



# Manu: A Cloud Native Vector Database Management System

Rentong Guo<sup>†\*</sup>, Xiaofan Luan<sup>†\*</sup>, Long Xiang<sup>‡\*</sup>, Xiao Yan<sup>‡\*</sup>, Xiaomeng Yi<sup>†\*</sup>, Jigao Luo<sup>†§</sup>  
Qianya Cheng<sup>†</sup>, Weizhi Xu<sup>†</sup>, Jiarui Luo<sup>‡</sup>, Frank Liu<sup>†</sup>, Zhenshan Cao<sup>†</sup>, Yanliang Qiao<sup>†</sup>, Ting Wang<sup>†</sup>

Bo Tang<sup>‡</sup>, Charles Xie<sup>†</sup>

<sup>†</sup>Zilliz    <sup>‡</sup>Southern University of Science and Technology    <sup>§</sup>Technical University of Munich

<sup>†</sup>{firstname.lastname}@zilliz.com

<sup>‡</sup>{xiangl3@mail., yanx@, 11911419@mail., tangb3@}ustech.edu.cn    <sup>§</sup>jigao.luo@tum.de

## ABSTRACT

With the development of learning-based embedding models, embedding vectors are widely used for analyzing and searching unstructured data. As vector collections exceed billion-scale, fully managed and horizontally scalable vector databases are necessary. In the past three years, through interaction with our 1200+ industry users, we have sketched a vision for the features that next-generation vector databases should have, which include long-term evolvability, tunable consistency, good elasticity, and high performance.

We present Manu, a cloud native vector database that implements these features. It is difficult to integrate all these features if we follow traditional DBMS design rules. As most vector data applications do not require complex data models and strong data consistency, our design philosophy is to relax the data model and consistency constraints in exchange for the aforementioned features. Specifically, Manu firstly exposes the write-ahead log (WAL) and binlog as backbone services. Secondly, write components are designed as log publishers while all read-only analytic and search components are designed as independent subscribers to the log services. Finally, we utilize multi-version concurrency control (MVCC) and a delta consistency model to simplify the communication and cooperation among the system components. These designs achieve a low coupling among the system components, which is essential for elasticity and evolution. We also extensively optimize Manu for performance and usability with hardware-aware implementations and support for complex search semantics. Manu has been used for many applications, including, but not limited to, recommendation, multimedia, language, medicine and security. We evaluated Manu in three typical application scenarios to demonstrate its efficiency, elasticity, and scalability.

## PVLDB Reference Format:

Rentong Guo, Xiaofan Luan, Long Xiang, Xiao Yan, Xiaomeng Yi, Jigao Luo, Qianya Cheng, Weizhi Xu, Jiarui Luo, Frank Liu, Zhenshan Cao, Yanliang Qiao, Ting Wang, Bo Tang, and Charles Xie. Manu: A Cloud Native Vector Database Management System. PVLDB, 15(12): 3548 - 3561, 2022. doi:10.14778/3554821.3554843

\*Co-first-authors are ordered alphabetically.

‡ Work done while working with Zilliz, correspondence to Bo Tang.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 15, No. 12 ISSN 2150-8097. doi:10.14778/3554821.3554843

## PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/milvus-io/milvus/tree/2.0>.

## 1 INTRODUCTION

According to IDC, unstructured data, such as text, images, and video, took up about 80% of the 40,000 exabytes of new data generated in 2020, their percentage keeps rising due to the increasing amount of human-generated rich media [47]. With the rise of learning-based embedding models, especially deep neural networks, using embedding vectors to manage unstructured data has become commonplace in many applications such as e-commerce, social media, and drug discovery [48, 62, 67]. A core feature of these applications is that they encode the semantics of unstructured data into a high-dimensional vector space. Given the representation power of embedding vectors, operations like recommendation, search, and analysis can be implemented via similarity-based vector search. To support these applications, many specialized vector databases are built to manage vector data [10, 12, 17–19, 80].

In 2019, we open sourced Milvus [80], our previous vector database, under the LF AI & Data Foundation. Since then, we collected feed-backs from more than 1200 industry users and found that some of the design principles adopted by Milvus are not suitable. Milvus followed the design principles of relational databases, which are optimized for either transaction [51] or analytical [80] workloads, and focused on functionality supports (e.g., attribute filtering and multi-vector search) and execution efficiency (e.g., SIMD and cache optimizations). However, vector database applications have different requirements in the following three aspects, which motivates us to restructure Manu from scratch with focuses on a cloud-native architecture.

- **Support for complex transactions is not necessary.** Instead of decomposing entity representations into different fields or tables, learning-based models encode complex and hybrid data semantics into a single vector. As a result, multi-row or multi-table transactions are not necessary; row-level ACID is sufficient for the majority of vector database applications.
- **A tunable performance-consistency trade-off is important.** Different users have different consistency requirements; some users prefer high throughput and eventual consistency, while others require some level of guaranteed consistency, i.e., newly inserted data should be visible to queries either immediately or within a pre-configured time. Traditional relational databases generally support either strong consistency or eventual consistency; there is little to no room for customization between these

two extremes. As such, *tunable* consistency is a crucial attribute for cloud-native vector databases.

- **High hardware cost calls for fine-grained elasticity.** Some vector database operations (e.g., vector search and index building) are computationally intensive, and hardware accelerators (e.g. GPUs or FPGAs) and/or a large working memory are required for good performance. However, depending on application types, workload differs amongst database functionalities. Thus, resources can be wasted or improperly allocated if the vector database does not have fine-grained elasticity. This necessitates a careful decoupling of functional and hardware layers; system-level decoupling such as separation of read and write logic is insufficient, elasticity and resource isolation should be managed at the functionalities level rather than the system level.

In summary, modern vector databases should have tunable consistency, functionality-level decoupling, and per-component scalability. Following the design principles of traditional relational databases makes achieving these design goals extremely difficult, if not impossible. A key opportunity for achieving these design goals lies in the potential for relaxing transaction complexity.

Manu follows the “log as data” paradigm. Specifically, Manu structures the entire system as a group of log publish/subscribe microservices. The write-ahead log (WAL) and inter-component messages are published as “logs”, i.e., durable data streams that can be subscribed. Read-side components, such as search and analytical engines, are all built as log subscribers. This architecture provides a simple yet effective way to decouple system functionalities; it enables the decoupling of read from write, stateless from stateful, and storage from computing. Each log entry is assigned a global unique timestamp, and special log entries called time-tick (similar to watermarks in Apache Flink [25]) are periodically inserted into each log channel signaling the progress of event-time for log subscribers. The timestamp and time-tick form the basis of the tunable consistency mechanism and multi-version consistency control (MVCC). To control the consistency level, a user can specify a tolerable time lag between a query’s timestamp and the latest time-tick consumed by a subscriber.

Additionally, we extensively optimize Manu for performance and usability. Manu supports various indexes for vector search, including vector quantization [21, 33, 36, 82], inverted index [23], and proximity graphs [32]. In particular, we tailor the implementations to better utilize the parallelization capabilities of modern CPUs and GPUs along with the improved read/write speeds of SSDs over HDDs. Manu also integrates refactored functionalities from Milvus [80], such as attribute filtering and multi-vector search. Moreover, build a visualization tool that allows users to track the performance of Manu in real time and include an auto-configuration tool that recommends indexing algorithm parameters using machine learning.

To summarize, this paper makes the following contributions:

- We summarize lessons learned from communicating with over 1200 industry users over three years. We shed light on typical application requirements of vector databases and show how they differ from those of traditional relational databases. We then outline the key design goals that vector databases should meet.

- We introduce Manu’s key architectural designs as a cloud native vector database, building around the core design philosophy of relaxing transaction complexity in exchange for tunable consistency and fine-grained elasticity.
- We present important usability and performance-related enhancements, e.g., high-level API, a GUI tool, automatic parameter configuration, and SSD support.

The rest of the paper is organized as follows. Section 2 provides background on the requirements and design goals of vector databases. Section 3 dives deep into Manu’s design. Section 4 highlights the key features for usability and performance. Section 5 discusses representative use cases for Manu. Section 6 review related works. Section 7 concludes the paper and outlines future work.

## 2 BACKGROUND AND MOTIVATION

Consider video recommendation as a typical use case of vector databases. The goal is to help users discover new videos based on their personal preferences and previous browsing history. Using machine learning models (especially deep neural networks), features of users and videos, such as search history, watch history, age, gender, video language, and tags are converted to embedding vectors. These models are carefully designed and trained to encode the similarity between user and video vectors into a common vector space. Recommendation is conducted by retrieving candidate videos from the collection of video vectors via similarity scores with respect to the specified user vector. The system also needs to handle updates to vectors when new videos are updated, some videos are deleted and the embedding model is changed.

Video recommendation and other applications of vector databases can involve hundreds of billions of vectors with daily growth at hundred-million scale, and serve million-scale queries per second (QPS). Existing DBMSs (e.g., relational databases [9, 11], NoSQL [75, 85], NewSQL [39, 73]) were not built to manage vector data on that scale. Moreover, the underlying data management requirements of their applications differ greatly from vector database applications.

First, when compared with relational databases, both the architecture and theory of vector databases are far from mature. A key reason for this is that AI- and data-driven applications are still in a state of constant evolution, thereby necessitating continued architectural and functionality changes to vector databases as well.

