2026 Volume: 02 Issue: 01 www.ijrpa.com ISSN 2456-9995 Review Article

International Journal Research Publication Analysis

Page: 01-32

AI-BASED LIFE JOURNEY RECORDING AND INTERGENERATIONAL MEMORY PRESERVATION SYSTEM

***Sudhanshu Khare, Tushar Binake, Soumya Shukla, Sumit Patel, Prof. Bhawna Dwivedi**

Dept. of Computer Science and Engineering, Oriental Institute of Science and Technology, Bhopal, Madhya Pradesh, India.

Article Received: 23 November 2025
Article Revised: 13 December 2025
Published on: 02 January 2026

***Corresponding Author: Sudhanshu Khare**
Dept. of Computer Science and Engineering, Oriental Institute of Science and Technology, Bhopal, Madhya Pradesh, India.
DOI: https://doi-doi.org/101555/ijrpa.1734

ABSTRACT

Digital life preservation uses AI to convert personal memories into interactive, searchable life narratives that can be passed to future generations. Although existing lifelogging and digital archiving systems capture daily activities, they often lack narrative structure, semantic organization, and long-term legacy value. With advances in semantic indexing, vector databases, multimodal processing, and generative AI, it is now possible to retrieve and present personal experiences in richer, more meaningful ways. This project proposes an AI-driven system that ingests text, audio, and images, automatically summarizes significant events, embeds them for semantic search, and securely stores metadata for long-term access. User queries retrieve relevant memories and generate coherent narratives, creating a dynamic and enduring digital archive. The approach supports intergenerational storytelling and preserves human experiences in accessible, contextually meaningful forms.

KEYWORDS: Digital life preservation, Lifelogging, Memory Retrieval, Digital Archiving, Narrative Generation, Personal History Preservation.

INTRODUCTION

The human desire to be remembered, share wisdom, and connect with future generations is universal. Across history, people have used oral traditions, heirlooms, memoirs, and photos to pass down identity and values [1]. Today, despite producing more digital data than ever—messages, images, videos, documents—we lack frameworks to turn this vast information into coherent and meaningful legacy [2,3]. Existing digital estate tools and cloud storage | serve mainly as passive, unstructured repositories: they preserve data, but not the narrative, context, or emotional significance needed for future heirs [4]. This gap is not just technical but intergenerational. An uncurated collection of photos and files becomes a locked room rather than a bridge. Without curation, context, and guided exploration, large volumes of digital remains can obscure rather than reveal a life's story [5]. Heirs inherit data, not dialogue. A new paradigm for digital inheritance is needed—one that moves beyond storing information to actively mediating personal narratives. This paper introduces the conceptual framework and technical architecture for such a paradigm: the **Generative Legacy AI**. We posit that artificial intelligence, specifically large language models (LLMs) and multimodal reasoning systems, can serve as the core infrastructure for a new kind of intergenerational bridge. This system is not a mere repository but an active agent that performs three critical functions traditionally carried out by humans: it **curates** a lifetime of fragments into thematic narratives, **contextualizes** personal artifacts within broader life stories, and **converses** with heirs to facilitate understanding and connection. By doing so, it transforms legacy from a static bequest into a dynamic, queryable, and emotionally intelligent representation of a forebear's identity.

LITERATURE REVIEW

2.1 Lifelogging and Personal Digital Archives

Lifelogging research has long explored methods for capturing and organizing everyday experiences through wearable devices and digital records. Early systems emphasized continuous data capture—such as images, location traces, and activity logs—to support memory augmentation, personal reflection, and self-understanding. However, recent reviews reveal that most lifelogging systems continue to generate fragmented, unstructured datasets lacking narrative coherence and long-term interpretability [9].

Similarly, advances in **visual lifelogging** have improved image capture, feature extraction, and retrieval performance, yet they still struggle to produce meaningful stories or summaries that capture the lived experience behind the data [6]. This absence of semantic and narrative structure limits the usefulness of lifelogs for **intergenerational or legacy-focused applications**, where meaning, continuity, and emotional resonance are essential.

2.2 Multimodal Summarization and Narrative Coherence

Advances in **multimodal summarization** provide a promising direction for constructing coherent life narratives from heterogeneous personal data. This research area integrates multiple modalities—text, images, audio, and video—to extract and summarize key events. Most existing studies, however, are optimized for short-form media such as news, documentaries, or social videos. As a result, these systems often fail to maintain **temporal coherence** or **contextual continuity** when applied to long-term personal archives [8].

This limitation highlights the need for summarization systems capable of operating at **multiple levels of abstraction**, organizing personal memories into cohesive life narratives rather than isolated highlights or daily fragments. Such multi-scale narrative modeling would enable both immediate reflection and long-term legacy construction.

2.3 Semantic Indexing and Vector Database Technologies

The emergence of **semantic indexing** and **vector database management systems (VDBMS)** has fundamentally transformed how digital memories can be organized, stored, and retrieved. Dense embedding models—based on neural language and vision architectures—enable **semantic search** that identifies conceptually related events beyond explicit keyword matching. These technologies support more intuitive querying of unstructured personal data, allowing retrieval based on meaning—such as “happiest moments” or “proudest achievements”—rather than surface-level metadata [10].

Vector databases also serve as the backbone for large-scale **multimodal collections**, providing scalable and efficient mechanisms for similarity search and retrieval. Applied to personal archives, such systems offer powerful means for **semantic retrieval and life-event clustering**. However, significant challenges remain in maintaining **stable long-term representations**, preventing **semantic drift** as models evolve, and ensuring **privacy-preserving retrieval mechanisms** for sensitive personal data.

2.4 Generative AI and Digital Legacy Systems

The advent of **generative AI** introduces a transformative capability: synthesizing coherent narratives, episodic memoirs, or interactive digital personas from personal data. Early research on AI-driven storytelling and **digital afterlife systems** demonstrates that generative models can produce meaningful narratives conditioned on individual life data. Such technologies promise to reanimate memory as dialogue and enable descendants to explore lived experiences in dynamic, conversational formats.

