

i-Octree: A Fast, Lightweight, and Dynamic Octree for Proximity Search

Jun Zhu¹, Hongyi Li¹, Zhepeng Wang², Shengjie Wang^{3,†}, and Tao Zhang^{1,†}, *Fellow, IEEE*

Abstract—Establishing the correspondences between newly acquired points and historically accumulated data (i.e., the map) through nearest neighbor search is crucial in numerous robotic applications. However, static tree data structures are inadequate to handle large and dynamically growing maps in real-time. To address this issue, we present the *i-Octree*, a dynamic octree data structure that supports both fast nearest neighbor search and real-time dynamic updates, such as point insertion, deletion, and on-tree down-sampling. The *i-Octree* is built upon a leaf-based octree and has two key features: a local spatially continuous storing strategy that allows for fast access to points while minimizing memory usage, and local on-tree updates that significantly reduce computation time compared to existing static or dynamic tree structures. The experiments show that the *i-Octree* outperforms contemporary state-of-the-art approaches by achieving, on average, a 19% reduction in runtime on real-world open datasets.

I. INTRODUCTION

Nearest neighbor search (NNS) is necessary in many robotic applications, such as real-time LiDAR-based simultaneous localization and mapping (SLAM) and motion planning, where the data is sampled sequentially and real-time mapping is essential. For example, in LiDAR-based SLAM, NNS is crucial to computing feature [1], [2], estimating normal [3], [4], and matching new points to the map [5]–[8]. Recent advances in LiDAR technology have not only significantly reduced the cost, size, weight, and power of LiDAR but also enabled high performance of LiDAR, making it an essential sensor for robots [9]. However, this also poses challenges to NNS. Current LiDAR sensor can produce a large amount of 3D points with centimeter accuracy at hundreds of meters measuring range per second. The large amount of sequential data requires being processed in real-time, which is a quite challenging task on robots with limited onboard computing resources. To guarantee the efficiency of NNS, maintaining a large map supporting high efficient inquiry and dynamic updates with newly arriving points in real-time comes to vital importance.

Although various static tree data structures have been proposed, they struggle to meet the demands. R-tree [10] and R*-tree [11] partition the data by grouping nearby points into their minimum bounding rectangle in the next higher level of the tree. k-d tree [12] is a well-known instance to split the space. The octree recursively splits the space equally into eight axis-aligned cubes, which form the volumes represented by eight child nodes. Although

k-d tree is a favored data structure in general k-nearest neighbor (KNN) search libraries, it is hard to draw any final conclusions as to whether the k-d tree is better suited for NNS compared to other data structures. Comparative studies [13], [14] show that the performance of different implementations of k-d trees can be diverse, and the octree is amongst the best performing algorithms especially for radius search due to its regular partitioning of the search space. Despite its performance, octree has not been fully exploited. Elseberg et al. [15] proposed an efficient octree to store and compress 3D data without loss of precision. Behley et al. [16] proposed an index-based octree that significantly improves the radius neighbor search in three-dimensional data, while the KNN search and dynamic updates are not enabled. When incorporating these static trees in real robotic applications, re-building the entire tree from scratch [17] repeatedly is inevitable, which is so time-consuming that the robots may fail to run.

In this paper, we propose a dynamic octree structure called *i-Octree*, which incrementally updates the octree with new points and enables fast NNS. In addition, our *i-Octree* boasts impressive efficiency in both time and memory, adaptable to various types of points, and allows for on-tree down-sampling and box-wise delete. We conduct validation experiments on both randomized data and real-world open datasets to assess the effectiveness of *i-Octree*. In the randomized data experiments, our *i-Octree* demonstrates significant improvements in runtime compared to the state-of-the-art incremental k-d tree (i.e., *ikd-Tree* [9] proposed recently). Specifically, it reduces run-time by 64% for building the tree, 66% for point insertion, 30% for KNN search, and 56% for radius neighbor search. Moreover, when applied to real-world data in LiDAR-based SLAM, *i-Octree* showcases remarkable time performance enhancements. It achieves over twice the speed of the original method while often maintaining higher accuracy levels. Furthermore, Our implementation of *i-Octree* is open-sourced on Github¹.

The remaining paper is organized as follows: the design of *i-Octree* is described in Section II. Experiments are shown in Section III, followed by conclusions in Section IV.

II. *i-Octree* DESIGN AND IMPLEMENTATION

The *i-Octree* takes sequential point clouds as input with two objectives: dynamically maintains a global map and performs fast NNS (i.e., KNN search and radius neighbor search) on the map. Fig. 1 illustrates a typical application

† Corresponding author: taozhang@tsinghua.edu.cn, wangsj23@mail.tsinghua.edu.cn.

