

# AI Powered Knowledge Graph Generator for Video Learning Using RAG

B Balaji<sup>1</sup>, M Venkata Sasidhar Kaushik<sup>2</sup>, N Raja Niketh Reddy<sup>3</sup>, V Lakshmi Gayathri<sup>4</sup>, B Siva Prasad<sup>5</sup>

<sup>1,2,3,4,5</sup>Department of Computer Science and Engineering, Raghu Engineering College, Visakhapatnam, India.

\*\*\*

**Abstract** - The rapid growth of educational video content has transformed modern learning, yet it introduces challenges such as knowledge fragmentation, poor retention, and inefficient retrieval. Existing tools such as video summarizers and note-taking systems fail to provide a persistent and interconnected learning experience. This paper presents an AI-powered knowledge graph generation system that transforms passive video consumption into an active and queryable knowledge base. The proposed system processes YouTube videos by extracting transcripts and performing multi-dimensional analysis using large language models, generating structured outputs such as summaries, quizzes, flashcards, and key insights. To enable efficient retrieval, the system encodes content into high-dimensional semantic embeddings stored in a vector database. A Retrieval-Augmented Generation (RAG) framework allows users to interact with stored knowledge through natural language queries, providing context-aware and source-grounded responses. Additionally, an automated knowledge graph captures relationships between concepts across multiple videos, enabling visual exploration of interconnected knowledge.

Experimental evaluation demonstrates improved knowledge retention and retrieval efficiency compared to traditional approaches. The proposed system provides a scalable and intelligent solution for long-term learning from video-based content.

**Key Words:** Artificial Intelligence, Knowledge Graph, Natural Language Processing, Named Entity Recognition, Relation Extraction, Retrieval-Augmented Generation, Graph Databases, Semantic Analysis

## 1 Introduction

The proliferation of digital learning platforms has significantly transformed the way individuals acquire knowledge in the modern era. Among these platforms, video-based learning—particularly through services such as YouTube—has emerged as one of the most dominant and accessible modes of education. Learners

across the globe increasingly rely on video content for understanding complex concepts, ranging from software engineering and data science to finance, medicine, and design. The flexibility, visual engagement, and on-demand nature of video learning have made it a cornerstone of self-directed education. Despite its widespread adoption, video-based learning suffers from fundamental limitations that hinder its effectiveness as a long-term knowledge acquisition medium. One of the primary challenges is the passive nature of content consumption. Unlike structured educational systems, where learners actively engage through problem-solving and revision, video consumption often lacks mechanisms that promote retention and recall. As a result, learners frequently experience rapid knowledge decay, commonly referred to as the “forgetting curve,” where information is lost shortly after initial exposure.

In addition to poor retention, knowledge fragmentation presents a significant obstacle. Learners typically consume content from multiple videos, often across different channels and timeframes. However, existing platforms do not provide mechanisms to connect concepts across these videos. For instance, a learner studying distributed systems may watch separate videos on caching, load balancing, and database consistency, but lacks tools to understand how these concepts interrelate. This fragmentation leads to isolated pockets of knowledge rather than a cohesive conceptual understanding. Another critical limitation is the inefficiency of knowledge retrieval. While platforms provide basic search capabilities within individual videos, they do not support semantic, cross-video querying. Consequently, users often spend considerable time re-watching content to locate specific information previously encountered. Traditional solutions such as bookmarks, timestamps, or manual notes are insufficient, as they rely heavily on user effort and lack scalability for large volumes of content. Existing tools, including video summarizers, note-taking applications, and browser extensions, attempt to address some of these challenges. However, they remain fundamentally limited. Most summarization tools provide only surface-level insights and do not preserve the deeper semantic

structure of the content. Note-taking applications depend on manual input, making them inconsistent and time-consuming. Similarly, standalone AI chat systems may allow querying of content but often lack persistent memory, structured storage, and reliable grounding mechanisms, leading to incomplete or hallucinated responses.

