

HARNESSING THE POWER OF VECTOR DATABASES: A NEW PARADIGM FOR VISUAL SEARCH SOLUTIONS

G.K. RAMANATHAN

Associate Specialist, BI Worldwide India Private Limited, Chennai, Tamil Nadu, India

Abstract: New trends in e-commerce play a significant part in the growth of technology through the internet, and the availability of modern devices and their sophisticated functions have triggered an increase in use for many. Choosing the proper product can be challenging due to the vast array of products showcased on websites, leaving customers feeling tired. These circumstances increase the rivalry between global commercial sites, which builds the need to work proficiently to increase financial profits. Simplifying the user experience is the goal of innovative technology. Visual search is the next innovation that will decouple users from the reliance on keyboards and open a new world of opportunities. Discover how the advancements in visual search technology are influencing future internet search strategies. Human and animal brains can easily detect objects, but computers struggle with this task. The latest technology is being developed, and we are witnessing these revolutionary innovations making our lives easier.

As visual content continues to dominate digital platforms, traditional search methods struggle to deliver accurate and intuitive results. This paper explores the transformative potential of vector databases in enabling efficient and intelligent visual search solutions. By converting images into high-dimensional vectors, these databases allow for similarity-based retrieval far beyond keyword matching. Leveraging machine learning and deep neural networks, visual features are encoded and compared using vector embeddings. This new paradigm not only enhances search relevance and speed but also supports scalable and real-time applications across eCommerce, healthcare, and media. The study highlights architecture, use cases, and performance benchmarks of vector-based systems.

Keywords: Visual Search, Vector Databases, eCommerce Innovation, Machine Learning, Deep Neural Networks

I. OBJECTIVES

- ✓ To study the vector database management systems (VDBMS)
- ✓ To reveal the issues in implementing vector database management systems
- ✓ To highlights the benefits of vector database management systems
- ✓ To know how vector database management systems behave in real world applications.
- ✓ To develop a vector database management system to solve the problems in the existing ecommerce system
- ✓ To create awareness of Win-Win Strategy

About Vector database management systems (VDBMS)

The vector database management systems are expected to further develop the e-commerce system performance by facilitating the clients to track down the suitable items as indicated by their search. Vector database management systems improve and increase the accuracy of products searched by the customers.

For example, if you want to look for a dress but don't know what brand or company it's from, you can't use keywords. In this case, you can just upload a photograph of the dress and retrieve all related facts and relevant results using our visual search technology.

Implementing vector database management systems (VDBMS) will result in a high conversion rate. Contextual relevance significantly enhances user experience by ensuring that the results returned by a vector database management system are not only accurate but also aligned with user intent and situational factors. This leads to more engaging interactions and increased satisfaction with the search process.

History and Evolution of Computer Vision

The first computer vision experiments took place in the 1950s, employing neural networks to recognize an object's edges and classify simple shapes like squares and circles.

Later in the 1970s, the commercial application of computer vision was introduced. It was the process of interpreting handwritten text with optical character recognition (OCR). This execution was used to understand written material for visually impaired individuals.


Facial recognition programs developed in the 1990s, as the Internet evolved. Later, in 2010 (and beyond), deep learning enabled computers to train and improve themselves over time.

Today, this technology is used in a variety of industries, including automotive, healthcare, retail, and smartphones. The AI in computer vision market was of worth USD 1.5 billion in 2023 and is expected to reach USD 4.3 billion by 2028, with a CAGR of 23.3%.

45% of retailers in the UK now use visual search.