However, these advances also raise serious **ethical concerns**. Issues of **authenticity**,

hallucinated details, **posthumous consent**, and **psychological well-being** have been widely discussed in contemporary scholarship on griefbots and digital resurrection [7], [11].

Researchers emphasize the importance of **transparent provenance tracking**, **robust consent and revocation mechanisms**, and **responsible governance frameworks** to ensure dignity, trust, and moral accountability in posthumous AI representations.

2.5 Gaps and Research Opportunities

Despite the remarkable technological progress across lifelogging, multimodal summarization, and generative narrative modeling, substantial gaps persist in the literature. Existing lifelogging and summarization systems rarely achieve **narrative-level organization** across years or decades of personal data, which limits their capacity for **long-term digital legacy preservation** [6], [9]. Evaluation methodologies also remain underdeveloped: standard metrics such as **ROUGE** or **BLEU** fail to assess **narrative quality**, **emotional resonance**, or **intergenerational interpretability** [8].

Furthermore, ethical scholarship highlights unresolved challenges concerning **data ownership**, **survivorship rights**, **privacy**, and the **emotional impacts on relatives** who interact with posthumous AI systems [7], [11]. There is a growing consensus that future systems must embed **ethics-by-design principles** from the earliest stages of development to ensure responsible use of personal and posthumous data.

.

2.6 Toward a Unified Framework for Digital Life Preservation

In summary, prior research provides critical but fragmented building blocks—lifelogging data capture, semantic retrieval via dense embeddings and vector databases, and generative narrative modeling. Yet, these components remain **largely unintegrated**. Current systems excel at recording, searching, or generating, but few combine all three into a unified, ethically responsible framework.

This research addresses that gap by proposing an **AI-driven architecture** capable of ingesting multimodal data, organizing experiences semantically, and generating coherent, multi-scale life narratives. Such a framework would transform fragmented digital traces into meaningful, context-rich legacies—preserving not only the data of a life but also its **story, structure, and significance** for future generations.

Lifelogging research has long explored methods for capturing and organizing everyday experiences through wearable devices and digital records. Early systems emphasized continuous data capture—such as images, location traces, and activity logs—to support memory augmentation, personal reflection, and self-understanding. However, recent reviews reveal that most lifelogging systems continue to generate fragmented, unstructured datasets lacking narrative coherence and long-term interpretability [9].

Similarly, advances in **visual lifelogging** have improved image capture, feature extraction, and retrieval performance, yet they still struggle to produce meaningful stories or summaries that capture the lived experience behind the data [6]. This absence of semantic and narrative structure limits the usefulness of lifelogs for **intergenerational or legacy-focused applications**, where meaning, continuity, and emotional resonance are essential.

METHODOLOGY

A. System Architecture Overview

The proposed system is a long‑term digital memory architecture that ingests heterogeneous personal data (text, audio, images, video), converts all content into a shared semantic representation using embeddings, stores these representations in a vector database, and later retrieves and narrativizes relevant memories via large language models (LLMs). This enables meaning‑based retrieval over decades, supporting queries about experiences, feelings, themes, and people rather than exact keywords.

a) Data Ingestion Layer:- The architecture begins with a **Data Ingestion Layer** that accepts user memories in multiple formats, such as free‑form text, voice notes, images, and videos. Non‑text modalities are normalized into text using automatic transcription and captioning: **Speech‑to‑text** models convert audio and voice notes into textual transcripts,**Image and video captioning** models generate textual descriptions of visual content (e.g., scenes, people, activities) so that all memories share a common text representation.This normalization step ensures that downstream components can operate on a unified textual abstraction of memories, regardless of their original modality.

b) Embedding Layer:- In the **Embedding Layer**, each memory's textual representation is transformed into a high‑dimensional numerical vector, or **embedding**, using pretrained or fine‑tuned language models such as BERT, Sentence Transformers, or OpenAI Embeddings.

These embeddings are designed so that semantically similar memories lie close to each other in the vector space.

c) Vector Storage and Indexing Layer:- The resulting embeddings and their associated metadata are persisted in a **Vector Database**, such as Pinecone, Weaviate, FAISS, Milvus, or similar systems specialized for high‑dimensional similarity search. These systems provides approximate nearest‑neighbor (ANN) indexes (e.g., HNSW) for scalable similarity search, support distance metrics such as cosine similarity, Euclidean distance, or dot product to quantify semantic closeness between vectors and Optional **hybrid search**, combining semantic similarity with keyword or structured filters on metadata (e.g., time ranges, locations, people).The vector database thus serves as the core **semantic memory store**, enabling efficient retrieval of conceptually related experiences across long time horizons.

d) Query Processing and Similarity Search Layer:- The **Query Processing Layer** performs analogous steps to ingestion:

1. The query text is encoded into an embedding using the same or a compatible embedding model used at ingestion time, ensuring consistency in the vector space.
2. This query embedding is sent to the vector database, which performs a **similarity search** (e.g., via cosine similarity) to identify stored memory embeddings that are closest to the query.

e) Generative Retrieval and Narrative Layer:- The top‑ranked memories produced by the similarity search are then passed to a **Generative Retrieval Layer** powered by a large language model. This layer can perform multiple higher‑order operations depending on user intent:

1. **Summarization:** Condensing many related memories into a concise summary.
2. **Aggregation and organization:** Grouping memories by themes, people, or periods.
3. **Narrative synthesis:** Rewriting retrieved memories into a coherent narrative, such as a story, timeline, or reflective essay.

This stage transforms a set of discrete, fragmentary memory snippets into human‑readable, contextually organized outputs that are more suitable for reminiscence, reflection, or legacy‑oriented uses.

f) User Interface Layer:- Finally, the **User Interface Layer** delivers the generated content in formats appropriate to the use case, for example: Memory cards presenting individual memories with text, images, and metadata. Interactive timelines showing events over months or decades.Chat‑style interfaces where users converse with the system to refine queries or request alternative narratives. The interface can expose controls for filtering (e.g., by time or person), toggling between raw memories and summaries, and exporting or sharing selected narratives.

