

---

# FlexFlood: Efficiently Updatable Learned Multi-dimensional Index

---

**Fuma Hidaka**  
The University of Tokyo  
hidaka@hal.t.u-tokyo.ac.jp

**Yusuke Matsui**  
The University of Tokyo  
matsui@hal.t.u-tokyo.ac.jp

## Abstract

A learned multi-dimensional index is a data structure that efficiently answers multi-dimensional orthogonal queries by understanding the data distribution using machine learning models. One of the existing problems is that the search performance significantly decreases when the distribution of data stored in the data structure becomes skewed due to update operations. To overcome this problem, we propose FlexFlood, a flexible variant of Flood. FlexFlood partially reconstructs the internal structure when the data distribution becomes skewed. Moreover, FlexFlood is the first learned multi-dimensional index that guarantees the time complexity of the update operation. Through experiments using both artificial and real-world data, we demonstrate that the search performance when the data distribution becomes skewed is up to 10 times faster than existing methods. We also found that partial reconstruction takes only about twice as much time as naive data updating.

## 1 Introduction

Filtering, scanning, and updating of data are fundamental operations for databases, and various data structures have been studied to perform these operations efficiently. In the real world, we often need to handle multi-dimensional data, and Kd-tree and its variants [5, 20, 31] are typical data structures for handling them. These data structures are widely used in real-world applications [3, 23, 22].

Recently, there has been active research to improve data structures by learning the distribution of data and queries with machine learning models. Such data structures are called learned index [24]. One of the significant challenges in learned multi-dimensional indexes [32, 12, 25, 19, 8] is that many do not support data update operations. Even if they do, none describe the time complexity for updating.

Therefore, we proposed a flexible variant of Flood (FlexFlood) that supports efficient data updating by adaptively modifying the internal structure of the existing learned multi-dimensional index, Flood [32]. We proved that the amortized time complexity of updating is  $O(D \log N)$  under two assumptions that the data increases at an approximately constant pace and that the training results of the ML model satisfy certain conditions. Here,  $D$  is the dimensionality of the data, and  $N$  is the total number of data. Furthermore, experiments using multiple artificial and real-world datasets confirm that the “certain conditions” can be sufficient. The original Flood’s search speed slows down as we update data, but FlexFlood remains fast. As a result, FlexFlood is up to 10 times faster than Flood after many updates. The source code for FlexFlood is available at <https://github.com/mti-lab/FlexFlood>.

## 2 Related Work

Many classical data structures have been proposed for handling multi-dimensional data. For example, tree-based structures include Kd-tree [5, 16], R-tree [20, 4], and Oct-tree [28]. Grid File [33] is also proposed as a space-partitioning structure. Additionally, there is the Z-order curve algorithm [31], which reduces multi-dimensional data into one dimension using a special sorting technique.

Learned index [24, 13, 2] is a data structure that incorporates machine learning models into classical data structures such as B-tree [1], Hash Map [10], and Bloom Filter [6]. Learned indexes improve performance by taking advantage of the distribution of data and queries. Recently, learned indexes have been actively researched. For example, several learned Bloom Filters have been proposed [29, 26, 30, 34, 35]. They achieved a better memory/accuracy trade-off than the original Bloom Filter.

Learned indexes are vulnerable to data update operations in general. This is because if the data distribution is distorted by updating operations, the accuracy of the machine learning model will decrease, and search performance will decrease, too. To address this problem, learned indexes that support efficient updating operations have been proposed [11, 15, 14, 37].

The above learned indexes can handle only one-dimensional data. On the other hand, learned indexes for multi-dimensional data have also been proposed. For example, Flood [32], Tsunami [12], Lisa [25], RLR-tree [19], Waffle [8], and so on [2]. Flood and Tsunami do not support data updating operations. Lisa, RLR-tree and Waffle support them, but there are no discussion of time complexity.

### 3 Preliminary

First, we define the problem setting.  $N$   $D$ -dimensional vectors  $\mathbf{v}_1, \dots, \mathbf{v}_N$  are given, where  $\mathbf{v}_n \in \mathbb{R}^D$ . We denote the  $d$ -th dimension of  $\mathbf{v}_n$  as  $v_n[d] \in \mathbb{R}$ . Initially, we construct an index based on the  $N$  vectors. Then, we process  $Q$  queries sequentially. Queries are provided in the following formats:

- Search Query:  $\mathbf{l} \in \mathbb{R}^D$  and  $\mathbf{r} \in \mathbb{R}^D$  are provided.  $\mathbf{l}$  and  $\mathbf{r}$  represent the endpoints of the diagonal of the search range hyper-rectangle. We enumerate all vectors  $\mathbf{v}_n$  contained within the data structure, such that for all  $d \in \{1, 2, \dots, D\}$ ,  $l[d] \leq v_n[d] \leq r[d]$  holds.
- Insert Query: A  $D$ -dimensional vector  $\mathbf{v} \in \mathbb{R}^D$  is provided. We add it if this vector is not in the data structure. We do nothing if the same vector already exists within the data structure.
- Erase Query: A  $D$ -dimensional vector  $\mathbf{v} \in \mathbb{R}^D$  is provided. We remove it if this vector exists within the data structure. We do nothing if the vector is not in the data structure.

Next, we explain the Flood algorithm [32], which is the basis of the proposed method. Flood divides the  $D$ -dimensional space into approximately equal parts by  $(D - 1)$ -dimensional grid cells. Given  $D$ -dimensional input vectors, Flood assigns them to the corresponding cell. Within each cell, Flood keeps vectors sorted using the value of the  $D$ -th dimension that is not used for the grid division. The parameters of Flood are the sort dimension and the number of cell partitions. Flood optimizes these parameters using a gradient descent method with a random forest regression model [7, 36].

Flood uses sorted arrays to hold data in cells, but updating data on a sorted array is very expensive. We can make Flood updatable by replacing sorted arrays with B-trees. (Appendix A discussed this overhead.) However, if the data distribution changes, the grid partitioning at the initialization becomes meaningless and the updatable Flood's search speed slows down greatly. Therefore, we propose an efficient data updating algorithm to solve this problem, using the updatable Flood as a baseline.