¹ Department of Automation, Tsinghua University. ² Lenovo Research. ³ Institute for Interdisciplinary Information Sciences (IIIS), Tsinghua University.

¹<https://sites.google.com/view/I-Octree>




Fig. 1. An example of using *i-Octree* in odometry. The *i-Octree* and odometry collaborate to estimate the poses of the 3D data obtained from the range sensors. The *i-Octree* provides a robust and efficient data structure for storing and querying the 3D data, while the odometry enables the estimation of the poses of the data points.




Fig. 2. Illustration of locations of an octant's points in memory. (a) scattered locations; (b) continual locations.

of the *i-Octree*. Range sensors continuously perceive their surroundings and generate sequential 3D range data periodically. The initial scan of the range data is utilized to construct the *i-Octree* and define the global coordinate frame. The *i-Octree* then facilitates the establishment of correspondences between the newly arriving data and the historical data through KNN search or radius neighbor search. Based on these correspondences, the poses of the new data can be estimated, and the 3D points with pose are added to the *i-Octree*. To prevent the map size in *i-Octree* from growing uncontrollably, only map points within a large local region (i.e., axis-aligned box) centered around the current position are maintained. In the following, we first describe the data structure and construction of the *i-Octree*, then we focus on dynamic updates and NNS.

A. Data Structure and Construction

An *i-Octree* node has up to eight child nodes, each corresponding to one octant of the overlying axis-aligned cube. An octant, starting with an axis-aligned bounding box with center $c \in \mathbb{R}^3$ and equal extents $e \in \mathbb{R}$, would be subdivided recursively into smaller octants of extent $\frac{1}{2}e$ until it contains less than a given number of points – the bucket size b or its extent is less than a minimal extent – e_{min} . For memory efficiency, octants without points are not created. In addition, we only keep indices and coordinates of points in each leaf octant and there is no points in non-leaf octants. To enable extremely fast access to each point in each leaf octant, we propose a local spatially continuous storing strategy (as shown in Fig. 2) which reallocates a segment of continuous memory for storing points information in the leaf octant after subdivision. Furthermore, the reallocation facilitates box-wise delete and incremental update, as it allows operating on a segment of memory without influencing others.

Based on above consideration, an octant C_o of our *i-Octree* contains the center $c_o \in \mathbb{R}^3$, the extent e_o , the points P_o storing coordinates and indices, and a pointer pointing to the address of the first child octant. The subscript “o” is used to




Fig. 3. Fig. (a) and (b) illustrate insertion of new points (red) out of the range to the *i-Octree*. In (a), the left yellow cube is the root octant as well as the leaf octant with points (black) in the beginning. After inserting, the root octant becomes the light purple cube. In (b), the purple node is the root octant and it is updated after inserting new octants (dashed black rectangle).

distinguish between different octants. Particularly, let C_r be the root octant, and P_r, e_r , and c_r are the points, the extent, and the center, respectively.

As for building an incremental octree, we firstly eliminate invalid points, calculate the axis-aligned bounding box of all valid points, and keep only indices and coordinates of valid points. Then, starting at the root, the *i-Octree* recursively splits the axis-aligned bounding box at the center into eight cubes indexed by Morton codes [18] and subdivides all the points in current octant into each cube according to their cube indices calculated. When a stopping criteria is satisfied, a leaf octant will be created and a segment of continuous memory will be allocated to store information in points of the leaf node.

B. Dynamic Updates

The dynamic updates include insertion of one or more points (i.e., incremental update) and deletion of all points in an axis-aligned box (i.e., box-wise delete). The insertion is integrated with down-sampling, which maintains the octree at a pre-determined resolution.

1) *Incremental Update*: When inserting new points, we have to consider the situation that some points may be beyond the boundary of the axis-aligned bounding box of the original tree. Once there are points out of the range of octree, we have to expand the bounding box by creating new root octant whose children contain current root octant. This process may be executed several times to ensure that all new points are within the range of the tree. Then, new points are added to the expanded octree (see Fig. 3).

In consideration of efficient points queries in robotic applications, *i-Octree* supports down-sampling that executes simultaneously with point insertion. The down-sampling focuses on the new points and deletes the points that satisfy a certain condition: they are subdivided into a leaf octant whose extent is less than $2e_{min}$ and size is larger than $b/8$.




Fig. 4. Illustration of box-wise delete. The blue box is the given box, and the numbers 0 to 7 (Morton codes) are the indices of the octants. Octants, with indices 0, 1, 4, and 5 are directly deleted. Octant 7 is also deleted due to having no points.