Recent advancements in artificial intelligence have opened new avenues for addressing these challenges. Large Language Models (LLMs) have demonstrated remarkable capabilities in natural language understanding and generation, enabling automated extraction of structured knowledge from unstructured text. Additionally, vector embeddings allow high-dimensional semantic representation of textual data, facilitating efficient similarity-based retrieval. Retrieval-Augmented Generation (RAG) systems further enhance these capabilities by combining retrieval mechanisms with generative models to produce context-aware and grounded responses. Furthermore, knowledge graphs provide a powerful framework for representing entities and relationships, enabling structured exploration of interconnected information.

Motivated by these advancements, this paper proposes an AI-powered knowledge graph generation system designed to transform passive video consumption into an active, persistent, and interconnected learning process. The system processes educational videos by extracting transcripts and performing multi-dimensional analysis using large language models. This analysis generates structured outputs, including summaries, key insights, quizzes, flashcards, action items, glossary terms, and misconceptions, thereby supporting multiple learning modalities.

## 2 Related Work

The rapid growth of digital learning platforms has led to extensive research in areas such as video summarization, knowledge management systems, semantic search, and intelligent tutoring systems. While significant progress has been made, existing approaches exhibit limitations in terms of persistence, cross-content reasoning, and structured knowledge representation. Early work in video-based learning primarily focused on video summarization techniques. These systems aim to condense long video content into shorter textual or visual summaries using natural language processing and machine learning methods. Extractive summarization approaches select key sentences from

transcripts, while abstractive methods generate concise summaries using deep learning models. Although these methods improve accessibility, they are inherently limited to one-time content reduction and do not support long-term knowledge retention or interaction. Furthermore, they fail to capture relationships between concepts across multiple videos. Note-taking and knowledge management systems, such as digital notebooks and bookmarking tools, have also been widely adopted to support learning from multimedia content. These systems allow users to manually store and organize information; however, they rely heavily on user effort and consistency. As the volume of content increases, manual note-taking becomes inefficient and unsalable. Additionally, such systems lack semantic understanding, making it difficult to retrieve information using natural language queries or to establish connections between different pieces of knowledge.

Recent advancements in artificial intelligence have led to the emergence of AI-powered assistants and conversational agents for learning. Large Language Models (LLMs) have been used to build systems capable of answering questions, summarizing documents, and generating explanations. While these systems provide powerful interaction capabilities, they often suffer from a lack of grounding and persistence. Many standalone AI chatbots do not retain user-specific knowledge over time, and their responses may rely on parametric memory rather than user-provided data, leading to potential inaccuracies or hallucinations. To address the limitations of standalone LLMs, Retrieval-Augmented Generation (RAG) frameworks have been proposed. RAG systems combine information retrieval techniques with generative models to produce context-aware responses based on external data sources. These systems typically use vector embeddings to represent textual data in high-dimensional space and perform similarity-based retrieval. While RAG significantly improves response accuracy and grounding, most existing implementations are designed for document-based corpora and do not specifically address the challenges of video-based learning, such as transcript processing, timestamp alignment, and multi-modal knowledge extraction.

Another relevant area of research is semantic search and vector databases. Technologies such as dense vector embeddings and approximate nearest neighbor (ANN) search algorithms have enabled efficient retrieval of semantically similar content. Systems leveraging vector

databases allow users to query information using natural language rather than exact keyword matching. However, these systems primarily focus on retrieval and do not incorporate higher-level knowledge structuring or visualization mechanisms.

