
Fast Exact Retrieval for Nearest-neighbor Lookup (FERN)

Richard Zhu¹

Abstract

Exact nearest neighbor search is a computationally intensive process, and even its simpler sibling — vector retrieval — can be computationally complex. This is exacerbated when retrieving vectors which have high-dimension d relative to the number of vectors, N , in the database. Exact nearest neighbor retrieval has been generally acknowledged to be a $O(Nd)$ problem with no sub-linear solutions. Attention has instead shifted towards Approximate Nearest-Neighbor (ANN) retrieval techniques, many of which have sub-linear or even logarithmic time complexities. However, if our intuition from binary search problems (e.g. $d = 1$ vector retrieval) carries, there ought to be a way to retrieve an organized representation of vectors without brute-forcing our way to a solution. For low dimension (e.g. $d = 2$ or $d = 3$ cases), `kd-trees` provide a $O(d \log N)$ algorithm for retrieval. Unfortunately the algorithm deteriorates rapidly to a $O(dN)$ solution at high dimensions, in practice. We propose a novel algorithm for logarithmic Fast Exact Retrieval for Nearest-neighbor lookup (FERN), inspired by `kd-trees`. The algorithm achieves $O(d \log N)$ look-up with 100% recall on 10 million $d = 128$ uniformly randomly generated vectors.

1. Introduction

Vector retrieval is pervasive, underlying search engines, transformers, and open-book language models. At heart, one of key attributes of computing systems lie in their ability to retrieve knowledge. Sometimes, when this knowledge is sufficiently broad — and a powerful enough retrieval architecture is built — these computing systems may even be so good at retrieving relevant knowledge that they appear to reason.

Given the power of knowledge retrieval for both commercial and academic pursuits, significant energy has been devoted towards effectively converting various types of data into vectors, from words (Mikolov et al., 2013), images (Radford et al., 2021), and audio (Radford et al., 2022) to documents, sentences, and paragraphs (Dai et al., 2015).

In this work, we differentiate between look-up and search. Look-up involves the retrieval of vectors guaranteed to be contained in the database, while search involves the retrieval of queries not necessarily contained in the database. Retrieving queries without exact matches may involve instead retrieving that vector’s nearest neighbors. The definition of nearest can be further disambiguated into Euclidean or cosine similarity measures, among others. Note that under the hood, a Euclidean distance-based nearest neighbor algorithm can be easily adapted to be cosine similarity-based simply by dividing each vector in the database by its magnitude during insertion. During look-up, the query vector is then also divided by magnitude. The resulting nearest neighbors we obtain are also such by cosine similarity.

The scalable and effective look-up of large numbers of high dimensional vectors is thus desired. While the vanilla hashmap algorithm provides $O(1)$ time complexity for scalars, extending to $O(d)$ for vectors in d -dimensional space, this holds only when the cardinality of the hash function range is large relative to the number of elements, n . When the number of elements becomes large relative to the number of bins, b , finding the key within each bin becomes a linear search problem. Since we expect each bin to have $\frac{n}{b}$ collisions, the time complexity for look-up is $O(\frac{nd}{b})$.

2. Related work

Prior work has taken 3 major directions to attain fast nearest neighbor search. The approaches involve bucketing, divide and conquer, or graph-based approaches. These techniques are specifically tuned to work well for vectors in high dimensional space, which should be unsurprising since each of these techniques is really just an extension of familiar 1-d concepts: hashmaps, binary search, and breadth first search. I believe there are key learnings that can be taken from both exact and approximate retrieval and search settings, even though we are interested in the exact variant. some new techniques such as Certified Cosine certificates (Francis-Landau & Van Durme, 2019) offer structure-agnostic tweaks to speed up performance.

Bucketing Locality Sensitive Hashing (Indyk & Motwani, 1998) and k-means (Lloyd, 1982)(MacQueen, 1967), including more recent variants like k-means++(Arthur & Vas-

silvitskii, 2007), are two techniques that use bucketing to group database vectors with their nearby peers. Since it is hard to deal with queries that fall on the boundaries of adjacent clusters, these algorithms are approximate search algorithms. Both of these algorithms in practice take linear time for sufficiently high dimensions and large numbers of elements.