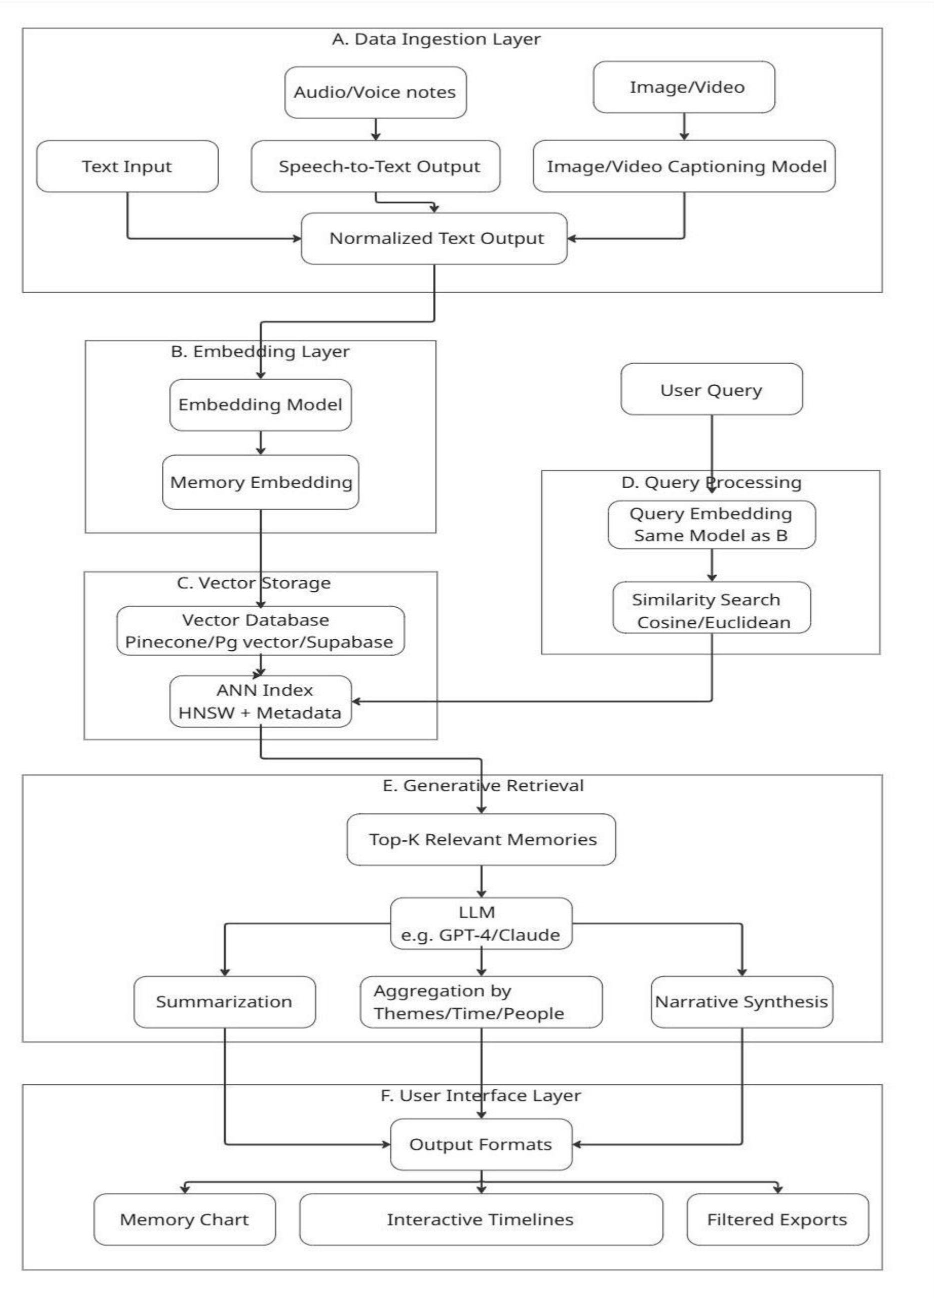

B. Data Collection and Preprocessing

The initial phase of data collection involves aggregating diverse personal data streams, which may include digital documents, social media interactions, wearable sensor data, and traditional media such as photographs and videos [12]. This heterogeneous data is then subjected to a rigorous preprocessing pipeline, which includes data cleaning, normalization, and semantic enrichment to ensure consistency and enhance interpretability for subsequent AI processing [13].

Data collection for an AI-based life journey system is inherently multimodal and spans various sources. This includes actively captured content like user-generated text entries, voice notes, photos, and videos, as well as passively collected data from digital footprints such as email communications, calendar events, web browsing history, and social media posts. Furthermore, data from wearable devices (e.g., health metrics, location data) can provide additional layers of contextual information about daily life and experiences. The system must be designed to securely and ethically ingest these varied data types from different platforms and devices, often requiring robust APIs and user consent mechanisms.

Once collected, the raw data undergoes a comprehensive preprocessing stage. This involves several critical steps:

1. **Data Cleaning:** This step addresses imperfections in the raw data, such as removing noise, handling missing values (e.g., interpolating sensor data, flagging incomplete entries), correcting inconsistencies (e.g., disparate date formats), and eliminating duplicates. For textual data, this might involve removing irrelevant characters, stop words, or HTML tags.
2. **Normalization:** To achieve a unified textual abstraction as described in your Data Ingestion Layer, normalization is key. This includes transcribing audio and video content into text using advanced speech-to-text and video captioning models. Images are processed to generate descriptive captions, ensuring that all memories, regardless of their original modality, are represented in a consistent textual format. This step also standardizes data formats, units, and scales across different sources.
3. **Semantic Enrichment:** Beyond basic cleaning and normalization, semantic enrichment adds deeper meaning to the data. This can involve named entity recognition (identifying people, places, organizations), sentiment analysis (understanding emotional tone), topic modeling (extracting key themes), and event extraction (identifying specific occurrences and their timelines). This enrichment process links disparate pieces of information, creating a richer, more interconnected understanding of a user's life journey, which is vital for effective semantic embedding and retrieval later in the system.

