb{X}$. If we solve the $n$ queries jointly we can share costs through the whole process. For instance, we can avoid some calculations by using the symmetry of $d$.
- We can upper bound some distances by computing shortest paths over the $k N N G$ under construction, maybe avoiding their actual computation. So, we can use the very $k \mathrm{NNG}$ in stepwise refinements to improve the second stage.


### 2.1 The ingredients of the recipe

The main data structure. Along all the algorithm, we use the Neighbor Heap Array (NHA) to store the $k$ NNG under construction. NHA can be regarded as the union of priority queues $N H A_{u}$, of size $k$, for all $u \in \mathbb{U}$. At any point in the process $N H A_{u}$ will contain the $k$ elements closest to $u$ known up to then, and their distances to $u$. Formally, $N H A_{u}=\left\{\left(x_{i_{1}}, d\left(u, x_{i_{1}}\right)\right), \ldots,\left(x_{i_{k}}, d\left(u, x_{i_{k}}\right)\right)\right\}$ sorted by decreasing $d\left(u, x_{i_{j}}\right)$ ( $i_{j}$ is the $j$-th neighbor identifier $)$.

For each $u \in \mathbb{U}$, we initialize $N H A_{u}=\{(\perp, \infty), \ldots,(\perp, \infty)\},\left|N H A_{u}\right|=k$. Let $\operatorname{cur} C R_{u}=d\left(u, x_{i_{1}}\right)$ be the current covering radius of $u$, that is, the distance from $u$ towards its current farthest neighbor candidate in $N H A_{u}$.

In the first stage, every distance computed to build the index $\mathcal{I}$ populates $N H A$. In the second, we refine $N H A$ with the following distance computations. We must ensure that $\left|N H A_{u}\right|=k$ upon successive additions. Hence, if we find some object $v$ such that $d(u, v)<\operatorname{curCR} R_{u}$, before adding $(v, d(u, v))$ to $N H A_{u}$ we extract the farthest candidate from $N H A_{u}$. This progressively reduces curCR $R_{u}$ from $\infty$ to the real covering radius. At the end, NHA stores the $k$ NNG of $\mathbb{U}$.

Using NHA as a graph. Once we calculate $d(u, v)$, if $d(u, v) \geq c u r C R_{u}$ we discard $v$ as a candidate for $N H A_{u}$, and moreover, we can discard all objects $w$ such that $d(v, w) \leq d(u, v)-\operatorname{cur} C R_{u}$ because of the triangle inequality. Unfortunately, we do not necessarily have stored $d(v, w)$. However, we can compute shortest paths over $N H A$ from $v$ to all $w \in \mathbb{U}$, and due to the triangle inequality, the sum of edge weights traversed in this path, $d_{N H A}(v, w)$, is an upper bound to $d(v, w)$. So, if $d(u, v) \geq c u r C R_{u}$, we also discard all objects $w$ such that $d_{N H A}(v, w) \leq d(u, v)-c u r C R_{u}$.
$\boldsymbol{d}$ is symmetric. Every time a distance $d(u, v)$ is computed, we check both $d(u, v)<c u r C R_{u}$ for adding $(v, d(u, v))$ to $N H A_{u}$, and $d(u, v)<c u r C R_{v}$ for adding $(u, d(u, v))$ to $N H A_{v}$. This can both reduce $\operatorname{curCR}$, , and cheapen the future query for $v$, even when we are solving neighbors for another object.
$\mathbb{U}$ is fixed. Assume we are solving query $N N_{k}(u)$, we have to check some already solved object $v$, and $\operatorname{curCR} R_{u} \leq \operatorname{curCR}$. Then, if $u \notin N N_{k}(v) \Rightarrow d(u, v) \geq$ $\operatorname{cur}^{\cos } R_{v}$, so $v \notin N N_{k}(u)$. Otherwise, if $u \in N N_{k}(v)$, then we already computed $d(u, v)$. Then, in those cases we avoid to compute $d(u, v)$. Fig. 1(a) illustrates.