Second, complex transactions are unnecessary for vector databases. In the above example, the recommendation system encodes all semantic features of users and videos into standalone vectors as opposed to multi-row or multi-column entity fields in a relational database. As a result, row-level ACID is sufficient; multi-table operations (such as joins) are inessential.

Third, vector database applications need a flexible performance-consistency trade-off. While some applications adopt a strong or eventual consistency model, there are others that fall between the two extremes. Users may wish to relax consistency constraints in exchange for better system throughput. In the video recommendation example, observing a newly uploaded video after several seconds is acceptable but keeping users waiting for recommendation harms user experience. Thus, the application can configure the

allowed maximal delay for the video updates in order to improve system throughput.

Fourth, vector databases have more stringent and diversified hardware requirements compared with traditional databases. This is attributed to three reasons. First, vector database operations are computation-intensive, and thus hardware accelerators such as GPUs are critical for computing functionalities such as search and indexing. Second, accesses to vector data (e.g., search or update) generally have poor locality, thereby requiring large RAM for good performance. Third, different applications vary significantly in their resource demands for the system functionalities. Core functionalities of a vector database include data insertion, indexing, filtering, and vector search. Applications such as video recommendation require online insertion and high concurrency vector search. In contrast, for interactive use cases such as drug discovery, offline data ingestion and indexing are generally acceptable. Although interactive applications usually require lower throughput than recommendation systems, they have high demands for real-time filtering, similarity-based vector search, and hybrid queries. The high hardware costs as well as diverse workload features call for fine-grained elasticity.

The key design goals of Manu are summarized below; these design goals not only fully encompass the above characteristics but also share some common goals with generic cloud-based databases.

- **Long-term evolvability:** Overall system complexity must be controlled for the continuous evolution of Manu’s functionalities. Without the need to support complex transactions, there lies an opportunity to model all the event sequences (such as WAL and inter-component messages) as message queues to cleanly decouple the entire system. In this way, individual components can evolve, be added, or be replaced easily with minimal interference to other components. This design echos large-scale data analytic platforms, which often rely on data streaming systems such as Kafka to connect system components.
- **Tunable consistency:** To enable flexible consistency-performance trade-off, Manu should introduce *delta consistency* that falls between strong consistency and eventual consistency, where a read operation returns the last value that was produced at most delta time units preceding itself. It’s worth noting that strong consistency and eventual consistency can be realized as special cases of this model, with delta being zero and infinity, respectively.
- **Good elasticity:** Workload fluctuations can cause different loads on individual system components. In order to dynamically allocate compute resources to high-load tasks, components must be carefully decoupled, taking both functionality and hardware dependencies into consideration. System elasticity and resource isolation should be managed at the component-level rather than at the system-level (e.g. decoupling indexing from querying versus decoupling read from write).
- **High availability:** Availability is a must-have for modern cloud-based applications; Manu must isolate system failures at the component level and make failure recovery transparent.
- **High performance:** Query processing performance is key to vector databases. For good performance, implementations to be extensively optimized for hardware. Moreover, the framework




Figure 1: An example of Manu’s schema.

should be carefully designed so as to minimize system overheads for query serving.

- **Strong adaptability:** Our customers use vector databases in a variety of environments, ranging from prototyping on laptops to large-scale deployments on the cloud. A vector database should provide consistent user experience and reduce code/data migration overhead across environments.

### 3 THE MANU SYSTEM

In this section, we begin by first introducing the basic concepts of Manu. Next, we present the system designs, including the overall system architecture, the log backbone, and how Manu conducts vector searches and builds vector search indexes.

#### 3.1 Schema, Collection, Shard, and Segment

**Schema:** The basic data types of Manu are vector, string, boolean, integer, and floating point. A schema example is given in Figure 1. Suppose each entity consists of five fields and corresponds to a product on an e-commerce platform. The *Primary key* is the ID of the entity. It can either be an integer or a string. If users do not specify this field, the system will automatically add an integer primary key for each entity. The *Feature vector* is the embedding of the product. The *Label* is the category of the product, such as food, book, and cloth. The *Numerical attribute* is a float or an integer associated with the product, such as price, weight, or production date. Manu supports multiple labels and numerical attributes in each entity. Note that these fields are used for filtering, rather than join or aggregation. The *Logical sequence number (LSN)* is a system field hidden from users.

**Collection:** A Collection is a set of entities similar to the concept of tables in relational databases. For example, a collection can contain all the products of an e-commerce platform. The key difference is that collections have no relations with each other; thus, relational algebra, such as join operations, are not supported.

**Shard:** The Shard correspondence to insertion/deletion channel. Entities are hashed into multiple shards based on their primary keys during insertion/deletion. Manu’s data placement unit is segment rather than shard.<sup>1</sup>

**Segment:** Entities from each shard are organized into segments. A segment can be in either a *growing* or *sealed* state. Sealed segments are read-only while growing segments can accept new entities. A growing segment will switch to sealed state when it reaches a predefined size (set to 512MB by default) or if a period of time has passed without an insertion (e.g., 10 seconds). As some segments may be

<sup>1</sup>Using segments for data placement is more flexible than shards, as the number of shards is static, while the number of segments grows as the volume of the collection increases.




Figure 2: The Architecture of Manu.

small (e.g., when insertion has a low arrival rate), Manu merges small segments into larger ones for search efficiency.

### 3.2 System Architecture

Manu adopts a service-oriented design [64] to achieve fine-grained decoupling among the system components. As shown in Figure 2, from top to bottom, Manu has four layers, i.e., *access layer*, *coordinator layer*, *worker layer*, and *storage layer*.

**Access layer** consists of stateless proxies that serve as the user endpoints. They work in parallel to receive requests from clients, distribute the requests to the corresponding processing components, and aggregate partial search results before returning to clients. Furthermore, the proxies cache a copy of the metadata for verifying the legitimacy of search requests (e.g., whether the collection to search exists). Search request verification is lightweight and moving it to the proxies has two key benefits. First, requests that fail verification are rejected early, thus lowering the load on other systems components. Second, it reduces the number of routing hops for requests, thus shortening request processing latency.

**Coordinator layer** manages system status, maintains metadata of the collections, and coordinates the system components for processing tasks. There are four coordinators, each responsible for different tasks. *Root coordinator* handles data definition requests, such as creating/deleting collections, and maintains meta-information of the collections. *Data coordinator* records detailed information about the collections (e.g., the routes of the segments on storage), and coordinates the data nodes to transform data update requests into binlogs [4]. *Query coordinator* manages the status of the query nodes, and adjusts the assignment of segments (along with related indexes) to query nodes for load balancing. *Index coordinator* maintains meta-information of the indexes (e.g., index types and storage routes), and coordinates index nodes in index building tasks. A coordinator can have multiple instances (e.g., one main and two backups) for reliability. As vector databases usually do not have the cross table operations that relational databases have, different collections can be served by separate coordinator instances for throughput.

**Worker layer** conducts the actual computation tasks. The worker nodes are stateless—they fetch read-only copies of data to conduct tasks and do not need to coordinate with each other. This ensures

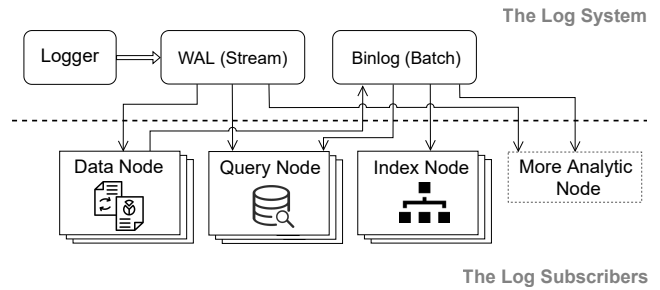


Figure 3: Overview of Manu's log system.

that computation intensive (thus expensive) worker nodes can be easily scaled on demand. We use different worker nodes for different tasks, i.e., *query nodes* for query processing, *index nodes* for index building, and *data nodes* for log archiving. Due to the fact that the workloads for different computation tasks vary significantly over time and across applications, each worker type can scale independently. This design also achieves resource isolation as different computation tasks have different QoS requirements.

**Storage layer** persists system status, metadata, the collections, and associated indexes. Manu uses etcd [7] (a key-value store) to host system status and metadata for the coordinators as etcd provides high availability with its leader election mechanism for failure recovery. When metadata is updated, the updated data is first written to etcd, and then synchronized to coordinators. Since the volume of other data (e.g., binlog, data, index) is large, Manu uses AWS S3 [13] (an object storage) for persistence due to its high availability and low cost. The API of many other object storage systems is compatible with AWS S3. This allows Manu to easily swap storage engines, if necessary. At present, storage engines including AWS S3, MinIO [8], and Linux file system are supported. Note that the high latency that comes with object storage is not a performance bottleneck as the worker nodes conduct computation tasks on in-memory, read-only copies of data.

### 3.3 The Log Backbone

The log system is the backbone of Manu, which connects the decoupled system components. As shown in Figure 3, Manu exposes the write-ahead log (WAL) and binlog as backbone services. The WAL is the incremental part of system log while the binlog is the base part; they complement each other in delay, capacity and cost. Loggers are entry points for publishing data onto the WAL. Data nodes subscribe to the WAL and convert the row-based WALs into column-based binlogs. All read-only components such as index nodes and query nodes are independent subscribers to the log service to keep themselves up-to-date. This architecture completely decouples the write and read components, thus allowing the components (e.g., WAL, binlog, data nodes, index nodes and query nodes) to scale independently.