### 4 Proposed Method: FlexFlood

Based on the updatable Flood, we propose FlexFlood, a data updating algorithm that ensures fast search even if data distribution changes. Figure 1 is an overview of FlexFlood. Remember Flood constructs cells so that for each axis  $d$ , the total number of data in the cell with the same  $d$ -dimensional value is  $\frac{N}{x_d}$  (an equal number of data for each cell). Here,  $x_d$  is the number of cell partitions on axis  $d$ . When the number of vectors in each cell deviates far from  $\frac{N}{x_d}$  due to data updating, the search of the updatable Flood slows down greatly. Therefore, FlexFlood re-partitions around the cells as follows. (Appendix B discussed thresholds for re-partitioning in more detail.)

- If the cell contains more than  $\frac{2N}{x_d}$  vectors, we split the cell into two.
- If the cell contains less than  $\frac{N}{3x_d}$  vectors and the neighboring cell contains less than  $\frac{7N}{6x_d}$  vectors, we merge the two cells.
- If the cell contains less than  $\frac{N}{3x_d}$  vectors and the neighboring cell contains more than  $\frac{7N}{6x_d}$  vectors, we equalize the number of vectors in the two cells.




Figure 1: Overview of our method ( $D = 2$ ): Even if  $N = 27$  data are equally divided into  $x_1 = 3$  cells at the initialization, the distribution can become skewed due to data updating, and the updatable Flood slows down. We aim to ensure the search performance by partially re-partitioning the cells.

Re-partitioning in this way keeps the number of vectors in the cell constant and keeps the search speed fast. However, since the above re-partition requires a large amount of movement of data within the cells, it is not obvious whether the update operation can be performed efficiently.

We therefore analyzed the time complexity of this algorithm. Based on the insight that re-partitions with very high computational costs occur only with sufficiently low frequency, we found that the computational complexity of the updating operation is not large when considered in terms of amortization. Under the two assumptions that (1) the data increases at a pace that can be regarded as constant and (2)  $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$ , we proved that the updating operation of FlexFlood is amortized at  $O(D \log N)$ . (See Appendix C for detailed analysis.)

The update operation of the updatable Flood is  $O(\log N)$ , which is  $D$  times faster than ours. Therefore, we can interpret FlexFlood as an algorithm that ensures high search performance even when the data distribution is skewed, instead of sacrificing up to  $D$  times the complexity of the update operations.

## 5 Experiment

We evaluate FlexFlood. The runtime environment is Intel Core i7-11800H, 8 cores, 2.3 GHz, 32 GB memory. The four data structures used for comparison and their implementations are shown below.

- Self-Balancing Kd-tree (SB-Kdtree) [5, 16]: We implemented it in C++.
- R-tree [20]: We used C++ boost::geometry::index::rtree [17].
- Updatable Flood: We implemented Flood in C++, then we replaced Flood’s sorted arrays with B-trees published by Google [18].
- FlexFlood: We added cell re-partitioning algorithm to the updatable Flood.

As a dataset, we used (1) normal distribution, (2) Stock Price dataset [21], and (3) Open Street Map dataset [9]. See Appendix D for details on data and query generation methods.

## 6 Result

Figures 2 illustrate the cumulative query processing time for each data structure for each dataset. The number of cell partitions and whether the condition  $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$  was satisfied or not in each dataset are shown in Table 1.

First, looking at Table 1, we can see that the assumption (2) was always satisfied within the range of the experiments conducted here. Therefore, we can assume that the amortized time complexity of the data updating operation of FlexFlood is  $O(D \log N)$  for practical purposes.

We then turn focus on Figures 2a, 2b, 2c, the results of the update queries. Comparing FlexFlood with the classical data structures such as SB-Kdtree and R-tree, FlexFlood processes update queries about 1.1 to 2.9 times faster. FlexFlood is better than the SB-Kdtree and R-tree for all datasets. Comparing FlexFlood with the updatable Flood, FlexFlood requires at most 2.0 times longer runtime. Remember that in theory, the cell re-partitioning algorithm takes about  $D$  times longer runtime than the updatable Flood’s data updating operation in the worst case. In light of this, we believe that

Table 1: Number of cell partitions and whether the conditional expression was satisfied or not

Dataset	$\{x_d\}_{d=1}^D$	$\prod_{d=1}^D x_d \sum_{d=1}^D x_d$	$DN \log N$	Assumption (2)
3D Normal Distribution	{21, 17, 1}	$1.4 \cdot 10^4$	$5.0 \cdot 10^6$	✓
4D Stock Price	{19, 19, 1, 39}	$1.1 \cdot 10^6$	$1.4 \cdot 10^7$	✓
5D Open Street Map	{17, 17, 13, 23, 1}	$6.1 \cdot 10^6$	$1.0 \cdot 10^8$	✓

