

RT-kNNS Unbound: Using RT Cores to Accelerate Unrestricted Neighbor Search

Vani Nagarajan
School of Electrical and Computer
Engineering
Purdue University
West Lafayette, IN, USA
nagara16@purdue.edu

Durga Mandarapu
Department of Computer Science
Purdue University
West Lafayette, IN, USA
dmandara@purdue.edu

Milind Kulkarni
School of Electrical and Computer
Engineering
Purdue University
West Lafayette, IN, USA
milind@purdue.edu

Abstract

The problem of identifying the k -Nearest Neighbors (kNNS) of a point has proven to be very useful both as a standalone application and as a subroutine in larger applications. Given its far-reaching applicability in areas such as machine learning and point clouds, extensive research has gone into leveraging GPU acceleration to solve this problem. Recent work has shown that using Ray Tracing cores in recent GPUs to accelerate kNNS is much more efficient compared to traditional acceleration using shader cores. However, the existing translation of kNNS to a ray tracing problem imposes a constraint on the search space for neighbors. Due to this, we can only use RT cores to accelerate *fixed-radius* kNNS, which requires the user to set a search radius *a priori* and hence can miss neighbors. In this work, we propose TrueKNN, the first *unbounded* RT-accelerated neighbor search. TrueKNN adopts an iterative approach where we incrementally grow the search space until all points have found their k neighbors. We show that our approach is orders of magnitude faster than existing approaches and can even be used to accelerate fixed-radius neighbor searches.

CCS Concepts: • **Computing methodologies** → **Ray tracing**; **Graphics processors**; • **Information systems** → **Nearest-neighbor search**; • **Theory of computation** → **Nearest neighbor algorithms**.

Keywords: Ray Tracing, k Nearest Neighbors, Neighbor Search

ACM Reference Format:

Vani Nagarajan, Durga Mandarapu, and Milind Kulkarni. 2023. RT-kNNS Unbound: Using RT Cores to Accelerate Unrestricted Neighbor Search. In *International Conference on Supercomputing (ICS '23)*,



This work is licensed under a Creative Commons Attribution International 4.0 License.

ICS '23, June 21–23, 2023, Orlando, FL, USA

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0056-9/23/06.

<https://doi.org/10.1145/3577193.3593738>

June 21–23, 2023, Orlando, FL, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3577193.3593738>

1 Introduction

The ability to leverage GPUs to accelerate general-purpose workloads has resulted in unprecedented performance gains in various applications. Re-purposing shader cores primarily achieved this in GPUs, which accelerate arithmetic computations in rendering algorithms, to accelerate other non-rendering applications. However, re-purposing the cores was not an easy task: non-rendering applications needed to be rewritten as rendering algorithms to be able to use the shader cores. As only skilled graphics programmers were able to do this translation, researchers began to work on creating programming models that would allow any programmer to leverage GPU acceleration. This led to the creation of popular programming models such as CUDA[20] and OpenCL[11]. They allowed programmers to write general-purpose code while abstracting the translations required to run it on the GPU. These programming models are responsible for cutting down execution times of machine learning applications from weeks to days and hours to seconds.

While the shader cores were ideal for performing arithmetic computation on regular structures, they performed poorly on irregular applications such as tree traversal. In such cases, the user was often better off executing the application on the CPU. The introduction of Ray Tracing (RT) cores in recent GPUs created an opportunity to accelerate a different class of applications: irregular applications. *Ray casting* is a popular rendering algorithm where a ray is launched through a pixel in the image plane and the interaction of the ray with objects in the scene is used to determine the color of the pixel. *Ray tracing* expands on ray casting by also accounting for reflected and refracted rays (secondary rays) that are generated due to the ray-object interaction and trace the interaction of these secondary rays with the objects in the scene. The RT cores on GPUs were designed to accelerate the ray tracing the process by offloading certain sections of the RT pipeline to hardware, namely the Bounding Volume Hierarchy traversal and ray-bounding volume intersection tests [21] (Details in Section 2.2.2).

Though the RT cores were designed to accelerate ray tracing applications, researchers are starting to look into exploiting these cores to perform general-purpose computations. Wald *et al.* introduced the idea of using RT cores to identify the tetrahedral mesh to which a point belongs [28]. The authors re-formulate the problem as a ray tracing problem by treating the query point as the ray origin and the tetrahedral meshes as objects in the scene. Now, the point-in-tetrahedron problem reduces to launching a ray from the query point and identifying the closest intersecting mesh. Later works by Zellman *et al.*, Evangelou *et al.*, and Zhu *et al.* looked at re-formulating the nearest neighbor search problem as a ray casting query [8, 32, 33]. The Nearest Neighbor Search (NNS) problem is the task of identifying the closest points to a query point. Its variant, k-Nearest Neighbor Search (kNNS), restricts the task of finding the k closest neighbors.

A major downside to RT-accelerated neighbor search is that the search space for nearest neighbors is constrained to a fixed-radius neighborhood due to the problem translation approach adopted by prior work (Details in Section 2.3). It is impossible to know the required radius to identify neighbors *a priori*, leading the approach to possibly find *less* than k neighbors. Prior work has suggested that this problem can be avoided by choosing a very large radius to ensure all neighbors are found [33], but we show that this approach is highly inefficient—trying to find all of the neighbors of query points using the existing approaches obviates the benefits of hardware acceleration entirely (Details in Section 5).

This paper presents an efficient solution to the problem of RT core acceleration of *unbounded* kNNS—ensuring that all query points will successfully find all k neighbors. We adopt an iterative solution: we start with a smaller search radius and keep track of points that have found their k nearest neighbors in each iteration. In the subsequent iterations, we incrementally increase the radius and *only* query the points that have not found their neighbors. While this approach seems like it should be significantly slower than choosing a single radius, we show that it is not. Since the number of points being queried decreases as the search radius increases, we find that our approach is significantly *faster* than choosing an arbitrarily large radius, and can even be faster than prior fixed-radius approaches even for smaller radii.

To summarize, the contributions of our paper are as follows:

- This paper introduces TrueKNN, the first RT-accelerated neighbor search algorithm that is not constrained to a fixed-radius.
- We show that TrueKNN outperforms fixed-radius, non-iterative approaches by large margins. This iterative solution gradually grows the neighbor search space while pruning query points that have already found their neighbors, leading to significantly fewer ray-object intersection tests

- We further show that, unlike prior approaches that require *a priori* selection of a query radius, we can adapt TrueKNN to find the appropriate radius dynamically, outperforming prior approaches even when those approaches select their query radius *a posteriori*.

2 Background

In this section, we introduce the k-Nearest Neighbor Search (kNNS) problem and explain how prior works have translated kNNS to a ray casting problem.