Manu records all the requests that change system state to the log, including data definition requests (e.g., create/delete collection), data manipulation requests (e.g., insert/delete a vector), and system coordination messages (e.g., load/dump a collection to/from memory). Note that vector search requests are not written to the log as they are read-only operations and do not change system

state. We use logical logs instead of physical logs, as logical logs focus on event recording, rather than describing the modifications to physical data pages. This allows the subscribers to consume the log data in different ways depending on their functions.

Figure 4 illustrates the detailed architecture of the log system. For the sake of clarity, we only illustrate the parts related to insert requests. The loggers are organized in a hash ring, and each logger handles one or more logical buckets in the hash ring based on consistent hashing. Each shard corresponds to a logical bucket in the hash ring and a WAL channel. Each entity in insert requests is hashed to a shard (and thus channel) based on their ID. When a logger receives a request, it will first verify the legibility of the request, assign an LSN for the logged entity by consulting the central time service oracle (TSO), determine the segment the entity should go to, and write the entity to WAL. The logger also writes the mapping of the new entity ID to segment ID into a local LSM tree and periodically flushes the incremental part of the LSM tree to object storage, which keeps the entity to segment mapping using the SSTable format of RocksDB. Each logger caches the segment mapping (e.g., for checking if the entity to delete exists) for the shards it manages by consulting the SSTable in object storage.

The WAL is row-based and read in a streaming manner for low delay and fine-grained log pub/sub. It is implemented via a cloud-based message queue such as Kafka or Pulsar. We use multiple logical channels for the WAL in order to prevent different types of requests from interfering with each other, thus achieving a high throughput. Data definition requests and system coordination messages use their own channels while data manipulation requests hashed across multiple channels to increase throughput.

Data nodes subscribe to the WAL and convert the row-based WALs into column-based binlogs. Specifically, values from the same field (e.g., attribute and vector) from the WALs are stored together in a column format in binlog files. The column-based nature of binlog makes it suitable for reading per field values in batches, thus increasing storage and IO efficiency. An example of this efficiency comes with the index nodes. Index nodes only read the required fields (e.g., attribute or vector) from the binlog for index building and thus are free from the read amplifications.

**System coordination:** Inter-component messages are also passed via log, e.g., data nodes announce when segments are written to storage and index nodes announce when indexes have been built. This is because the log system provides a simple and reliable mechanism for broadcasting system events. Moreover, the time semantics of the log system provide a deterministic order for coordination messages. For example, when a collection should be released from memory, the query coordinator publishes the request to log, and does not need to confirm whether the query nodes receive the message or handle query node failure. The query nodes independently subscribe to the log and asynchronously release segments of the collection.

### 3.4 Tunable Consistency

We adopt a delta consistency model to enable flexible performance-consistency trade-offs, which guarantees a bounded staleness of data seen by search queries. Specifically, the data seen by a query can be stale for up to delta time units, with respect to time of the last

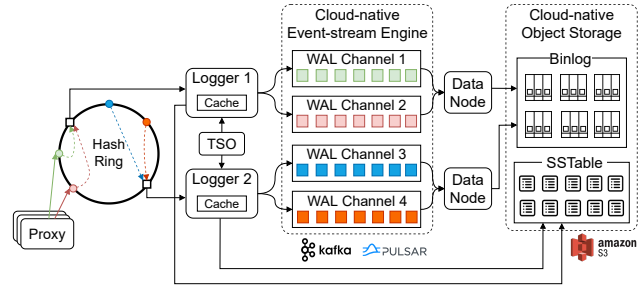


Figure 4: Detailed structure of Manu’s log system.

Table 1: Major indexes in Manu

<b>Vector Quantization</b>	PQ, OPQ, RQ, SQ
<b>Inverted Index</b>	IVF-Flat, IVF-PQ, IVF-SQ, IVF-HNSW, IMI
<b>Proximity Graph</b>	HNSW, NSG, NGT
<b>Numerical Attribute</b>	B-Tree, Sorted List

data update, where delta is an user-specified “staleness tolerance” given in virtual time.

In practice, users prefer to define temporal tolerance as physical time, e.g., 10 seconds. Manu achieves this by making the LSN assigned to each request extremely close to physical time. Manu uses a hybrid logical clock in the TSO to generate timestamps. Each timestamp has two components: a physical component that tracks physical time, and a logical component that tracks event order. The logical component is needed since multiple events may happen at the same physical time unit. Since a timestamp is used as a request’s LSN, the value of the physical component indicates the physical time when the request was received by Manu.

For a log subscriber, e.g., a query node, to run the delta consistency model, it needs to know three things: (1) the user-specified staleness tolerance  $\tau$ , (2) the time of the last data update, and (3) the issue time of the search request. In order to let each log subscriber know (2), we introduce a time-tick mechanism. Special control messages called time-ticks (similar to watermarks in Apache Flink [25]) are periodically inserted into each log channel (for example, WAL channel) signaling the progress of data synchronization. Denote the latest time-tick a subscriber consumed as  $L_s$  and the issue time of a query as  $L_r$ , if  $L_r - L_s < \tau$  is not satisfied, the query node will wait for the next time-tick before executing the query.

Note that strong consistency and eventual consistency are two special cases of delta consistency, where delta equals to 0 and infinity, respectively. To the best of our knowledge, our work is the first to support delta consistency in a vector database.

### 3.5 Index Building

Searching similar vectors in large collections by brute-force, i.e., scanning the whole dataset, usually yields unacceptably long delays. Numerous indexes have been proposed to accelerate vector search and Manu automatically builds user specified indexes. Table 1 summarizes the indexes currently supported by Manu, and we are continuously adding new indexes following the latest indexing algorithms. These indexes differ in their properties and use cases. Vector quantization (VQ) [33, 44] methods compress vectors



to reduce memory footprint and the costs for vector distance/similarity computation. For example, scalar quantization (SQ) [90] maps each dimension of vector (data types typically are int32 and float) to a single byte. Inverted indexes [68] group vectors into clusters, and only scan the most promising clusters for a query. Proximity graphs [32, 41, 60] connect similar vectors to form a graph, and achieve high accuracy and low latency at the cost of high memory consumption [53]. Besides vector indexes, Manu also supports indexes on the attribute field of the entities to accelerate attribute-based filtering.

There are two index building scenarios in Manu, i.e., *batch indexing* and *stream indexing*. Batch indexing occurs when the user builds an index for an entire collection (e.g., when all vectors are updated with a new embedding model). In this case, the index coordinator obtains the paths of all segments in the collection from the data coordinator, and instructs index nodes to build indexes for each segment. Stream indexing happens when users continuously insert new entities, and indexes are built asynchronously on-the-fly without stopping search services. Specifically, after a segment accumulates a sufficient number of vectors, its resident data node seals the segment and writes it to object storage as a binlog. The data coordinator then notifies the index coordinator, which instructs a index node to build index for the segment. The index node loads only the required column (e.g., vector or attribute) of the segment from object storage for indexing building to avoid read amplification. For entity deletions, Manu uses a bitmap to record the deleted vectors and rebuilds the index for a segment when a sufficient number of its entities have been deleted. In both batch and stream indexing scenarios, after the required index is built for a segment, the index node persists it in the object storage and sends the path to the index coordinator, which notifies the query coordinator so that query nodes can load the index for processing queries. The index coordinator also monitors the status of the index nodes and shuts down idle index nodes to save costs. As vector indexes generally have sub-linear search complexity w.r.t. the number of vectors, searching a large segment is cheaper than several small segments, Manu builds joint indexes on multiple segments when appropriate.

### 3.6 Vector Search

Manu supports *classical vector search*, *attribute filtering*, and *multi-vector search*. For classical vector search, the distance/similarity function can be Euclidean distance, inner product or angular distance. Attribute filtering is useful when searching vectors similar to the query subject to some attribute constraints. For example, an e-commerce platform may want to find products that interest the customer and cost less than 100\$. Manu supports three strategies for attribute filtering and uses a cost-based model to choose the most suitable strategy for each segment. Multi-vector search is required when an entity is encoded by multiple vectors, for example, a product can be described by both embeddings of its image and embeddings of its text description. In this case, the similarity function between entities is defined as a composition of similarity functions on the constituting vectors. Manu supports two strategies for multi-vector search and chooses the one to use according to the

entity similarity function. For more details about how Manu handles attribute filtering and multi-vector search, interested readers can refer to Milvus [80].

For vector search, Manu partitions a collection into segments and distributes the segments among query nodes for parallel execution.<sup>2</sup> The proxies cache a copy of the distribution of segments on query nodes by inquiring the query coordinator, and dispatch search requests to query nodes that hold segments of the searched collection. The query nodes perform vector searches on their local segments without coordination using a *two-phase reduce* procedure. For a top- $k$  vector search request, the query nodes search their local segments to obtain the segment-wise top- $k$  results. These results are merged by each query node to form the node-wise top- $k$  results. Then, the node-wise top- $k$  results are aggregated by the proxy for the global top- $k$  results and returned to the application. To handle the deletion of vectors, the query nodes use a bitmap to record the deleted vectors in each segment and filter the deleted vectors from the segment-wise search results. Users can configure Manu to batch search requests to improve efficiency. In this case, the proxies organize cache search requests if results of the previous batches have not been returned yet. In the cache, requests of the same type (i.e., target the same collection and use the same similarity function) are organized into the one batch and handled by Manu together. Manu also allows maintaining multiple hot replicas of a collection to serve queries for availability and throughput.