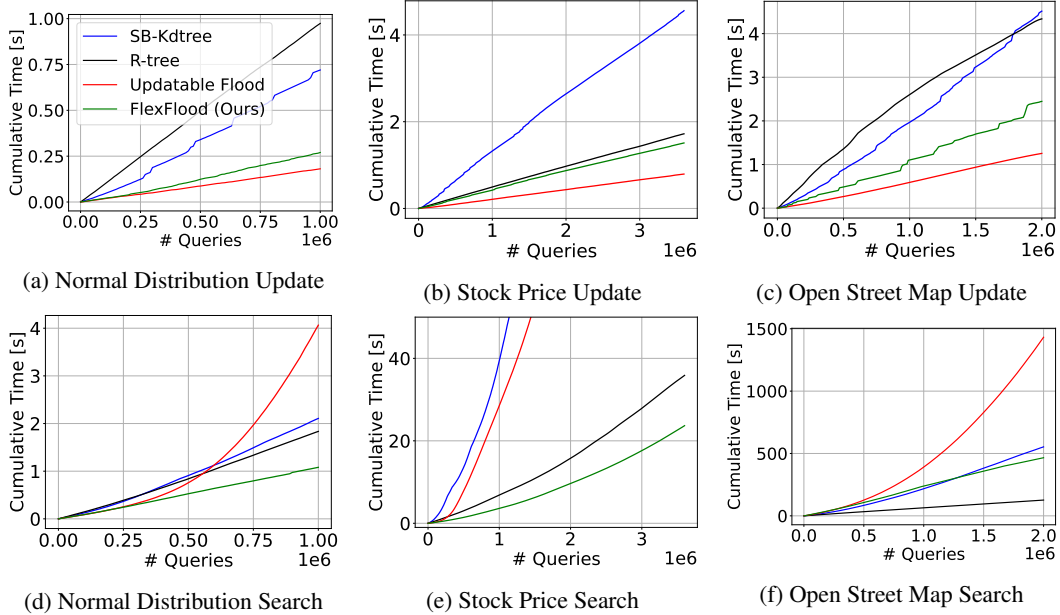


Figure 2: Experimental results: The upper panel shows the update queries, and the lower panel shows the results for the search queries. (Lower is better.)

FlexFlood is not only theoretically guaranteed to be computationally feasible but also fast enough for practical use (e.g., FlexFlood could be  $D = 5$  times slower than the updatable Flood for 5D Open Street Map dataset, but Figure 2c shows that FlexFlood is only about 2.0 times slower in practice).

Finally, we refer to Figures 2d, 2e, 2f, the results of the search queries. Comparing FlexFlood with the updatable Flood, FlexFlood processes search queries about 3.3 to 10 times faster, outperforming the updatable Flood on all datasets. Comparing FlexFlood with SB-Kdtree and R-tree, FlexFlood processes queries on the normal distribution dataset and the Stock Price dataset about 1.2 to 12 times faster. On the Open Street Map dataset, however, FlexFlood is slower than the R-tree although FlexFlood is faster than SB-Kdtree. Regarding this result, the benchmark paper [27] pointed out that the learned indexes perform poorly compared to classical data structures on the Open Street Map dataset because it lacks local structure, making them difficult to learn. Figure 2f also shows that the slope of the curve near the origin for the updatable Flood is steeper than that of SB-Kdtree and R-tree. In light of this fact, we interpret this result as consistent with the results of previous studies.

## 7 Conclusion

By adaptively reconstructing the internal structure of Flood, we proposed FlexFlood, which supports efficient data updating. Experimental results show that FlexFlood does not reduce the search speed and has advantages over classical data structures. Furthermore, we proved that the amortized time complexity of data updating is  $O(D \log N)$  under two experimentally valid assumptions.

On the other hand, FlexFlood loses the optimality guarantees regarding the sort dimension and the number of cell divisions after the data update. Therefore, it may be possible to ensure even faster search by periodically relearning the distribution. (Appendix E discusses this in more detail.)

## References

- [1] Database architects blog: The case for b-tree index structures. <http://databasearchitects.blogspot.com/2017/12/the-case-for-b-tree-index-structures.html>.
- [2] Abdullah Al-Mamun, Hao Wu, Qiyang He, Jianguo Wang, and Walid G. Aref. A survey of learned indexes for the multi-dimensional space, 2024.
- [3] Amazon AWS. Amazon redshift engineering’s advanced table design playbook: Compound and interleaved sort keys. 2016. <https://aws.amazon.com/jp/blogs/big-data/amazon-redshift-engineerings-advanced-table-design-playbook-compound-and-interleaved-sort-keys/>.
- [4] Norbert Beckmann, Hans-Peter Kriegel, Ralf Schneider, and Bernhard 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*, pages 322–331, 1990.
- [5] Jon Louis Bentley. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9):509–517, 1975.
- [6] Burton H Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [7] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.
- [8] Dalsu Choi, Hyunsik Yoon, Hyubjin Lee, and Yon Dohn Chung. Waffle: in-memory grid index for moving objects with reinforcement learning-based configuration tuning system. *Proceedings of the VLDB Endowment*, 15(11):2375–2388, 2022.
- [9] Contributors. Open street map. <https://download.geofabrik.de/>. Viewed 10 January 2024, Licence: ODbL 1.0.
- [10] Martin Dietzfelbinger, Anna Karlin, Kurt Mehlhorn, Friedhelm Meyer Auf Der Heide, Hans Rohnert, and Robert E Tarjan. Dynamic perfect hashing: Upper and lower bounds. *SIAM Journal on Computing*, 23(4):738–761, 1994.
- [11] Jialin Ding, Umar Farooq Minhas, Jia Yu, Chi Wang, Jaeyoung Do, Yinan Li, Hantian Zhang, Badrish Chandramouli, Johannes Gehrke, Donald Kossmann, et al. Alex: an updatable adaptive learned index. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 969–984, 2020.
- [12] Jialin Ding, Vikram Nathan, Mohammad Alizadeh, and Tim Kraska. Tsunami: A learned multi-dimensional index for correlated data and skewed workloads. *Proceedings of the VLDB Endowment*, 14(2):74–86, 2020.
- [13] Paolo Ferragina and Giorgio Vinciguerra. Learned data structures. In *Recent Trends in Learning From Data: Tutorials from the INNS Big Data and Deep Learning Conference (INNSBDDL2019)*, pages 5–41. Springer, 2020.
- [14] Paolo Ferragina and Giorgio Vinciguerra. The pgm-index: a fully-dynamic compressed learned index with provable worst-case bounds. *Proceedings of the VLDB Endowment*, 13(8):1162–1175, 2020.
- [15] Alex Galakatos, Michael Markovitch, Carsten Binnig, Rodrigo Fonseca, and Tim Kraska. Fiting-tree: A data-aware index structure. In *Proceedings of the 2019 ACM SIGMOD International Conference on Management of Data*, pages 1189–1206, 2019.
- [16] Igal Galperin and Ronald L Rivest. Scapegoat trees. In *Proceedings of the fourth annual ACM-SIAM Symposium on Discrete algorithms*, pages 165–174, 1993.
- [17] Barend Gehrels, Bruno Lalande, Mateusz Loskot, Adam Wulkiewicz, and Oracle and/or its affiliates. boost::geometry::index::rtree. [https://beta.boost.org/doc/libs/1\\_82\\_0/libs/geometry/doc/html/geometry/reference/spatial\\_indexes/boost\\_\\_geometry\\_\\_index\\_\\_rtree.html](https://beta.boost.org/doc/libs/1_82_0/libs/geometry/doc/html/geometry/reference/spatial_indexes/boost__geometry__index__rtree.html). Viewed 18 September 2024, License: Boost Software License 1.0.