The process of adding new points to an octant C_o is similar to the construction of an octant. If C_o is a leaf node and it satisfies the subdivision criteria, all points (i.e., old points and new added points) in C_o will be recursively subdivided into child octants. If down-sampling is enabled, $e_o \leq 2e_{min}$, and $|P_o| > b/8$, new points will be deleted later instead of being added to C_o . Otherwise, a segment of continuous memory will be allocated for the updated points. If C_o has child octants, the problem becomes assigning the newly added points to various children, and only the new points are to be subdivided. This process is similar to the one mentioned above, except for the recursive updating of octants.

2) *Box-wise Delete*: In certain robotic applications, such as SLAM, only the points near the agent are required to estimate the state. Consequently, the points located far away from the agent in the *i-Octree* are not essential and can be removed for efficiency concerns.

When it comes to removing unnecessary points in an axis-aligned cube, instead of directly searching for points within the given space and deleting them, the *i-Octree* firstly checks whether the octants are inside the given box. All octants inside the given box will be directly deleted without searching for points in them, which significantly reduces the deletion time. The deletion of octants have no influence on others thanks to the local spatially continuous storing strategy. For the leaf octants overlapping with the given box, we delete the points within the box and allocate a new segment of memory for the remaining. If a leaf octant contains no points after deletion, it will be deleted as shown in Fig. 4.

C. K-Nearest Neighbor Search

Using *i-Octree*, we can retrieve k nearest neighbors for an arbitrary query point $\mathbf{q} \in \mathbb{R}^3$. The nearest search on the *i-Octree* is an accurate nearest search [19] instead of an approximate one as [20]. We maintain a priority queue \mathbf{h} with a maximal length of k to store the k -nearest neighbors so far encountered and their distance to the query point \mathbf{q} . The last element of \mathbf{h} always has the largest distance regardless of pushing or popping. The axis-aligned box of each octant is well utilized effectively to accelerate the nearest neighbor search using a “bounds-overlap-ball” test [19] and a proposed priority search order pre-computed according to the fixed indices of child octants.

Firstly, we recursively search down the *i-Octree* from its root node until reach the leaf node closest to \mathbf{q} . Then the distances from \mathbf{q} to all points in the leaf node and

corresponding indices will be pushed to priority queue \mathbf{h} . All leaf nodes so far encountered will be searched before \mathbf{h} is full. If \mathbf{h} is full and the search ball $\mathcal{S}(\mathbf{q}, d_{max})$ defined by \mathbf{q} and the largest distance d_{max} in \mathbf{h} is inside the axis-aligned box of current octant, the searching is over. If an octant C_k doesn't contain the search ball $\mathcal{S}(\mathbf{q}, d_{max})$, then one of the following three conditions must be satisfied:

$$e_k - |q_x - c_{k,x}| < d_{max}, \quad (1)$$

$$e_k - |q_y - c_{k,y}| < d_{max}, \quad (2)$$

$$e_k - |q_z - c_{k,z}| < d_{max}, \quad (3)$$

where $\mathbf{q} = (q_x, q_y, q_z)^T$, $\mathbf{c}_k = (c_{k,x}, c_{k,y}, c_{k,z})^T$. If none of the above conditions hold, the search ball is inside the octant.

We update \mathbf{h} by investigating octants overlapping the search ball $\mathcal{S}(\mathbf{q}, d_{max})$, since only these could potentially contain points that are also inside the desired neighborhood. We define the distance d between \mathbf{q} and C_k as below:

$$d = \|\sigma(|\mathbf{q} - \mathbf{c}_k| - \mathbf{1}e_k)\|_2, \quad (4)$$

where $\mathbf{1} = (1, 1, 1)^T$ and $\sigma(x) = x$ if $x > 0$, otherwise $\sigma(x) = 0$. $d < d_{max}$ indicates that C_k overlaps $\mathcal{S}(\mathbf{q}, d_{max})$. Let C_o be the parent node of C_k . In order to speed up the process, we sort the candidate child octants of C_o according to their distances to C_k and get 8 different sequences in \mathbf{I}_{order} . Such that the closer octants are earlier to be searched and the search reaches its end early.

D. Radius Neighbor Search

For an arbitrary query point $\mathbf{q} \in \mathbb{R}^3$ and radius r , the radius neighbor search method finds every point \mathbf{p} satisfying $\|\mathbf{p} - \mathbf{q}\|_2 < r$. The process is similar to KNN search except for a fixed radius and an unlimited k . We adopt the pruning strategy proposed by Behley et al. [16] with improvements to reduce computation cost. Before test whether $\mathcal{S}(\mathbf{q}, r)$ completely contains an octant C_k , we make a simple test whether r^2 is larger than $3e_k^2$. The simple test is a necessary condition for $\mathcal{S}(\mathbf{q}, r)$ containing C_k with small cost. Besides, we try to avoid the extraction of a root in the algorithm. These two tricks together with our local spatially continuous storing strategy play key roles in accelerating the searching.