Divide and conquer kd-trees provide strong algorithms for pruning, with new variants attempting to decrease constant factors in the time complexity. (Zhang et al., 2012) A recent work (Ram & Sinha, 2019) attempts to create partitions based on random rotations of the dataset achieve the same search accuracy guarantees as RPTree (Dasgupta & Sinha, 2013) but with $O(d \log d + \log n)$ time complexity for approximate search.

Graph-based approaches Recent graph-based approaches to search, such as Navigable Small World (NSW) graphs (Malkov et al., 2014) and a later variant involving layers of NSWs providing increasing granularity. Further optimizations have been presented recently (Fu et al., 2019) (Jayaram Subramanya et al., 2019).

3. FERN

Goal We aim to build an algorithm that satisfies two primary goals: we must be able to perform quick look-up on vectors guaranteed to be contained in the database and we must be able to quickly insert vectors. When designing the algorithm, we assume our database contains n vectors, each spanning d dimensional space. The i -th vector v_i is defined as follows

$$\mathbf{v}_i = [v_{i,1}, v_{i,2}, \dots, v_{i,d}]^\top$$

where $v_i \sim \mathcal{U}(-1, 1)$

This process effectively generates vectors lying within a d dimensional "ball"-like shape of radius 1 in each direction. The directions in which the vectors point are also evenly distributed direction-wise.

In terms of time complexity, we define quick as anything taking logarithmic time — this means that lookup in a database of ten billion vectors should only take seven times longer than lookup in a database of a thousand vectors. That is remarkable, because a naive linear search would take ten million times longer, a nearly intractable time difference at scale. Since a logarithmic lookup time and linear space complexity is state of the art (SoTA), we believe a key contribution of our work is an alternative data structure and algorithm that achieves SoTA while simultaneously providing the capacity to be extended to logarithmic-time nearest neighbor search, given its unique approach to dividing the

vectors by hyperplanes defined based on the vectors in the database rather than measuring along a specific direction like the traditional kd-trees process. Each node is an object that stores a vector, pointers to the left and right children, and a pointer to that node's parent node. This results in a binary tree with undirected edges. While we ultimately implement retrieval using a queue structure, this bidirectional edge only adds marginal complexity to the algorithm and underlying data structure while enabling a backtracking-based traversal method. The queue-based method emulates a level-order traversal of candidate nodes while a stack-based backtracking-based traversal method that full explores a specific path before backtracking to explore each sibling node that could not be pruned without potentially missing the nearest neighbor, emulates a depth-first search.

Methodology We design a novel algorithm that is a variant of kd-trees, but has the capacity to perform better at higher dimensions. Broadly, the structure is a binary tree. Each node that has both left and right children defines a hyperplane using the vectors of its left and right children as support vectors.

The tree is constructed so that all vectors in each child's subtree are on the same side of the hyperplane as that child. This allows us to perform vector look-up in logarithmic time, provided that vectors are added to the database in a sufficiently random manner. It is feasible, however, for an adversarial insertion process to result in a heavily imbalanced tree and consequently worst-case linear look-up time. While we don't observe this as an issue in practice, we can resolve the issue by implementing a slightly more complicated variant of FERN - using a variant of the Red-Black Tree technique (Guibas & Sedgwick, 1978) to guarantee balanced trees, logarithmic depth, and thus logarithmic retrieval time complexity.

There are two key components to our algorithm: one method for insertion (Algorithm 1) and another for lookup (Algorithm 2).

The insertion algorithm (Algorithm 1) is fairly concise. When inserting a vector into the tree, it is placed at the root if the tree has not been initialized yet. Otherwise, if the current node is missing a left or right child, we insert the vector as a child node. If the node is a leaf node — that is, missing both left and right children — then the left child is always inserted first, before the right child.