- [18] Google. `cpp-btree`. <https://code.google.com/archive/p/cpp-btree/wikis/UsageInstructions.wiki>. Viewed 30 January 2024, License: Apache License 2.0.
- [19] Tu Gu, Kaiyu Feng, Gao Cong, Cheng Long, Zheng Wang, and Sheng Wang. The rlr-tree: A reinforcement learning based r-tree for spatial data. *Proceedings of the ACM on Management of Data*, 1(1):1–26, 2023.
- [20] Antonin Guttman. R-trees: A dynamic index structure for spatial searching. In *Proceedings of the 1984 ACM SIGMOD International Conference on Management of Data*, pages 47–57, 1984.
- [21] Evan Hallmark. Daily historical stock prices (1970 - 2018). <https://www.kaggle.com/datasets/ehallmar/daily-historical-stock-prices-1970-2018>. Viewed 10 January 2024, License: Unknown.
- [22] IBM. The spatial index. <https://www.ibm.com/docs/en/informix-servers/12.10?topic=data-spatial-index>.
- [23] Adrian Ionescu. Processing petabytes of data in seconds with databricks delta. <https://www.databricks.com/blog/2018/07/31/processing-petabytes-of-data-in-seconds-with-databricks-delta.html>.
- [24] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 ACM SIGMOD International Conference on Management of Data*, pages 489–504, 2018.
- [25] Pengfei Li, Hua Lu, Qian Zheng, Long Yang, and Gang Pan. Lisa: A learned index structure for spatial data. In *Proceedings of the 2020 ACM SIGMOD international Conference on Management of Data*, pages 2119–2133, 2020.
- [26] Qiyu Liu, Libin Zheng, Yanyan Shen, and Lei Chen. Stable learned bloom filters for data streams. *Proceedings of the VLDB Endowment*, 13(12):2355–2367, 2020.
- [27] Ryan Marcus, Andreas Kipf, Alexander van Renen, Mihail Stoian, Sanchit Misra, Alfons Kemper, Thomas Neumann, and Tim Kraska. Benchmarking learned indexes. *Proceedings of the VLDB Endowment*, 14(1):1–13, 2020.
- [28] Donald Meagher. Octree encoding: A new technique for the representation, manipulation and display of arbitrary 3-d objects by computer. 1980.
- [29] Michael Mitzenmacher. A model for learned bloom filters and optimizing by sandwiching. *Proceedings of the NeurIPS*, pages 462–471, 2018.
- [30] Michael Mitzenmacher. Partitioned learned bloom filters. *Proceedings of the ICLR*, 2021.
- [31] Guy M Morton. A computer oriented geodetic data base and a new technique in file sequencing. 1966.
- [32] Vikram Nathan, Jialin Ding, Mohammad Alizadeh, and Tim Kraska. Learning multi-dimensional indexes. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*, pages 985–1000, 2020.
- [33] Jürg Nievergelt, Hans Hinterberger, and Kenneth C Sevcik. The grid file: An adaptable, symmetric multikey file structure. *ACM Transactions on Database Systems (TODS)*, 9(1):38–71, 1984.
- [34] Atsuki Sato and Yusuke Matsui. Fast partitioned learned bloom filters. *Proceedings of the NeurIPS*, 2023.
- [35] Atsuki Sato and Yusuke Matsui. Fast construction of partitioned learned bloom filter with theoretical guarantees. *arXiv preprint arXiv:2410.13278*, 2024.
- [36] scikit-learn developers. `Randomforestregressor`. <https://scikit-learn.org/dev/modules/generated/sklearn.ensemble.RandomForestRegressor.html>. Viewed 28 October 2024, License: BSD License.
- [37] Jiacheng Wu, Yong Zhang, Shimin Chen, Jin Wang, Yu Chen, and Chunxiao Xing. Updatable learned index with precise positions. *Proceedings of the VLDB Endowment*, 14(8):1276–1288, 2021.




Figure 3: Processing time per search query.

## A Sorted Array vs B-tree

The updatable Flood replaces the Flood’s sorted arrays with B-trees, which slows down the search. Therefore, we conducted comparative experiments on workloads where no update queries existed. The data structures used for the comparison are (1) Self-Balancing Kd-tree (SB-Kdtree), (2) R-tree, (3) Flood, and (4) updatable Flood. Note that FlexFlood is exactly the same as the updatable Flood if there are no update queries. The dataset used for the experiments are (1) normal distribution, (2) Stock Price dataset, and (3) Open Street Map dataset. We initialized each data structure, and performed  $10^4$  search queries to measure the processing time per search query.

The experimental results are in Figure 3. The search speed of the updatable Flood is about 1.5 to 2.7 times slower than that of Flood. Therefore, we should use regular Flood for workloads where it is known in advance that there will be no update queries at all. However, for many datasets, updatable Flood achieves faster search than classical data structures. Therefore, we believe that updatable Flood is worth using for workloads where update queries are likely to be present.

## B Threshold for Re-partitioning

The thresholds for the re-partitioning algorithm introduced in Section 4 are hyperparameters. In this section, we discuss the intuitive reasons for determining the thresholds as in Section 4 and how the performance changes when the thresholds are varied.