Check Order Heap ( $\mathbf{C O H}$ ). We create the priority queue $C O H=\left\{\left(u, c u r C R_{u}\right)\right.$, $u \in \mathbb{U}\}$ to complete $N N_{k}(u)$ queries in increasing $\operatorname{curCR} R_{u}$ order, because a small radius query has larger discriminative power and produces candidates that are closer to the query $u$. This reduces the CPU time and increases the chance of
$K \mathbf{N N}$ (Integer $k$, ObjectSet $\mathbb{U}$ )
Stage 1: Initialize NHA and construct the index $\mathcal{I}$

1. For each $u \in \mathbb{U}$ Do $N H A_{u} \leftarrow\{(\perp, \infty), \ldots,(\perp, \infty)\} / / k$ pairs
2. Create $\mathcal{I}$, all computed distances populate symmetrically $N H A$

Stage 2: Complete the $N N_{k}(u)$ for all $u \in \mathbb{U}$
3. $C O H \leftarrow\left\{\left(u, \operatorname{cur} C R_{u}\right), u \in \mathbb{U}\right\}$

For each $\left(u, \mathrm{curCR}_{u}\right) \in \mathrm{COH}$, in increasing $\mathrm{curCR}_{u}$ order Do
Create the candidate set $\mathcal{C}$ according to $\mathcal{I} / /$ exclude $N H A_{u}$
While $\mathcal{C} \neq \emptyset$ Do
$c \leftarrow$ extract a candidate from $\mathcal{C}$
If "U is fixed" does not apply for $u$ and $c$ Then
$d_{u, c} \leftarrow d(u, c)$, try to insert $c$ into $N H A_{u}$
try to insert $u$ into $N H A_{c}$, update $c$ in COH (symmetry)
use NHA as a graph and $\mathcal{I}$ to discard objects from $\mathcal{C}$
(a) $\mathbb{U}$ is fixed.
11.
12. Return $N H A$ as a graph
(b) Sketch of the methodology.

Fig. 1. In 1(a), assume we are solving $u, v$ is already solved, and $c u r C R_{u} \leq c u r C R_{v}$. On the top, if $u \notin N N_{k}(v) \Rightarrow d(u, v) \geq \operatorname{cur} C R_{v} \geq \operatorname{cur} C R_{u}$. On the bottom, if $u \in N N_{k}(v)$, we already computed $d(u, v)$. Then, in those cases we avoid computing $d(u, v)$. In 1(b), we sketch the methodology.
improving candidate sets in $N H A$ for other objects $v$, using the symmetry of $d$ as explained to reduce $\operatorname{cur}_{2} R_{v}$. This, in turn, changes the position of $v$ in COH .

The recipe. We split the process into two stages. The first is to build $\mathcal{I}$ to preindex the objects. The second is to use $\mathcal{I}$ and all the ingredients to solve the $N N_{k}(u)$ queries for all $u \in \mathbb{U}$. Fig. 1(b) depicts the methodology.

For practical reasons, we allow that our algorithms use at most $O(n(\log n+k))$ memory both to index $\mathbb{U}$ and to store the $k$ NNG under construction.

### 2.2 The resulting algorithms

Based on our methodology, we propose two $k$ NNG construction algorithms focused on decreasing the total number of distance computations. They are:

1. Recursive partition based algorithm: In the first stage, we build a preindex by performing a recursive partitioning of the space. In the second stage, we complete the $N N_{k}(u)$ queries using the order induced by the partitioning.
2. Pivot based algorithm: In the preindexing stage, we build the pivot index. Later, we complete the $N N_{k}(u)$ queries by performing range-optimal queries.

The experiments confirm that these algorithms are efficient. For instance, in the string space, the pivot-based algorithm requires CPU time of the empirical form $O\left(n^{1.85}\right)$, and $O\left(n^{1.26}\right)$ in distance computations. In the high-dimensional document space, the recursive partition-based algorithm requires empirically $O\left(n^{1.955}\right)$ both in distance computations and CPU time.