During insertion, if a node has both left and right children, then we set the current node instead to the child node that is closest to the vector we are inserting. That is, if we form a hyperplane from the set of points equidistant to both left and right children, then we set the current node to the left child if the vector to be inserted lies on the same side of the hyperplane as the left child, otherwise we set the current

Algorithm 1 FERN Insertion

```

1: function INSERT(vector)
2:   if root not initialized then
3:     root  $\leftarrow$  VectorNode(null)
4:   end if
5:   node  $\leftarrow$  root
6:   while True do
7:     if no left child then
8:       set left node to vector
9:       break
10:    else if no right child then
11:      set right node to vector
12:      break
13:    else if vector closer to left child then
14:      node  $\leftarrow$  node.left
15:    else
16:      node  $\leftarrow$  node.right
17:    end if
18:  end while
19: end function

```

node to the right child.

The result of this insertion algorithm is that — for a balanced binary tree — we get a maximum tree depth of $O(\log n)$ where n is the number of elements in the database. Insertion time per vector is thus $O(\log n)$ since we only visit one node per depth level.

When looking up vectors from the data structure, we demonstrate a method (Algorithm 2) that has a per-vector retrieval time proportional to the maximum depth of the tree, since we only look at one node per depth level. However, when extended to search settings where the query is not known to be contained in the database, retrieval time becomes proportional to the number of elements in the tree. We can no longer automatically prune any queries that lie close to the hyperplane boundary since there is a possibility that the nearest neighbor and query may lie on different sides of the hyperplane.

Ostensibly, we would expect a non-negligible proportion of vectors to be sufficiently far from the hyperplane to be pruned. However, we notice that in practice, as the dimensionality of the vectors increase, so too does the proportion of vectors lying close to the boundary. This makes sense intuitively since we are trying to project increasingly higher dimensions of vectors onto a 1-d line (the line normal to the hyperplane and passing through both support vectors). During retrieval, we effectively perform the depth-first or level-order search, as described previously. For a balanced tree with strong boundaries (that is, most queries lie far away from the hyperplane), per-vector time should be logarithmic with respect to the number of elements already present in

the database (another word for our proposed data structure) at insertion-time. However, it becomes linear otherwise. We first define `mip` and `mip_vector`, the distance to the nearest neighbor found thus far and the vector representing the nearest vector retrieved thus far. We then create a queue and insert the root node.

Algorithm 2 FERN Lookup

```

1: function RETRIEVE(query)
2:   mip, mip_vector  $\leftarrow$   $\infty$ , None
3:   curr  $\leftarrow$  None
4:   queue  $\leftarrow$  [self.root]
5:   while queue not empty do
6:     curr  $\leftarrow$  oldest element in queue
7:     update mip, mip_vector if curr closer to query
8:     if curr has both left and right children then
9:       if query is closer to left child then
10:        queue.append(curr.left)
11:        if query close to boundary then
12:          queue.append(curr.right)
13:        end if
14:      else
15:        queue.append(curr.right)
16:        if query close to boundary then
17:          queue.append(curr.left)
18:        end if
19:      end if
20:     else if curr has left child only then
21:       queue.append(curr.left)
22:     else if curr has right child only then
23:       queue.append(curr.right)
24:     end if
25:   end while
26:   return mip_vector
27: end function

```

We then continuously pop a node from the head of the queue until the queue is empty. Each time we pop a node, we check whether its vector is closer to the query vector than the current best candidate for nearest neighbor, `mip_vector`, which is a Euclidean distance `mip` away from the query. For lookup, we are looking for an exact match, so we are seeking an `mip` of 0. If the node has a left child only or a right child only (the latter should never happen, but we have it as a redundancy against exceptions) then we add that node to the queue. Otherwise, if both children exist then we add to the queue the node that shares the same side of the hyperplane as the query. For lookup cases, we consider any query to be sufficiently far from the boundary that only one child node needs to be added to the queue per node. After all, whether a query is close to the boundary is somewhat arbitrary and the exact function definition depends on whether we are performing lookup or search.

For mapping applications, we can add an additional variable, `data` (a byte array), to the `Node` class.