III. APPLICATION EXPERIMENTS

We compare our *i-Octree* to publicly available implementations of static k-d trees (i.e., k-d tree used in Point Cloud Library (PCL)), incremental k-d tree (i.e., *ikd-Tree*), and PCL octree. We conduct the experiments on randomized data and various open real-world datasets. We first evaluate our *i-Octree* against the PCL octree and the state-of-the-art *ikd-Tree* across various operations: tree construction, point insertion, KNN search, radius neighbor search, and box-wise delete, using randomized 3D point data of different sizes. Then, we validate the *i-Octree* in actual robotic applications on real-world datasets by replacing the static k-d tree in the LiDAR-based SLAM without any refinement and evaluate the time performance and accuracy. All experiments are performed on a PC with Intel i7-13700K CPU at 3.40GHz.




Fig. 5. Dynamic data structure comparison over different tree sizes.

A. Randomized Data Experiments

The efficiency of our *i-Octree* is fully investigated by the following experiments on randomized data.

1) *Performance Comparison*: In this experiment, we investigate the time performance of three implementations of dynamic data structures (i.e., the *i-Octree*, the *ikd-Tree*, and the PCL octree). Both *ikd-Tree* and PCL octree are state-of-the-art implementations with high efficiency and support point insertion, which is the key consideration why we choose them to be compared with our *i-Octree* for a fair comparison. In this experiment, 200,000 points are generated randomly in a $10m \times 10m \times 10m$ space (i.e., the workspace) to build the trees. Then 100 test operations are conducted. In each test operation, 2,000 new points randomly sampled in the workspace are inserted into the trees. Besides, we perform KNN search with $k = 5$ on each of 200 new points randomly sampled and radius neighbor search with $r = 0.3m$ on each of another 200 ones. Down-sampling is not enabled on *i-Octree* and *ikd-Tree* in this experiment. We fix $e_{min} = 0.01m$ for *i-Octree* and PCL octree and $b = 32$ for *i-Octree* in the following experiments. We record the time for tree construction and point insertion, the total time of KNN search for 200 points, and the total time of radius neighbor search for another 200 points at each step.

The dynamic data structure comparison over different tree sizes is shown in Fig. 5. The Table I shows the comparison of average time consumption and peak memory usage. When the tree size increases from 200,000 to 400,000, the time for point insertion (without down-sampling) of *i-Octree* and PCL octree remains stable at $0.8ms$ while that for the *ikd-Tree* is 3 times larger and grows linearly with the tree size. For KNN

TABLE I
THE COMPARISON OF AVERAGE TIME CONSUMPTION AND PEAK MEMORY USAGE

	<i>i-Octree</i>	<i>ikd-Tree</i>	PCL octree
Construction[ms] ¹	16.63	46.26	52.66
Point Insertion[ms]	0.83	2.45	0.85
KNN Search[ms]	0.41	0.59	2.53
Radius Neighbor Search[ms]	0.81	1.85	3.96
Peak Memory Usage[Mb]	21.53	62.91	136.82

¹ The construction of the tree is conducted only once at the first time.




Fig. 6. The comparison of time consumption.

search, our method demonstrates a performance that is more than twice as fast as the PCL octree and achieves a 20% reduction in runtime compared to the *ikd-Tree*. Regarding the execution time for radius neighbor search, the *i-Octree* demonstrates a performance that is more than twice as fast as both the *ikd-Tree* and the PCL octree. Furthermore, the construction time of the *i-Octree* is less than 36% compared to the *ikd-Tree*. Additionally, the peak memory usage of the *i-Octree* is the lowest, being less than 35% of that observed for the *ikd-Tree*.

2) *Box-wise Delete*: The *ikd-Tree* is chosen to be compared with the *i-Octree* since both support the box-wise delete. This experiment investigates the time performance of deleting all points in an axis-aligned box and the time performance of KNN search and radius neighbor search after deletion. In the experiment, we sample 400,000 points randomly in a $10m \times 10m \times 10m$ space (i.e., the workspace) to initialize both trees. Then 100 test operations are conducted on the trees. In each test operation, we randomly sample 200 points in the workspace for KNN search and another 200 points for radius neighbor search. For every 20 test operations, one axis-aligned box is sampled in the workspace with side length of $1.0m$ and points contained in the box are deleted from the trees. We record the tree size after each deletion, the time of box-wise delete, the total time of KNN search for 200 points, and the total time of radius neighbor search for another 200 points at each step.