### 2.3 Recursive partition based algorithm

This algorithm is based on using a preindex slightly different to the Bisector Tree $(B S T)$ [17]. We call our modified $B S T$ the Division Control Tree ( $D C T$ ), which is a binary tree representing the shape of the partitioning. The $D C T$ node structure is $\{p, l, r, p r\}$, which represents the parent, left and right children, and partition radius of the node, respectively. The partition radius is the distance from the node towards the farthest node of its partition. (With respect to the $B S T$ structure, we have added the pointer $p$ to easily navigate trough the tree.)

For simplicity we use the same name for the node and for its representative in the $D C T$. Then, given a node $u \in \mathbb{U}, u_{p}, u_{l}$, and $u_{r}$, refer to nodes that are the parent, left child, and right child of $u$ in the $D C T$, respectively, and also to their representative nodes in $\mathbb{U}$. Finally, $u_{p r}$ refers to the partition radius of $u$.

In this algorithm, we use $O(k n)$ space to store the $N H A$ and $O(n)$ to store the $D C T$. The remaining memory is used as a cache of computed distances, $C D$, whose size is limited to $O(n \log n)$. Thus, every time we need to compute a distance, we check if it is present in $C D$, in which case we return the stored value. Note that the $C D \subset \mathbb{U}^{2} \times \mathbb{R}^{+}$can also be seen as graph of all stored distances. The criterion to insert distances into $C D$ depends on the stage (see later). Once we complete the $N N_{k}(u)$, we remove its adjacency list from $C D$.

First stage: construction of $\boldsymbol{D C T}$. We partition the space recursively to construct the $D C T$, and populate symmetrically $N H A$ and $C D$ with all the computed distances. The $D C T$ is built as follows. Given the node root and the set $S$, we choose children objects $l$ and $r$ from $S$. Then we generate two subsets: $S_{l}$, objects nearest to $l$, and $S_{r}$, objects nearest to $r$. Finally, we compute both partition radii. The recursion follows with $\left(l, S_{l}\right)$ and $\left(r, S_{r}\right)$, finishing when $|S|<2$. Once we finish the division, leaves in the $D C T$ have partition radii 0 . The $D C T$ root is fictitious, having no equivalent in $\mathbb{U}$, and partition radius $\infty$.

Since the $D C T$ has $n$ nodes, its expected height is $2 \ln n$ (the $D C T$ construction is statistically identical to populating a binary search tree). For each $D C T$ level, each node computes two distances towards the splitting nodes, which accounts for $2 n$ distances per level. So, we expect to compute $4 n \ln n$ distances in the partitioning. As we store 2 edges per distance, we need to store $8 n \ln n$ in $C D$. Hence, we fix the maximum space of $C D$ as $8 n \ln n=O(n \log n)$.

Solving $\boldsymbol{N} \boldsymbol{N}_{\boldsymbol{k}}(\boldsymbol{u})$ queries with $\boldsymbol{D C T}$. The construction of $D C T$ ensures that every node has already computed distances to all of its ancestors, its ancestor's siblings, and its parent descent. Then, to finish the $N N_{k}(u)$ query, it is enough to check whether there are relevant objects in all the descendants of $u$ 's ancestors' siblings. This corresponds to white nodes and subtrees in Fig. 2(a). Nevertheless, the $D C T$ allows us to avoid some work. Assume we are checking whether $v$ is relevant to $u$, and the balls ( $u, c u r C R_{u}$ ) and ( $v, v_{p r}$ ) do not intersect each other, then we discard $v$ and its partition. Otherwise, we recursively check children $v_{l}$ and $v_{r}$. Fig. 2(b) illustrates this.

