

Parallel Nearest Neighbors in Low Dimensions with Batch Updates

Guy E. Blelloch & Magdalen Dobson

Authors




Guy Blelloch



Magdalen Dobson

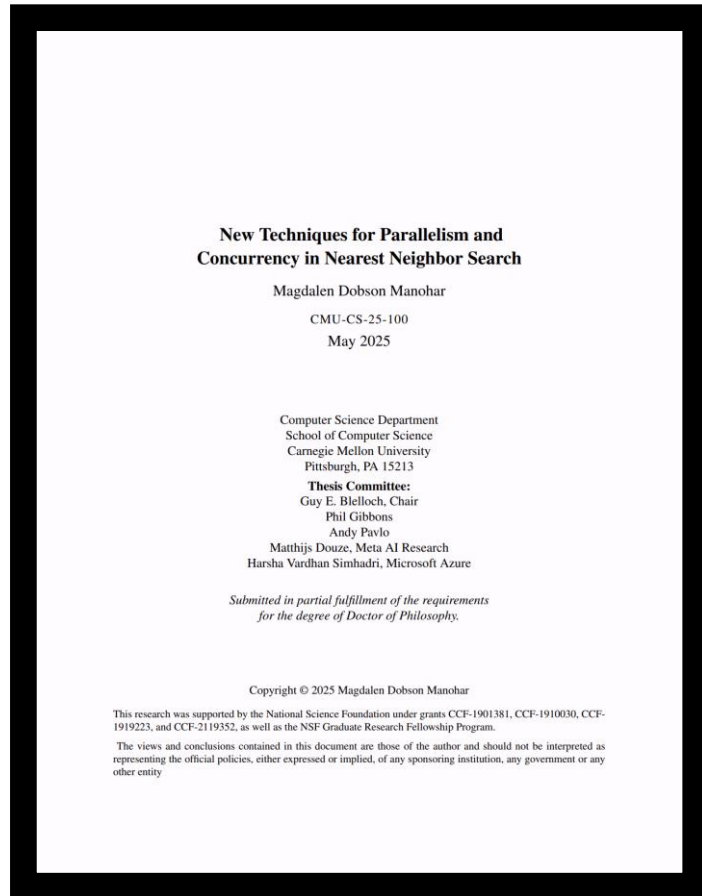
Authors



Magdalen Dobson

<https://www.csd.cs.cmu.edu/academics/doctoral/degrees-conferred/magdalen-dobson-manohar>

Authors



- Introduces the ideas from this paper (zd-tree)
- Extends zd-tree to support concurrent queries and updates
- Modify existing high-dimensional algorithms to be lock-free, deterministic, and scalable
 - ParlayANN

<https://www.csd.cs.cmu.edu/academics/doctorsal/degrees-conferred/magdalen-dobson-manohar>

Definitions

K-nearest neighbors:


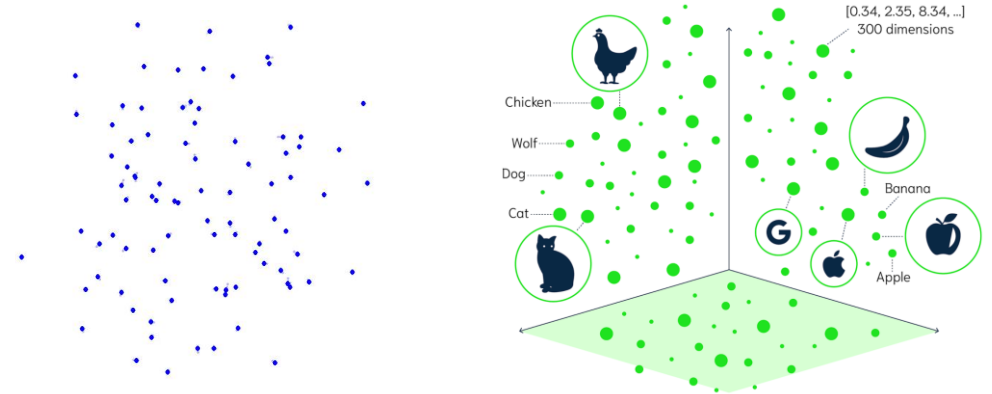
- Given a point x and set of points P , determine k nearest points to x in P