### B.1 Threshold Selection Criteria

We explain why the threshold for “split” in Section 4 is set to  $\frac{2N}{x_d}$  and the threshold for “merge” and “equalize” to  $\frac{N}{3x_d}$ . The purpose of cell re-partitioning is to maintain the number of data in cells at a baseline value  $\frac{N}{x_d}$  when the number of data in a particular cell increases or decreases too much. Since “split” halves the number of data in a cell, it is efficient to perform a “split” when the number of data in a cell is  $\frac{2N}{x_d}$ . For the same reason, it seems intuitive that “merge” and “equalize” should be performed when the number of data in a cell reaches  $\frac{N}{2x_d}$ . However, the number of data in the neighboring cell is greater than  $\frac{N}{2x_d}$  and less than  $\frac{2N}{x_d}$ . Therefore, the number of data in the cell after “merge” and “equalize” is considered to be more than  $\frac{N}{x_d}$ . For these reasons, we selected  $\frac{N}{3x_d}$  as the threshold for “merge” and “equalize”, which is slightly smaller than  $\frac{N}{2x_d}$ . This is expected to bring the number of data in the cell after “merge” and “equalize” closer to  $\frac{N}{x_d}$ .

Also, we explain why the threshold for switching between “merge” and “equalize” is set to  $\frac{7N}{6x_d}$ . This is because  $\frac{7N}{6x_d} = \frac{\frac{N}{3x_d} + \frac{2N}{x_d}}{2}$ . When “merge” or “equalize” is performed, the number of data in the neighboring cell is between  $\frac{N}{3x_d}$  and  $\frac{2N}{x_d}$ . Therefore, we switch between “merge” and “equalize” at the intermediate value.

## B.2 Performance Variation with Thresholds

In order to confirm the appropriateness of the thresholds in Section 4, we describe the results of the experiments in which FlexFlood’s performance changed when the thresholds were changed. We conducted the experiments by varying the coefficients of  $\frac{N}{x_d}$  on the threshold of “split” in the range  $\{1.5, 1.6, 1.7, 1.8, 1.9, 2.0, 2.1, 2.2, 2.3, 2.4, 2.5\}$  and on the threshold of “merge” and “equalize” in the range  $\{0.1, 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6\}$ . The threshold for switching between “merge” and “equalize” was set to a value halfway between the two thresholds for re-partitioning. The experimental settings are the same as in Section 5.

Figure 4 shows the results of the experiments in which search time and update time were measured separately. Red squares in the heatmap indicate a long search/update time, while blue squares indicate a short time. That is, blue squares mean better performance. In more detail, we calculated the percentage change compared to the search/update time when the threshold was set as section 4. Figure 4 shows that there are few thresholds that are blue squares in the heatmaps for both search and update. Therefore, we believe that the thresholds of Section 4 are one of the best practical values of hyperparameters.

## C Amortized Time Complexity Analysis

We prove that the amortized time complexity of FlexFlood’s update operation is  $O(D \log N)$  under two assumptions. Let  $X$  denote the total number of cells, that is,  $\prod_{d=1}^D x_d = X$ .

### C.1 Worst Time Complexity of Vector Insertion or Erasion

First, we show that the worst time complexity of vector insertion or erasion is  $O(\log N)$ . Insertion or erasion of a vector  $\mathbf{v} \in \mathbb{R}^D$  consists of two steps: (1) identifying the cell that should contain  $\mathbf{v}$  and (2) inserting or erasing  $\mathbf{v}$  for a B-tree within the cell. For (1), for each dimension  $d \in \{1, 2, \dots, D\}$ , we can perform a binary search on the set of cell boundary coordinates to determine where  $v[d]$  should be placed. This takes  $O(\sum_{d=1}^D \log x_d) = O(\log \prod_{d=1}^D x_d) = O(\log X)$ . We can do (2) with worst  $O(\log N)$  because we only insert or erase  $\mathbf{v}$  into the B-tree that holds the data in the identified cell. The total computational cost of inserting or erasing vectors is  $O(\log X + \log N)$ . Here,  $X < N$  is considered to be valid unless the learning of the distribution is very unsuccessful. This is because  $X \geq N$  means that the number of cells is larger than the total number of data, which is obviously wasteful. Table 1 also shows that this is valid. Therefore, the worst time complexity of vector insertion or erasion is  $O(\log N)$ .

### C.2 Worst Time Complexity of Cell Re-partition

Next, we discuss the worst time complexity of cell re-partition. Since cell re-partition occurs independently for each axis  $d \in \{1, 2, \dots, D\}$ , we consider each axis  $d$  independently and sum up later. We assume the following two conditions.

- The number of data  $N$  increases at an approximately constant rate.
- $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$

We define the first assumption as increasing the number of data by  $\Delta \in [0, 1]$  per updating query. From a micro perspective, of course, each update operation either increases or decreases the data count by one. However, from a macro perspective, we assume that after performing  $k$  update operations, the number of data  $N$  can be approximated as  $N = N_0 + k\Delta$  using the initial number of data  $N_0$ .

We discuss the worst-case time complexity of the “split” operation. An overview of the “split” is shown in Figure 5. “Split” first inserts empty cells (B-trees) in the appropriate location, and then distributes the data equally between the old and new cells. We can implement the insertion of empty cells by simultaneously inserting new  $\frac{X}{x_d}$  cells while sliding (at most)  $X$  cells that already exist. Thus, the insertion of empty cells costs  $O(\frac{X}{x_d} + X) = O(X)$ . For the data distribution, the erase from old B-trees and insert into new B-trees are performed  $\frac{2N}{x_d}$  times in total (Remember the definition of the “split”). Since  $\Delta \geq 0$  holds, the “split” occurs when the number of data of the target cells is exactly