Hence, in the candidate set $\mathcal{C}$, it suffices to manage the set of ancestors' siblings, and if it is not possible to discard the whole sibling's partition we add


Fig. 2. Using the $D C T$ to solve $N N_{k}(q)$ queries. In 2(a), $u$ has been compared with all black nodes and all the descent of its parent. To finish the query, we just process white nodes and subtrees. In $2(\mathrm{~b})$, as $d(u, v) \leq c u r C R_{u}+v_{p r}$ the partition of $v$ intersects the ball ( $u$, cur $C R_{u}$ ), so we recursively check children $v_{l}$ and $v_{r}$. As $v_{r}$ 's partition does not intersect the ball ( $u, \operatorname{curCR} R_{u}$ ), we discard $v_{r}$ and its partition. However, we continue the checking on $v_{l}$ 's partition as it intersects the ball ( $u$, cur $C R_{u}$ ).
its children to $\mathcal{C}$. Since it is more likely to discard small partitions, we process $\mathcal{C}$ in order of increasing radius. This agrees with the intuition that the partition radius of $u$ 's parent's sibling is likely the smallest of $\mathcal{C}$, and that some of its descendants could be relevant to $u$.

Second stage: Completing the queries. As $C D$ can be seen as a graph, we use $N H A \cup C D$ to upper bound distances: when $d(u, v) \geq c u r C R_{u}$, we discard objects $w$ such that their shortest path $d_{N H A \cup C D}(v, w) \leq d(u, v)-c u r C R_{u}$. We do this by adding them to $\mathcal{C}$ marked as EXTRACTED.

In this stage, if we have available space in $C D$, we cache all the computed distances small enough so as to get into their respective queues in $N H A$, since these distances can be used in future symmetric queries. Note that adding distances to $C D$ without considering the space limitation could increase its size beyond control, as it is shown by the following average case analysis. With probability $\frac{n-k}{n}$, a random distance is greater than the $k$-th shortest one (thus, not stored), and with probability $\frac{k}{n}$ it is lower, then it is stored in $C D$ using one cell. The base case uses $k$ cells for the first distances. Then, the recurrence for the average case of edge insertions for each $N H A_{u}$ is: $T(n, k)=T(n-1, k)+\frac{k}{n}, T(k, k)=k$. We obtain $T(n, k)=k\left(H_{n}-H_{k}+1\right)=O\left(k \log \frac{n}{k}\right)$. As we have $n$ priority queues, if we do not consider the limitation, we could use $O\left(n k \log \frac{n}{k}\right)$ memory cells, which can be an unpractical memory requirement.

Finally, we combine all of these ideas to complete the $N N_{k}(u)$ queries for all nodes in $\mathbb{U}$. We begin by creating the priority queue COH . Then, for each node $u$ picked from $C O H$ we do the following. We add the edges of $N H A_{u}$ to $C D_{u}$, where $C D_{u}$ refers to the adjacency list of $u$ in $C D$. (Due to the size limitation it is likely that some of the $u$ 's current neighbors do not belong $C D_{u}$.) Then, we compute shortest paths from all $u$ 's ancestors discarding as many objects as we can. Then, we finish the query, and finally, delete $C D_{u}$. Given node $u$, to finish its query, we first add all $u$ 's ancestors to $\mathcal{C}$. Later, we take objects $c$ from $\mathcal{C}$ in increasing $c_{p r}$ order, and process $c$ according one of the following rules:

1. If $c$ was already marked as EXTRACTED, we add children $\left\{c_{l}, c_{r}\right\}$ to $\mathcal{C}$;
2. If " $\mathbb{U}$ is fixed" applies for $c$ and $u$, and $d(u, c) \notin C D$, we add $\left\{c_{l}, c_{r}\right\}$ to $\mathcal{C}$; or
3. If we have $d(u, c)$ stored in $C D$, we retrieve it, else we compute it, and use the symmetry of $d$. Then, if $d(u, c)<c u r C R_{u}+c_{p r}$, we have region intersection, so we add $\left\{c_{l}, c_{r}\right\}$ to $\mathcal{C}$. Next, we use $N H A \cup C D$ as a graph computing shortest paths from $c$ to discard as many object as we can.