K-nearest neighbor graph:

- Calculate k -nearest-neighbors for all points in P


Motivation

- Nearest neighbors is a fundamental problem with many applications
 - Classification
 - Computer Graphics
 - Physics Simulations
 - Databases and IR
 - Robotics
 - Computational Biology
 - Etc.
- Existing parallel implementations:
 - Aren't performant
 - Do not support batch updates




Example (Dynamic Query)

K-nearest neighbors for **new** point ($k = 3$)




Example (Dynamic Query)

K-nearest neighbors for **new** point ($k = 3$)




Example (Dynamic Query)

K-nearest neighbors for **new** point ($k = 3$)




Example (Dynamic Query)

K-nearest neighbors for **new** point ($k = 3$)




Example (Dynamic Query)

K-nearest neighbors for **new** point




Example (K-Nearest Neighbor Graph)

K-nearest neighbors for every point ($k=2$)




Example (K-Nearest Neighbor Graph)

K-nearest neighbors for every point (k=2)




Example (K-Nearest Neighbor Graph)

K-nearest neighbors for every point (k=2)




Example (K-Nearest Neighbor Graph)

K-nearest neighbors for every point ($k=2$)



Example (K-Nearest Neighbor Graph)

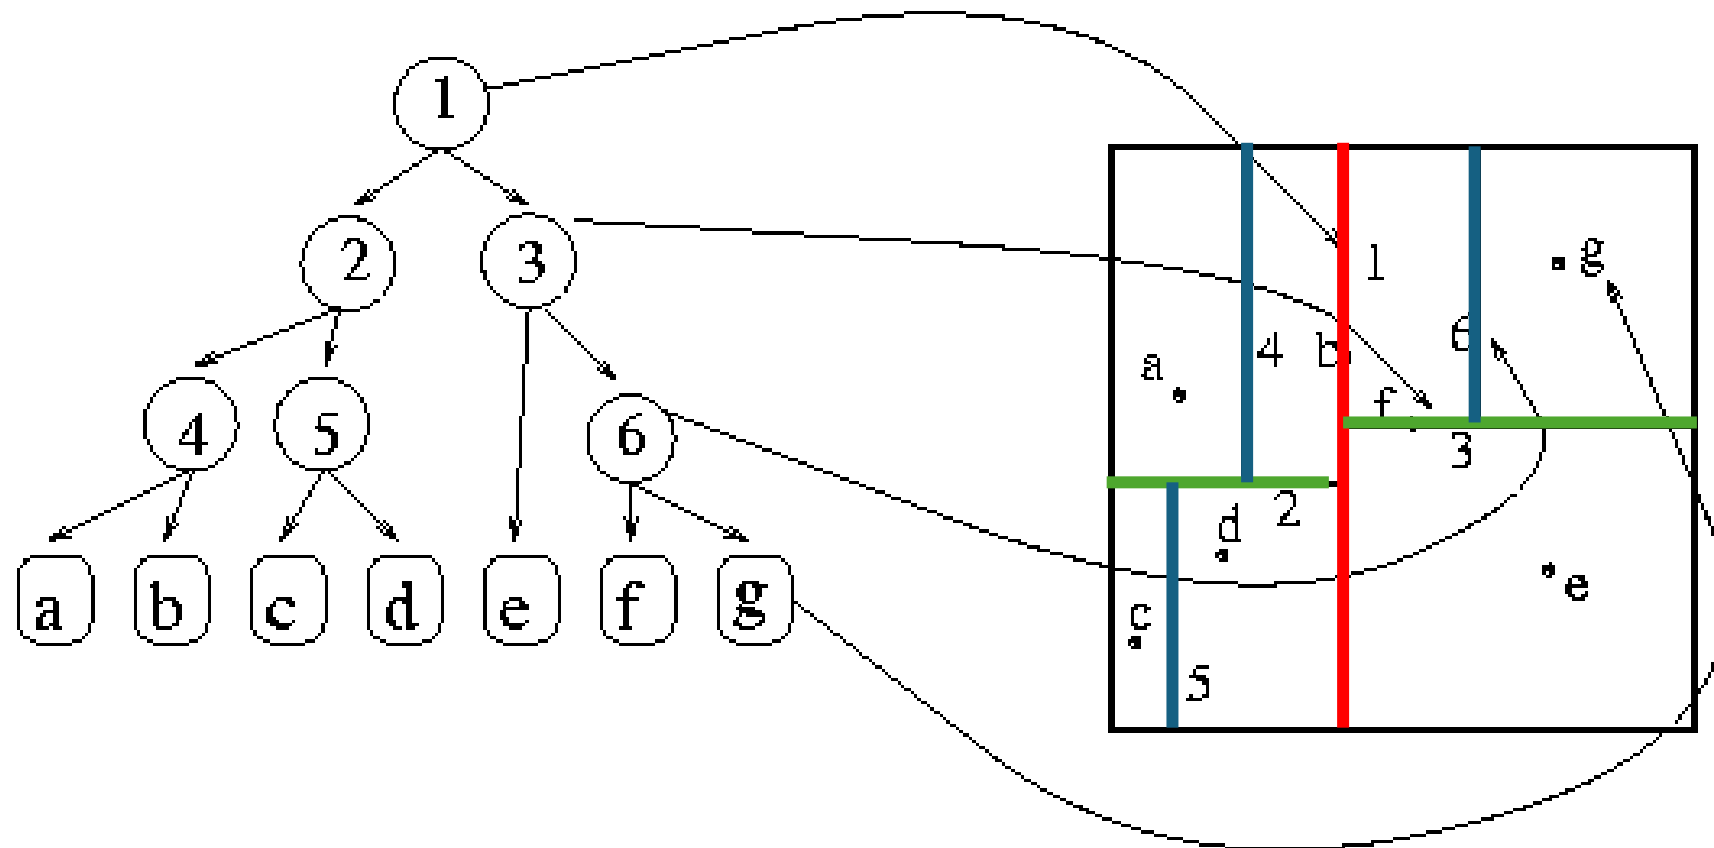


Relevant Preliminaries

Kd-Trees

- Tree data structure where each node represents a bounding box in d-dimensional space
 - Children split parent region into smaller bounding boxes at each level
- Common Splitting Rules:
 - Largest Dimension
 - Equal points on each side (median point)
- {quad,oct}-tree: 2^d equal sized children in d-dimensional space

Kd-Trees




Morton Ordering

- Maps points from multidimensional space to one-dimensional sequence while preserving spatial locality
- Definition: Interleave the binary form of each integer coordinate

Morton Ordering

- Maps points from multidimensional space to one-dimensional sequence while preserving spatial locality
- Definition: Interleave the binary form of each integer coordinate



Existing Approaches to Nearest Neighbors

- Kd-Tree based algorithms
 - Construct kd-tree to break up d-dimensional space and explicitly explore
 - Slow construction, Fast Queries
- Morton based algorithms
 - Sort input by Morton order and traverse implicit tree
 - Fast construction, Slow Queries

Novel Idea

Zd-tree based algorithm

- Definition:
 - kd-tree whose splitting rule uses the Morton ordering
 - Root - Entire bounding box
 - Children - Split points at level i based on whether the bit at place i is 0 or 1
- Implementation:
 - Internal nodes stores corners of bounding box, 2 children, parent
 - Leaf nodes store constant number of points instead of children

Algorithm 1: buildTree(P, b)

Input: A set of randomly shifted points sorted according to their Morton ordering and an integer b representing the bit we are working on, starting with the highest bit.