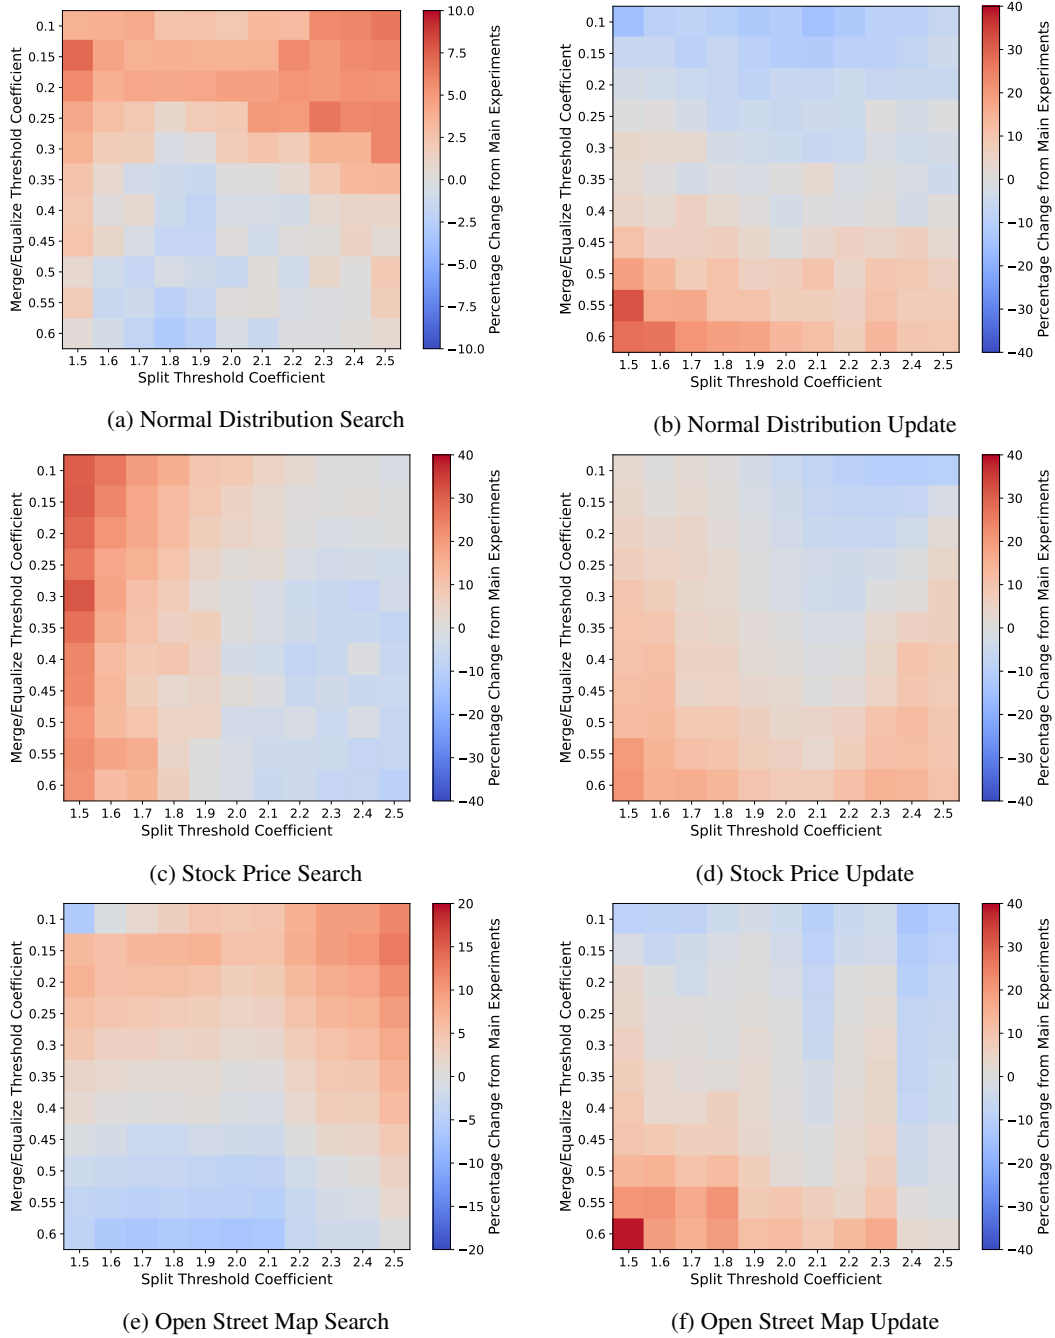


Figure 4: Heatmaps show percentage changes compared to experimental results in Section 6. Blue squares mean better performance.




Figure 5: Overview of “split” ( $D = 3, x_1 = 8, x_2 = 3, X = 24, N = 48$ ): The cells outlined in red are subject to “split” because they contain 13 data in total. This value is larger than the “split” condition:  $2 \cdot \frac{N}{x_1} = 2 \cdot \frac{48}{8} = 12$ . We insert  $\frac{X}{x_1} = \frac{24}{8} = 3$  new cells (B-trees), and distribute the 12 data points evenly between the old and new cells, with 6 points in each.

$\frac{2N}{x_d}$ , and the number of B-tree operations cannot be more than  $\frac{2N}{x_d}$ ). Therefore, the time complexity of operating B-trees for data distribution is  $O(\frac{N}{x_d} \log N)$ . Thus, the worst time complexity of the “split” is  $O(X + \frac{N}{x_d} \log N)$ . Considering the “merge” and the “equalize” in the same way, the worst time complexities are  $O(X + \frac{N}{x_d} \log N)$ ,  $O(\frac{N}{x_d} \log N)$ , respectively.