2.1 k-Nearest Neighbor Search

The Nearest Neighbor Search (NNS) problem was first introduced by Fix *et al.* [9] and expanded by Cover *et al.* [7]. It is defined as follows:

Definition 2.1. For a dataset D and query point $q \in D$, find the set of k nearest points to q .

The nearest neighbors are typically identified using a distance metric, with Euclidean distance being the most popular choice. kNNS is primarily used as a subroutine in k-Nearest Neighbor classification and regression algorithms. The basic idea behind these algorithms is that a property of a query point can be determined by observing its nearest neighbors. For example, a query point can be classified into the same class as a majority of its neighbors in classification problems. Similarly, the properties of a query point can be averaged using its neighbors in regression problems.

kNN classifier and regression models are widely used in point cloud applications to compute surface normals [26], recommendation systems to assign recommendations based on similar users [1], healthcare to classify patients [31] and pattern recognition, to name a few.

2.2 Ray Tracing Hardware

As a part of NVIDIA’s Turing architecture, each Streaming Multiprocessor (SM) has an RT core to accelerate BVH traversal and perform ray-triangle intersection tests, allowing the SM to perform other computations in the meantime. Both the SM and RT core share the same memory, allowing us to use both units in parallel. We direct the reader to [21] for more details.

2.2.1 Ray Casting. The RT cores are able to accelerate the ray casting process by reducing the *number* of intersection tests performed, and this is facilitated by the Bounding Volume Hierarchy (BVH) structure used for object representation. The ray-casting process involves launching rays from a source through each pixel in the image plane, recording their interactions with objects in the scene, and using that information to determine the color of the pixels. It would seem that *each ray* would have to be tested for intersection against *each object* to check whether an intersection could

affect the color of the pixel. However, this would be very inefficient since rays may not intersect a large subset of objects and intersection tests are computationally very expensive.

2.2.2 Bounding Volume Hierarchy. Bounding Volume Hierarchies are an acceleration structure used to reduce the number of required ray-object intersection tests. The general idea is as follows: if we group objects that are spatially close to each other, we can test for intersection against *groups* of objects rather than *individual* objects, reducing the number of intersection tests performed. The idea is to enclose objects in bounding volumes and recursively enclose these volumes in larger bounding volumes until we create a volume that is large enough to enclose the entire scene. These bounding volumes are represented hierarchically using a tree structure called the Bounding Volume Hierarchy (BVH). Each node in the tree represents a bounding volume that encloses all its descendent nodes. The most commonly used bounding volume is an Axis-Aligned Bounding Box (AABB).

Figure 1 shows how objects in a scene are used to construct the BVH. On the left, each object A, B, C and D are enclosed in their corresponding bounding volumes BV_A, BV_B, BV_C and BV_D. As A and B are spatially close together, they are combined into a larger bounding volume BV_AB that encloses the bounding volumes of A and B. BV_CD is constructed similarly and BV_ABCD encloses both BV_AB and BV_CD to capture the entire scene. The hierarchical relationship between the bounding volumes is captured on the right. For example, BV_AB encloses BV_A and BV_B, and this relationship is captured in the BVH with BV_AB as the parent node and BV_A and BV_B as the children node.

It is possible to reduce the number of ray-object intersection tests by performing ray-AABB intersection tests and only performing the ray-object intersection test if the previous test succeeds. For example, in Fig 1, if the ray does not intersect BV_AB, it is *guaranteed* to not intersect BV_A and BV_B. Since A and B are contained in these bounding volumes, we do not have to perform the corresponding ray-object intersection tests. This approach allows us to prune large parts of the search space since if the ray does not intersect the AABB, it is guaranteed to not intersect any of the objects or other bounding volumes contained in the AABB.

2.2.3 Optix API. The Optix API handles offloading Bounding Volume Hierarchy build and traversal to the RT core, while allowing the user to write custom CUDA kernels that use the GPU's shader cores.

Fig 2 shows the working of the Optix API. The *RayGen* program is responsible for creating rays with the specified origin (\vec{o}), direction (\vec{d}), and length (t)

$$\vec{r} = \vec{o} + t\vec{d}, t \in [t_{min}, t_{max}]$$

The ray traverses the BVH and checks for intersection with the Axis-aligned Bounding Boxes (AABBs) in hardware (Highlighted in yellow in Fig 2). The ray-object intersection test

can happen in software or hardware, depending on the object. If the object is a triangle, the test happens in hardware. If not, the user provides a custom intersection test written as a software CUDA *Intersection* program. Optix allows us to perform the ray-AABB intersection test in hardware and the ray-object intersection test in software/hardware.

The user can evoke the *AnyHit* program to record intersections, decide whether to continue or terminate the BVH traversal for a particular ray and launch subsequent rays. After the BVH traversal completes entirely, the user can specify a *ClosestHit* to record the closest intersected object to the ray. The user can also specify what to do in case of no intersections using the *Miss* program.

2.3 RT-accelerated k-Nearest Neighbor Search (RT-kNNS)

We use the reduction proposed by Zellman *et. al.* to translate fixed-radius nearest neighbor queries to ray tracing queries [32]. The translation relies on a key observation: to find neighbors within a radius r of point p , we can expand a sphere of radius r around *all* points and check how many spheres contain the point p . The centers of the k -closest spheres are the neighboring points of p .

Algorithm 1: RT-kNNS

Input : Dataset D , radius r
Output : Neighbors within radius r

```

1 spheres  $\leftarrow$  createSpheres( $D, r$ )
2 boundedSpheres  $\leftarrow$  createAABB(spheres)
3 constructBVH(boundedSpheres)
4 for Point  $p \in D$  do
5   ray  $\leftarrow$  RayGen( $\vec{p}, \vec{d}, 0, \text{FLOAT\_MIN}$ )
6   while traverseBVH(ray) do
7     if Intersect(ray, AABB) then
8       if Intersect(ray, sphere) then
9         p.neighbor  $\leftarrow$ 
10        p.neighbor  $\cup$  ray.center
11       end
12     end
13 end

```

RT-accelerated kNNS is outlined in Algorithm 1 and is based on the Optix API pipeline from Section 2.2.3. In Line 1, we create spheres for all points with $p \in D$ (D is the input dataset). The centers of the sphere are the points p and the radius is supplied by the user. In Line 2, we specify a *BoundingBox* program to create Axis-Aligned Bounding Boxes (AABBs) to enclose the spheres. We proceed to construct a Bounding Volume Hierarchy (BVH) by recursively combining bounding volumes (See Section 2.2.2).