Output: The leaf or internal node that contains P 's bounding box

```
1 if  $b == 0$  or  $\text{size}(P) < \text{sizeCutoff}$  then
2   |   return createLeaf( $P$ )
3 else
4   |    $i = \text{splitUsingBit}(P, b)$  ;
5   |   do in parallel
6   |   |    $L = \text{buildTree}(P[1 : i], b - 1)$  ;
7   |   |    $R = \text{buildTree}(P[i : n], b - 1)$  ;
8   |   return createInternalNode( $L, R$ ) ;
```

Search Routines

- Goal: Add nearest neighbor candidates to min heap until certain all possible points have been explored
- Downward Search – Start at node and recurse downward into children
 - root-based algorithm starts at root with empty heap (single query)
- Upward Search – Start at leaf containing query point and recurse upward, performing downward search into sibling subtrees
 - bit-based algorithm finds leaf then calls searchUp (single query)
 - leaf-based algorithm runs searchUp from every leaf (full NN graph)

Downward Search

- Search if the furthest current candidate vertex is farther than the boxes bounding box from p
 - If leaf:
 - Check and add points to heap
 - Else:
 - Recurse on closer child bounding box
 - Recurse on further child bounding box

Algorithm 2: $\text{searchDown}(T, p, N)$

needed subroutines:

distance (p, N, k) returns infinity if N has fewer than k points and otherwise returns the distance from p to the furthest point in N ; $k = 1$ if not specified.

insert (N, p, k) adds p to the set N keeping only the k closest points to p .

withinBox (T, p, r) returns true if p is within a distance r of the bounding box for T .

Input: A pointer to a tree node T , the query point p , and a current set of up to k nearest neighbors N .

Output: The k -nearest neighbors of p .

```
1  $r \leftarrow \text{distance}(p, N, k)$  ;
2 if  $\text{withinBox}(T, p, r)$  then
3   if  $T = \text{Leaf}$  then
4      $Q \leftarrow$  set of points contained in  $T$  ;
5     for  $q \in Q$  do
6       if  $q \neq p$  then
7         if  $\text{distance}(q, p) < \text{distance}(p, N, k)$ 
8           then
9              $\text{insert}(N, p, k)$ ;
10  else
11     $R \leftarrow T.\text{Right}()$  ;
12     $L \leftarrow T.\text{Left}()$  ;
13     $\ell \leftarrow \text{distance}(p, L.\text{center}())$ ;
14     $r \leftarrow \text{distance}(p, R.\text{center}())$  ;
15    if  $\ell < r$  then
16       $N' = \text{searchDown}(L, p, N)$ ;
17      return  $\text{searchDown}(R, p, N')$ ;
18    else
19       $N' = \text{searchDown}(R, p, N)$ ;
20      return  $\text{searchDown}(L, p, N')$ ;
```

Upward Search

- Find leaf node p belongs to
- Add points in leaf to heap
- If furthest point is closer than edge of bounding box, return
- Otherwise, search down the sibling bounding box
- Recheck condition for parent bounding box

Algorithm 3: $\text{searchUp}(C, p)$
 $\text{withinBox}(T, p, r)$ with negative r returns true if p is in within the bounding box of T and at least r from the boundary.

Input: A leaf C of the kd-tree and a point p within the bounding box of C .

```
1  $N = \emptyset$ 
2  $Q \leftarrow$  set of points contained in  $C$  ;
3 for  $q \in Q$  do
4   if  $q \neq p$  then
5     if  $d(q, p) < \text{distance}(p, N, k)$  then
6        $\text{insert}(N, p, k)$ ;
7  $r \leftarrow \text{distance}(p, N, k)$ ;
8  $P \leftarrow C.\text{Parent}()$  ;
9 while not  $\text{withinBox}(C, p, -r)$  and  $P \neq \top$  do
10   if  $P.\text{Left}() = C$  then
11      $N = \text{searchDown}(P.\text{Right}(), p, N)$ ;
12   else
13      $N = \text{searchDown}(P.\text{Left}(), p, N)$ ;
14    $C = P$ ;
15    $r \leftarrow \text{distance}(p, N, k)$ ;
16    $P \leftarrow C.\text{Parent}()$  ;
17 return  $N$ 
```

Batch Dynamic Updates

- Recursively insert into children in parallel
- Split leaves if insertion would exceed leaf size restriction

Algorithm 4: batchInsert(T, P)