### C.3 Amortized Time Complexity of Cell Re-partition

However, under the assumption that  $N = N_0 + k\Delta$ , we can show that a re-partition occurs at most once every  $O(\frac{N}{x_d})$ . When  $N = N_0 + k\Delta$ , we insert  $\frac{(1+\Delta)k}{2}$  times and erase  $\frac{(1-\Delta)k}{2}$  times. “Splits” due to intensive insertions occur only  $\frac{(1+\Delta)k}{2} \cdot \frac{x_d - 2\Delta}{N}$  times at most. This is because the “split” occurs when the number of data in the cell reaches exactly  $\frac{2N}{x_d}$  (because  $\Delta \geq 0$  holds), and once the “split” is executed, the number of data in the cell is reset to  $\frac{N}{x_d}$ . So, we need  $\frac{N}{x_d - 2\Delta}$  consecutive intensive insertions to execute the “split” again (Check  $\frac{N}{x_d} + \frac{N}{x_d - 2\Delta} = 2 \cdot \frac{N + \frac{N - 2\Delta}{x_d} \Delta}{x_d}$ ). In the same way, we can say that the “merge” and the “equalize” by intensive erasion occur only  $\frac{(1-\Delta)k}{2} \cdot \frac{3x_d + \Delta}{2N}$  times at most. Moreover, even for cells that are not erased at all, the number of data in the cell may reach the threshold of “merge” or “equalize” because  $N$  increases. About this case, we can prove that the number of computations is maximized when the “merge” or “equalize” occurs immediately when the number of data reaches the threshold. Therefore, the “merge” and the “equalize” due to increasing data can occur only  $x_d \cdot \frac{k}{\frac{2N}{x_d}}$  times at most. From the above, the upper bound on the number of cell re-partitions that occur after  $k$  update operations is  $\frac{k((5+\Delta)x_d - 3\Delta - 5\Delta^2)}{4N}$ . That is, the cell re-partitioning occurs at most once every  $\frac{4N}{(5+\Delta)x_d - 3\Delta - 5\Delta^2} = O(\frac{N}{x_d})$  times.

Remember the worst-case complexity of the cell re-partitioning is  $O(X + \frac{N}{x_d} \log N)$ . This happens once every  $O(\frac{N}{x_d})$  times, so the amortized time complexity of the cell re-partitioning is  $O(\frac{X}{N} x_d + \log N)$ . Since we have considered each axis  $d$  independently, the amortized time complexity of the overall cell re-partitioning is  $O(\frac{X}{N} \sum_{d=1}^D x_d + D \log N)$  by summing them up. Here, from the

second assumption  $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$ ,  $\frac{X}{N} \sum_{d=1}^D x_d < D \log N$  holds. Therefore, it is proved that the amortized time complexity of the cell re-partitioning is  $O(D \log N)$ .

#### C.4 Limitation

We proved that the amortized time complexity of the FlexFlood’s update operation is  $O(D \log N)$  when  $N \simeq N_0 + k\Delta$  and  $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$  are assumed to hold. It remains to be discussed how realistic these assumptions are. In the real world, data basically tends to increase. However, for example, there may be cases where periods of rapid data growth alternate with periods of slower growth on an annual cycle. It is very important to observe how FlexFlood performs in such cases. If the performance drops significantly, we should develop a hybrid approach that balances these periods. Furthermore, although  $\prod_{d=1}^D x_d \sum_{d=1}^D x_d \leq DN \log N$  always held in our experiment, it may not hold for some datasets. We should observe how much the update speed decreases due to the breakdown of this assumption. Furthermore, it would be interesting to identify common characteristics among datasets where this assumption breaks down.

## D Dataset Details

We describe the details of the three dataset used in the experiments. In addition, we also introduce the method for generating queries.

### D.1 Normal Distribution Dataset

Dataset 1 is a Normal Distribution Dataset. We tested  $D = 3$  as the number of dimensions of the normal distribution. As the initial data, we independently generated  $10^5$  data. Each data was generated independently for each axis according to a normal distribution with  $\mu = 3 \cdot 10^8$ ,  $\sigma = 10^8$ . Note that  $\mu$  is the mean and  $\sigma$  is the standard deviation of the normal distribution.

The number of queries we generated is  $2 \cdot 10^6$ , with an update and search queries alternating every  $10^4$  queries. The  $i$ -th query is an update query if  $\lfloor \frac{i}{10^4} \rfloor \equiv 0 \pmod{2}$ , and a search query otherwise. (Note that  $\lfloor x \rfloor$  denotes the largest integer less than or equal to  $x$ .)

When we generated search queries, we first generated a hyper-rectangle, a search region. The length of one side of the hyper-rectangle was independently set to a random value less than  $3 \cdot 10^8$ . We placed that hyper-rectangle uniformly at random inside the hyper-rectangle whose diagonals are  $(0, 0, 0)$  and  $(10^9, 10^9, 10^9)$ , and we used this as a search query.

The update query has a 50% chance of being selected as an insert query and a 50% chance of being selected as an erase query. We generated the insertion queries according to a normal distribution with  $\mu = 3 \cdot 10^8 + 4 \cdot 10^8 \cdot \frac{i}{2 \cdot 10^6}$ ,  $\sigma = 10^8$  independently for each axis when the query was the  $i$ -th from the first. This simulates the gradual change of the data distribution. We generate deletion queries by randomly selecting one of the data currently in the data structure.

### D.2 Stock Price Dataset

Dataset 2 is the Stock Price Dataset [21] used by [12]. we tested  $D = 4$  as the number of dimensions. Of the approximately  $2 \cdot 10^7$  of data included in the dataset, we randomly chose  $2 \cdot 10^6$  data and used them. Each data has four attributes: lowest price, highest price, volume, and date.

Out of  $2 \cdot 10^6$  data points, we took  $2 \cdot 10^5$  with the oldest dates as the initial data. We generated  $7.2 \cdot 10^6$  queries, with the search and update query order matching Dataset 1. Hyper-rectangles were generated randomly for search queries, searching approximately 0.1% of the data per query. Insert queries were generated by selecting the oldest data not yet in the structure, and erase queries by selecting the oldest data in the structure. This simulated keeping the last  $2 \cdot 10^5$  data points and evaluated each data structure’s robustness to real-world data distribution shifts over time.