With the BVH constructed, we can begin to generate rays that traverse the BVH and test for ray-AABB and ray-object

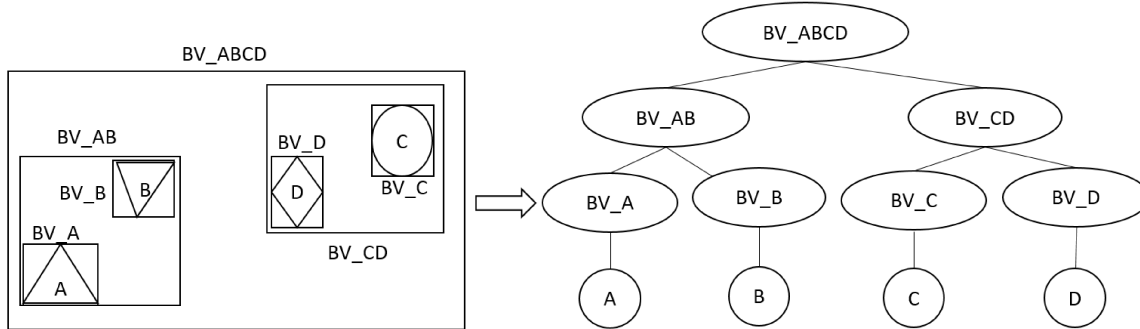


Figure 1. Bounding Volume Hierarchy Construction. The scene with objects enclosed in bounding volumes is shown on the left. On the right, the BVH corresponding to the objects and bounding volumes is shown.

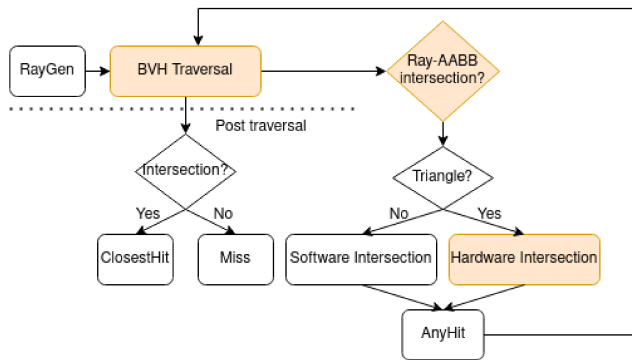


Figure 2. Ray Tracing Pipeline in Optix API

intersections. We create rays for all points $p \in D$ in Line 5. The *RayGen* program accepts origin, direction, and ray interval (t_{min}, t_{max}) as inputs. We set point \vec{p} as center, $(0,0,1)$ as the direction (\vec{d}) and t_{max} as `FLOAT_MIN`. Since a ray of infinitesimal length is sufficient to intersect neighbors, t_{max} can be set to a very small value. In Lines 6-12, each ray traverses the BVH and performs intersection tests against AABBs. If the ray intersects the AABB, ray-object intersection tests are performed for *each* object contained in the intersected AABB. If the ray intersects the sphere, we add the sphere’s center to the list of point p ’s neighbors.

3 Design

In this section, we describe how we use the fixed-radius k -Nearest Neighbor Search (kNNS) reduction from Section 2.3 to implement our *unbounded* RT-accelerated kNNS algorithm, TrueKNN. TrueKNN incrementally increases the search space till we find each point’s k nearest neighbors.

3.1 TrueKNN Overview

We define TrueKNN as the task of finding the k nearest neighbors of all points in a dataset. TrueKNN differs from fixed-radius kNNS in that we do not constrain our neighbor search space to a particular neighborhood. Section 2.3 shows how to reduce fixed-radius nearest neighbor queries to ray

tracing queries by expanding spheres with a user-specified radius around all points. In TrueKNN, the user does not specify the radius.

Since the user does not have to specify the radius, we need to figure out a radius such that each point in the input dataset can find all its neighbors. A solution proposed by Zhu is to use an arbitrarily large radius for the sphere expansion phase [33]. This would guarantee that each point finds its k nearest neighbors. However, choosing this appropriately-large radius is difficult *a priori*: without performing kNNS, to begin with, we do not know how large a radius will guarantee that all points will find their k nearest neighbors. The only radius that is guaranteed to work *a priori* is one that encompasses *all* the points. But in such a scenario, the volumes around each point are guaranteed to overlap as the bounding boxes become larger when the search radius is large. As the bounding volume hierarchy can no longer separate AABBs, the algorithm is definitive $O(n^2)$.

So *a priori* radius setting cannot work. Instead, our approach is to incrementally identify the radius needed to identify all of the neighbors. We propose a multi-round fixed-radius kNNS approach that implements TrueKNN. We begin by performing fixed-radius kNNS with a small radius and determine the starting radius for the first round of our neighbor search by selecting the minimum distance between points in a subset of the input dataset (Section 3.2). Because this radius is small, the search completes quickly, but many points may not find their k neighbors. In each subsequent round, we increment the radius and re-run fixed-radius kNNS on any points that have not yet found their k neighbors until all points find their k nearest neighbors (Section 3.3).

3.2 Determining start radius

In Algorithm 1, the user specifies the radius for sphere expansion. However, in the case of TrueKNN, there is no user input and we need to determine an appropriate start radius to begin the neighbor search. The selection of a start radius is crucial to performance. If the radius is too small, we will have many rounds where no points find any neighbors and we will pay for the cost of a context switch between device

and host, BVH refit, ray launch, and intersection tests. On the other hand, if the radius is too large¹, the number of intersection tests performed will increase by several orders of magnitude, leading to poor performance.

Algorithm 2: RandomSample

Input :Dataset D

Output:Radius $start_radius$

```

1 random_sample  $\leftarrow$  sample( $D$ , 100)
2 neighbors, distances  $\leftarrow$   $k$ NearestNeighbors( $D$ ,  $k=4$ )
3 start_radius  $\leftarrow$  min(distances)

```

It was evident that we needed to have some idea about the input dataset to select a good start radius. Without this information, it would be impossible to know whether our chosen radius is too small or large. To incorporate dataset information, we propose a random sampling approach (Algorithm 2) to find the minimum distance between neighbors of a subset of the input dataset. In Line 1, we choose 100 random points from the input dataset. We then use Python scikit-learn’s built-in, ball-tree based k nearest neighbors algorithm to find the 4 nearest neighbors of these randomly sampled points. We empirically chose $k = 4$ as it worked well in our experimental evaluation and had negligible execution time (5 to 8 ms). We note that increasing k and/or sample size *could* result in a better start radius. As shown in Line 3, we then find the minimum distance between a point and its neighbors and set that as our starting radius.

We experimented with different starting radii and found that the cost of choosing a larger radius was much higher than starting off with a smaller radius. When starting off with a very small radius, we observe that some points find few to no neighbors in the initial rounds. However, we found that the time taken by these initial rounds was insignificant compared to the total execution time. On the other hand, starting off with a larger radius leads to fewer rounds but more intersection tests, justifying our decision to use the minimum distance as the start radius. We show that our random sampling approach produces useful starting radii in Section 5.4.2.