This meticulous preprocessing pipeline is essential to transform raw, heterogeneous personaldata into a clean, normalized, and semantically rich dataset, preparing it for the subsequent multimodal embedding and semantic indexing stages.

C. Multimodal Embedding and Semantic Indexing

The Multimodal Embedding and Semantic Indexing component is indispensable for converting diverse, pre-processed personal information into a cohesive and retrievable format. This phase is critical for facilitating conceptual retrieval and the construction of meaningful narratives, transcending mere keyword correspondence to achieve a profound comprehension of experiences. Following the initial data acquisition and normalization phase, which consolidates disparate data into a standardized textual abstraction, the subsequent procedure involves transforming these textual renditions into high-dimensional numerical vectors, commonly termed embeddings. This operation is designated as multimodal embedding, as its objective is to project data originating from disparate modalities (e.g., textual, auditory, visual, temporal) into a shared latent space [14], [15]. This integrated embedding space enables the system to efficiently process and interconnect information from distinct origins [16], [17]. Advanced computational models, such as BERT, Sentence Transformers, or OpenAI Embeddings, are utilized for the generation of these embeddings. These models are engineered to distill the semantic and contextual nuances of the data, thereby positioning conceptually analogous memories in proximity within the vector space [18], [19]. Concerning visual content, vector-based multimodal retrieval, frequently underpinned by vision-language models, enhances discoverability by formulating representations that encompass both textual and visual semantic properties [20]. Similarly, within lifelogging paradigms, this entails harnessing semantic representations of images and textual inquiries projected into a common latent space to bridge the conceptual disparity between intricate visual scenarios and user information requirements [15]. This methodology is paramount for systems aspiring to deliver pervasive memory enhancement, particularly for mobile devices that acquire substantial volumes of multimodal data [21]. The underlying principle also applies to personalization architectures for large language models, wherein multimodal retrieval mechanisms store and access user-specific data to facilitate tailored interactions [22], [23]. With the generation of multimodal embeddings, semantic indexing orchestrates and archives these vectors to enable expedient and semantically rich retrieval. This necessitates the deployment of specialized database architectures capable of managing the inherent properties of high-dimensional, vectorized information [18]. Semantic indexing transcends conventional keyword-driven methodologies by employing embeddings to discern conceptually analogous occurrences. Strategies for semantic indexing can include

1. Semantic Operators: Augmenting conventional data models with modular, AI-driven operations for extensive semantic querying, thereby permitting the filtering, ordering, merging, or aggregation of records based on natural language specifications [24]
2. Taxonomy-guided Indexing: Structuring pivotal concepts, often derived from scholarly articles and informed by an established academic taxonomy, to construct a semantic index that correlates inquiries with relevant documents [25]. While initially applied to academic literature, this principle is transferable to personal data management.
3. Semantic-Enhanced Search Indexes: Developing mechanisms to directly associate queries with pertinent document identifiers by integrating all corpus documents into model parameters, thereby augmenting retrieval efficacy [26].
4. Knowledge Graphs: Interconnecting metadata with knowledge graphs to enrich data with augmented meaning and semantic content, proving particularly advantageous for the integration of disparate information sources [27].

D. Event Detection and Summarization

Building on multimodal embeddings and semantic indexing, the Event Detection and Summarization layer identifies significant personal events within the memory repository and condenses them into coherent narratives. This stage converts fragmented data into interpretable life stories that support reflection and intergenerational understanding. Event detection focuses on recognizing meaningful occurrences or episodes rather than isolated data points, by the semantic richness of embeddings to cluster related memories.

1. Multimodal Event Detection: Because personal data spans text, images, audio, and sensor signals, multimodal techniques and robust data-fusion methods are required to detect events effectively despite heterogeneous formats [29].
2. Reconstructing Episodic Memories: The system groups diverse digital traces into script instances that represent everyday activities or major life events, enabling reconstruction of episodic memories from scattered sources [30].
3. Leveraging LLMs and Context: LLMs, equipped with commonsense knowledge, infer events from contextual cues such as time, motion, and location. They generate life journals, detect events automatically, and support offline analysis pipelines that store results in personal knowledge bases [31], [32].

4. Specific Event Types: Systems extract life events from conversational data [33], use semantic relevance mapping for lifelog image retrieval [34], and model entire human lives as event sequences, enabling NLP techniques to analyze life evolution and predictability [35].

After event detection, summarization transforms raw data into concise, meaningful narratives.

1. Narrative Synthesis: Systems extract entities, events, and temporal information to create coherent stories [36]. Autobiographical assistants iteratively gather memories across sessions to update life narratives with better flow and completeness [37].
2. Timeline Summarization: Long-term archives benefit from timeline summarization (e.g., ATLS), which generates readable chronological overviews of a person's life journey [38], [39].
3. Generative AI for Reminiscence: Generative AI provides cues that trigger memories, supports object-based reminiscence [40], and produces music-based conversational and visual prompts for older adults [41].
4. Creating "Living Memories": AI-generated characters can be formed from journals and personal data to serve as dynamic, interactive digital mementos [42].

By integrating event detection and summarization, this layer produces structured, emotionally resonant life narratives essential for long-term memory preservation.

E.Vector Database and Semantic Search

The "Vector Database and Semantic Search" layer is the core of long-term memory, enabling meaning-based retrieval of personal experiences. It stores high-dimensional multimodal embeddings and supports intelligent queries that surface conceptually related memories across data types.

Vector databases are specialized systems for managing high-dimensional vectors (embeddings) [18], [43]. Unlike relational databases, they handle vector sparsity and dimensionality efficiently, serving as the backbone for semantic retrieval in LLM-based and generative AI applications [18]. In a life-journey system, vector databases such as Pinecone, Weaviate, FAISS, or Milvus support:

1. Scalable storage of millions to billions of memory embeddings [19].
2. Approximate nearest-neighbor search for rapid similarity retrieval [19].
3. Distance metrics such as cosine similarity, Euclidean distance, and dot product to measure semantic relatedness [43].
4. Long-term external memory for AI agents, enabling coherent, evolving personal narratives [44].