### 2.4 Pivot-based algorithm

Pivot-based algorithms have good performance in low dimensional spaces, but worsen quickly as the dimension grows. However, our methodology compensates this failure in medium and high dimensions. In this algorithm we use $O(k n)$ space in NHA and $O(n \log n)$ space to store the pivot index.

First stage: construction of the pivot index. We select at random a set of pivots $\mathcal{P}=\left\{p_{1}, \ldots, p_{|\mathcal{P}|}\right\} \subseteq \mathbb{U}$, and store a table of $|\mathcal{P}| n$ distances $d\left(p_{j}, u\right), j \in$ $\{1, \ldots,|\mathcal{P}|\}, u \in \mathbb{U}$. We give the same space in bytes to the table as that of the cache of distances and the division control tree of the recursive based algorithm. Then, in our implementation we use $|\mathcal{P}|=12 \ln n+2.5=O(\log n)$.

Solving $N N_{k}(u)$ queries with the pivot table. To perform a range-optimal query for $u$ we use $\mathcal{C}$ as an array to store maximum lower bounds of distances from $u$ to other objects. Because of the triangle inequality, for each $v \in \mathbb{U}$ and $p \in \mathcal{P}$, $|d(v, p)-d(u, p)|$ is a lower bound of $d(u, v)$. Let $\mathcal{C}_{v}=\max _{p \in \mathcal{P}}\{|d(v, p)-d(u, p)|\}$. So, we can discard non-relevant objects $v$ such that $\mathcal{C}_{v} \geq c u r C R_{u}$.

Then, we store $\mathcal{C}$ values in a priority queue $\operatorname{SDiff}=\left\{\left(v, \mathcal{C}_{v}\right), v \in \mathbb{U}-(\mathcal{P} \cup\right.$ $\left.\left.N H A_{u} \cup\{u\}\right)\right\}$. For each object $v$ picked from SDiff by ascending $\mathcal{C}_{v}$, we check if $\mathcal{C}_{v}<\operatorname{curCR}$. In such case, when " $\mathbb{U}$ is fixed" applies for $u$ and $v$ we avoid the distance computation and process the next node, else we compute the distance $d(u, v)$. So, if $d(u, v)<c u r C R_{u}$ we add $v$ to $N H A_{u}$ (this could reduce $c u r C R_{u}$ ). Moreover, using the symmetry of $d$, we can refine $N H A_{v}$ and consequently update $v$ in $C O H$. Finally, we use $N H A$ as a graph computing shortest path from $v$ to extract from SDiff as many object as we can. Each $N N_{k}(u)$ query finishes when we reach a node $v$ such that $\mathcal{C}_{v} \geq c u r C R_{u}$, or SDiff gets empty.

Second stage: Completing the queries. Since pivots $p \in \mathcal{P}$ compute distances towards all objects, once we compute the table, they have already solved their $k$ nearest neighbors. So, we only have to complete $n-|\mathcal{P}|$ range-optimal queries for objects $u \in \mathbb{U}-\mathcal{P}$. Notice that because of the symmetry of $d$, these objects already have candidates in their respective queues in $N H A$.

## 3 Experimental results

We have tested our algorithms on synthetic and real-world metric spaces. The first synthetic set is formed by 65,536 points uniformly distributed in the metric space $\left([0,1]^{D}, L_{2}\right)$ (the unitary real $D$-dimensional cube with Euclidean distance). This space allows us to measure the effect of the space dimension $D$ on our
algorithms. The second set is formed by 65,536 points in a 20 -dimensional space with Gaussian distribution forming 256 clusters randomly placed in $\left([0,1]^{20}, L_{2}\right)$. We consider three standard deviations to make more crisp or more fuzzy clusters $(\sigma=0.1,0.2,0.3)$. Of course, we have not used the fact that vectors have coordinates, but have treated them as abstract objects.