Query nodes obtain data from three sources, i.e., the WAL, the index files, and the binlog. For data in the growing segments, query nodes subscribe to the WAL and conduct searches using brute force scan so that updates can be searched within a short delay. A dilemma for segment size is that larger size yields better search efficiency once the index is built but brute force scan on growing segment is also more costly. To tackle this problem, we divide each segment into *slices* (each containing 10,000 vectors by default). New data are inserted into the slices sequentially, and after a slice is full, a light-weight temporary index (e.g., IVF-FLAT) is built for it. Empirically, we observed that the temporary index brings up to 10X speedup for searching growing segments. When a segment changes from growing state to sealed state, its index will be built by an index node and then stored in object storage. After that, query nodes are notified to load the index and replace the temporary index.

Query nodes access the binlog for data when the distribution of segments among the query nodes changes, which may happen during scaling, load-balancing, query node failure and recovery. Specifically, the query coordinator manages the segment distribution and monitors the query nodes for liveness and workload to coordinate failure recovery and scaling. On failure recovery, the segments and their corresponding indexes (if they exist) handled by failed query nodes are loaded to the healthy ones.<sup>3</sup> In the case of scaling down, a query node can be removed once other query nodes load the indexes for the segments it handles from the object storage. When scaling up, the query coordinator assigns some of

<sup>2</sup>Manu loads all data to the query nodes as different queries may access different parts of the data, and a hot compute side cache is necessary for low latency. This is different from general cloud DBMSs that decouple compute and storage (e.g., Snowflake [28]), which only fetch the required data to compute side upon request.

<sup>3</sup>The WAL channels subscribed to by failed query nodes are also assigned to healthy ones.

**Table 2: Main commands of Python-based PyManu API**

Methods	Description
Collection( <i>name, schema</i> )	Create collection with name <i>str</i> and schema <i>schema</i>
Collection.insert( <i>vec</i> )	Insert vector <i>vec</i> into collection
Collection.delete( <i>expr</i> )	Delete vectors satisfying boolean expression <i>expr</i> from collection
Collection.create_index( <i>field, params</i> )	Create index on a <i>field</i> of the vectors using parameters <i>params</i>
Collection.search( <i>vec, params</i> )	Vector search for <i>vec</i> with parameters <i>params</i>
Collection.query( <i>vec, params, expr</i> )	Vector search for <i>vec</i> with boolean expression <i>expr</i> as filters

the segments to the newly added nodes. A new query node can join after it loads the assigned segments, and existing query nodes can release the segments no longer handled by them. The query coordinator also balances the workloads (and memory consumption) of the query nodes by migrating segments. Note that Manu does not ensure that segment redistribution is atomic, and a segment can reside on more than one query node. This does not affect correctness as the proxies remove duplicate result vectors for a query.

## 4 FEATURE HIGHLIGHTS

In this part, we introduce several key features of Manu for usability and performance.

### 4.1 Cloud Native and Adaptive

The primary design goal of Manu is to be a cloud native vector database such that it fits well into cloud-based data pipelines. To this end, Manu decouples system functionalities into storage, coordinators, and workers in the overall design. For storage, Manu uses a transaction KV for metadata, message queues for logs, and an object KV for data, which are all general storage services provided by major cloud vendors and thus enables easy deployment. For coordinators that manage system functionalities, Manu uses the standard one main plus two hot backups configuration for high availability. For workers, Manu decouples vector search, log archiving and index building tasks for component-wise scaling, a model suitable for cloud-based on-demand resource provisioning. The log backbone allows the system components to interact by writing/reading logs in their own ways. This enables the system components to evolve independently and makes it easy to add new components. The log backbone also provides consistent time semantics in the system, which are crucial for deterministic execution and failure recovery.

Our customers use vector databases in the entire life-cycle of their applications. For example, an application usually starts with data scientists conducting proof of concept (PoC) on their personal computers. Then, it is migrated to dedicated clusters for testing and finally deployed on the cloud. Thus, to reduce migration costs, our customers expect vector databases to adapt to different deployment scenarios while providing a consistent set of APIs. To this end, Manu defines unified interface for the system components but provides different invocation methods and implementations for different platforms. For example, on cloud, local cluster and personal computer, Manu uses cloud service APIs, remote procedure call (RPC) and direct function calls to invoke system functionalities, respectively. The object KV can be the local file system (e.g., MinIO [8]) on personal computers, and S3 on AWS. Thus, Manu

applications can migrate with little or no change across different deployment scenarios.

### 4.2 Good Usability

Data pipelines interact with Manu in simple ways: vector collections, updates for vector data and search requests are fed to Manu, and Manu returns the identifiers of the search results for each search request, which can be used to retrieve objects (e.g., images, advertisements, movies) in other systems. Because different users adopt different programming languages and development environments, Manu provides APIs in popular languages including Python, Java, Go, C++, along with RESTful APIs. As an example, we show key commands of the Python-based PyManu API in Table 2, which uses the object-relational mapping (ORM) model and most commands are related to the *collection* class. As shown in Table 2, PyManu allows users to manage collections and indexes, update collections, and conduct vector searches. The *search* command is used for similarity-based vector search while the *query* command is mainly used for attribute filtering. We show an example of conducting top-*k* vector search by specifying the parameters in *params* in as follows.

```
query_param = {
    "vec": [[0.6, 0.3, ..., 0.8]],
    "field": "vector",
    "param": {"metric_type": "Euclidean"},
    "limit": 2,
    "expr": "product_count > 0",
}
res = collection.search(**query_param)
```

In the above example, the search request provides a high dimensional vector [0.6, 0.3, ..., 0.8] as query and searches the *feature vector* field of the collection. The similarity function is Euclidean distance and the targets are the top-2 similar vectors in the collection (i.e., with limit = 2).

For easy system management, Manu provides a GUI tool called Attu, for which a screen shot is shown in Figure 5. In the *system view*, users can observe overall system status including queries processed per second (QPS), average query latency, and memory consumption on the top of screen. By clicking a specific service (e.g., data service), users can view detailed information of the worker nodes for the service on the side. We also allow users to add and drop worker nodes with mouse clicks. In the *collection view*, users can check the collections in the system, load/dump collections to/from memory, delete/import collections, check the index built for the collections, and build new indexes. In the *vector search view*, users can check the search traffic and performance on each collection, configure the




Figure 5: A screenshot of Attu, the GUI tool of Manu.

index and search parameters to use for each collection. The vector search view also allows to issue queries for functionality test.

For vector search, using different parameters for the indexes (e.g., neighbor size  $M$  and queue size  $L$  for HNSW [60]) yields different trade-offs among cost, accuracy, and performance. However, even experts find it difficult to set proper index parameters as the parameters are interdependent and their influences vary across collections. Manu adopts a Bayesian Optimization with Hyperband (BOHB) [31] method to automatically explore good index parameter configurations. Users provide a utility function to score the configurations (e.g., according to search recall, query throughput) and set a budget to limit the costs of parameter search. BOHB starts with a group of initial configurations and evaluates their utilities. Then, Bayesian Optimization is used to generate new candidate configurations according to historical trials and Hyperband is used to allocate budgets to different areas in the configuration space. The idea is to prioritize the exploration of areas close to high utility configurations to find even better configurations. Manu also supports sampling a subset of the collection for the trails to reduce search costs. We are still improving the automatic parameter search module and plan to extend it to searching system configurations (e.g., the number and type of query nodes).

### 4.3 Time Travel

Users often need to rollback the database to fix corrupted data or code bugs. Manu allows users to specify a target physical time  $T$  for database restore, and jointly uses checkpoint and log replay for rollback. We mark each segment with its progress  $L$  and periodically checkpoints the segment map for a collection, which contains information (such a route, rather than data) of all its segments. To restore the database at time  $T$ , we read the closest checkpoint before  $T$ , load all segments in the segment map and replay the WAL log for each segment from its local progress  $L$ . This design reduces storage consumption as we do not write entire collection for each checkpoint. Instead, segments that have no changes are shared among checkpoints. The replay overhead is also reduced as each segment has its own progress. Users can also specify a expiration period to delete outdated log and segments to reduce storage consumption.

### 4.4 Hardware Optimizations

Manu comes with extensively optimized implementations for CPU, GPU and SSD for efficiency. For more details about our CPU and GPU optimizations, interested readers can refer to Milvus [80].




Figure 6: Manu and Milvus for mixed workloads, numbers behind legends (e.g., 1k) indicate insertion rate.