The time of box-wise delete and the tree size are shown in Table II. The run time performance of *i-Octree* and *ikd-*

TABLE II
THE TIME CONSUMPTION OF BOX-WISE DELETE AND TREE SIZE

Counter	Box-wise Delete[ms]		Tree Size		
	<i>i-Octree</i>	<i>ikd-Tree</i>	GT ¹	<i>i-Octree</i>	<i>ikd-Tree</i>
0	–	–	400,000	400,000	400,000
20	0.017	0.222	399,614	399,614	399,623
40	0.013	0.343	399,205	399,205	399,236
60	0.017	0.250	398,783	398,783	398,826
80	0.016	0.320	398,375	398,375	398,437
100	0.006	0.126	398,185	398,185	398,247

¹ GT denotes the exact size of the point cloud at each test.

Tree on KNN search and radius neighbor search is shown in Fig. 6. In this experiment, it appears that the box-wise delete operation on the *ikd-Tree* does not remove all points within the given box as shown in Table II. The results show that our *i-Octree* runs over 10 times faster than the *ikd-Tree* on

the box-wise delete and approximately 3 times faster on the radius neighbor search. Besides, the *i-Octree* reduces the run time of KNN search by 58% over the *ikd-Tree*.

B. Real-world Data Experiments

We test our developed *i-Octree* in actual robotic applications (e.g., LiDAR-inertial SLAM [8], [9], [21], [22] and pure LiDAR SLAM [5], [6], [23]). In these experiments, we directly replace the static k-d tree in the LiDAR-based SLAM by our proposed *i-Octree* without any refinement and evaluate the time performance and accuracy. To ensure a fair and complete comparison, we test on three datasets: M2DGR [24], Newer College Dataset [25] and NCLT [26] with different sensor setups. The M2DGR is a low-beam LiDAR data, the Newer College Dataset is a high-beam LiDAR data, and the NCLT is a long-term dataset. They all have ground truth trajectories such that we can evaluate the accuracy of the LiDAR trajectories using absolute trajectory error (ATE) [27].

For the sake of simplicity, we will focus on three representative LiDAR-based SLAM algorithms: FAST-LIO2 [9] and LIO-SAM [21] that combine LiDAR and inertial measurements (LiDAR-inertial SLAM), and FLOAM [23] that relies solely on LiDAR data (pure LiDAR SLAM). LIO-SAM and FLOAM are both the state-of-the-art algorithms and use the static k-d tree for the nearest points search, which is crucial to match a point in a new LiDAR scan to its correspondences in the map (or the previous scan). For the LIO-SAM, we choose the LIO.SAM.6AXIS¹, which is modified from LIO-SAM to support a wider range of sensors (e.g., 6-axis IMU and low-cost GNSS) without accuracy loss. Besides, the loop closure module of LIO.SAM.6AXIS is deactivated. The parameter setups for every sequence in a dataset are the same, as shown in Table III. Both algorithms build two independent maps: surface map and edge map, where two k-d trees or *i-Octree* will be built. The SurfRes and EdgeRes determine the minimal extent of the two *i-Octree* structures, respectively. Moreover, the performance of *i-Octree* and *ikd-Tree*, once integrated into FAST-LIO2, is evaluated using the Newer College Dataset.

TABLE III
THE PARAMETERS SETUP FOR EACH DATASET

Dataset	FLOAM			LIO.SAM.6AXIS		
	SurfRes	EdgeRes	MinRange	SurfRes	EdgeRes	MinRange
M2DGR	0.1m	0.05m	0.2m	0.4m	0.2m	1.0m
Newer College	0.2m	0.1m	0.2m	0.4m	0.2m	1.0m
NCLT	0.4m	0.2m	0.2m	0.4m	0.2m	1.0m

* SurfRes and EdgeRes denote the map resolution of surface map and edge map, respectively; MinRange means that any points within the LiDAR range defined by MinRange will not be considered.

¹https://github.com/JokerJohn/LIO_SAM_6AXIS

TABLE IV

THE COMPARISON OF MAXIMAL AND AVERAGE TIME CONSUMPTION PER SCAN ON INCREMENTAL UPDATE

Sequence Name	FLOAM[ms]				LIO.SAM.6AXIS[ms]			
	Incre_max		Incre_mean		Incre_max		Incre_mean	
	ioct	orig	ioct	orig	ioct	orig	ioct	orig
<i>gate_01</i>	7.20	969.98	5.57	354.73	5.16	33.99	2.82	17.56
<i>walk_01</i>	6.92	271.49	5.17	141.08	4.66	15.39	2.53	11.18
<i>street_02</i>	7.07	793.63	4.20	286.17	6.95	19.88	2.94	10.84
<i>street_03</i>	9.22	1983.50	6.52	893.20	7.27	30.05	3.58	19.57
<i>room_01</i>	2.22	17.41	0.64	5.67	0.42	1.70	0.29	0.48
<i>room_02</i>	3.00	24.00	0.65	6.30	0.43	0.73	0.28	0.46
<i>room_dark_01</i>	2.48	13.57	0.71	6.13	0.52	0.74	0.29	0.44

* Note: Incre_max and Incre_mean denote the maximal and average incremental update time of each scan respectively; ioct means *i-Octree* and orig means the static k-d tree.