Input: A pointer to a node T of the kd-tree and a set of points P contained in its bounding box and sorted according to their Morton order

```
1 if  $T = \text{Leaf}$  then
2   | if  $\text{size}(P) + \text{size}(T) \geq \text{leafCutoff}$  then
3   |   | Insert  $p \in P$  into  $T$  ;
4   | else
5   |   | Split  $T$  into multiple leaf nodes ;
6 else
7   |  $b = T \rightarrow \text{bit}$  ;
8   |  $i = \text{splitUsingBit}(P, b)$  ;
9   | do in parallel
10  |   | batchInsert( $T \rightarrow \text{Right}, P[1 : i]$ ) ;
11  |   | batchInsert( $T \rightarrow \text{Left}, P[i : n]$ ) ;
```

Theorem 2.1

For a point set P of size n with bounded ratio, the zd-tree can be built using $O(n)$ work with $O(n^\epsilon)$ span and resulting tree height $O(\log n)$

Proof

- Bounding cube of P has side length within constant factor of max distance between two points
- Must be divided until two points minimum distance from one another are in separate cubes (assuming no coarsening)
- $d_{\max} / d_{\min} = \text{poly}(n)$ due to bounded ratio assumption
- Tree has $O(\log n)$ depth
- This implies radix sort can be used since we only need compare $O(\log n)$ bits to construct the tree

Theorem 2.2

For a z_d -tree representing a point set P of size n with bounded expansion, finding the k -nearest neighbors of a point $p \in P$ requires expected $O(k \log k)$ work.

Proof

- Work separated into two parts:
 - Work of searching through the points in leaves of the zd-tree
 - Work of finding candidate leaves

Lemma 2.1

- Searching for the k -nearest neighbor of a point p
 - $O(k)$ candidate points will be considered
 - Resulting in $O(k \log k)$ work to evaluate candidate points

Lemma 2.1

- Searching for the k -nearest neighbor of a point p
 - $O(k)$ candidate points will be considered
 - Resulting in $O(k \log k)$ work to evaluate candidate points

Proof

- Find ancestor box B with $O(k)$ points
- Expand radius to $4r$
- All candidate points contained in box
- Constant number of “overlapping” leaves

Lemma 2.2

- Traversing zd-tree
 - Expected number of tree edges to find nearest neighbors is $O(k)$

Lemma 2.2

- Traversing zd-tree
 - Expected number of tree edges to find nearest neighbors is $O(k)$

Proof

- Lemma 2.1 established search space contains $O(k)$ points
- Each edge we traverse to visit a leaf can be amortized against the total number of leaves
- Length of longest path is $O(1)$ in expectation

Theorem 2.3

Let T be a pruned z_d -tree representing point set P , and let Q be a point set of size k , such that $|P| + |Q| = n$. Then if $P \cup Q$ and Q both have bounded expansion and bounded ratio in the same hypercube X , Q can be inserted into T in $O(k \log(n/k))$ work and $O(k^c + \text{polylog}(n))$ span.

Implementation

- C++ using ParlayLib
- Optimizations:
 - Squared distances
 - Vector for small k and C++ stl priority queue for large k
 - Sort queries according to Morton order to improve cache utilization