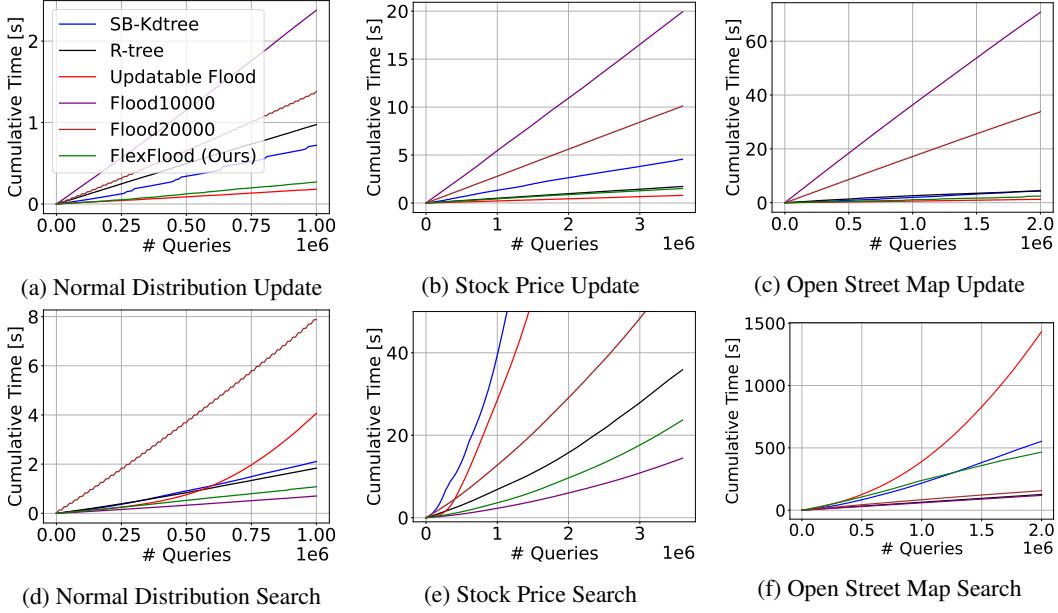


Figure 6: Experimental results: The upper panel shows the update queries, and the lower panel shows the results for the search queries. (Lower is better.)

### D.3 Open Street Map Dataset

Dataset 3 is the Open Street Map Dataset [9] used by [32]. We tested  $D = 5$  as the number of dimensions. Of the approximately  $10^8$  of data, we randomly extracted  $2 \cdot 10^6$  data and used them. Each data set has five attributes: ID, version, date, latitude, and longitude.

First, we clustered the dataset into two groups of  $10^6$  each using K-means, and used one cluster as the initial data. We generated  $2 \cdot 10^6$  queries, with the search and update query order matching Dataset 1. We generated hyper-rectangles randomly for search queries so that approximately 0.3% of the data were searched per query. Insert queries were generated by randomly selecting data from the cluster not selected as the initial data, while erase queries were generated by randomly selecting data from the initial cluster.

## E Re-initialization

There is a delta-buffer method according to the survey paper on learned multi-dimensional indexes [2]. It stores data updating in a small array and periodically merges them with the data structure. We compare the performance of FlexFlood with that of a delta-buffered version of Flood.

In designing the delta-buffer Flood, it is computationally too expensive to incorporate re-learning of parameters (sort dimension and the number of cell partitions). In our implementation, learning the distribution took 32 seconds for the normal distribution dataset, 46 seconds for the Stock Price dataset, and 138 seconds for the Open Street Map dataset. Since there are about  $10^6$  update queries in our workload, we need to re-learn more than  $10^2$  times even if we re-learn every  $10^4$  times. The total re-learning time would then be about  $10^4$  seconds  $\simeq$  3 hours, which is too long as shown in Figures 2. Therefore, we adopt the method of re-initializing the data structure without changing the parameters.

In this experiment, we used Flood10000, which re-initializes the entire structure after every  $10^4$  data updates, and used Flood20000 with re-initialization after every  $2 \cdot 10^4$  updates, in addition to the four data structures described in Section 5. As noted in Appendix D, search queries and update queries come alternately,  $10^4$  each. Thus, Flood10000 is re-initialized for each set of update queries, and Flood20000 stores the first set of update queries in the buffer and re-initializes after the second set.

Figures 6 illustrate the cumulative query processing time for each data structure for each dataset. Looking at the update query results (Figures 6a, 6b, 6c), Flood10000 and Flood20000 are slower

than all the compared data structures. Especially, both are more than five times slower than FlexFlood. From this result, the delta buffer Flood is undesirable when there are many update queries.

We then turn our attention to the search query results (Figures 6d, 6e, 6f). Flood10000 has the fastest search speed of all data structures. This is because Flood10000 re-initializes every update query set, so Flood10000 can achieve virtually the same performance as Flood, and the original Flood is extremely fast, as discussed in Appendix A. On the other hand, Flood20000 is slower than FlexFlood for normal distribution dataset and Stock Prices dataset, but achieves a very fast search speed for Open Street Map dataset. This is likely due to the large amount of data being searched in the Open Street Map dataset (See Appendix D). The time spent searching the internal structure of the Flood has become the rate-limiting step, making the time spent reading from the buffer relatively negligible.

Therefore, delta-buffer Flood is worth considering in situations where an extremely large amount of data is being searched, or when spending a relatively long time on updates is not an issue. On the other hand, we believe that FlexFlood excels in scenarios where high update speed is required or when the amount of data being searched is relatively small.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In the abstract and Section 1, we describe that FlexFlood is the first learned multi-dimensional index that guarantees that the computational complexity of the update operation is  $O(D \log N)$ , and that experimentally it outperforms Flood in certain situations.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss the limitation in Section 7 and Appendix E.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We discuss the computational complexity of the proposed method in Section 4 and Appendix C.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We introduce Flood, the baseline of the FlexFlood, in Section 3. We also describe how we modified it to obtain the FlexFlood in Section 3, 4. In addition, we describe the details of the dataset in Section 5 and Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We do not provide the open access to the data and code. We believe that everyone can reproduce the source code and dataset by reading the paper.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Section 5 and Appendix D provide sufficient experimental detail to understand the results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: Error bars are not reported because it would be too computationally expensive.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.



- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Section 5 describes the execution environment. You can also see the time required for the experiment in Section 6 and Appendix E.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We read the NeurIPS Code of Ethics and conducted our research accordingly.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This is a foundational study focusing on the data structures, so this paper does not address the social implications.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.

- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This is a foundational study focusing on the data structures, so this paper does not involve data or models with a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [No]

Justification: In Section 5, most of the assets were properly introduced. However, the license for the Stock Price Dataset [21] was “unknown”.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: We do not release the source code.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.