SSD is 100x cheaper than dram and offers 10x larger bandwidth than HDD. thus, Manu supports using SSD to store large vector collections on cheap query nodes with limited dram capacity. the challenge is that SSD bandwidth is still much smaller than dram, which may lead to low query processing throughput and thus necessitates careful designs for storage layout and index structure. as SSD reads are conducted with 4kb blocks (i.e., reading less than 4kb has the same cost as reading 4kb), Manu organizes the vectors into buckets whose sizes are close to but smaller than 4kb. <sup>4</sup> this is achieved by conducting hierarchical k-means for the vectors and controlling the sizes of the clusters. each bucket is stored on 4kb aligned blocks on SSD for efficient read and represented by its k-means center in dram. these centers are organized using existing vector search indexes (e.g., ivf-flat, hnsw).

vector search with SSD is conducted in two stages. first, we search the cluster centers in dram for the ones that are most similar to the query. then, the corresponding buckets are loaded from SSD for scan. to reduce the amount of data fetched from SSD, we compress the vectors using scalar quantization, which has negligible influence on the quality of search results according to our trials. another problem is that k-means can put vectors similar to a query into several buckets but the centers of some buckets may not be similar to the query, which leads to a low recall. to tackle this problem, Manu uses a strategy similar to multiple hash tables in locality sensitive hashing [40]. hierarchical k-means is conducted by multiple times, each time assigning a vector to a bucket. this means that a vector is replicated multiple times in SSD and we index all cluster centers for bucket search in dram. Manu’s SSD solution wins track 2 (search with SSD) of the billion-scale approximate nearest neighbor search challenge at neurips’2021 [3]. tests results show that Manu’s solution improves the recall of the competition baseline by up to 60% at the same query processing throughput. <sup>5</sup> we notice that another work adopts similar designs for SSD-based vector search [26].

## 5 USE CASES AND EVALUATION

Before introducing the use cases of Manu, we first compare Manu with Milvus, our previous vector database. Milvus adopts an eventual consistency model and thus does not support the tunable consistency of Manu. To show the advantages brought by Manu’s

<sup>4</sup>we set the bucket size to a few times (e.g., 4 and 8) of 4kb if the size of an individual vector is large.

<sup>5</sup>for more details about results please refer to [71].



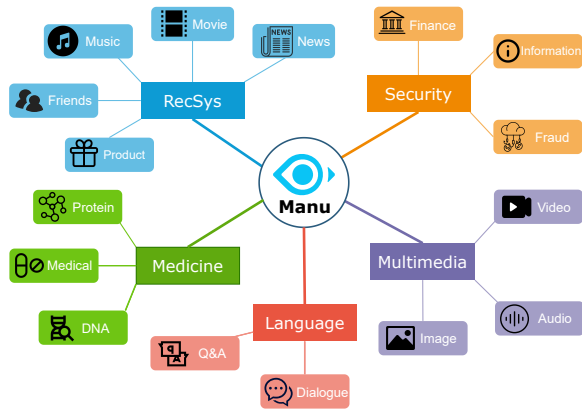


Figure 7: The use cases of Manu.

fine-grained functionality decomposition, we create a mixed workload. Specifically, we start with an empty collection, insert vectors at a fixed rate (e.g., 2k vectors per second), and measure the latency for search requests over time. Both Manu and Milvus use 6 nodes and are properly configured for good performance. The results in Figure 6 show that the search latency of Milvus is significantly longer than Manu, especially when insertion rates are high (e.g., at 3k and 4k). Milvus has multiple read nodes, but only one write node, to ensure eventual consistency. The write node responsible for data insertion and index construction, and thus write tasks and index building tasks contend for resource. As a result, the index building latency is long and brute force search is used for a large amount of data. In contrast, with dedicated index nodes, Manu finishes index building quickly and thus search latency remains low over the entire period.

## 5.1 Overview of Use Cases

We classify our customers into 5 application domains in Figure 7 and briefly elaborate them as follows.

**Recommendation:** Platforms for e-commerce [78], music [76], news [54], video [27], and social network [43] record user-content interactions, and use the data to map users and contents to embedding vectors with techniques such as ALS [74] and deep learning [49]. Finding contents of interest for a user is conducted by searching content vectors having large similarity scores (typically inner product) with user vector.

**Multimedia:** Multimedia contents (e.g., images, video and audio) are becoming increasingly popular, and searches for multimedia contents from large corpus are common online. The general practice is to embed both user query and corpus contents into vectors using tools such as CNN [48] and RNN [86]. Searching multimedia contents is conducted by finding vectors similar to the user query.

**Language:** Automatic questioning answering and machine-based dialogue attract much attention recently with products such as Siri [14] and Xiaoice [20], and searches for text contents is a general need. With models such as Word2Vec [62] and BERT [30], language sequences are embedded into vectors such that retrieving language




Figure 8: Recall vs. throughput comparison.

contents boils down to finding content vectors that are similar to user query.


**Security:** Blocking spams and scanning viruses are important for security. The common practice is to map spams and viruses into vectors using hashing [58] or tailored algorithms [38]. After that, suspicious spams and viruses can be checked by finding the most similar candidates in the corpus for further check.

**Medicine:** Many medical applications search for certain chemical structures and gene sequences for drug discovery or health risk identification. With tools such as GNN [67] and LSTM [83], chemical structures and gene sequences can be embedded into vectors and their search tasks are cast into vector search.

Full-fledged vector databases are necessary for the forgoing domains as they require much more complex functionality support in addition to vector search. Specifically, as the vector datasets are large and applications have high requirements for throughput, they need distributed computing with multiple nodes for scalability. The vectors are also continuously updated when new user/content comes, user behavior changes or the embedding model is updated. Since most of these applications serve end users, they require high availability and durability. Some of our customers have deployed Manu in their production environment, and they found Manu satisfactory in terms of usability, performance, elasticity, and adaptability. In what follows, we simulate some typical application scenarios of our customers to demonstrate the advantages of Manu.

## 5.2 Example Use Cases

Due to business security, the names of the customers are anonymous. For the experiments, we use two datasets widely used for vector search research, i.e., SIFT [5] (with 128-dim vectors) and DEEP [2] (with 96-dim vectors), and extract sub-datasets with the required sizes. By default, we use two query nodes, one data node and one index node for Manu. Each worker node is an EC2 m5.4xlarge instance running on Amazon Linux AMI version 5.4.129. For index, we experiment with IVF-FLAT [45] and HNSW [60], which are widely used in practice. When comparing Manu with other systems, we always ensure that the systems use the same resource and are properly configured. Due to time and expense limits, we are only able to compare with some vector databases in a subset of the experiments. We search the top-50 most similar vectors for each query, and ensure that average search recall is above 0.8 if recall is not reported explicitly.




**Figure 9: Search workload, query latency, and number of query nodes used by Manu over time. Different colors indicate different number of query nodes are used.**

**E-commerce recommendation.** Company A is a leading online shopping platform in China that mainly sells clothing and make-ups. They use Manu for recommendation, and products are recommended to a user according to their similarity scores with the user embedding vector. They have three main requirements for vector database: (1) *high throughput* as they need to handle the requests of many concurrent costumers; (2) *high quality* search results for good recommendation effect; (3) *good elasticity* for low costs as their search requests have large fluctuations over time (peaks in evening but very low in midnight, very high at promotion events).


In Figure 8, we compare the recall-throughput performance of Manu with four popular open-source vector search systems (represented as System A, System B, System C and System D), when using a single node. We use Euclidean distance for SIFT and inner product for DEEP to test different similarity functions. As System A only supports its own graph-based index, System B and System D only supports the HNSW index [60], we have only a single curve for them in each plot. The results show that Manu consistently outperforms the baselines across different datasets and similarity functions. System C and System D achieve significantly lower query processing throughput than Manu at the same recall. This is because that System C’s complex aggregation procedure for search results introduces high overhead and System D is a disk-based solution. The performances of System A and System B are much better than System C and System D but still inferior compared with Manu. We conjecture this is because Manu has better implementations with optimizations for CPU cache and SIMD.

To test the elasticity of Manu, we use the search traffic of an e-commerce platform over one day period [16], which is plotted as the purple curve in Figure 9. The results show that search workload fluctuates violently over time, and the peak is much higher than the valley. We use SIFT100M as the dataset and Euclidean distance as the similarity function. Manu is configured to reduce query nodes by 0.5x when search latency is shorter than 100ms and add query nodes to 2x when search latency is over 150ms. The colors in Figure 9 indicate the number of query nodes used by Manu, which shows that Manu has good elasticity to adapt to query workload. The black line reports the search latency and shows that Manu can keep search latency within the target range via scaling.

**Video deduplication.** Company B is a video sharing website in Europe, on which users can upload videos and share with others.



**Figure 10: Scalability of Manu w.r.t. query nodes.**



**Figure 11: Scalability of Manu w.r.t. data volume.**

They find that there are many duplicate videos that result in high management costs and thus conduct deduplication before archiving the videos. They model a video as a set of its critical frames and encode each frame into a vector. They use vector search to find videos in the corpus that are most similar to a new video and conduct further checking on the shortlisted videos to determine if the new video is a duplicate. They also use vector search to find videos similar to those viewed by users for recommendation. They require vector DBMS to have *good scalability* with respect to both *data volume* and *computing resource* as their corpus grows quickly.

In Figure 10 and Figure 11, we test the scalability of Manu when changing the number of query nodes and the size of dataset, respectively. The results show that query processing throughput scales almost linearly with the number of query nodes and the reciprocal of dataset size. The observation is consistent for different datasets, indexes and similarity functions. This is because Manu uses segments to distribute search tasks among query nodes. With segment size fixed, each query node handles more segments when data volume increases, and fewer segments when the number of query nodes increases. Note that better scalability w.r.t. data volume can be achieved by configuring Manu to use larger segments when dataset size increases. This is because similarity search indexes usually have sub-linear complexity w.r.t. dataset size.