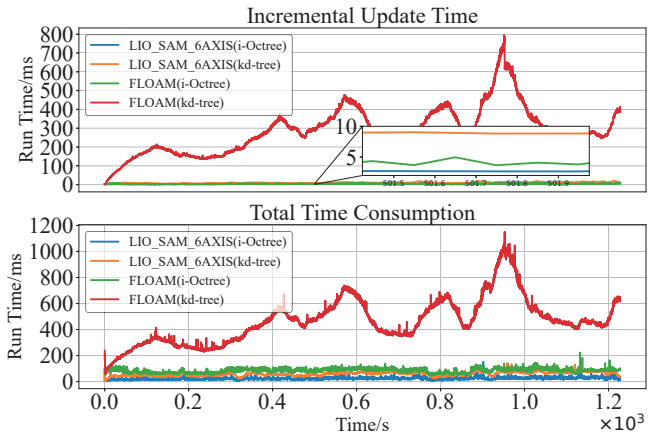


Fig. 7. The Comparison of time consumption on the *street_02* sequence.

1) *Low-beam LiDAR Data*: The first dataset is M2DGR, which is a large-scale dataset collected by an unmanned ground vehicle (UGV) with a full sensor-suite including a Velodyne VLP-32C LiDAR sampled at 10 Hz, a 9-axis Handsfree A9 inertial measurement unit (IMU) sampled at 150 Hz, and other sensors. The dataset comprises 36 sequences captured in diverse scenarios including both indoor and outdoor environments on the university campus. The ground truth trajectories of the outdoor sequences are obtained by the Xsens MTi 680G GNSS-RTK suite while for the indoor environment, a motion-capture system named Vicon Vero 2.2 whose localization accuracy is 1mm is used to collect the ground truth. We choose several sequences from this dataset and the detailed information is summarized in Table V.

The maximal and average time consumption of incremental update for each scan are listed in Table IV. For the FLOAM running in the outdoor environment, the size of map grows gradually making it computational costly when using the static k-d tree. The rebuilding process may cost a lot when the size of map is large, which deteriorates the real-time performance sharply. The time for rebuilding the entire k-d tree fills a large percentage of the total time per scan, as shown in the Fig. 7. After replacing the k-d tree by *i-Octree*, the incremental time reduces to less than 10ms and the real-time performance is approximately guaranteed. The FLOAM with *i-Octree* runs over 5 times faster than

TABLE V

THE COMPARISON OF AVERAGE TIME CONSUMPTION PER SCAN AND ABSOLUTE TRANSLATIONAL ERRORS (RMSE, METERS)

Dataset Name	Sequence Name	Attributes		Average Time Consumption per Scan[ms]				RMSE[m]			
				FLOAM		LIO_SAM_6AXIS		FLOAM		LIO_SAM_6AXIS	
		Duration	Length	<i>i-Octree</i>	static k-d tree	<i>i-Octree</i>	static k-d tree	<i>i-Octree</i>	static k-d tree	<i>i-Octree</i>	static k-d tree
M2DGR	<i>gate_01</i>	172s	-	101.70	570.62	40.69	100.50	8.667	9.206	0.167	0.340
	<i>walk_01</i>	291s	-	91.21	273.72	18.98	52.42	0.108	0.087	0.078	0.079
	<i>street_02</i>	1227s	-	82.85	462.21	24.26	58.0	133.602	80.935	2.296	2.942
	<i>street_03</i>	354s	-	117.55	1318.47	39.11	114.55	15.781	16.923	0.137	0.138
	<i>room_01</i>	72s	-	22.42	25.91	3.11	12.32	0.134	0.135	0.137	0.209
	<i>room_02</i>	75s	-	22.52	26.60	3.02	6.32	0.128	0.129	0.123	0.121
	<i>room_dark.01</i>	111s	-	22.85	26.16	2.70	5.73	0.146	0.148	0.151	0.155
Newer College	<i>Quad-Easy</i>	198s	247m	73.87	225.31	38.03	100.73	0.795	0.399	0.083	0.079
	<i>Quad-Medium</i>	190s	260m	58.19	430.78	43.02	96.86	15.626	18.358	0.100	0.095
	<i>Quad-Hard</i>	187s	234m	44.86	300.38	32.85	70.35	17.133	7.326	0.130	0.152
	<i>Cloister</i>	278s	429m	41.06	156.80	20.07	53.77	15.100	7.685	0.115	0.169
	<i>Maths-Easy</i>	216s	264m	62.53	180.89	31.34	74.46	1.026	0.260	0.088	0.125
	<i>Underground-Easy</i>	141s	162m	28.16	41.27	9.37	30.85	0.143	0.078	0.077	0.073
NCLT	<i>2012-01-08</i>	5633s	6.4km	21.88	70.62	24.65	53.54	187.588	62.188	1.890	9.291