4. Experimental results

During experiments, we typically utilize $d = 128$ with the same uniform distribution previously assumed. While this may not be characteristic of all data distributions, we note that our architecture is actually agnostic to the distribution of the vectors being inserted. What matters (in terms of potential effects on performance) is the order in which vectors are inserted based on their relative positions.

To run our experiments, we use the Intel Xeon Platinum 8380 CPU (2.30 GHz), the same processor used for running the popular `ann-benchmark`. For values of n we use 10^4 , $5 * 10^4$, 10^5 , $5 * 10^5$, and 10^6 , $5 * 10^6$, 10^7 . The last setting, equivalent to look-up on 10 million vectors, has comparable values of n and d to many of the Euclidean distance based benchmarks in `ann-benchmark`. We notice a nearly perfect logarithmic time complexity, and at $n = 10^7$ we run approximately 3000 retrievals/second without additional optimizations.

5. Discussion

During our experiments, we noticed that system dynamics can change drastically based on the dimensionality of the vectors. The bulk of the time spent on this paper (80 or so hours) was spent playing around with different ways of defining the hyperplanes and various algorithms that would balance or repair the tree to try to guarantee logarithmic retrieval for large d and n . Unfortunately, while these algorithms almost universally gave logarithmic time complexity for 100% recall, the performance broke down drastically beyond $d = 2$ or $d = 3$. In particular, we note the importance of having well-defined hyperplane boundaries.

Boundary sharpness We want to be able to maximize pruning since we achieve $O(\log_2 n)$ time complexity when we prune 50% of the nodes in the database each time we measure the distance between a node and the query). Indeed, note that since a k -means based nearest neighbors search allows us to prune up to n/k nodes per comparison, we might wonder why 2-means search doesn't achieve $O(\log n)$ complexity. It's because k -means has an $O(n/k + k)$ time complexity, the n/k term means that we would need to do a linear number of searches regardless of the number of clusters.

We observe empirically that the proportion of nodes in the database that are visited increases sharply when there are more nodes that are closer to the hyperplane than the support vectors that define the plane. That is to say, during the insertion process, there will be vectors that lie between




Figure 1. FERN lookup with vectors where $d = 128$ and look-up time is averaged over 1000 vectors randomly sampled from the database

a support vector and the hyperplane. Now when we're retrieving, the query may lie on one side of the hyperplane but its nearest neighbor may be one of these "in-between" nodes on the other side of the hyperplane. This means we now must be much more prudent when pruning which decreases the proportion of vectors that are pruned and thus increases the time complexity.

6. Conclusion

We are able to achieve our goal of creating a novel vector database structure that achieves state of the art look-up time complexity that is logarithmic in the number of vectors. The algorithm presented here, FERN, maintains 100% recall while performing lookup on vectors in high-dimensional space (e.g. $d = 128$) with n varying from 10^4 to 10^7 , and presents a potential path towards a data structure and algorithm that will allow for the first sub-linear exact nearest neighbor retrieval process.

We believe that the exact process for attempting to perform binary search on a vector database requires carefully defined hyperplanes, which presents an area for further work. We find the "fixing" step of the Red-Black tree algorithm to be particularly inspirational as a direction of future work.

We further also believe that an alternative for hyperplanes is to use a graph based approach, similar to the approach taken in many recent works (Malkov et al., 2014) (Fu et al., 2019) (Jayaram Subramanya et al., 2019), since this could allow us to more easily divide the database in a well-defined and easy to update way. Overall we are excited by the potential and hope to further develop this algorithm in pursuit of sub-linear exact search.