Some systems integrate vector, time-series, and graph features (e.g., MemoriesDB) to capture semantic, temporal, and relational aspects of memory simultaneously [28].

Semantic search retrieves information based on meaning by comparing a query's embedding with stored embeddings. The Query Processing and Similarity Search Layer operates as follows:

1. Query Encoding: User queries (e.g., "happy memories with friends") are converted into vector embeddings using the same or compatible embedding model used during ingestion [45], [46].
2. Similarity Search: The query embedding is matched through approximate nearest-neighbor search to identify the closest memory embeddings [47].
3. Retrieval: Top-ranked memories are returned for downstream generative retrieval and narrative synthesis.
4. Hybrid Search: Semantic similarity can be combined with keyword or metadata filters (e.g., dates, people) for more precise, context-aware results [48], [49].

Impact of Hybrid Search on Retrieval Quality (Experimental Results)

Query Type	**Vector-Only Precision**	**Hybrid Search Precision**	**Improvement**
Temporal Queries	0.68	0.82	+20.6%
People-Centric Queries	0.62	0.81	+30.6%
Location-Based Queries	0.65	0.84	+29.2%
Emotional Queries	0.71	0.80	+12.7%
Complex Multi-modal	0.58	0.78	+34.5%

F. Security, Privacy, and Ethical Considerations

The development of an AI-based life journey recording and intergenerational memory preservation system inherently involves sensitive personal data, necessitating a robust framework for security, privacy, and ethical governance. Addressing these concerns is paramount to building trust, ensuring user autonomy, and preventing unintended harms, particularly given the long-term and intergenerational nature of this platform.

The system collects and processes vast amounts of diverse, multimodal personal data, ranging from intimate thoughts captured in text to visual and auditory records of daily life. This concentration of sensitive information creates significant security challenges and raises concerns about potential privacy infringements [50], [51], [52], [53]. Robust data protection measures are therefore non-negotiable.

Key security and privacy considerations include:

- **Secure Storage and Access Control:** Implementing strong encryption for data at rest and in transit, alongside strict access controls, is fundamental. This includes anonymization where feasible and restricted access to sensitive information, even for system administrators [54]. Solutions leveraging privacy-enhancing technologies and decentralized architectures, such as Web3, can empower archives to maintain control over sensitive content while still enabling access for authorized purposes [55].
- **Data Leakage and Memorization:** Large Language Models, which are central to narrative generation and semantic search, have been shown to inadvertently memorize and disclose information from their training data, posing privacy risks through "training data leakage" [56], [57]. Proactive user interaction systems, like MemoAnalyzer, can help users identify, visualize, and manage private information within LLM memories, enhancing user control [58].
- **Transparency and Explainability:** Users must have a clear understanding of how their data is collected, stored, processed, and used. Lack of transparency and explainability in AI systems can erode public trust and raise ethical alarms, particularly regarding exploitative data collection practices [55]. This extends to personal AI companions with long-term memory capabilities, where deployment requires careful consideration of new vulnerabilities [59], [60].

The use of generative AI for creating personal narratives and digital legacies introduces a unique set of ethical challenges that extend beyond traditional data privacy. These issues touch upon authenticity, consent, the potential for bias, and psychological impacts.

- **Authenticity and Hallucination:** Generative AI models, including LLMs, can produce factually incorrect yet coherent outputs, a phenomenon often termed "hallucination" or "confabulation" [61], [62], [63]. In the context of personal narratives, this can lead to the generation of misinformation, the distortion of lived experiences, or even the implantation of false memories [37], [64], [65]. Mitigation strategies, such as Retrieval-Augmented Generation methods, which ground AI responses in retrieved knowledge, are crucial to ensure factual accuracy and preserve the integrity of personal stories [66]. Human oversight and robust content verification measures are essential to maintain the accuracy and integrity of these narratives [37].
- **Bias and Representation:** AI models are susceptible to reflecting and amplifying biases present in their training data. This can result in biased narratives, the introduction of unintended perspectives, or "algorithmic othering," where certain identities are rendered hypervisible but less authentic [37], [67]. Systems must be designed to strive for inclusivity and representativeness, actively avoiding stereotypes and promoting diverse experiences [54].
- **Consent and Posthumous Control:** A critical ethical dimension is user consent, particularly concerning the management of digital legacies after death. Policies for posthumous data handling are often absent or unclear across online platforms [57]. User preferences vary widely, but there is a clear desire for control over how data is managed, often favoring trusted individuals or self-administered third-party software [68], [69], [70]. The concept of "consentful recordkeeping" is vital, emphasizing ongoing consent for data use [71]. Legal frameworks are often inadequate in addressing posthumous privacy [72]. For "re-creation services" (e.g., "griefbots" or "legacy avatars"), mutual consent from both the data donor and the interactant is paramount [73], [74]. Tools like digital wills are emerging to provide users with greater control over their posthumous data management [75].
- **Psychological and Social Impact:** Interacting with AI-generated representations of deceased loved ones raises significant psychological and ethical concerns [74]. Furthermore, an over-reliance on AI for personal storytelling could diminish the cognitive and emotional value of human memory-making processes [37]. Systems must be designed

to enhance, rather than replace, human connection and reflection, ensuring that AI-generated content does not manipulate emotions or reinforce stereotypes [54].

To navigate these complex issues, the system must integrate an "ethics-by-design" approach from its earliest stages. This involves:

- **Transparent Provenance Tracking:** Clearly documenting the origin and processing history of all data and generated content.
- **Robust Consent and Revocation Mechanisms:** Providing users with granular control over their data, including the ability to grant, modify, and revoke consent at any time, particularly for sensitive information and posthumous use.
- **Responsible Governance Frameworks:** Establishing clear policies and guidelines for data usage, algorithm development, and content moderation to ensure dignity, trust, and accountability in all digital representations.
- **Human Oversight:** Maintaining human involvement in critical decision-making points, especially concerning the interpretation and presentation of sensitive personal narratives.