Knowledge graphs have been extensively studied as a means of representing structured relationships between entities. They are widely used in domains such as search engines, recommendation systems, and enterprise knowledge management. Knowledge graphs enable the modeling of complex relationships and support reasoning over interconnected data. Despite their advantages, most existing knowledge graph systems require manual curation or are limited to specific domains. Automated knowledge graph construction from unstructured video content remains a relatively underexplored area. In the context of educational technology, some platforms have attempted to integrate multiple features such as summarization, quizzes, and flashcards to enhance learning. However, these systems are typically limited to isolated functionalities and do not provide a unified framework that combines semantic memory, interactive querying, and graph-based knowledge representation. Moreover, they often lack scalability and fail to support cross-content reasoning across large collections of learning materials.

The proposed system distinguishes itself from existing work by integrating multiple advanced technologies into a single cohesive framework. Unlike traditional summarization tools, it generates multi-dimensional knowledge artefacts that support diverse learning modalities. In contrast to note-taking systems, it provides automated semantic storage and retrieval using high-dimensional embeddings. Compared to standalone RAG systems, it incorporates persistent user-specific memory and timestamp-grounded responses. Furthermore, the integration of an automated knowledge graph enables cross-video conceptual linking, addressing the problem of knowledge fragmentation. By combining video processing, semantic embeddings, retrieval-augmented generation, and knowledge graph construction, the proposed approach provides a comprehensive solution for transforming passive video consumption into an active and interconnected learning experience. This holistic integration represents a significant advancement over existing approaches and highlights the novelty and contribution of the proposed system.

### 3 Proposed Methodology

This section presents the design and implementation of the proposed AI-powered knowledge graph generation system. The system is designed to transform unstructured video content into a structured, queryable, and interconnected knowledge base. It achieves this through a combination of multi-dimensional analysis, semantic embedding, retrieval-augmented generation, and knowledge graph construction.

#### A. System Overview

The proposed system follows a three-tier architecture consisting of a frontend layer, a backend processing layer, and a database layer. The frontend is responsible for user interaction and visualization, providing an intuitive interface for input and graph exploration. The backend handles computationally intensive AI tasks such as natural language processing, entity extraction, and relationship identification. The database layer manages persistent storage, including structured relational data as well as vector embeddings for efficient semantic retrieval.

This modular architecture ensures scalability, maintainability, and efficient data processing, allowing the system to seamlessly integrate user inputs with intelligent knowledge graph generation and visualization.




Figure1: System Architecture Diagram

## B. Video Processing Pipeline

The video processing pipeline begins when a learner submits a YouTube URL. The system validates the source, retrieves the transcript, and forwards normalized text to a multi-dimensional analysis layer for downstream processing. After ingestion, the transcript is analyzed through parallel Large Language Model (LLM) tasks. This concurrent workflow produces structured learning artifacts, including summaries, key insights, quizzes, flashcards, action items, glossary entries, and misconception corrections.

| Analysis Type  | Output           | Purpose             |
|----------------|------------------|---------------------|
| Summary        | Short + Detailed | Quick understanding |
| Key Insights   | Bullet points    | Core learning       |
| Quizzes        | MCQs             | Active recall       |
| Flashcards     | Q/A pairs        | Memory retention    |
| Action Items   | Steps            | Practical use       |
| Glossary       | Terms            | Vocabulary          |
| Misconceptions | Corrections      | Concept clarity     |

Table1: Multi Dimensional Analysis Outline

Parallel execution reduces end-to-end latency while improving compute utilization, enabling responsive interaction even for long-form videos. Table 1 presents the core outputs and their instructional purpose within the learning workflow.




Figure2: Video Processing Pipeline

## C. Semantic Embedding and Vector Storage

To enable efficient retrieval and long-term memory, the system converts textual content into high-dimensional semantic embeddings. The transcript is divided into smaller chunks, each of which is encoded into a vector representation. Similarity between a query and stored content is computed using cosine similarity:

$$\text{sim}(\mathbf{q}, \mathbf{d}) = \mathbf{q} \cdot \mathbf{d} / (\|\mathbf{q}\| \cdot \|\mathbf{d}\|)$$

where  $\mathbf{q}$  represents the query embedding and  $\mathbf{d}$  represents the document embedding. This formulation allows the system to retrieve semantically relevant information even when exact keyword matches are absent. The embeddings are stored in a vector-enabled database using efficient indexing techniques, enabling sub-millisecond retrieval performance. This semantic memory layer forms the foundation for intelligent querying and knowledge persistence.