The first real-world set is the string metric space under the edit distance, a discrete function that measures the minimum number of character insertions, deletions and replacements needed to make the strings equal. We index a random subset of 65,536 words taken from an English dictionary. The second set is the document space under the cosine distance, a function that measures the angle between two documents when they are represented as vectors in a highdimensional vector model. We index a random subset of 1,215 English documents taken from the TREC-3 collection.

Experiments were run on an Intel Pentium IV of 2 GHz and 512 MB of RAM. We measure distance evaluations and CPU time. For shortness we have called the basic $k$ NNG construction algorithm $K \mathrm{NNb}$, the recursive partition based algorithm KNNrp, and the pivot based algorithm KNNpiv. We are not aware of any published $k$ NNG practical implementation for general metric spaces.

We summarize our experimental results in Fig. 3, where we show distance computations per element, and Table 1 for the least square fittings. The dependence on $k$ turns out to be so mild that we neglect $k$ in the fittings. From them, we induce empirical complexities of the form $O\left(n^{\alpha}\right)$, however these are not mean to constitute asymptotic complexity figures, but rather performance predictors.

Figs. 3(a), 3(b) and 3(c) show experimental results for $\mathbb{R}^{D}$. Fig. 3(c) shows that, as $D$ grows, the performance of our algorithms degrade, phenomenon known as the curse of dimensionality. For instance, for $D=4, K N N p i v ~ u s e s$ $O\left(n^{1.10}\right)$ distance evaluations, but for $D=24$, it is $O\left(n^{1.96}\right)$ distance evaluations. Fig. 3(a) shows that for all dimensions our algorithms are subquadratic in distance evaluations, instead of $K \mathrm{NNb}$ which is always $O\left(n^{2}\right)$. For low and medium dimensions ( $D \leq 16$ ) ours have better performance than $K \mathrm{NNb}$, being $K \mathrm{NNpiv}$ the best of ours. Moreover, for lower dimensions $(D \leq 8)$ ours are only slightly superlinear. Fig. 3(b) shows a sublinear dependence on $k$ for all dimensions, however, KNNpiv is more sensitive to $k$ than $K N N r p$. Also, the dependence on $k$ diminishes as long as $D$ grows, although it is always monotonically increasing on $k$. Finally, it is shown that for $k \leq 4$, our algorithms behave better than $K \mathrm{NNb}$, even in high dimensional spaces ( $K$ NNpiv in $D=20$ ).

Figs. 3(e) and 3(f) show results in Gaussian space. For crisp clusters ( $\sigma=0.1$ ) the performance of our algorithms improves significantly, even for high values of $k$. It is interesting to note that for $k \leq 8$ our algorithms are more efficient than $K \mathrm{NNb}$ for the three variances. Again, $K \mathrm{NNpiv}$ has the best performance.

Figs. 3(g) and 3(h) show results for strings. The plots show that both KNNrp and $K N N p i v$ are subquadratic for all $k \in[2,128]$. For instance, for $n=65,536$, $K$ NNrp costs $28 \%$, and $K$ NNpiv just $8 \%$, of $K$ NNb to build the 32 NNG.

Finally, Fig. 3(d) shows that our methodology save lots of work in the highdimensional document space. For instance, for $n=1,215, K$ NNrp costs $63 \%$,


Fig. 3. Distance evaluations per node during $k$ NNG construction. Fig. 3(b)/3(f) follows the legend of Fig. 3(a)/3(e).