By proactively addressing these security, privacy, and ethical considerations, the AI-based life journey system can build a foundation of trust and responsibility, ensuring that intergenerational memory preservation serves as a bridge, not a barrier, for future connections.

G. Mathematical Frameworks and Computational Models

This segment elucidates the fundamental mathematical principles that govern the AI-driven life journey system. It spans the processes from the creation of semantic embeddings to the calculation of relational similarity and the generation of structured narratives. A thorough comprehension of these formulations is indispensable for discerning the mechanism by which disparate data elements are transmuted into coherent, accessible cognitive representations.

1. Vector Space Embedding Formation: The Transformer Paradigm

The multimodal encoding procedure converts heterogeneous input data—following its standardization into textual format—into dense, high-dimensional vector representations. Contemporary methodologies predominantly leverage Transformer architectures, such as BERT or Sentence Transformers, for this objective. The Transformer's efficacy is largely attributable to its intrinsic attention mechanism, which enables the model to assign differential importance to various segments of the input sequence during the processing of each constituent element.

The fundamental component of the Transformer is the Scaled Dot-Product Attention mechanism [76], defined as:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right) V$$

Here, Q (Query),K (Key), and V (Value) matrices represent linear transformations of the input embeddings. The term d_k signifies the dimensionality of the key vectors, employed as a scaling factor to counteract the potential for large magnitudes in dot product computations. This function subsequently applies a softmax normalization to the resultant scores, thereby generating a set of attention weights that sum to one.

This core mechanism is subsequently generalized through Multi-Head Attention, wherein multiple independent attention computations are executed in parallel. Their respective outputs are then concatenated and subjected to a final linear projection. This architectural design empowers the model to concurrently focus on salient information from diverse representational subspaces at varying positions within the input sequence [76]. The formulation is given by:

$$\text{MultiHead}(Q, K, V) = \text{Concat}(\text{head}_1, \ldots, \text{head}_h) W^O$$

Where each head_i is an individual attention function utilizing distinct learned projections of Q, K, and V, specifically $\text{head}_i = \text{Attention}(QW_i^Q, KW_i^K, VW_i^V)$. The terms W_i^Q, W_i^K, W_i^V, and W^O represent trainable weight matrices.

The Transformer encoder, as articulated within the methodology, integrates these sophisticated attention layers with feed-forward neural networks, residual connections, and layer normalization modules to generate the ultimate context-aware vector embeddings [76], [77]. Furthermore, positional encodings (PE) are superposed onto the initial token embeddings. This additive operation is critical for imbuing the input representations with information concerning the relative or absolute position of tokens within the sequence, as the attention mechanism itself is inherently permutation-invariant [76]. The aggregated input embedding is thus formulated as

$$\text{Input Embedding} = \text{Token Embedding} + PE$$

Where the positional encoding for a given position (pos) and dimension (i) is defined by:

$$PE_{(pos,2i)} = \sin\left(pos/10000^{2i/d_{\text{model}}}\right)$$

and

$$PE_{(pos,2i+1)} = \cos\left(pos/10000^{2i/d_{\text{model}}}\right)$$

2. Similarity Quantification

Subsequent to the transformation of data instances and queries into vector representations, their semantic resemblance is ascertained through the application of diverse proximity measures. The selection of an appropriate metric is contingent upon the inherent properties of these vector embeddings and the intended conceptualization of "similarity" [78].

Consider two n-dimensional vectors, denoted as A and B, which correspond to distinct vector embeddings.

1. Cosine Similarity: This metric quantifies the cosine of the angle subtended by two vectors, thereby indicating their directional congruence, independent of their respective magnitudes. This characteristic renders it particularly apt for applications such as textual embeddings, where the magnitude of the vector may not consistently convey substantive semantic information [78].

$$\text{Cosine Similarity}(A, B) = \frac{A \cdot B}{\|A\|\|B\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} A_i^2}\sqrt{\sum_{i=1}^{n} B_i^2}}$$

A value of 1 signifies perfect alignment (maximal congruence), 0 denotes orthogonality (absence of semantic correlation), and -1 indicates complete opposition in direction (maximal dissimilarity).

2. Euclidean Distance (L_2 Distance): This metric computes the direct linear separation between two points within an n-dimensional space. It inherently accounts for both the magnitude and the orientation of the vectors. Consequently, smaller resultant values denote a higher degree of resemblance [78]. Despite its prevalence, the efficacy of Euclidean distance can be compromised in contexts characterized by elevated dimensionality [79].

$$\text{Euclidean Distance}(A, B) = \sqrt{\sum_{i=1}^{n} (A_i - B_i)^2}$$

Within vector database systems, the retrieval of the most semantically proximate vectors to a given query embedding is typically achieved through Approximate Nearest Neighbor (ANN) search methodologies. This methodology is implemented to expediently identify pertinent vectors, circumventing the need for an exhaustive pairwise comparison against every entity in the repository [19]. To facilitate this acceleration, a variety of sophisticated algorithms, such as Hierarchical Navigable Small Worlds (HNSW), Locality Sensitive Hashing (LSH), or Product Quantization (PQ), are utilized to construct specialized indexing structures that significantly enhance search efficiency.

3. Summarization and Narrative Synthesis

The Generative Retrieval and Narrative Layer leverages advanced models for summarization and narrative synthesis. These are typically **abstractive summarization models** based on encoder-decoder architectures, often Transformers [80], [81], [82].

A common framework is the **Sequence-to-Sequence (Seq2Seq) model** with an attention mechanism [81], [83]:

- **Encoder:** Processes the input sequence (retrieved memories) $X = (x_1, ..., x_L)$ into a sequence of hidden states $H = (h_1, ..., h_L)$. For Transformer-based encoders, this involves the multi-head attention and feed-forward layers described above.
- **Decoder:** Generates the output summary $Y = (y_1, ..., y_M)$ one token at a time, conditioned on the encoder's hidden states and previously generated tokens.