## D. Retrieval Augmented Generate Framework

The system employs a Retrieval-Augmented Generation (RAG) framework to provide context-aware responses to user queries. When a user submits a query, the system first retrieves relevant chunks from the vector database and then generates a response using a language model conditioned on the retrieved context.

The RAG process can be formalized as:

$$y = f(q, R(q))$$

where  $q$  is the input query,  $R(q)$  represents the retrieved context, and  $y$  is the generated response.




Figure3: RAG Pipeline Diagram

This approach ensures that generated responses are grounded in actual data rather than relying solely on the model's internal knowledge. Additionally, the system provides timestamp-based citations, improving transparency and reliability.

### E. Knowledge Graph Construction

To address the problem of knowledge fragmentation, the system constructs a knowledge graph that represents relationships between concepts extracted from video content. The graph is defined as:

$$G = (V, E)$$

where  $V$  represents the set of nodes (concepts such as topics, technologies, or entities) and  $E$  represents the relationships between them.

The knowledge graph is generated automatically using LLM-based extraction techniques. Concepts are identified from transcripts and linked based on semantic relationships. These relationships are aggregated across multiple videos to form a unified graph.




Figure4: Knowledge Graph Visualization

The resulting graph enables users to explore connections between concepts visually, facilitating deeper understanding and cross-domain learning.

### F. Asynchronous Processing Design

The system leverages asynchronous processing to handle computationally intensive tasks efficiently. Operations such as LLM analysis and embedding generation are executed in parallel using asynchronous programming techniques. This design ensures that the system remains responsive to user

interactions while processing tasks in the background. It also improves scalability by allowing multiple requests to be handled concurrently.

### G. Data Flow and Interaction

The interaction between system components follows a structured data flow. User actions trigger backend processes, which perform analysis and store results in the database. The frontend periodically updates the interface to reflect the latest state. This workflow ensures seamless integration between components and enables real-time interaction with the knowledge base.

### H. Summary of Methodology

The proposed methodology integrates multiple advanced techniques into a unified framework. By combining multi-dimensional analysis, semantic embeddings, retrieval-augmented generation, and knowledge graph construction, the system provides a comprehensive solution for intelligent knowledge management.

The modular design of the system allows for easy extensibility and future enhancements, making it adaptable to evolving requirements in AI-powered learning systems.

## 4 Implementation and System Design

This section describes the practical implementation of the proposed system, including the technology stack, system components, and integration strategies. The system is designed as a scalable, modular, and production-ready architecture that efficiently handles AI-driven processing and user interactions.

### A. Technology Stack

The implementation leverages a modern full-stack architecture combining frontend, backend, and database technologies. Each component is selected to ensure scalability, performance, and ease of integration.

| Layer    | Technology            | Purpose                 |
|----------|-----------------------|-------------------------|
| Frontend | Next.js               | UI + SSR                |
| Backend  | FastAPI               | Async processing        |
| DB       | PostgreSQL + pgvector | Storage + vector search |
| AI       | LLM Embeddings +      | Analysis retrieval +    |

Table2: Technology Stack

The frontend is developed using Next.js, enabling server-side rendering and efficient client-side interactions. The backend is implemented using FastAPI, which provides asynchronous request

handling and high-performance API execution. The database layer utilizes PostgreSQL extended with vector capabilities to support both relational data storage and semantic search operations.

## B. Frontend Design

The frontend layer is responsible for providing an interactive and user-friendly interface. It is designed to support multiple functionalities, including video submission, visualization of analysis results, querying, and knowledge graph exploration.