TABLE VI

PERFORMANCE COMPARISON OF *i-Octree* AND *ikd-Tree*

Sequence Name	AveTim[ms]		RMSE[m]		PeaMem[MB]	
	<i>i-Octree</i>	<i>ikd-Tree</i>	<i>i-Octree</i>	<i>ikd-Tree</i>	<i>i-Octree</i>	<i>ikd-Tree</i>
<i>Quad-Easy</i>	14.58	19.70	0.069	0.071	242.06	252.26
<i>Quad-Medium</i>	14.24	16.21	0.059	0.059	239.04	259.13
<i>Quad-Hard</i>	10.35	12.08	0.051	0.047	240.58	257.86
<i>Cloister</i>	7.30	9.08	0.053	0.063	255.93	261.62
<i>Maths-Easy</i>	12.45	16.41	0.093	0.093	252.77	278.04
<i>Underground-Easy</i>	4.82	5.96	0.029	0.036	235.12	239.61

* Note: AveTim denotes the average time required for updating the tree and executing KNN search for each scan ; PeaMem means the peak memory usage of ROS node.

the original, with sometimes a slight accuracy loss as shown in Table V. The LIO_SAM_6AXIS with *i-Octree* runs over 2 times faster than the original on almost all sequences and the accuracy is usually improved. Besides, the *i-Octree* reduces the peak time usage which degrades the real-time ability.

2) *High-beam LiDAR Data*: The second dataset is the Newer College Dataset, which is a multi-camera LiDAR inertial dataset of 4.5 km walking distance. The dataset is collected by a handheld device equipped with an Ouster OS0-128 LiDAR sampled at 10 Hz with embedded IMU sampled at 100 Hz. Precise centimetre accurate ground truth is provided and calculated by registering each undistorted LiDAR frame to a prior map scanned by 3D imaging laser scanner, Leica BLK360. Table V shows the 6 sequences selected from the dataset. These sequences contain small and narrow passages, large scale open spaces, as well as vegetated areas. Besides, challenging situations such as aggressive motion are presented.

In this test, we record average total time consumption of each scan and calculate the RMSE of the absolute translational errors, as shown in Table V. We find that the *i-Octree* shows promising compatibility with LIO_SAM_6AXIS since it not only reduces the time consumption a lot but also improves the accuracy. For the LOAM, the *i-Octree* significantly accelerates the processing speed per scan, but sometimes at a cost of degrading accuracy, as it directly replaces the static k-d tree without any refinement. Furthermore, the performance comparison between *i-Octree* and *ikd-Tree*, as integrated into FAST-LIO2 without parallelization, is evaluated using this dataset, as presented in Table VI. In this experiment, we set $e_{min} = 0.5m$ for both *i-Octree* and *ikd-Tree*, and $b = 1$ specifically for *i-Octree*. The other *ikd-Tree* parameters are kept at their default settings. The *i-Octree* method exhibits superior memory efficiency and speed compared to the *ikd-Tree*, achieving an average reduction in runtime of 19% without sacrificing accuracy, consistent with findings from experiments on randomized

data.

3) *Long-term LiDAR Data*: Finally, we test on a long sequence called *2012-01-08* from a large scale dataset named NCLT captured by a Velodyne HDL-32E LiDAR sampled at 10 Hz, a Microstrain 3DM-GX3-45 IMU sampled at 100 Hz, and other sensors. The length of the sequence is 6.4km and the duration is 5633s. The average run time of each scan and the absolute translational errors are added to Table V. The FLOAM drifts severely on the large scale sequence while LIO_SAM_6AXIS (*i-Octree*) shows surprising performance, except for a slight drift along the z-axis, even without enabling loop closure. LIO_SAM_6AXIS (static k-d tree) can only use the very recent key frames to estimate the poses, whereas LIO_SAM_6AXIS (*i-Octree*) can take advantage of point information from the start, which contributes to its better performance.

IV. CONCLUSION

In this paper, we propose a novel dynamic octree data structure, *i-Octree*, which supports incremental point insertion with on-tree down-sampling, box-wise delete, and fast NNS. Besides, a large amount of experiments on both randomized data and open datasets show that our *i-Octree* can achieve the best overall performance among the state-of-the-art tree data structures.