| Space | KNNrp <br> Dist. evals. | $K$ NNrp <br> CPU time | KNNpiv <br> Dist. evals. | $K$ NNpiv <br> CPU time |
| :---: | :---: | :---: | :---: | :---: |
| $[0,1]^{4}$ | $O\left(n^{1.32}\right)$ | $O\left(n^{2.24}\right)$ | $O\left(n^{1.09}\right)$ | $O\left(n^{2.01}\right)$ |
| $[0,1]^{8}$ | $O\left(n^{1.38}\right)$ | $O\left(n^{2.11}\right)$ | $O\left(n^{1.06}\right)$ | $O\left(n^{1.69}\right)$ |
| $[0,1]^{12}$ | $O\left(n^{1.59}\right)$ | $O\left(n^{2.03}\right)$ | $O\left(n^{1.27}\right)$ | $O\left(n^{1.79}\right)$ |
| $[0,1]^{16}$ | $O\left(n^{1.77}\right)$ | $O\left(n^{2.14}\right)$ | $O\left(n^{1.64}\right)$ | $O\left(n^{1.97}\right)$ |
| $[0,1]^{20}$ | $O\left(n^{1.88}\right)$ | $O\left(n^{2.18}\right)$ | $O\left(n^{1.87}\right)$ | $O\left(n^{2.10}\right)$ |
| $[0,1]^{24}$ | $O\left(n^{1.96}\right)$ | $O\left(n^{2.16}\right)$ | $O\left(n^{1.96}\right)$ | $O\left(n^{2.16}\right)$ |
| Gaussian $\sigma=0.1$ | $O\left(n^{1.33}\right)$ | $O\left(n^{2.07}\right)$ | $O\left(n^{0.91}\right)$ | $O\left(n^{1.63}\right)$ |
| Gaussian $\sigma=0.2$ | $O\left(n^{1.71}\right)$ | $O\left(n^{2.09}\right)$ | $O\left(n^{1.60}\right)$ | $O\left(n^{1.94}\right)$ |
| Gaussian $\sigma=0.3$ | $O\left(n^{1.85}\right)$ | $O\left(n^{2.17}\right)$ | $O\left(n^{1.81}\right)$ | $O\left(n^{2.06}\right)$ |
| String | $O\left(n^{1.54}\right)$ | $O\left(n^{2.09}\right)$ | $O\left(n^{1.26}\right)$ | $O\left(n^{1.85}\right)$ |
| Document | $O\left(n^{1.95}\right)$ | $O\left(n^{1.96}\right)$ | $O\left(n^{1.86}\right)$ | $O\left(n^{1.87}\right)$ |

Table 1. KNNrp and KNNpiv empirical models for distance evaluations and CPU time for all the metric spaces.
and $K$ NNpiv costs $67 \%$, of $K \mathrm{NNb}$ to build the 8 NNG . These two last results show that our methodology is also practical in real-world situations.

All of these conclusions are confirmed in Table 1. We remark that in some practical conditions (vectors in $[0,1]^{D}$ with $D \leq 8$ and $k \leq 32$ and Gaussian
 $K \mathrm{NNb}$ in CPU time. This is important since the Euclidean distance is very cheap to compute. Note that in the metric space context, superquadratic CPU time in side computations is not as important as a subquadratic number of computed distances. In fact, in the document space, $K$ NNrp and $K$ NNpiv perform better in CPU time that $K \mathrm{NNb}$, showing that in practice the leading complexity (computing distances) is several orders of magnitude larger than other side computations such as traversing pointers or scanning the pivot table.

## 4 Conclusions

We have presented a general methodology to construct the $k$-nearest neighbor graph ( $k \mathrm{NNG}$ ) in general metric spaces. Based on our methodology we give two algorithms. The first is based on a recursive partitioning of the space ( $K \mathrm{NNrp}$ ), and the second on the classic pivot technique (KNNpiv). Our methodology considers two stages: the first indexes the space, and the second completes the $k \mathrm{NNG}$ using the index and some metric and graph optimizations.

Experimental results confirm the practical efficiency of our approach in vectorial metric spaces of wide dimensional spectrum $(D \leq 20)$, and real-world metric spaces. For instance, in the string space, our algorithms achieve empirical CPU time of the form $O\left(n^{1.85}\right)$, and $O\left(n^{1.26}\right)$ in distance computations; and in the high-dimensional document space, they reach empirical $O\left(n^{1.87}\right)$ both in distance computations and CPU time. In low dimensional metric spaces, our algorithms behave even better. KNNpiv is in general better than $K$ NNrp for small and moderate $k$ values, yet $K$ NNrp is less sensitive to larger $k$ values or higher dimensional spaces.