Figure5: Frontend Interface Overview

The user interface allows users to input YouTube URLs, view generated summaries and insights, attempt quizzes, and interact with flashcards. Additionally, a chat interface is integrated to enable natural language querying of stored knowledge using the RAG framework. The frontend communicates with the backend through API endpoints and server actions, ensuring

## C. Backend Architecture

The backend serves as the core processing unit of the system, handling video processing, AI analysis, embedding generation, and query handling. It is built using an asynchronous framework to efficiently manage multiple concurrent tasks. Key backend modules include:

1. Video Processing Module: Handles video validation, transcript extraction, and preprocessing.
2. AI Analysis Module: Executes multiple LLM-based tasks in parallel to generate structured outputs such as summaries, quizzes, and insights.
3. Embedding Module: Converts textual content into vector embeddings for semantic storage.
4. RAG Module: Retrieves relevant content and generates context-aware responses.
5. Graph Generation Module: Extracts entities and relationships to construct the knowledge graph.

The use of asynchronous execution ensures that these modules operate efficiently without causing delays in user interactions.

## D. Database Design

The system uses PostgreSQL as the primary database, extended with vector storage capabilities to support semantic search. The database is designed to handle both structured and unstructured data efficiently. The relational schema includes tables for users, videos, transcripts, analysis outputs, and query logs. Each table is normalized to ensure data consistency and integrity.

To support semantic retrieval, embeddings are stored as high-dimensional vectors. These vectors are indexed using efficient similarity search techniques, enabling fast retrieval of relevant content.

## E. API and Integration

The system exposes RESTful API endpoints for communication between the frontend and backend. These APIs handle operations such as video submission, data retrieval, query processing, and graph generation. The integration of relational and vector data within a single database simplifies system design and reduces the need for external services.




Figure6: Overview of Video Summarization

The APIs are designed to be stateless, ensuring scalability and ease of deployment. Authentication and authorization mechanisms are implemented to secure user data and maintain isolation between different users' knowledge bases.

## F. Asynchronous Task Handling

Given the computational complexity of LLM-based processing, the system employs asynchronous task handling mechanisms. Tasks such as transcript analysis and embedding generation are executed in parallel using background workers. This approach

improves system responsiveness and allows users to continue interacting with the application while processing is ongoing. It also enables efficient resource utilization and scalability.

## G. System Integration Workflow

The complete system workflow integrates all components into a seamless pipeline. When a user submits a video, the frontend sends a request to the backend, which processes the video and stores the results in the database. The frontend then retrieves and displays the processed data. For query operations, the frontend sends a query to the backend, which retrieves relevant embeddings, generates a response using the RAG framework, and returns the result to the user.

## H. Summary of Implementation

The implementation demonstrates the feasibility of integrating advanced AI techniques into a scalable and user-friendly system. By combining modern web technologies with AI-driven processing, the system provides an efficient solution for intelligent knowledge management.



Figure7: Overview of Knowledge Graph




Figure8: Overview of Memory Interaction

The modular design ensures that individual components can be updated or replaced without affecting the overall system, making it adaptable to future advancements in AI and software engineering.

## 5 Results and Evaluation

This section evaluates the performance and effectiveness of the proposed AI-powered knowledge graph generation system. The evaluation focuses on key aspects such as retrieval accuracy, response quality, processing efficiency, and overall learning effectiveness.

### A. Evaluation Objectives

The primary objective of the evaluation is to assess how effectively the system improves video-based learning compared to traditional approaches. Specifically, the evaluation addresses the following research questions:

1. Does the system improve knowledge retrieval efficiency?
2. Does the use of Retrieval-Augmented Generation (RAG) enhance response accuracy?
3. Does multi-dimensional analysis improve knowledge retention?
4. How effectively does the knowledge graph reduce knowledge fragmentation?

### B. Evaluation Setup

The system was tested using a collection of educational YouTube videos across domains such as computer science, data structures, and machine learning. Each video was processed through the proposed pipeline to generate structured outputs and semantic embeddings.