3.3 Multi-round kNNS

Now that we have a starting radius, Algorithm 3 describes our multi-round approach for TrueKNN.

We use the random sampling approach outlined in Section 3.2 to determine the radius for sphere expansion in the first round of fixed-radius RT-kNNS. We use the minimum neighbor distance returned by RandomSample as our start radius ($radius$) in Line 1. In Line 3, we call RT-kNNS with our dataset (D) and radius ($radius$) to find all points within

¹We define the notion of small and large radius in comparison to the maximum distance between a point and any of its k nearest neighbors in the input dataset.

Algorithm 3: TrueKNN

```

1 radius  $\leftarrow$  RandomSample( $D$ )
2 while  $D \neq \{\emptyset\}$  do
3   neigh, dist  $\leftarrow$  RT-kNNS( $D$ , radius)
4   for  $p \in D$  do
5     if  $|p.neigh| == k$  then
6        $D - = \{p\}$ 
7     end
8   end
9   if  $D \neq \{\emptyset\}$  then
10    radius  $\leftarrow$  radius * 2
11    REFIT_BVH( $D$ , radius)
12  end
13 end

```

a fixed neighborhood of p . We then check if the previously chosen radius was sufficient for all points to find k neighbors in Line 5. In Line 6, we remove all the points that have found their k neighbors from our dataset so that the next iteration will only consider points that have not found their neighbors yet. For points that have not found all k neighbors, we expand the neighbor search space by incrementing the radius of spheres and re-fitting the bounding boxes around the spheres for the BVH, as shown in Lines 10 and 11. We continue this process till $D = \{\emptyset\}$ and all points have found their k nearest neighbors.

3.4 TrueKNN Discussion

It may seem curious that this algorithm is more efficient than prior fixed-radius approaches—after all, each iteration of TrueKNN performs an *entire fixed-radius* nearest neighbor search. The key is in the interplay between radius size, BVH-based search speed, and sorting time. When the radius is small, the spheres around the points are small and well-separated. As a result, the BVH is extremely effective at accelerating ray tracing, and query points very quickly identify nearby points. By starting with a small radius, this first iteration is *faster* than “normal” fixed-radius kNNS, with the trade-off that many points are unable to find their needed k neighbors.

In subsequent iterations, the radius is increased, but because *some* points in earlier iterations have already found their needed neighbors, those points do not need to be re-queried. So although the ray tracing process is slower, there are fewer query points, and the overall search is faster. In the final iterations, when the radius is quite large and BVH acceleration is essentially useless, there are only a few query points remaining (the outliers).

As mentioned earlier, setting a large starting radius for fixed-radius kNNS leads to inefficient filtering, and a $O(n^2)$ runtime. By incrementally increasing the radius, TrueKNN avoids this problem—while outlier points have to do an $O(n)$

search to find their candidate neighbors, most points are resolved with smaller radii and require closer to a $O(\log n)$ search.

There is another beneficial effect of TrueKNN’s iterative approach. With a large starting radius (e.g., enough to find at least k neighbors for 99% of the points), *many* points will find *too many* candidate neighbors in their initial search, and hence will waste time sorting those candidates to find the k closest. By resolving those points with smaller radii, TrueKNN also reduces sorting time. Indeed, as we show in Section 5.5.1, TrueKNN’s benefits are not only because outliers force large query radii for fixed-radius search. Even in settings where we do not attempt to resolve outlier points, TrueKNN outperforms fixed-radius kNNS.

4 Implementation Details

In this section, we discuss the implementation of TrueKNN using the Optix API. We used the Optix Wrapper Library (OWL) [27], which is built on top of Optix 7, to implement TrueKNN and our baseline. OWL allows the user to define custom shader programs to test for the ray-object intersection. Though one would typically use the *AnyHit* program (Section 2.2.3) to collect multiple hits, we implemented the TrueKNN logic in the *Intersection* program to avoid incurring overhead costs associated with calling the *AnyHit* program. In fact, we disable both the *AnyHit* and *ClosestHit* program invocations to avoid performance penalties.

Since TrueKNN (Algorithm 3) increases the radius of spheres in every iteration to expand the search space, the BVH corresponding to the objects also needs to change every iteration. One way to handle this is to *re-build* the BVH in every iteration. However, Optix² provides the option of BVH *re-fit*. This allows us to *re-fit* the bounding volumes to accommodate the expanded objects without having to explicitly re-build the BVH. We found that re-fitting was between 10-25% faster than re-building.

5 Evaluation

In this section, we evaluate TrueKNN’s performance by analyzing the effect of varying parameters such as dataset size and k . We also look at how outliers in the dataset affect both TrueKNN and the baseline.

5.1 Datasets

We used a mixture of 2D and 3D real-world datasets (3DRoad, Porto, KITTI, and 3DIono) and a 3D synthetic dataset (UniformDist) to evaluate TrueKNN.

3DRoad The 3DRoad dataset captures the road network of North Jutland, Denmark [14]. The dataset consists of 430K points. We use this as a 2D dataset, using only the latitude and longitude parameters.

²We use OWL and Optix interchangeably as every feature in OWL is also present in Optix

Porto The Taxi Service Trajectory - Prediction Challenge 2015 dataset captures vehicle movement trajectory data of 442 taxis in the city of Porto, Portugal [15]. The dataset has just over 81M points and we use Porto as a 2D dataset, using only the latitude and longitude parameters.

KITTI The KITTI vision benchmark dataset captures data from the movement of an autonomous vehicle around the city of Karlsruhe, Germany [10]. The dataset has just over 1M points and we use KITTI as a 3D dataset.

3DIono The 3D Ionosphere dataset captures the behavior of electrons in the ionosphere [22]. The dataset has just over 1M points and we use it as a 3D dataset.

UniformDist We create a synthetic 3D dataset of 1M points that is uniformly distributed on $[0,1]$ to study the impact of outliers on our algorithm.

5.2 Experimental Setup

We ran our experiments on an NVIDIA GeForce RTX 2060 GPU with 6 GB device memory, CUDA version 10.1, and Optix 7.1. As Optix is a graphics rendering API, it accepts only 3D input data. As a workaround, we set the z-dimension to 0 for 2D datasets.

5.2.1 Baseline. We use RT-kNNS (Algorithm 1) as our baseline by setting the radius for sphere expansion as the maximum distance (*maxDist*) between a point and any of its k nearest neighbors. This way, the baseline is guaranteed to find all k nearest neighbors of each point in the dataset. We chose this baseline as prior work has shown that RT-accelerated neighbor search is consistently faster than other GPU-based implementations [8, 28, 33]. We also note that our baseline represents the *best case* scenario since our neighbor search is constrained exactly to a *maxDist*-neighborhood. In practice, the user would probably select some arbitrary d -neighborhood, where $d \gg \text{maxDist}$.