**Virus scan.** Company C is a world leading software security service provider and one of its main service is scanning viruses for smart phones. They have a virus base that continuously collects new viruses and develop specialized algorithms to map virus and user APK to vector embedding. To conduct a virus scan, they find viruses in their base that have embedding similar to the query APK and then compare the search results with the APK in more detail. They have two requirements for vector DBMS: (1) *short delay* for streaming




Figure 12: The relation between search latency and grace time, legends stand for the time tick interval.




Figure 13: Index construction time of Manu vs. data volume.

update as new viruses (vectors) are continuously added to their virus base and vector search needs to observe the latest viruses with a short delay; (2) *fast index building* as they frequently adjust their embedding algorithm to fix problems, which leads to update of the entire dataset and requires to rebuild index.

In Figure 12, we show the average delay of search requests for Manu. Recall that *grace time* (i.e.,  $\tau$ ) means that a search request must observe updates that happen time  $\tau$  before it, and is configurable by users. The legends correspond to different *time tick interval*, with which the loggers write time tick to WAL. The results show that search latency decreases quickly with grace time, and shorter time tick interval results in shorter search latency. This is because with longer grace time, search requests can tolerate longer update delay and are less likely to wait for updates. When the time tick interval reduces, each segment can confirm that all updates have been received more quickly, thus the search requests wait for a shorter time. In Figure 13, we report the index building time of Manu when changing data volume. The results show that index building time scales linearly with data volume. This is because Manu builds index for each segment and larger data volume leads to more segments.

## 6 RELATED WORK

**Vector search algorithms.** Vector search algorithms have a long research history, and most works focus on efficient approximate search on large-scale datasets. Existing algorithms can be roughly classified into four categories, i.e., *space partitioning tree* (SPT), *locality sensitive hashing* (LSH), *vector quantization* (VQ) and *proximity graph* (PG). SPT algorithms divide the space into areas, and

use tree structures to quickly narrow down search results to some areas [29, 57, 63, 70, 79]. LSH algorithms design hash functions such that similar vectors are hashed to the same bucket with high probability, and examples include [34, 35, 40, 42, 52, 55, 56, 59, 69, 89]. VQ algorithms compress vectors and accelerate similarity computation by quantizing the vectors using a set of vector codebooks, and well-known VQ algorithms include [22, 33, 41, 44, 90]. PG algorithms form a graph by connecting a vector with those most similar to it in the dataset, and conduct vector search by graph walk [32, 60, 72, 88]. Different algorithms have different trade-offs, e.g., LSH is cheap in indexing building but poor in result quality, VQ reduces memory and computation but also harms result quality, PG has high efficiency but requires large memory. Manu supports a comprehensive set of search algorithms such that users can trade off between different factors.

**Vector databases.** Vector data management solutions have gone through two stages of development. Solutions in the first stage are libraries (e.g., Facebook Faiss [45], Microsoft SPTAG [15], HNSWlib [60] and Annoy [1]) and plugins (e.g., ES plugins [6], Postgres plugins [11]) for vector search. They are insufficient for current applications as full-fledged management functionalities are required, e.g., distributed execution for scalability, online data update, and failure recovery. Two OLAP database systems, AnalyticDB-V [81] and PASE [84] support vector data by adding a table column to store them but lacks optimizations tailored for vector data.

The second stage solutions are full-fledged vector databases such as Vearch [51], Vespa [18], Weaviate [19], Vald [17], Qdrant [12], Pinecone [10], and our previous effort Milvus [80].<sup>6</sup> Vearch uses Faiss as the underlying search engine and adopts a three-layer aggregation procedure to conduct distributed search. Similarly, Vespa distributes data over nodes for scalability. A modified version of the HNSW algorithm is used to support online updates for vector data, and Vespa also allows attribute filtering during search and learning-based inference on search results (e.g., for re-ranking). Weaviate adopts a GraphQL interface and allows storing objects (e.g., texts, images), properties, and vectors. Users can directly import vectors or customize embedding models to map objects to vectors, and Weaviate can retrieve objects based on vector search results. Vald supports horizontal scalability by partitioning a vector dataset into segments and builds indexes without stopping search services. Qdrant is a single-machine vector search engine with extensive support for attribute filtering. It allows filtering with various data types and query conditions (e.g., string matching, numerical ranges, geo-locations), and uses a tailored optimizer to determine the filtering strategy. Note that Vespa, Weaviate and Vald only support proximity graph index.

We can observe that these vector databases focus on different functionalities, e.g., learning-based inference, embedding generation, object retrieval, and attribute filtering. Thus, we treat evolvability as first class priority when design Manu such that new functionalities can be easily introduced. Manu also differs from these vector databases in important perspectives. First, the log backbone of Manu provides time semantics and allows tunable consistency.

<sup>6</sup>Pinecone is offered as SaaS and closed source. Thus, we do not know its design details.

Second, Manu decomposes system functionalities with fine granularity and instantiates them as cloud services for performance and failure isolation, and thus is more suitable for cloud deployment. Third, Manu comes with more comprehensive optimizations for usability and performance, e.g., support various indexes, hardware tailored implementations, and GUI tools.

**Cloud native databases.** Many OLAP databases are designed to run on the cloud recently and examples include Redshift [37], BigQuery [61], Snowflake [28] and AnalyticDB [87]. Redshift is a data warehouse system offered as a service on Amazon Web Service and adopts a *shared-nothing* architecture. It scales by adding or removing EC2 instances, and data is redistributed in the granularity of columns. Snowflake uses a *shared-data* architecture by delegating data storage to Amazon S3. Compute nodes are stateless and fetch read-only copies of data for tasks, and thus can be easily scaled. For efficiency, high-performance local disk is used to cache hot data.

Aurora [77] and PolarDB Serverless [50] are two cloud native OLTP databases. Aurora uses a *shared-disk* architecture and proposes the “log is database” principle by pushing transaction processing down to the storage engine. It observes that the bottleneck of cloud-based platforms has shifted from computation and storage IO to network IO. Thus, it only persists redo log for transaction processing and commits transaction by processing log according to LSN. PolarDB Serverless adopts a *disaggregation* architecture, which uses high-speed RDMA network to decouple hardware resources (e.g., compute, memory and storage) as resource pools.

Our Manu follows the general design principles of cloud native databases to decouple the system functionalities at fine granularity for high elasticity, fast evolution and failure isolation. However, we also consider the unique design opportunities of vector databases to trade the simple data model and weak consistency requirement for performance, cost and flexibility. Specifically, complex transactions are not supported and the log backbone is utilized to support tunable consistency-performance trade-off. Moreover, vector search, index building and log archiving tasks are further decoupled as their workloads may have significant variations.

## 7 CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we introduce the designs of Manu as a cloud native vector database. To ensure that Manu suits vector data applications, we set ambitious design goals, which include good evolvability, tunable consistency, high elasticity, good efficiency, and etc. To meet these design goals, Manu trades the simple data model of vectors and weak consistency requirement of applications for performance, costs and flexibility. Specifically, Manu conducts fine-grained decoupling of the system functionalities for component-wise scaling and evolution, and uses the log backbone to connect the system components while providing time semantics and simplifying inter-component interaction. We also introduce important features such as high-level API, GUI tool, hardware optimizations, and complex search. We think Manu is still far from perfect and some of our future directions include:

- **Multi-way search:** Many applications jointly search multiple types of contents, e.g., vector and primary key, vector and text. The log system of Manu allows to add search engines for other

contents (e.g., primary key and text) as co-processors by subscribing to the log stream. We will explore how multiple search engines can interact efficiently and how to flexibly coordinate different search engines to meet application requirements.

- **Modularized algorithms:** We think vector search algorithms can be distilled into independent components, e.g., *compression* for memory reduction and efficient computation, *indexing* for limiting computation to a small portion of vectors, and *bucketing* for grouping similar vectors. Existing vector search algorithms only explore some combinations of techniques for different components. We will provide a unified framework for vector search such that users can flexibly combine different techniques according to their desired trade-off between cost and performance.
- **Hierarchical storage aware index:** Current vector search index assumes a single type of storage, e.g., GPU memory, main memory or disk. We will explore indexes that can jointly utilize all devices on the storage hierarchy. For example, most applications have some hot vectors (e.g., popular products in e-commerce) that are frequently accessed by search requests, which can be placed in fast storage. As a query accesses only a portion of the vectors and a node processes many concurrent queries, the storage swap latency may be hidden by pipelining.
- **Advanced hardware:** NVM [66] costs about one-third of DRAM for unit capacity but provides comparable read bandwidth and latency comparable, which makes it a good choice for replacing expensive DRAM when storing large datasets. RDMA [24, 46] significantly reduces the communication latency among nodes, and NVLink [65] directly connects GPUs with much larger bandwidth than PCIe. By exploiting these fast connections, we will explore indexes and search algorithms that jointly use multiple devices. We are also working with hardware vendors to apply FPGA and MLU for vector search and index building.
- **Embedding generation toolbox:** For better application level-integration, we plan to incorporate a application-oriented toolbox for generating embedding vectors. This toolbox would incorporate model fine-tuning in addition to providing a number of pre-trained models that can be used out-of-the-box, allowing for rapid prototyping.