References

- Arthur, D. and Vassilvitskii, S. k-means++: the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '07, pp. 1027–1035, USA, January 2007. Society for Industrial and Applied Mathematics. ISBN 978-0-89871-624-5.
- Dai, A. M., Olah, C., and Le, Q. V. Document Embedding with Paragraph Vectors, July 2015. URL <http://arxiv.org/abs/1507.07998>. arXiv:1507.07998 [cs].
- Dasgupta, S. and Sinha, K. Randomized partition trees for exact nearest neighbor search. *ArXiv*, February 2013. URL <https://www.semanticscholar.org/paper/dfa274737dbad5b73cf3c1c76c7dd701c0f3f29d>.
- Francis-Landau, M. and Van Durme, B. Exact and/or Fast Nearest Neighbors, November 2019. URL <http://arxiv.org/abs/1910.02478>. arXiv:1910.02478 [cs].
- Fu, C., Xiang, C., Wang, C., and Cai, D. Fast approximate nearest neighbor search with the navigating spreading-out graph. *Proceedings of the VLDB Endowment*, 12(5):461–474, January 2019. ISSN 2150-8097. doi: 10.14778/3303753.3303754. URL <https://doi.org/10.14778/3303753.3303754>.
- Guibas, L. J. and Sedgewick, R. A dichromatic framework for balanced trees. In *19th Annual Symposium on Foundations of Computer Science (sfcs 1978)*, pp. 8–21, October 1978. doi: 10.1109/SFCS.1978.3. URL <https://ieeexplore.ieee.org/document/4567957>. ISSN: 0272-5428.
- Indyk, P. and Motwani, R. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing - STOC '98*, pp. 604–613, Dallas, Texas, United States, 1998. ACM Press. ISBN 978-0-89791-962-3. doi: 10.1145/276698.276876. URL <http://portal.acm.org/citation.cfm?doid=276698.276876>.
- Jayaram Subramanya, S., Devvrit, F., Simhadri, H. V., Krishnawamy, R., and Kadekodi, R. DiskANN: Fast Accurate Billion-point Nearest Neighbor Search on a Single Node. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.nips.cc/paper_files/paper/2019/hash/09853c7fb1d3f8ee67a61b6bf4a7f8e6-Abstract.html.
- Lloyd, S. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, March 1982. ISSN 1557-9654. doi: 10.1109/TIT.1982.1056489. URL <https://ieeexplore.ieee.org/document/1056489>. Conference Name: IEEE Transactions on Information Theory.
- MacQueen, J. Some methods for classification and analysis of multivariate observations. 1967. URL <https://www.semanticscholar.org/paper/Some-methods-for-classification-and-analysis-of-MacQueen/Some-methods-for-classification-and-analysis-of-MacQueen/Some-methods-for-classification-and-analysis-of-MacQueen/Some-methods-for-classification-and-analysis-of-MacQueen/ac8ab51a86f1a9ae74dd0e4576d1a019f5e654ed>.
- Malkov, Y., Ponomarenko, A., Logvinov, A., and Krylov, V. Approximate nearest neighbor algorithm based on navigable small world graphs. *Information Systems*, 45:61–68, September 2014. ISSN 0306-4379. doi: 10.1016/j.is.2013.10.006. URL <https://www.sciencedirect.com/science/article/pii/S0306437913001300>.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient Estimation of Word Representations in Vector Space, September 2013. URL <http://arxiv.org/abs/1301.3781>. arXiv:1301.3781 [cs].
- Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., and Sutskever, I. Learning Transferable Visual Models From Natural Language Supervision, February 2021. URL <http://arxiv.org/abs/2103.00020>. arXiv:2103.00020 [cs].
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. Robust Speech Recognition via Large-Scale Weak Supervision, December 2022. URL <http://arxiv.org/abs/2212.04356>. arXiv:2212.04356 [cs, eess].
- Ram, P. and Sinha, K. Revisiting kd-tree for Nearest Neighbor Search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 1378–1388, Anchorage AK USA, July 2019. ACM. ISBN 978-1-4503-6201-6. doi: 10.1145/3292500.3330875. URL <https://dl.acm.org/doi/10.1145/3292500.3330875>.
- Zhang, Y., Zha, L., Liu, J., and Xu, Z. A Large-Scale Online Search System of High-Dimensional Vectors Based on Key-Value Store. In *2012 Eighth International Conference on Semantics, Knowledge and Grids*, pp. 233–236, October 2012. doi: 10.1109/SKG.2012.3. URL <https://ieeexplore.ieee.org/abstract/document/6391840>.