II. REVIEW OF LITERATURE

Datta, R., Joshi, D., Li, J., & Wang, J. Z. (2008). Provides an extensive overview of image retrieval systems and highlights the shift from keyword-based to content-based image search, which laid the foundation for visual search technologies. **Johnson, J., Douze, M., & Jégou, H.** (2019). Discusses scalable similarity search using high-dimensional vectors with GPU acceleration, relevant for understanding how vector databases like FAISS enable high-performance visual search. **Krizhevsky, A., Sutskever, I., & Hinton, G. E.** (2012). This landmark paper introduced convolutional neural networks (CNNs) for image recognition and feature extraction, forming the backbone of AI-powered visual search technologies. **Zhou, J., Li, L., & Wang, S.** (2022).: Reviews the structure, indexing methods, and scalability of modern vector databases such as Milvus, FAISS, and Annoy, essential tools for implementing real-time visual search. **Bell, S., Upchurch, P., Snavely, N., & Bala, K.** (2015). Highlights how visual search and recognition technologies are used in real-world applications such as product identification in eCommerce, demonstrating their commercial value.



III. INTRODUCTION

Buyers enjoy a style but are unsure how to characterize it in a keyword search. Customers can order matching products on the moment using visual search technology, which finds exactly what they're seeking for based on an image. This technology helps clients find the things they desire by searching by image or photo rather than a keyword. Visual search helps to connect clients with similar products, increasing their buying potential and cross-selling.

In the digital age, the surge in visual content has transformed how users interact with information online. Traditional keyword-based search methods often fall short in accurately identifying and retrieving image-based content. Vector databases offer a groundbreaking solution by converting visual data into high-dimensional vectors, enabling efficient and accurate similarity-based searches. Combined with machine learning and deep neural networks, these systems can analyze, compare, and retrieve images with unprecedented precision. This new paradigm is revolutionizing search capabilities across sectors like eCommerce, healthcare, and digital media. This paper explores the fundamentals, applications, and future impact of vector databases in visual search technologies.

Suggestions for similar or equally relevant product recommendations to reduce bounce rates, enhance basket size, and boost sales.

Currently data contains

- ✓ Structured data
- ✓ Semi Structured data
- ✓ Unstructured data

In future 80% of the data will only be unstructured data.

How to analysis unstructured data?

Unstructured data cannot be stored in the pre-defined format. Not fit for the existing model.

Example of unstructured data


- ✓ Image
- ✓ Video
- ✓ Audio
- ✓ Text files

How can we search data in unstructured data?

Need database for the AI. VDBMS is a special type of database.

Uses of vector database (VDBMS)

VECTOR databases store data in high-dimensional vectors, which are mathematical representations of attributes.



Vector databases offer various advantages over regular databases. Quick and accurate similarity search and retrieval. Vector databases can identify the most similar or relevant data based on vector distance or similarity, which is a critical feature for many applications such as natural language processing, computer vision, recommendation systems, and so on.

Why are vector databases important?

Vectors are generated using an embedding model, which accepts unstructured data as input and returns a vector comprising a multidimensional array of numbers that captures the semantics, context, and meaning of the incoming data object.

Vectors are numerical representations of data items, also called vector embeddings.

Who uses vector databases?

To build embeddings and populate a vector database, an application developer can leverage open-source models, automated machine learning (ML) tools, and foundational model services. This requires only basic machine learning skills.

A team of data scientists and engineers can create expertly tuned embeddings and implement them using a vector database. This allows them to deploy artificial intelligence (AI) solutions faster.

Categorization based on data type

- ✓ Text vector databases
- ✓ Image vector databases
- ✓ Multimedia vector databases
- ✓ Graph vector databases

Categorization based on indexing technique

- ✓ Tree-based indexing databases
- ✓ Hashing-based indexing databases
- ✓ Quantization-based databases

How Vector Databases Work

- ✓ Data Ingestion
- ✓ Indexing
- ✓ Querying


Index Types

IVFFLAT

- Divides vectors into **lists**
- **Faster build** times
- Uses **less memory**
- **Lower query performance** (speed-recall tradeoff)
- Create index **after the table has some data**


HNSW

- Creates a **multilayer graph**
- **Slower build** times
- Uses **more memory**
- **Better query performance**
- Index can be created **without any data in the table** (no training step)



Use-cases


- ✓ Similarity search in general
- ✓ image and video similarity search
- ✓ Voice recognition
- ✓ Recommendation Systems
- ✓ Fraud Detection
- ✓ Chatbots and long-term memory



Vector search

- K-nearest neighbors (KNN)
- Approximate nearest neighbors (ANN)

Vector databases




PgVector is an open-source vector database that extends PostgreSQL, specifically addressing vector similarity search. Supporting exact and approximate nearest neighbor search, PgVector excels with its comprehensive support for L2 distance, inner product, and cosine distance.