## ACKNOWLEDGMENTS

Manu is a multi-year project open sourced by Zilliz. The development of Manu involves many engineers in its community. In particular, we thank Bingyi Sun, Weida Zhu, Yudong Cai, Yihua Mo, Xi Ge, Yihao Dai, Jiquan Long, Cai Zhang, Congqi Xia, Xuan Yang, Binbin Lv, Xiaoyun Liu, Wenxing Zhu, Yufen Zong, Jie Zeng, Shaoyue Chen, Jing Li, Zizhao Chen, Jialian Ji, Min Tian, Yan Wang and all the other contributors in the community for their contributions. We also thank Filip Haltmayer for proofreading the paper and valuable suggestions to improve paper quality.

## REFERENCES

- [1] 2021. Annoy: Approximate Nearest Neighbors Oh Yeah. <https://github.com/spotify/annoy>.
- [2] 2021. Benchmarks for Billion-Scale Similarity Search. <https://research.yandex.com/datasets/biganns>.

- [3] 2021. Billion-Scale Approximate Nearest Neighbor Search Challenge. <https://big-ann-benchmarks.com>.
- [4] 2021. binlog. <https://hevodata.com/learn/using-mysql-binlog/>.
- [5] 2021. Datasets for approximate nearest neighbor search. <http://corpus-texmex.irisa.fr/>.
- [6] 2021. ElasticSearch: Open Source, Distributed, RESTful Search Engine. <https://github.com/elastic/elasticsearch>.
- [7] 2021. etcd. <https://etcd.io/>.
- [8] 2021. MinIO. <https://min.io/>.
- [9] 2021. MySQL. <https://www.mysql.com/>.
- [10] 2021. Pinecone. <https://www.pinecone.io/>.
- [11] 2021. PostgreSQL: The World's Most Advanced Open Source Relational Database. <https://www.postgresql.org/>.
- [12] 2021. Qdrant. <https://qdrant.tech/>.
- [13] 2021. S3. <https://aws.amazon.com/cn/s3/>.
- [14] 2021. siri. <https://www.apple.com/siri/>.
- [15] 2021. SPTAG: A library for fast approximate nearest neighbor search. <https://github.com/microsoft/SPTAG>.
- [16] 2021. User Behavior Data from Taobao for Recommendation. <https://tianchi.aliyun.com/dataset/dataDetail?dataId=649>.
- [17] 2021. Vald. <https://github.com/vdaas/vald>.
- [18] 2021. Vespa. <https://vespa.ai/>.
- [19] 2021. Weaviate. <https://github.com/semi-technologies/weaviate>.
- [20] 2021. Xaoice. <https://en.wikipedia.org/wiki/Xaoice>.
- [21] Reza Akbarinia, Esther Pacitti, and Patrick Valduriez. 2007. Best Position Algorithms for Top-k Queries. In *Proceedings of the 33rd International Conference on Very Large Data Bases (Vienna, Austria)*. VLDB Endowment, 495–506.
- [22] Fabien André, Anne-Marie Kermarrec, and Nicolas Le Scouarnec. 2015. Cache Locality is Not Enough: High-Performance Nearest Neighbor Search with Product Quantization Fast Scan. *Proc. VLDB Endow.* 9, 4, 288–299.
- [23] Artem Babenko and Victor S. Lempitsky. 2015. The Inverted Multi-Index. *IEEE Trans. Pattern Anal. Mach. Intell.* 37, 6 (2015), 1247–1260. <https://doi.org/10.1109/TPAMI.2014.2361319>
- [24] Wei Cao, Yingqiang Zhang, Xinjun Yang, Feifei Li, Sheng Wang, Qingda Hu, Xuntao Cheng, Zongzhi Chen, Zhenjun Liu, Jing Fang, et al. 2021. PolarDB Serverless: A Cloud Native Database for Disaggregated Data Centers. In *Proceedings of the 2021 International Conference on Management of Data*. 2477–2489.
- [25] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache flink: Stream and batch processing in a single engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 36, 4 (2015).
- [26] Qi Chen, Bing Zhao, Haidong Wang, Mingqin Li, Chuanjie Liu, Zhiyong Zheng, Mao Yang, and Jingdong Wang. 2021. SPANN: Highly-efficient Billion-scale Approximate Nearest Neighborhood Search. *Advances in Neural Information Processing Systems* 34 (2021).
- [27] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [28] Benoît Dageville, Thierry Cruanes, Marcin Zukowski, Vadim Antonov, Artin Avanes, Jon Bock, Jonathan Claybaugh, Daniel Engovatov, Martin Hentschel, Jiansheng Huang, Allison W. Lee, Ashish Motivala, Abdul Q. Munir, Steven Pelley, Peter Povinec, Greg Rahn, Spyridon Triantafyllis, and Philipp Unterbrunner. 2016. The Snowflake Elastic Data Warehouse. In *Proceedings of the 2016 International Conference on Management of Data, SIGMOD Conference 2016, San Francisco, CA, USA, June 26 - July 01, 2016*. ACM, 215–226.
- [29] Sanjoy Dasgupta and Yoav Freund. 2008. Random projection trees and low dimensional manifolds. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*. 537–546.
- [30] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, 4171–4186. <https://doi.org/10.18653/v1/n19-1423>
- [31] Stefan Falkner, Aaron Klein, and Frank Hutter. 2017. Combining hyperband and bayesian optimization. In *NIPS 2017 Bayesian Optimization Workshop (Dec 2017)*.
- [32] Cong Fu, Chao Xiang, Changxu Wang, and Deng Cai. 2019. Fast approximate nearest neighbor search with the navigating spreading-out graph. *Proceedings of the VLDB Endowment* 12, 5 (2019), 461–474.
- [33] Tiezheng Ge, Kaiming He, Qifa Ke, and Jian Sun. 2013. Optimized product quantization for approximate nearest neighbor search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2946–2953.
- [34] Aristides Gionis, Piotr Indyk, Rajeve Motwani, et al. 1999. Similarity search in high dimensions via hashing. In *Vldb*, Vol. 99. 518–529.
- [35] Long Gong, Huayi Wang, Mitsunori Ogihara, and Jun Xu. 2020. iDEC: indexable distance estimating codes for approximate nearest neighbor search. *Proceedings of the VLDB Endowment* 13, 9 (2020).
- [36] Ruiqi Guo, Sanjiv Kumar, Krzysztof Choromanski, and David Simcha. 2016. Quantization based fast inner product search. In *Artificial Intelligence and Statistics*. PMLR, 482–490.
- [37] Anurag Gupta, Deepak Agarwal, Derek Tan, Jakub Kulesza, Rahul Pathak, Stefano Stefani, and Vidhya Srinivasan. 2015. Amazon Redshift and the Case for Simpler Data Warehouses. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*. ACM, 1917–1923.
- [38] Michael Hersovici, Michal Jacovi, Yoelle S Maarek, Dan Pelleg, Menachem Shtalham, and Sigalit Ur. 1998. The shark-search algorithm. an application: tailored web site mapping. *Computer Networks and ISDN Systems* 30, 1-7 (1998), 317–326.
- [39] Dongxu Huang, Qi Liu, Qiu Cui, Zhuhe Fang, Xiaoyu Ma, Fei Xu, Li Shen, Liu Tang, Yuxing Zhou, Menglong Huang, et al. 2020. TiDB: a Raft-based HTAP database. *Proceedings of the VLDB Endowment* 13, 12 (2020), 3072–3084.
- [40] Piotr Indyk and Rajeve Motwani. 1998. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*. 604–613.
- [41] Masajiro Iwasaki and Daisuke Miyazaki. 2018. Optimization of indexing based on k-nearest neighbor graph for proximity search in high-dimensional data. *arXiv preprint arXiv:1810.07355* (2018).
- [42] Omid Jafari, Parth Nagarkar, and Jonathan Montañó. 2020. mmlsh: A Practical and Efficient Technique for Processing Approximate Nearest Neighbor Queries on Multimedia Data. In *International Conference on Similarity Search and Applications*. Springer, 47–61.
- [43] Mohsen Jamali and Martin Ester. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In *Proceedings of the fourth ACM conference on Recommender systems*. 135–142.
- [44] Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence* 33, 1 (2010), 117–128.
- [45] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with gpus. *IEEE Transactions on Big Data* (2019).
- [46] Anuj Kalia, Michael Kaminsky, and David G Andersen. 2014. Using RDMA efficiently for key-value services. In *Proceedings of the 2014 ACM Conference on SIGCOMM*. 295–306.
- [47] Timothy King. 2019. 80 Percent of Your Data Will Be Unstructured in Five Years. <https://solutionsreview.com/data-management/80-percent-of-your-data-will-be-unstructured-in-five-years/>.
- [48] Yann LeCun, Yoshua Bengio, et al. 1995. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks* 3361, 10 (1995), 1995.
- [49] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436–444.
- [50] Feifei Li. 2019. Cloud native database systems at Alibaba: Opportunities and Challenges. *Proc. VLDB Endow.* 12, 12 (2019), 2263–2272.
- [51] Jie Li, Haifeng Liu, Chuanghua Gui, Jianyu Chen, Zhenyuan Ni, Ning Wang, and Yuan Chen. 2018. The design and implementation of a real time visual search system on jd e-commerce platform. In *Proceedings of the 19th International Middleware Conference Industry*. 9–16.
- [52] Mingjie Li, Ying Zhang, Yifang Sun, Wei Wang, Ivor W Tsang, and Xuemin Lin. 2020. I/O efficient approximate nearest neighbour search based on learned functions. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 289–300.
- [53] Wen Li, Ying Zhang, Yifang Sun, Wei Wang, Mingjie Li, Wenjie Zhang, and Xuemin Lin. 2019. Approximate nearest neighbor search on high dimensional data—experiments, analyses, and improvement. *IEEE Transactions on Knowledge and Data Engineering* 32, 8 (2019), 1475–1488.
- [54] Jiahui Liu, Peter Dolan, and Elin Rønby Pedersen. 2010. Personalized news recommendation based on click behavior. In *Proceedings of the 15th international conference on Intelligent user interfaces*. 31–40.
- [55] Wanqi Liu, Hanchen Wang, Ying Zhang, Wei Wang, and Lu Qin. 2019. I-LSH: I/O efficient c-approximate nearest neighbor search in high-dimensional space. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 1670–1673.
- [56] Kejing Lu and Mineichi Kudo. 2020. R2LSH: A Nearest Neighbor Search Scheme Based on Two-dimensional Projected Spaces. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 1045–1056.
- [57] Kejing Lu, Hongya Wang, Wei Wang, and Mineichi Kudo. 2020. VHP: approximate nearest neighbor search via virtual hypersphere partitioning. *Proceedings of the VLDB Endowment* 13, 9 (2020), 1443–1455.
- [58] Lailong Luo, Deke Guo, Richard TB Ma, Ori Rottenstreich, and Xueshan Luo. 2018. Optimizing bloom filter: Challenges, solutions, and comparisons. *IEEE Communications Surveys & Tutorials* 21, 2 (2018), 1912–1949.
- [59] Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. 2017. Intelligent probing for locality sensitive hashing: Multi-probe LSH and beyond. (2017).
- [60] Yu A Malkov and Dmitry A Yashunin. 2018. Efficient and robust approximate nearest neighbor search using hierarchical navigable small world graphs. *IEEE*