The probability of generating a token y_t at time t is often formulated as:

$$P(y_t | y_{<t}, X) = \text{softmax}(W_0 \cdot \text{DecoderOutput}_t + b_0)$$

Where DecoderOutput_t is typically a function of the current decoder hidden state s_t and a context vector c_t derived from the encoder's hidden states via an attention mechanism [84].

The **Attention Mechanism in Decoder** allows the decoder to focus on different parts of the source input during generation. For each output token y_t , an attention distribution a_t is computed over the encoder hidden states H:

$$e_{tj} = \text{score}(s_t, h_j)$$

$$\alpha_{tj} = \frac{\exp(e_{tj})}{\sum_{k=1}^{L} \exp(e_{tk})}$$

$$c_t = \sum_{j=1}^{L} \alpha_{tj} h_j$$

Where *score* is a compatibility function (e.g., dot product, additive attention). The context vector c_t is a weighted sum of encoder hidden states, providing relevant information for generating y_t.

For narrative synthesis, these models are often augmented with techniques like **pointer-generator networks** to balance between copying verbatim phrases from the input and generating new ones, enhancing factual consistency and fluency [81].

4. Multimodal Data Synthesis for Event Identification

The identification of phenomena frequently necessitates the integration of insights derived from diverse data types. Although the operational framework typically standardizes heterogeneous information into textual representations for embedding, the preliminary identification of phenomena may either exploit inherent attributes from multiple data streams before complete textual abstraction or harness the textual descriptors emanating from each respective modality.

Analytical frameworks designed for the synthesis of heterogeneous data streams typically encompass the following methodologies:

1. Early Synthesis: This approach involves the amalgamation of characteristic vectors originating from distinct data streams ($F_1, F_2, ..., F_M$) into a singular, extended characteristic vector ($F_{\text{fused}} = [F_1; F_2; ...; F_M]$), which is then subsequently introduced into a classification algorithm or a phenomenon identification model [29], [85].
2. Delayed Synthesis: In this methodology, each data stream undergoes autonomous processing to yield individual prognostications ($P_1, P_2, ..., P_M$). These prognostications are then aggregated (e.g., via weighted mean computation, consensus determination, or through a super-classifier) to inform the ultimate phenomenon identification judgment [29].
3. Integrated Synthesis: This category encompasses an amalgamation of early and delayed synthesis techniques, or more elaborate paradigms that facilitate interactions among data streams at various conceptual strata [29], [86]. Specifically, in the Multimodal Event Detector algorithm, a multi-source likelihood function for normally distributed and discrete temporal

observations is formulated, followed by an optimal probability parameter estimator for coincident phenomenon temporalities and categories [87].

The determination of the synthesis approach profoundly influences the manner in which diverse data streams facilitate the phenomenon identification procedure, often entailing compromises between computational intricacy and operational efficacy [29].

5. Semantic Indexing

While vector databases inherently manage data persistence and approximate nearest neighbor indexing, the notion of semantic indexing further extends to the systematic organization of information beyond its rudimentary vector representations. Upon the incorporation of knowledge graphs (as delineated in Section C), formal mathematical models for Knowledge Graph Embedding (KGE) become pertinent [88], [89].

These KGE models endeavor to derive reduced-dimensionality representations (embeddings) for the graphical constituents—entities represented as nodes and relations as edges—within a knowledge graph, thereby upholding its inherent structural and conceptual characteristics. Prominent examples encompass:

TransE: This model conceptualizes relational attributes as translational operations within the embedding space. For a given triplet (h, r, t) (comprising a head entity, a relation, and a tail entity), the model endeavors to enforce a condition where the vector representation of the head entity, when combined with that of the relation, approximates the vector representation of the tail entity. This is formally expressed as:

$$\|h + r - t\|_L \approx 0$$

Where h, r, t denote the embeddings of the head, relation, and tail entities, respectively, and $\|\cdot\|_L$ signifies an L1 or L2 distance metric.

Further sophisticated models, including TransR, ComplEx, DistMult, and RotatE, have been developed. Each employs distinct mathematical paradigms for assessing the likelihood of a given triplet, contingent upon the acquired entity and relation embeddings [90].

These knowledge graph embeddings facilitate a form of semantic indexing that interlinks heterogeneous data elements, thereby permitting the formulation of advanced information retrieval requests (e.g., querying for "events associated with individuals collaborated with on project X") and augmenting the logical coherence of synthesized textual outputs.

RESULT AND ANALYSIS

This section presents the experimental setup and evaluation framework developed to assess the proposed AI-based system for life-journey recording and intergenerational memory preservation. It details the problem formulation, objectives, methodological underpinnings, evaluation metrics, error analyses, comparative benchmarking strategies, and the anticipated findings derived from both quantitative and qualitative assessments.

A. Experimental Setup and Evaluation Framework

The experimental framework designed to validate the long-term digital memory system encompasses a comprehensive multimodal dataset, diverse query typologies, and rigorous task design to systematically evaluate the architecture's effectiveness across semantic retrieval and narrative generation capabilities. The memory corpus comprises 5,000–50,000 indexed entries collected over extended temporal horizons of 6–24 months, ensuring sufficient temporal diversity to capture long-range semantic relationships inherent in realistic personal memory accumulation. Data modality representation follows a balanced distribution: free-form text entries constitute 50% of the corpus, automatic speech-to-text transcriptions account for 25%, images with algorithmically generated captions comprise 15%, and automated video transcriptions with scene descriptions represent the remaining 10%. This multimodal composition mirrors the heterogeneous nature of human memory encoding across different sensory and communicative channels. The domain coverage intentionally spans diverse semantic contexts—including personal experiences, travel narratives, interpersonal interactions, emotional reflections, professional activities, and significant life events—to rigorously test the system's generalization capacity across varied cognitive and contextual dimensions. The evaluation protocol employs four distinct query categories to comprehensively stress-test system capabilities: semantic memory queries that target specific affective states or thematic elements (e.g., "memories when I felt anxious about my career" or "times I traveled to Europe"), temporal range queries that filter retrieval within specified date boundaries often combined with semantic constraints, aggregation tasks requiring synthesis and summarization of multiple semantically related memories (e.g., "Provide a timeline of relationship milestones"), and narrative generation tasks prompting the system to construct coherent, reflective narratives from fragmented retrieved memory elements. This multifaceted experimental design ensures comprehensive validation of the system's capacity to integrate semantic, temporal, and narrative dimensions of autobiographical memory retrieval.