5.3 Performance Evaluation

We compare the performance of TrueKNN against RT-kNNS as described in Section 5.2.1. We study the performance impact of dataset size by varying the dataset size between 100K and 1M. For each dataset size, we chose $k = 5$ and $k = \sqrt{\text{DatasetSize}}$ to study the impact of varying k . We chose $k = \sqrt{\text{DatasetSize}}$ as it is the commonly used k value for KNN classifier and regression models [18]. We only vary dataset size up to 1M since we run out of memory to store neighbors in our GPU when $k = \sqrt{\text{DatasetSize}}$. We also perform an experiment where we modify the kNNS problem to one of finding 99th percentile neighbors and evaluate our performance against the baseline to understand the effect of outliers in the dataset.

For different dataset sizes (d), we always used the first d points in our experiments and averaged results over 5 runs. The reported execution time for all our experiments includes

BVH build and refit times, context switching overheads and neighbor search time. We assume that the input is already on the device and do not include data transfer time.

5.3.1 Impact of Dataset size. We study the effect of varying the number of query points for a fixed number of neighbors ($k = \sqrt{\text{DatasetSize}}$). From Fig 3, we see that TrueKNN outperforms the baseline on all datasets³. Table 1 also shows the raw execution time for all datasets. A major advantage of TrueKNN is that points that have found their k nearest neighbors in previous rounds are no longer queried in subsequent rounds with larger search radii. As the search radius increases, the size of the Axis-Aligned Bounding Boxes (AABBs) enclosing the spheres also increases. This leads to an increase in ray-AABB intersection tests, which, in turn, leads to an increase in ray-object intersection tests. Since TrueKNN can correctly identify the k nearest neighbors, any additional intersection tests (compared to TrueKNN) performed by the baseline are unnecessary. As the number of query points increases, we see that the number of these unnecessary intersection tests also increases.




Figure 3. TrueKNN's speedup compared to baseline while varying dataset size

Since the ray-AABB intersection tests are performed in hardware, we do not have any information on the number of tests performed. However, since the ray-object intersection tests are performed in software, we can compare the number of intersection tests performed by TrueKNN and the baseline. Table 2 shows the number of ray-sphere intersection tests performed (in billions) on the Porto dataset. We notice that the baseline performs 9x the tests performed by TrueKNN for 100K points and this increases to 32.1x for 1M points. As TrueKNN incrementally increases the search radius and reduces the number of query points launched every round, we perform far fewer intersection tests compared to the baseline, leading to significant performance improvements.

³3DRoad does not have speedup bars for 800K and 1M as it only has 400K points in total

Though this gives us an idea of where our speedup comes from, speedup does not always increase proportionally to the decrease in the number of intersection tests. For example, the baseline performs 16.8x and 24.7x more intersection tests compared to TrueKNN for 400K and 800K points, respectively. However, our speedup falls from 99.9x to 65x. There are possibly other RT Core-related aspects (ray-AABB intersection tests, BVH traversal) that contribute to this⁴.

Comparison with RTNN. RTNN was proposed by Zhu to optimize nearest neighbor search using RT cores [33]. RTNN proposes two main optimizations (1) query reordering to improve ray coherence (2) query partitioning to minimize intersection tests. To demonstrate the effectiveness of our approach, we compare an unoptimized TrueKNN (does not use query sorting or partitioning optimizations) against a fully optimized RTNN on the Porto dataset. We found that TrueKNN was between 1.5x to 8x faster than RTNN, showing that our ability to perform fewer computations results in huge performance improvements.

Comparison with non-RT Baseline. We evaluate TrueKNN's performance against cuML's kNNS implementation, which is a purely CUDA-based implementation [24]. As cuML's kNNS ran out of memory when we set $k = \sqrt{\text{DatasetSize}}$, where $\text{DatasetSize} > 200K$, we chose $k = 5$ for our experiment.

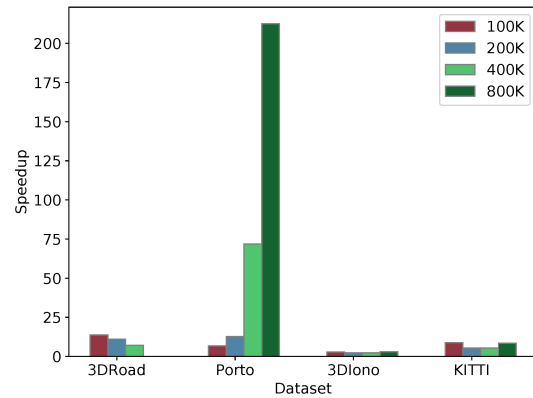


Figure 4. TrueKNN's speedup compared to cuML's kNNS implementation on varying dataset size

From Fig 4, we see that TrueKNN outperforms cuML's kNNS across all the different datasets and dataset sizes. We attribute this to TrueKNN's ability to leverage hardware acceleration while reducing the computation footprint. In general, we found that TrueKNN's speedup increased on increasing the dataset size.

⁴At the time of writing this paper, there were no tools available to profile the RT Cores


Dataset size	3DRoad		Porto		3DIono		KITTI	
	TrueKNN	Baseline	TrueKNN	Baseline	TrueKNN	Baseline	TrueKNN	Baseline
100K	6.28	44.71	25.27	722.93	9.48	308.16	13.8	520.01
200K	23.97	146.83	89.72	3965.7	38.16	2086.13	51.6	2769.96
400K	120.41	753.81	456.78	45639.7	149.6	11433.5	175.38	12503.4
800K	-	-	1128.67	73423.5	776.96	57184.7	641.96	776562.4
1M	-	-	1052.95	127720.1	973.3	1862527.9	892.21	121256.2

Table 1. Execution time (in seconds) for TrueKNN and baseline for all datasets



Dataset size	TrueKNN	Baseline
100K	1.09	9.9
200K	2.85	39.9
400K	7.16	189.1
800K	25.86	639.31
1M	31.12	999.19

Table 2. Number of ray-object intersection tests performed (in billions) for the Porto dataset

5.3.2 Impact of k . We set the k parameter as 5 and $\sqrt{\text{DatasetSize}}$ for 400K points from the 3DRoad, Porto, 3DIono, and KITTI datasets. In both cases, we consistently outperform the baseline, as shown in Fig 5. We also notice that the extent of our speedup is larger when $k = 5$ compared to $k = 660$. We mainly attribute this to the overhead of sorting and maintaining the list of k nearest neighbors over multiple rounds in the case of large k values, compared to the one-time cost incurred by the baseline.

**Figure 5.** Impact of k

Additionally, we notice that as k increases, the number of ray-object intersection tests also increases as we need to identify more neighbors. For example, on the Porto dataset, the baseline performs 839x the intersection tests performed by TrueKNN when $k = 5$ but this number decreases to 17x when $k = 660$. However, we are still over 70x faster than the baseline for most datasets when $k = 660$.