Feature	pgvector
Database type	Extension for PostgreSQL
Deployment	Self-hosted
Scalability	Limited by PostgreSQL
Integration	Works with existing PostgreSQL stack
Cost	Free, open-source


Cosine Similarity

Cosine similarity measures the cosine of the angle between two vector embeddings, and it's often used as a distance metric in text analysis and other domains where the magnitude of the vector is less important than the direction.




Euclidean Distance

Euclidean distance measures the straight-line distance between two points in Euclidean space.



Manhattan Distance

Manhattan distance sums the absolute differences of their coordinates.



Augmented functionality to enrich visual search experiences

The ability to crop or zoom in on specific things inside an image enhances visual search capabilities. Customers frequently draw inspiration from photos that include numerous products. Customers can easily find what they want by isolating a certain item, such as a shirt, footwear, or accessory, and searching for similar-looking products. Sell more today!

IV. CONCLUSION

Vector database is a growing data model intended for storing vectors which describe rich data in high-dimensional vectors. This study provided an overview of fundamental concepts behind vector databases and vector database management systems, such as different types of vector similarity comparison types, different vector index types, and the principal software components in a VDBMS. The integration of vector databases into visual search technology marks a significant shift in how digital content is retrieved and experienced. Unlike traditional keyword-based systems, vector-based search leverages high-dimensional embeddings and machine learning models to understand and compare visual data with greater accuracy and efficiency. This approach not only enhances the relevance of search results but also dramatically improves user experience by providing intuitive and intelligent content discovery. As industries like eCommerce, healthcare, and digital media increasingly rely on image-centric platforms, the demand for scalable and real-time visual search solutions will continue to rise. Vector databases meet this demand by supporting rapid image retrieval, high concurrency, and seamless integration with AI technologies. With continued advancements in deep learning and neural networks, visual search powered by vector databases is poised to become a core component of next-generation search engines, fundamentally reshaping how users interact with digital information in a visually driven world.

REFERENCES

- [1]. <https://www.oracle.com/a/ocom/docs/database/cube-research-oracle-db-vector-search.pdf>
- [2]. <https://arxiv.org/pdf/2309.11322>
- [3]. <https://arxiv.org/pdf/2310.11703>
- [4]. https://www.cs.purdue.edu/homes/csjpgwang/pubs/SIGMOD24_VecDB_Tutorial.pdf
- [5]. https://cs.purdue.edu/homes/csjpgwang/pubs/VLDB24_SingleStoreVec.pdf
- [6]. <https://projector-video-pdf-converter.datacamp.com/33674/chapter3.pdf>
- [7]. https://www.cs.ru.nl/bachelors-theses/2024/Jaap_Aarts___1056714___Implementing_in-Database_Similarity_Search.pdf
- [8]. Bell, S., Upchurch, P., Snavely, N., & Bala, K. (2015). Material recognition in the wild with the Materials in Context Database (MINC). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 3479–3487). <https://doi.org/10.1109/CVPR.2015.7298948>
- [9]. Datta, R., Joshi, D., Li, J., & Wang, J. Z. (2008). Image retrieval: Ideas, influences, and trends of the new age. *ACM Computing Surveys*, 40(2), 1–60. <https://doi.org/10.1145/1348246.1348248>
- [10]. Zhou, J., Li, L., & Wang, S. (2022). Towards efficient similarity search in vector databases: Techniques and trends. *Information Systems*, 104, 101868. <https://doi.org/10.1016/j.is.2021.101868>
- [11]. Johnson, J., Douze, M., & Jégou, H. (2019). Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3), 535–547. <https://doi.org/10.1109/TBDATA.2019.2921572>
- [12]. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (pp. 1097–1105). <https://doi.org/10.1145/3065386>