User queries were simulated to evaluate retrieval and response generation. These queries ranged from factual questions to conceptual explanations, ensuring comprehensive testing of the system.

### C. Performance Metrics

To evaluate the system, the following metrics were considered:

1. Retrieval Accuracy: Measures how relevant the retrieved content is to the user query.
2. Response Relevance: Evaluates the correctness and contextual alignment of generated answers.
3. Processing Time: Measures the time required for video analysis and query response.
4. User Effort Reduction: Estimates the reduction in manual effort compared to traditional learning methods.

### D. Retrieval Performance

The semantic embedding-based retrieval mechanism demonstrated high accuracy in identifying relevant content. By using cosine similarity, the system was able to retrieve contextually similar information even when queries did not contain exact keywords.

$$\text{sim}(q, d) = q \cdot d / |q| \cdot |d|$$

Experimental observations show embedding-based retrieval outperforms keyword search for contextual relevance.

## E. Response Quality Evaluation

The integration of the Retrieval-Augmented Generation (RAG) framework improved the quality of generated responses. Unlike standalone language models, the system generates answers grounded in retrieved content, reducing hallucinations and improving factual accuracy.

The response generation process is defined as:

$$y = f(q, R(q))$$

Responses were evaluated based on correctness, completeness, and interpretability. The inclusion of timestamp-based references further enhanced the reliability of responses.

## F. Comparative Analysis

To demonstrate the effectiveness of the proposed system, it was compared with existing approaches.

| Feature             | Traditional Notes | Video Summarizers | AI Chatbots | Proposed System |
|---------------------|-------------------|-------------------|-------------|-----------------|
| Persistent Memory   | No                | No                | Limited     | Yes             |
| Semantic Search     | No                | No                | Yes         | Strong          |
| Cross-Video Linking | No                | No                | No          | Yes             |
| Knowledge Graph     | No                | No                | No          | Yes             |
| Context-Aware QA    | No                | Limited           | Yes         | Strong          |

Table2: Comparison of existing systems

The results show that the proposed system provides a more comprehensive and integrated solution compared to existing tools.

## G. Processing Efficiency

The use of asynchronous processing significantly improved system performance. Multi-dimensional analysis tasks were executed in parallel, reducing overall processing time. For query handling, the system achieved near real-time response generation due to efficient vector search and optimized backend processing.

## H. Knowledge Retention and Learning Impact

The system enhances knowledge retention by combining multiple learning modalities such as summaries, quizzes, and flashcards. These features promote active recall and spaced repetition, which are known to improve long-term memory. Additionally, the knowledge graph provides a visual representation of relationships between concepts, enabling users to develop a deeper understanding of interconnected topics.

## I. Discussion

The evaluation demonstrates that the proposed system effectively addresses the limitations of traditional video-based learning. By integrating semantic search, RAG-based querying, and knowledge graph visualization, the system provides a unified platform for intelligent learning.

While the results are promising, certain limitations remain. The quality of outputs depends on the accuracy of transcript extraction and the performance of language models. Additionally, large-scale deployment may require further optimization of computational resources.

## J. Summary of Results

Overall, the proposed system achieves significant improvements in retrieval accuracy, response quality, and learning efficiency. The combination of advanced AI techniques enables a more interactive and effective learning experience compared to existing approaches.

## 6. Conclusion and Future Work

This paper presented an AI-powered knowledge graph generation system designed to transform passive video-based learning into an active, structured, and persistent knowledge experience. By addressing key challenges such as knowledge fragmentation, poor retention, and inefficient retrieval, the proposed system provides a comprehensive solution for modern digital learning environments.