**(a)** Time taken by each round**(b)** Number of points left after each round**Figure 6.** Breakdown of 3DRoad execution time

5.4 Performance Analysis

In this section, we justify our design decisions by analyzing the run time of the various rounds in our TrueKNN approach and showing that our choice of start radius yields good results.

5.4.1 Runtime Analysis. To understand why TrueKNN is able to largely outperform the baseline, we evaluated TrueKNN on the first 400K points of the 3DRoad dataset, with a start radius of 0.001 and $k = 5$. Fig 6a shows the

time taken for each round of TrueKNN and Fig 6b shows the number of query points that *have not* found their k nearest neighbors after each round.

From Fig 6a, it is evident that rounds 7 and 8 take significantly more time compared to the other rounds. This is interesting, since these rounds query far fewer points compared to earlier rounds, as most points have already found their k neighbors. Indeed, from Fig 6b, we see that only 3 query points remain in these last few rounds. Though one would expect that fewer query points would lead to faster execution times, we find that it is not true because the search radius is much larger in the latter rounds.

As the search radius increases in latter rounds, even if we have only 3 query points, the BVH traversal will result in a large number of intersection tests as the size of the bounding volumes enclosing the spheres also increases to accommodate the larger sphere. As we double the radius in each round, we see a huge uptick in the number of intersection tests. This shows us why TrueKNN is able to outperform the baseline: we expand *fewer* points to larger radii in latter rounds, allowing us to drastically reduce the number of intersection tests performed.




Figure 7. Impact of start radius selection

5.4.2 Start Radius. In Section 3.2, we proposed a method to choose a useful start radius for sphere expansion in the first stage of the TrueKNN algorithm. In particular, we decided to always choose the minimum distance between the 4 nearest neighbors of 100 points as our starting radius. In this section, we show that this approach typically produces useful start radii.

We ran Algorithm 2 multiple times to generate different start radii for 400K points of the Porto dataset. Fig 7 shows the execution time (in seconds) of TrueKNN, where $k = \sqrt{DatasetSize}$, for the various start radii. In most cases, we found that there was an insignificant difference in the execution times for the different start radii. We found similar results on the other datasets and dataset sizes.

5.5 Impact of outliers

From our analysis in Section 5.4.1, we see that using larger radii for neighbor search results in higher execution times even for a few query points. Since our baseline uses this large radius to query *all* points, it is unsurprising that it is much slower than TrueKNN.

In this section, we perform an experiment where we only identify neighbors upto the 99th percentile distance to eliminate the influence of these outliers. We also create a synthetic dataset of points uniformly distributed over $[0,1]$ to study the impact of outliers. In both experiments, we find that TrueKNN almost always outperforms the baseline.

5.5.1 99th Percentile. We calculated the 99th percentile search radius by computing the distance between query points and their neighbors and selecting the 99th percentile distance. We perform this experiment to show that our approach does not unfairly benefit from the presence of outliers in the dataset. As the 99th percentile radius is much smaller, the baseline greatly benefits from performing fewer intersection tests.




Figure 8. TrueKNN's speedup compared to baseline for 99th percentile neighbor search on different datasets

We evaluated performance on the Porto, 3DIono, and KITTI datasets, with the dataset size varying between 100K and 800K and $k = \sqrt{DatasetSize}$. We modified TrueKNN to terminate the execution when we reached the 99th percentile radius. We note that this is only a thought experiment and that it is not possible to know the 99th percentile without *actually* computing the neighbors of each query point. From Fig 8, it is evident that TrueKNN still outperformed the baseline in all cases. This showed us that even with a drastic reduction in the search radius for the baseline ($\approx 30x$ on average), our approach still manages to reduce the number of software intersection tests. *In fact, our original TrueKNN, which identifies all neighbors, is also faster than the 99th percentile-modified baseline in all cases!*

5.5.2 Uniformly Distributed Dataset. We created a synthetic dataset of 1M points that are uniformly distributed on $[0,1]$ to see whether our approach would work well in the absence of blatant outliers. We studied the effects of varying dataset size and also repeated our 99th percentile experiment on this dataset for $k = \sqrt{\text{DatasetSize}}$.

In both experiments, TrueKNN outperformed the baseline, as shown in Table 3. This was surprising since this would be the worst-case input to TrueKNN, due to the absence of outliers. Though our speedup margin is not as high as the previous experiments, we still manage to perform fewer intersection tests, resulting in speedups of up to 4.2x on the regular kNNS experiment and 1.7x on the 99th percentile experiment.

Dataset size	kNNS	99 th percentile kNNS
100K	3.5x	1.5x
200K	3.25x	1.23x
400K	4.28x	1.7x
800K	4.15x	1.78x

Table 3. TrueKNN’s speedup over baseline on kNNS problems for UniformDist Dataset

6 Discussion

In this section, we discuss TrueKNN’s ability to accelerate *fixed-radius* kNNS and look at the hardware limitations that affect the efficiency of RT-accelerated general purpose applications.

6.1 Fixed-radius kNNS

Fixed-radius kNNS is a variant of kNNS, where the search space for neighbors is restricted to a particular radius. From our experiments in Section 5, we believe that, in addition to unbounded kNNS, our approach could *also* work well for fixed-radius kNNS applications.

In the 99th percentile experiment in Section 5.5.1, which is a fixed-radius kNNS problem, we found that we were consistently faster than the fixed-radius baseline. This observation gives us an important result: our multi-round approach can be effectively used even for *fixed* radius neighbor searches, as our approach minimizes the number of intersection tests performed. For large datasets and large k values, TrueKNN’s overhead costs are amortized, as we will see in Section 6.2.1. For smaller datasets and k values, we find that TrueKNN is sometimes slower than the baseline as these overhead costs are not amortized. We experimented with 100K-800K points from the Porto, 3DRoad, 3DIono and KITTI datasets and set $k = 5$. Though we were consistently faster on the Porto, 3DRoad and KITTI datasets, we were upto 1.6x slower on the 3DIono dataset, as shown in Fig 9. This is mainly due to the overhead of switching between host and device contexts

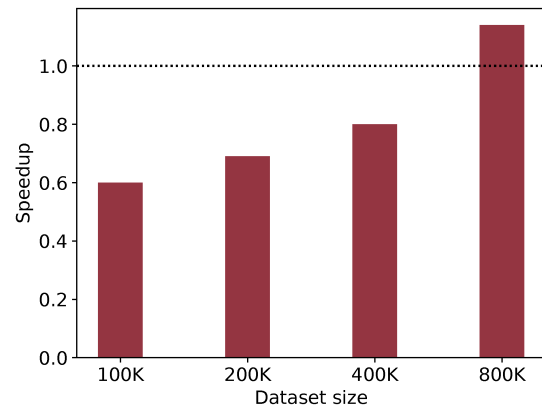


Figure 9. TrueKNN’s speedup compared to baseline for 99th percentile neighbor search on 3DIono dataset

and the cost is not amortized across the iterations needed to find all neighbors.