REFERENCES

- [1] X. Xiong, D. Munoz, J. A. Bagnell, and M. Hebert, "3-d scene analysis via sequenced predictions over points and regions," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 2609–2616.
- [2] J. Behley, V. Steinhage, and A. B. Cremers, "Performance of histogram descriptors for the classification of 3d laser range data in urban environments," in *2012 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2012, pp. 4391–4398.
- [3] N. J. Mitra and A. Nguyen, "Estimating surface normals in noisy point cloud data," in *Proceedings of the Nineteenth Annual Symposium on Computational Geometry*, ser. SCG '03. New York, NY, USA: Association for Computing Machinery, 2003, p. 322–328. [Online]. Available: <https://doi.org/10.1145/777792.777840>
- [4] K. Klasing, D. Wollherr, and M. Buss, "A clustering method for efficient segmentation of 3d laser data," in *2008 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2008, pp. 4043–4048.
- [5] J. Zhang and S. Singh, "Loam: Lidar odometry and mapping in real-time," in *Robotics: Science and Systems*, vol. 2, no. 9, 2014, pp. 1–9.
- [6] T. Shan and B. Englot, "Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4758–4765.
- [7] J. Lin and F. Zhang, "Loam livox: A fast, robust, high-precision lidar odometry and mapping package for lidars of small fov," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 3126–3131.
- [8] W. Xu and F. Zhang, "Fast-lidar: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3317–3324, 2021.
- [9] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang, "Fast-lidar2: Fast direct lidar-inertial odometry," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2053–2073, 2022.
- [10] A. Guttman, "R-trees: A dynamic index structure for spatial searching," in *Proceedings of the 1984 ACM SIGMOD international conference on Management of data*, 1984, pp. 47–57.
- [11] N. Beckmann, H.-P. Kriegel, R. Schneider, and B. Seeger, "The r*-tree: An efficient and robust access method for points and rectangles," in *Proceedings of the 1990 ACM SIGMOD international conference on Management of data*, 1990, pp. 322–331.
- [12] J. L. Bentley, "Multidimensional binary search trees used for associative searching," *Communications of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.
- [13] J. L. Vermeulen, A. Hillebrand, and R. Geraerts, "A comparative study of k-nearest neighbour techniques in crowd simulation," *Computer Animation and Virtual Worlds*, vol. 28, no. 3-4, p. e1775, 2017.
- [14] J. Elseberg, S. Magnenat, R. Siegwart, and A. Nüchter, "Comparison of nearest-neighbor-search strategies and implementations for efficient shape registration," *Journal of Software Engineering for Robotics*, vol. 3, no. 1, pp. 2–12, 2012.
- [15] J. Elseberg, D. Borrmann, and A. Nüchter, "One billion points in the cloud—an octree for efficient processing of 3d laser scans," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 76, pp. 76–88, 2013.
- [16] J. Behley, V. Steinhage, and A. B. Cremers, "Efficient radius neighbor search in three-dimensional point clouds," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3625–3630.
- [17] R. B. Rusu and S. Cousins, "3d is here: Point cloud library (pcl)," in *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2011, pp. 1–4.
- [18] G. M. Morton, "A computer oriented geodetic data base and a new technique in file sequencing," International Business Machines Company, New York, Tech. Rep., 1996.
- [19] J. H. Friedman, J. L. Bentley, and R. A. Finkel, "An algorithm for finding best matches in logarithmic expected time," *ACM Trans. Math. Softw.*, vol. 3, no. 3, p. 209–226, sep 1977. [Online]. Available: <https://doi.org/10.1145/355744.355745>
- [20] M. Muja and D. G. Lowe, "Fast approximate nearest neighbors with automatic algorithm configuration," in *VISAPP 2009: Proceedings of the Fourth International Conference on Computer Vision Theory and Applications, VOL 1*, A. Ranchordas and H. Araujo, Eds. Inst Syst & Technologies Informat, Control & Commun; ACM SIGGRAPH, 2009, pp. 331–340.
- [21] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping," in *2020 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2020, pp. 5135–5142.
- [22] K. Li, M. Li, and U. D. Hanebeck, "Towards high-performance solid-state-lidar-inertial odometry and mapping," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5167–5174, 2021.
- [23] H. Wang, C. Wang, C.-L. Chen, and L. Xie, "F-loam: Fast lidar odometry and mapping," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 4390–4396.
- [24] J. Yin, A. Li, T. Li, W. Yu, and D. Zou, "M2dgr: A multi-sensor and multi-scenario slam dataset for ground robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 2266–2273, 2021.
- [25] L. Zhang, M. Camurri, and M. F. Fallon, "Multi-camera lidar inertial extension to the newer college dataset," *CoRR*, vol. abs/2112.08854, 2021. [Online]. Available: <https://arxiv.org/abs/2112.08854>
- [26] N. Carlevaris-Bianco, A. K. Ushani, and R. M. Eustice, "University of Michigan North Campus long-term vision and lidar dataset," *International Journal of Robotics Research*, vol. 35, no. 9, pp. 1023–1035, 2015.
- [27] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 573–580.