The system integrates multiple advanced technologies, including large language models for multi-dimensional content analysis, semantic embeddings for efficient knowledge representation, and a Retrieval-Augmented Generation (RAG) framework for context-aware querying. Additionally, the incorporation of automated knowledge graph construction enables the visualization of relationships between concepts, allowing users to explore interconnected knowledge across multiple videos. The implementation demonstrates the feasibility of combining these technologies into a scalable and modular architecture. Through asynchronous processing and efficient database design, the system achieves high performance while maintaining responsiveness and usability. The evaluation results indicate significant improvements in retrieval accuracy, response quality, and overall learning efficiency when compared to traditional approaches such as manual note-taking and basic video summarization tools.

One of the key strengths of the proposed system lies in its ability to provide persistent and personalized knowledge storage. Unlike conventional learning tools, the system allows users to build a continuously evolving knowledge base that can be queried and expanded over time. This capability supports long-term learning and reduces the need for repetitive

content consumption. Despite its advantages, the system has certain limitations. The quality of generated outputs is dependent on the accuracy of video transcripts and the capabilities of underlying language models. Additionally, large-scale deployment may require further optimization in terms of computational resources and latency management. The automatic construction of knowledge graphs, while effective, may also benefit from refinement to improve the precision of extracted relationships.

Future work will focus on enhancing the system in several directions. First, improvements in transcript processing and noise reduction can increase the reliability of extracted information. Second, incorporating user feedback mechanisms can enable adaptive learning and personalization. Third, advanced graph algorithms and reasoning techniques can be integrated to enable deeper insights and automated knowledge discovery. Furthermore, expanding support to include multi-modal data such as images, slides, and audio cues can provide a richer learning experience.

In conclusion, the proposed system represents a significant step toward intelligent knowledge management for video-based learning. By combining semantic understanding, interactive querying, and structured knowledge representation, it enables a more effective and engaging learning process. As digital learning continues to evolve, such systems will play a crucial role in bridging the gap between content consumption and knowledge acquisition.

## 7. References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova, "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding," in Proceedings of NAACL-HLT, 2019.
- [2] Tom B. Brown et al., "Language Models are Few-Shot Learners," in Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [3] Patrick Lewis, Ethan Perez, Aleksandra Piktus et al., "Retrieval-Augmented Generation for

Knowledge-Intensive NLP Tasks,” in NeurIPS, 2020.

[4] Tomas Mikolov et al., “Efficient Estimation of Word Representations in Vector Space,” in ICLR Workshops, 2013.

[5] Jeffrey Pennington, Richard Socher, and Christopher D. Manning, “GloVe: Global Vectors for Word Representation,” in EMNLP, 2014.

[6] Ashish Vaswani et al., “Attention Is All You Need,” in NeurIPS, 2017.

[7] Zhiyuan Liu and Maosong Sun, “Knowledge Graph Embedding: A Survey of Approaches and Applications,” *IEEE Transactions on Knowledge and Data Engineering*, 2016.

[8] Jeff Johnson, Matthijs Douze, and Hervé Jégou, “Billion-scale Similarity Search with GPUs,” *IEEE Transactions on Big Data*, 2019.

[9] Oren Etzioni et al., “Open Information Extraction: The Second Generation,” in *IJCAI*, 2011.

[10] YouTube, “YouTube for Learning,” Available: <https://www.youtube.com/learning>

[11] Danqi Chen et al., “Dense Passage Retrieval for Open-Domain Question Answering,” in *EMNLP*, 2020.

[12] Kurt VanLehn, “The Relative Effectiveness of Human Tutoring, Intelligent Tutoring Systems, and Other Tutoring Systems,” *Educational Psychologist*, 2011.

[13] Reimers Nils and Iryna Gurevych, “Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks,” in *Proceedings of EMNLP*, 2019.

[14] James Thorne et al., “FEVER: A Large-Scale Dataset for Fact Extraction and Verification,” in *NAACL-HLT*, 2018.

[15] Petar Veličković et al., “Graph Attention Networks,” in *International Conference on Learning Representations (ICLR)*, 2018.