6.2 Limitations of RT Hardware and API

The fundamental limitation of using the RT hardware is the restriction to 3D datasets. As the hardware was created to accelerate graphics rendering, it was designed to work *only* with 3D data. We can get the hardware to work for 2D and 1D datasets by setting the corresponding dimensions to 0. For example, we set the z-dimension to 0 for 2D datasets, allowing the hardware to treat it as a 3D dataset.

Though it is not possible to express higher dimensional datasets using this hardware, we can use dimensionality reduction techniques such as Principal Component Analysis (PCA) [12], Linear Discriminant Analysis (LDA) and Generalized Discriminant Analysis (GDA) [4] to reduce the multi-dimensional dataset to just 3 dimensions. We also note that there are, in fact, many important 2D and 3D datasets, such as point clouds, geospatial data and geometries. These datasets are widely used in applications such as clustering and surface normal computation, both of which use kNNS as a subroutine.

6.2.1 Optix Overhead Costs. As the OWL API was intended to serve ray tracing applications, its setup is not ideal for general-purpose computations. We notice this in particular when we implement our multi-stage TrueKNN algorithm (Section 3.1). At the end of each round, we re-fit the bounding boxes around the spheres to expand our neighbor search space. For this operation, we need to switch the context from the device to the host to increment the size of the bounding boxes in the BVH. We then call the RayGen kernel (Fig 2) to begin the ray casting process, which requires a context switch from host to device. As the Optix API does not allow the BVH re-fit to happen on the device side, we incur this cost every round.

7 Related Work

Leveraging RT cores for non-Ray-Tracing Applications. The idea of using RT cores to accelerate applications other than ray tracing was first introduced by Wald *et al.* [28]. They formulated the problem of identifying a point's location in a tetrahedral mesh as a ray tracing problem by declaring the meshes as 3D objects in a scene and tracing rays originating at the query point. They show how leveraging both hardware-accelerated BVH traversal and ray-triangle intersection tests resulted in up to 6.5x speedup over other CUDA implementations. Morrical *et al.* used RT cores to accelerate the problem of finding a point's location in unstructured elements with both planar and bilinear faces [17] and rendering of unstructured meshes [16]. Zellmann *et al.* proposed a mapping of the fixed-radius nearest neighbor query to ray tracing queries by expanding spheres over points in the dataset and launching a small ray to record intersections [32]. They used the nearest neighbor query as a subroutine for the Spring Embedders force-directed graph drawing algorithm and show performance improvement between 4x to 13x over purely CUDA-based implementations. Evangelou *et al.* used the nearest neighbor mapping to solve the k -Nearest Neighbors problem, which returns the k closest neighbors of a point, and also perform photon mapping [8]. Zhu proposed z-order sorting and query partitioning of input data points to reduce control flow and memory divergence in RT-accelerated neighbor searches [33].

k -Nearest Neighbor Search. The k -Nearest Neighbor Search (kNNS) algorithm is used to find similarities within the feature space of input vectors. Researchers proposed using index structures to optimize the neighbor search, but soon found that this approach was useful only when the dataset had less than 10 dimensions [29]. Nagarkar *et al.* provides an overview of popular indexing techniques for lower and higher dimensional data [19]. Some of the techniques used in lower dimensions include M Tree [6], R/R* tree [25], and k -d Tree [5]. Hashing and quantization algorithms were used to find Approximate Nearest Neighbors (ANN) in higher dimensions [2, 3].

Researchers have worked on leveraging GPU acceleration for kNNS as it is often the computational bottleneck in applications. Qui *et al.* implemented kNNS on the GPU to accelerate point cloud registration and show that it is 88x faster than its CPU counterpart [23]. Johnson *et al.* created a suite of GPU-accelerated approximate neighbor searches called FAISS that used quantization optimizations [13]. Wieschollek *et al.* also uses a quantization-based approach in the form of Product Quantization Trees (PQT) to accelerate ANN [30] on GPUs.

8 Conclusion

In this work, we implemented TrueKNN, the first RT-accelerated k -Nearest Neighbor Search algorithm that *does not* restrict

the neighbor search space to a pre-defined fixed radius. TrueKNN uses an iterative approach where we initially sample the input dataset to guess a good search radius, and then incrementally increase the search space such that each point finds its k nearest neighbors. We found that TrueKNN was orders of magnitude faster than existing algorithms on the *unbounded* neighbor search task and significantly faster even on the *fixed-radius* neighbor search task.

Acknowledgments

We thank the anonymous ICS reviewers for their valuable feedback. We are grateful to Kirshanthan Sundararajah for his comments that helped improve the paper. This work was funded by NSF grants CCF-1908504, CCF-1919197 and CCF-2216978.

References

- [1] D.A. Adeniyi, Z. Wei, and Y. Yongquan. 2016. Automated web usage data mining and recommendation system using K-Nearest Neighbor (KNN) classification method. *Applied Computing and Informatics* 12, 1 (2016), 90–108. <https://doi.org/10.1016/j.aci.2014.10.001>
- [2] Alexandr Andoni and Piotr Indyk. 2006. Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 459–468. <https://doi.org/10.1109/FOCS.2006.49>
- [3] Martin Aumüller, Erik Bernhardtsson, and Alexander John Faithfull. 2018. ANN-Benchmarks: A Benchmarking Tool for Approximate Nearest Neighbor Algorithms. *CoRR abs/1807.05614* (2018). arXiv:1807.05614 <http://arxiv.org/abs/1807.05614>
- [4] G. Baudat and F. Anouar. 2000. Generalized Discriminant Analysis Using a Kernel Approach. *Neural Comput.* 12, 10 (oct 2000), 2385–2404. <https://doi.org/10.1162/089976600300014980>
- [5] Jon Louis Bentley. 1975. Multidimensional Binary Search Trees Used for Associative Searching. *Commun. ACM* 18, 9 (sep 1975), 509–517. <https://doi.org/10.1145/361002.361007>
- [6] Paolo Ciaccia, Marco Patella, and Pavel Zezula. 1997. M-Tree: An Efficient Access Method for Similarity Search in Metric Spaces. In *Proceedings of the 23rd International Conference on Very Large Data Bases (VLDB '97)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 426–435.
- [7] T. Cover and P. Hart. 1967. Nearest neighbor pattern classification. *IEEE Transactions on Information Theory* 13, 1 (1967), 21–27. <https://doi.org/10.1109/TIT.1967.1053964>
- [8] I. Evangelou, G. Papaioannou, K. Vardis, and A. A. Vasilakis. 2021. Fast Radius Search Exploiting Ray Tracing Frameworks. *Journal of Computer Graphics Techniques (JCGT)* 10, 1 (5 February 2021), 25–48. <http://jcgt.org/published/0010/01/02/>
- [9] Evelyn Fix and Joseph L. Hodges. 1989. Discriminatory Analysis - Nonparametric Discrimination: Consistency Properties. *International Statistical Review* 57 (1989), 238.
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. 2013. Vision meets Robotics: The KITTI Dataset. *International Journal of Robotics Research (IJRR)* (2013).
- [11] Khronos Group. 2009. OpenCL. <https://www.khronos.org/opencv/>
- [12] Harold Hotelling. 1933. Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology* 24 (1933), 498–520.
- [13] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data* 7, 3 (2019), 535–547.