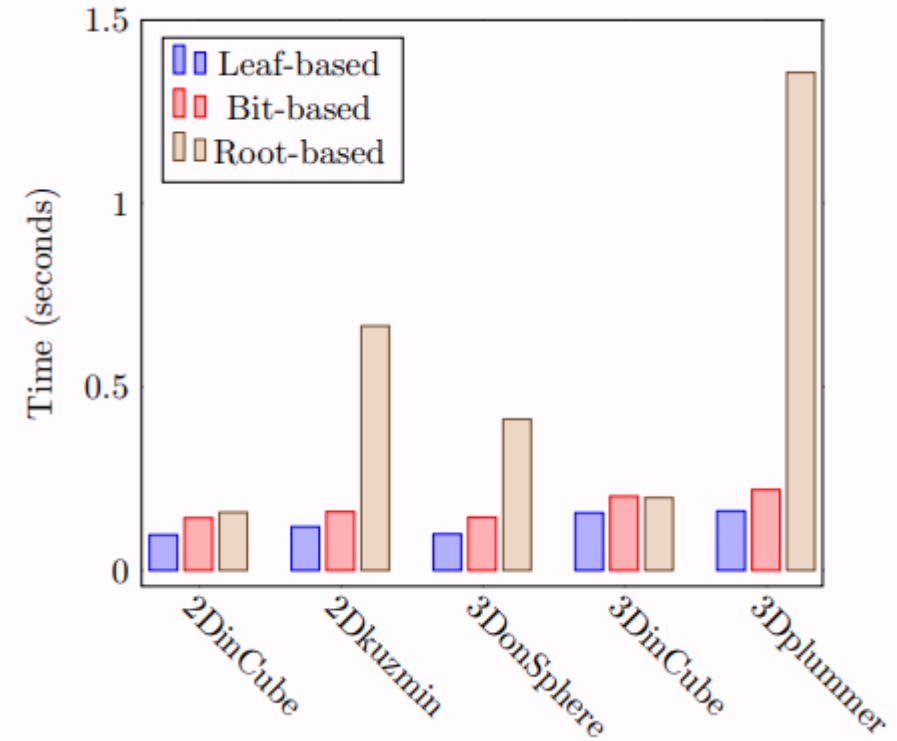
Empirical Results

Benchmarked against 3 existing implementations

- Chan: Naïve search from root
 - STANN: k-nearest neighbors and graph functions
 - CGAL: Widely adopted parallel implementation
-
- Note: Authors updated Chan and STANN to utilize modern parallel primitives


Search Strategies

- Compared three search strategies on four datasets of 10 million points
- Root based algorithms perform poorly
- Leaf based algorithm barely beats bit-based despite $O(n)$ vs $O(n \log n)$ work because constant factor is much larger




Comparison to State-of-the-Art


- Both implementations faster (often by order of magnitude)
- Robust performance across data sets and k
- Work efficiency
- Good scalability
 - 75x speedup with 72 cores and 144 hyperthreads




(a) Time required to calculate nearest neighbors as the size of the dataset increases. Calculated by dividing the total time by the number of points queried.




(b) The same as (a) but with a 2D dataset drawn randomly from a square instead of a 3D dataset.



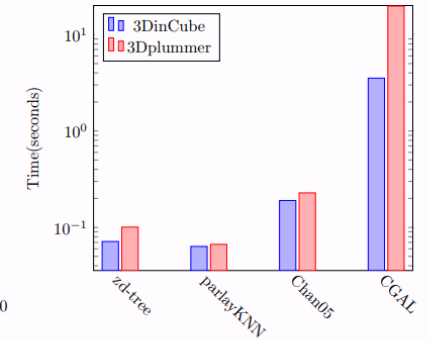
(c) Total work (threads \times time) required to build a tree of 10 million points, then build the nearest neighbor graph of the point set. Shown as the number of threads vary.



(d) The same measurements as (c), but on the 32-core AMD machine.



(e) Time required to calculate a neighbor as the number of neighbors k increases. Calculated by dividing the total time by k times the number of queries.



(f) The bars represent the total time each algorithm takes to build the data structure, for points drawn randomly from a 3D cube, and points drawn from a Plummer distribution.

Figure 4: Statistics related to non-dynamic queries. Unless otherwise stated, the size of the dataset is 10 million, the number of nearest neighbors $k = 1$, experiments were performed on 144 threads on a 72-core Dell R930, and data points are drawn randomly from a 3D cube.

Strengths

- Solid performance and scalability
- Elegant idea bridging existing strategies
- Support for parallel batch dynamic insertions and deletions

Weaknesses

- Requires bounding box of all data to be specified before hand
- Proofs are difficult to follow
- Tradeoff between work and span for sorting strategies

Future Directions

- Extensions to higher dimensions
- Relaxing theoretical assumptions
- Concurrent updates and queries
- External memory

Discussion Questions

- How realistic are theoretical assumptions for real world data?
- Would this implementation scale well on GPUs?
- What is the cost of preprocessing? How does it compare to existing implementations?