- transactions on pattern analysis and machine intelligence* 42, 4 (2018), 824–836.
- [61] Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. 2010. Dremel: Interactive Analysis of Web-Scale Datasets. *Proc. VLDB Endow.* 3, 1 (2010), 330–339.
- [62] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*. 3111–3119.
- [63] Marius Muja and David G Lowe. 2014. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence* 36, 11 (2014), 2227–2240.
- [64] Michael P Papazoglou and Willem-Jan Van Den Heuvel. 2006. Service-oriented design and development methodology. *International Journal of Web Engineering and Technology* 2, 4 (2006), 412–442.
- [65] Carl Pearson, I-Hsin Chung, Zehra Sura, Wen-Mei Hwu, and Jinjun Xiong. 2018. NUMA-aware data-transfer measurements for power/NVLink multi-GPU systems. In *International Conference on High Performance Computing*. Springer, 448–454.
- [66] Jie Ren, Minjia Zhang, and Dong Li. 2020. Hm-ann: Efficient billion-point nearest neighbor search on heterogeneous memory. *Advances in Neural Information Processing Systems* 33 (2020).
- [67] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. 2008. The graph neural network model. *IEEE transactions on neural networks* 20, 1 (2008), 61–80.
- [68] Falk Scholer, Hugh E Williams, John Yiannis, and Justin Zobel. 2002. Compression of inverted indexes for fast query evaluation. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*. 222–229.
- [69] Anshumali Shrivastava and Ping Li. 2014. Asymmetric LSH (ALSH) for sublinear time maximum inner product search (MIPS). *Advances in neural information processing systems* 27 (2014).
- [70] Chanop Silpa-Anan and Richard Hartley. 2008. Optimised KD-trees for fast image descriptor matching. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 1–8.
- [71] Harsha Vardhan Simhadri, George Williams, Martin Aumüller, Matthijs Douze, Artem Babenko, Dmitry Baranchuk, Qi Chen, Lucas Hosseini, Ravishankar Krishnaswamy, Gopal Srinivasa, et al. 2022. Results of the NeurIPS’21 Challenge on Billion-Scale Approximate Nearest Neighbor Search. *arXiv preprint arXiv:2205.03763* (2022).
- [72] Suhas Jayaram Subramanya, Fnu Devvrit, Harsha Vardhan Simhadri, Ravishankar Krishnaswamy, and Rohan Kadekodi. 2019. Rand-NSG: Fast Accurate Billion-point Nearest Neighbor Search on a Single Node. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. 13748–13758.
- [73] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis, Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, et al. 2020. Cockroachdb: The resilient geo-distributed sql database. In *Proceedings of the 2020 ACM SIGMOD International Conference on Management of Data*. 1493–1509.
- [74] Gábor Takács and Domonkos Tikk. 2012. Alternating least squares for personalized ranking. In *Proceedings of the sixth ACM conference on Recommender systems*. 83–90.
- [75] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghotham Murthy. 2009. Hive: a warehousing solution over a map-reduce framework. *Proceedings of the VLDB Endowment* 2, 2 (2009), 1626–1629.
- [76] Aäron Van Den Oord, Sander Dieleman, and Benjamin Schrauwen. 2013. Deep content-based music recommendation. In *Neural Information Processing Systems Conference (NIPS 2013)*, Vol. 26. Neural Information Processing Systems Foundation (NIPS).
- [77] Alexandre Verbitski, Anurag Gupta, Debanjan Saha, Murali Brahmadesam, Kamal Gupta, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, Tengiz Kharatishvili, and Xiaofeng Bao. 2017. Amazon Aurora: Design Considerations for High Throughput Cloud-Native Relational Databases. In *Proceedings of the 2017 ACM International Conference on Management of Data, SIGMOD Conference 2017, Chicago, IL, USA, May 14-19, 2017*. ACM, 1041–1052.
- [78] Jizhe Wang, Pipei Huang, Huan Zhao, Zhibo Zhang, Binqiang Zhao, and Dik Lun Lee. 2018. Billion-scale commodity embedding for e-commerce recommendation in alibaba. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 839–848.
- [79] Jingdong Wang, Naiyan Wang, You Jia, Jian Li, Gang Zeng, Hongbin Zha, and Xian-Sheng Hua. 2014. Trinary-Projection Trees for Approximate Nearest Neighbor Search. *IEEE Trans. Pattern Anal. Mach. Intell.* 36, 2 (2014), 388–403.
- [80] Jianguo Wang, Xiaomeng Yi, Rentong Guo, Hai Jin, Peng Xu, Shengjun Li, Xi-angyu Wang, Xiangzhou Guo, Chengming Li, Xiaohai Xu, et al. 2021. Milvus: A Purpose-Built Vector Data Management System. In *Proceedings of the 2021 International Conference on Management of Data*. 2614–2627.
- [81] Chuangxian Wei, Bin Wu, Sheng Wang, Renjie Lou, Chaoqun Zhan, Feifei Li, and Yuanzhe Cai. 2020. AnalyticDB-V: A Hybrid Analytical Engine Towards Query Fusion for Structured and Unstructured Data. *Proc. VLDB Endow.* 13, 12 (2020), 3152–3165.
- [82] Xiang Wu, Ruiqi Guo, Ananda Theertha Suresh, Sanjiv Kumar, Daniel N Holtmann-Rice, David Simcha, and Felix Yu. 2017. Multiscale quantization for fast similarity search. *Advances in Neural Information Processing Systems* 30 (2017), 5745–5755.
- [83] SHI Xingjian, Zhouong Chen, Hao Wang, Dit-Yan Yeung, Wai-Kin Wong, and Wang-chun Woo. 2015. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *Advances in neural information processing systems*. 802–810.
- [84] Wen Yang, Tao Li, Gai Fang, and Hong Wei. 2020. PASE: PostgreSQL Ultra-High-Dimensional Approximate Nearest Neighbor Search Extension. In *Proceedings of the 2020 International Conference on Management of Data, SIGMOD Conference 2020, online conference [Portland, OR, USA], June 14-19, 2020*. ACM, 2241–2253.
- [85] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, Ion Stoica, et al. 2010. Spark: Cluster computing with working sets. *HotCloud* 10, 10-10 (2010), 95.
- [86] Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. Recurrent neural network regularization. *arXiv preprint arXiv:1409.2329* (2014).
- [87] Chaoqun Zhan, Maomeng Su, Chuangxian Wei, Xiaoqiang Peng, Liang Lin, Sheng Wang, Zhe Chen, Feifei Li, Yue Pan, Fang Zheng, et al. 2019. Analyticdb: Real-time olap database system at alibaba cloud. *Proceedings of the VLDB Endowment* 12, 12 (2019), 2059–2070.
- [88] Weijie Zhao, Shulong Tan, and Ping Li. 2020. SONG: Approximate Nearest Neighbor Search on GPU. In *36th IEEE International Conference on Data Engineering, ICDE 2020, Dallas, TX, USA, April 20-24, 2020*. IEEE, 1033–1044.
- [89] Bolong Zheng, Zhao Xi, Lianggui Weng, Nguyen Quoc Viet Hung, Hang Liu, and Christian S Jensen. 2020. PM-LSH: A fast and accurate LSH framework for high-dimensional approximate NN search. *Proceedings of the VLDB Endowment* 13, 5 (2020), 643–655.
- [90] Wengang Zhou, Yijuan Lu, Houqiang Li, and Qi Tian. 2012. Scalar quantization for large scale image search. In *Proceedings of the 20th ACM international conference on Multimedia*. 169–178.