- [14] Manohar Kaul, Bin Yang, and Christian S. Jensen. 2013. Building Accurate 3D Spatial Networks to Enable Next Generation Intelligent Transportation Systems. In *2013 IEEE 14th International Conference on Mobile Data Management*, Vol. 1. 137–146. <https://doi.org/10.1109/MDM.2013.24>
- [15] Luis Moreira-Matias, João Gama, Michel Ferreira, João Mendes-Moreira, and Luis Damas. 2013. Predicting Taxi–Passenger Demand Using Streaming Data. *IEEE Transactions on Intelligent Transportation Systems* 14, 3 (2013), 1393–1402. <https://doi.org/10.1109/TITS.2013.2262376>
- [16] Nate Morrical, Will Usher, Ingo Wald, and Valerio Pascucci. 2019. Efficient Space Skipping and Adaptive Sampling of Unstructured Volumes Using Hardware Accelerated Ray Tracing. *2019 IEEE Visualization Conference (VIS) (2019)*, 256–260.
- [17] Nate Morrical, Ingo Wald, Will Usher, and Valerio Pascucci. 2020. Accelerating Unstructured Mesh Point Location with RT Cores. *IEEE transactions on visualization and computer graphics (2020)*.
- [18] Prakash Nadkarni. 2016. Chapter 10 - Core Technologies: Data Mining and “Big Data”. In *Clinical Research Computing*, Prakash Nadkarni (Ed.). Academic Press, 187–204. <https://doi.org/10.1016/B978-0-12-803130-8.00010-5>
- [19] Parth Nagarkar, Arnab Bhattacharya, and Omid Jafari. 2021. Exploring State-of-the-Art Nearest Neighbor (NN) Search Techniques. In *Proceedings of the 3rd ACM India Joint International Conference on Data Science; Management of Data (8th ACM IKDD CODS; 26th COMAD) (Bangalore, India) (CODS-COMAD '21)*. Association for Computing Machinery, New York, NY, USA, 443–446. <https://doi.org/10.1145/3430984.3431968>
- [20] NVIDIA. 2007. CUDA. <https://developer.nvidia.com/cuda-zone>
- [21] NVIDIA. 2021. NVIDIA Turing Architecture Whitepaper. <https://gpltech.com/wp-content/uploads/2018/11/NVIDIA-Turing-Architecture-Whitepaper.pdf>
- [22] Victor Pankratius, A. Coster, Juha Vierinen, Philip Erickson, and Bill Rideout. 2015. GPS Data Processing for Scientific Studies of the Earth’s Atmosphere and Near-Space Environment. (01 2015), 1–12. https://doi.org/10.1007/978-3-319-23519-6_1651-1
- [23] Deyuan Qiu, Stefan May, and Andreas Nüchter. 2009. GPU-Accelerated Nearest Neighbor Search for 3D Registration. In *Computer Vision Systems*, Mario Fritz, Bernt Schiele, and Justus H. Piater (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 194–203.
- [24] Sebastian Raschka, Joshua Patterson, and Corey Nolet. 2020. Machine Learning in Python: Main developments and technology trends in data science, machine learning, and artificial intelligence. *arXiv preprint arXiv:2002.04803 (2020)*.
- [25] Nick Roussopoulos, Stephen Kelley, and Frédéric Vincent. 1995. Nearest Neighbor Queries. *SIGMOD Rec.* 24, 2 (may 1995), 71–79. <https://doi.org/10.1145/568271.223794>
- [26] Jagan Sankaranarayanan, Hanan Samet, and Amitabh Varshney. 2006. A Fast k-Neighborhood Algorithm for Large Point-Clouds. In *Symposium on Point-Based Graphics*, Mario Botsch, Baoquan Chen, Mark Pauly, and Matthias Zwicker (Eds.). The Eurographics Association. <https://doi.org/10.2312/SPBG/SPBG06/075-084>
- [27] Ingo Wald, Nathan Morrical, and Haines E. 2020. OWL-The Optix 7 Wrapper Library.
- [28] Ingo Wald, Will Usher, Nathan Morrical, Laura Lediaev, and Valerio Pascucci. 2019. RTX Beyond Ray Tracing: Exploring the Use of Hardware Ray Tracing Cores for Tet-Mesh Point Location. In *High-Performance Graphics - Short Papers*, Markus Steinberger and Tim Foley (Eds.). The Eurographics Association. <https://doi.org/10.2312/hpg.20191189>
- [29] Roger Weber, Hans-Jörg Schek, and Stephen Blott. 1998. A Quantitative Analysis and Performance Study for Similarity-Search Methods in High-Dimensional Spaces. In *Proceedings of the 24th International Conference on Very Large Data Bases (VLDB '98)*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 194–205.
- [30] Patrick Wieschollek, Oliver Wang, Alexander Sorkine-Hornung, and Hendrik P. A. Lensch. 2016. Efficient Large-Scale Approximate Nearest Neighbor Search on the GPU. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2027–2035. <https://doi.org/10.1109/CVPR.2016.223>
- [31] Wenchao Xing and Yilin Bei. 2020. Medical Health Big Data Classification Based on KNN Classification Algorithm. *IEEE Access* 8 (2020), 28808–28819. <https://doi.org/10.1109/ACCESS.2019.2955754>
- [32] Stefan Zellmann, Martin Weier, and Ingo Wald. 2020. Accelerating Force-Directed Graph Drawing with RT Cores. In *2020 IEEE Visualization Conference (VIS)*. 96–100. <https://doi.org/10.1109/VIS47514.2020.00026>
- [33] Yuhao Zhu. 2022. RTNN: Accelerating Neighbor Search Using Hardware Ray Tracing. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (Seoul, Republic of Korea) (PPoPP '22)*. Association for Computing Machinery, New York, NY, USA, 76–89. <https://doi.org/10.1145/3503221.3508409>