Future work involves developing another $k$ NNG constructing algorithm based on the list of clusters [7] so that we can also obtain good performance in higher di-
mensional metric spaces. We are also researching how to enhance the data structure to allow dynamic insertions/deletions in reasonable time, so as to maintain an up-to-date set of $k$-nearest neighbors for each element in the database.

## References

1. S. Arya, D. Mount, N. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimension. In Proc. SODA'94, pages 573-583, 1994.
2. F. Aurenhammer. Voronoi diagrams - a survey of a fundamental geometric data structure. ACM Computing Surveys, 23(3), 1991.
3. R. Baeza-Yates, C. Hurtado, and M. Mendoza. Query clustering for boosting web page ranking. In Proc. AWIC'04, LNCS 3034, pages 164-175, 2004.
4. M. Brito, E. Chávez, A. Quiroz, and J. Yukich. Connectivity of the mutual $k$ nearest neighbor graph in clustering and outlier detection. Statistics \& Probability Letters, 35:33-42, 1996.
5. P. Callahan. Optimal parallel all-nearest-neighbors using the well-separated pair decomposition. In Proc. FOCS'93, pages 332-340, 1993.
6. P. Callahan and R. Kosaraju. A decomposition of multidimensional point sets with applications to $k$ nearest neighbors and $n$ body potential fields. JACM, 42(1):6790, 1995.
7. E. Chávez and G. Navarro. A compact space decomposition for effective metric indexing. Pattern Recognition Letters, 26(9):1363-1376, 2005.
8. E. Chávez, G. Navarro, R. Baeza-Yates, and J.L. Marroquin. Searching in metric spaces. ACM Computing Surveys, 33(3):273-321, September 2001.
9. K. Clarkson. Fast algorithms for the all-nearest-neighbors problem. In Proc. FOCS'83, pages 226-232, 1983.
10. K. Clarkson. Nearest neighbor queries in metric spaces. Discrete Computational Geometry, 22(1):63-93, 1999.
11. M. Dickerson and D. Eppstein. Algorithms for proximity problems in higher dimensions. Computational Geometry Theory and Applications, 5:277-291, 1996.
12. R. Duda and P. Hart. Pattern Classification and Scene Analysis. Wiley, 1973.
13. H. Edelsbrunner. Algorithms in Combinatorial Geometry. Springer-Verlag, 1987.
14. D. Eppstein and J. Erickson. Iterated nearest neighbors and finding minimal polytopes. Discrete E Computational Geometry, 11:321-350, 1994.
15. K. Figueroa. An efficient algorithm to all $k$ nearest neighbor problem in metric spaces. Master's thesis, Universidad Michoacana, Mexico, 2000. In Spanish.
16. G. Hjaltason and H. Samet. Incremental similarity search in multimedia databases. Technical Report TR 4199, Dept. of Comp. Sci. Univ. of Maryland, Nov 2000.
17. I. Kalantari and G. McDonald. A data structure and an algorithm for the nearest point problem. IEEE Trans. Software Eng., 9(5):631-634, 1983.
18. D. R. Karger and M. Ruhl. Finding nearest neighbors in growth-restricted metrics. In Proc. STOC'02, pages 741-750, 2002.
19. R. Krauthgamer and J. Lee. Navigating nets: simple algorithms for proximity search. In Proc. SODA'04, pages 798-807, 2004.
20. R. Paredes and E. Chávez. Using the $k$-nearest neighbor graph for proximity searching in metric spaces. In Proc. SPIRE'05, LNCS 3772, pages 127-138, 2005.
21. P. Vaidya. An $O(n \log n)$ algorithm for the all-nearest-neighbor problem. Discrete \& Computational Geometry, 4:101-115, 1989.