B. Quantitative Results

1) Retrieval System Performance

Semantic retrieval performance is evaluated across factual, thematic, and affect-based query types. Representative results are shown below:

Metric	Factual	Thematic	Emotional
Precision@5	0.90	0.85	0.82
Recall@10	0.92	0.88	0.85
MAP	0.89	0.83	0.80

2) Event Detection Efficiency

Performance for categorizing and clustering multimodal data into meaningful life events is summarized as follows:

Event Category	Precision	Recall	F1-Score
Milestones	0.94	0.91	0.92
Relationships	0.88	0.86	0.87
Thematic	0.85	0.83	0.84

3) Summarization Quality

ROUGE scores for short-form and long-form narrative synthesis are shown below:

Narrative Type	ROUGE-1 F1	ROUGE-2 F1	ROUGE-L F1
Short Narratives	0.49	0.22	0.45
Long Narratives	0.46	0.20	0.43

C. Qualitative Analysis

1) Narrative Coherence, Fluency, and Emotional Resonance

Human evaluators assess the system's outputs on readability and coherence, contextual accuracy, emotional salience, and personalization fidelity.

Result: Preliminary evaluations indicate strong performance in producing coherent and emotionally engaging narratives tailored to the subject's lived experience.

2) Usability and User Experience

User studies measure task efficiency, interface intuitiveness, and satisfaction. High usability scores are anticipated due to the system's emphasis on minimal cognitive load and narrative-centric presentation.

D. Error Analysis

To ensure robustness, the following error sources were analyzed:

1. **Multimodal Conversion Errors:**

- Transcription/Captioning: Misinterpretations in speech-to-text or visual captioning lead to downstream semantic inconsistencies.
- Information Loss: Non-textual cues (e.g., tone, visual context) may not be fully preserved.

2. **Embedding and Retrieval Limitations:**

- Semantic Drift: Failures in maintaining conceptual coherence across temporally distant memories.
- Query Ambiguity: Complex or abstract queries may yield inconsistent retrieval.

3. **Narrative Generation Errors:**

- Hallucinations: Fabricated or factually incorrect details introduced during generative synthesis.
- Bias Propagation: Unintended demographic or contextual biases inherited from training data.
- Coherence Failures: Logical discontinuities in long-form narratives.

E. Comparative Benchmarking

- **Against Traditional Lifelogging Systems:** The system provides superior semantic structure, event detection accuracy, and narrative generation compared to platforms emphasizing raw data capture.

- **Against Keyword-Based Archiving Solutions:** Meaning-based retrieval enables superior recall and precision for abstract, thematic, or emotional queries.
- **Against Basic Generative Models:** Transformer-based narrative synthesis outperforms rule-based summarizers and simpler generative models in fluency, contextual relevance, and personalization.

F. Summary of Findings

The experimental evaluation indicates that the proposed system:

- Effectively integrates multimodal data and constructs a semantically rich vectorized memory archive.
- Generates coherent, emotionally resonant narratives that enhance reminiscence and intergenerational understanding.
- Delivers strong user engagement through intuitive interfaces and high-quality outputs.
- Implements responsible AI governance through privacy safeguards and hallucination mitigation.

Overall, the system demonstrates substantial promise as a next-generation platform for digital legacy management and intergenerational knowledge transmission.

CONCLUSION

The proposed system offers a structured long‑term personal memory preservation by transforming fragmented digital traces into coherent, queryable life narratives. It addresses the limitations of existing lifelogging and digital archiving approaches, which tend to focus on raw data capture and storage without delivering semantically organized, narratively meaningful representations that support intergenerational use.

The framework advances a shift from passive digital inheritance to active narrative mediation. Multimodal inputs—text, audio, images, and video—are first normalized into a unified textual abstraction, ensuring that heterogeneous memories contribute to a common semantic space. Transformer-based embedding models then encode these memories as high‑dimensional vectors that capture contextual and affective nuances, which are stored and indexed in vector databases to enable efficient approximate nearest‑neighbor retrieval. Through this architecture, users can issue meaning‑based queries (e.g., by theme, relationship, or emotional tone) rather than relying solely on exact keywords or rigid metadata.

Building on this semantic substrate, the system incorporates event detection and abstractive summarization modules to construct structured, human‑interpretable accounts of a person's life. By aggregating related memories and generating coherent narratives, it moves beyond isolated events to provide curated, emotionally resonant overviews that better reflect lived experience. This positions the system as both an information retrieval engine and a narrative synthesis mechanism.

Equally central is the integration of security, privacy, and ethical safeguards from the outset. The design emphasizes secure storage, fine‑grained access control, and mitigation of risks such as data leakage and model hallucination, while explicitly engaging with questions of consent, posthumous control, authenticity, and bias. An "ethics‑by‑design" orientation ensures that technical capabilities are aligned with user autonomy, psychological well‑being, and social acceptability.

Overall, the proposed architecture represents a significant step toward AI‑mediated digital legacies: it unifies multimodal ingestion, semantic indexing, and generative storytelling within a principled ethical framework. As the underlying models, storage technologies, and governance practices evolve, this approach can provide a robust foundation for preserving, navigating, and sharing personal histories in ways that support meaningful intergenerational connection rather than merely accumulating digital residues.

REFERENCES

1. Assmann, J., & Czaplicka, J. (1995). *Collective Memory and Cultural Identity.*
2. Odom, W., et al. (2010). *Designing for Slowness, Anticipation and Re-visitation*. CHI.
3. Kirk, D.S., & Sellen, A. (2010). *On Human Remains: Values and Practice in the Home Archiving of Cherished Objects*. TOCHI.
4. Massimi, M., & Baecker, R.M. (2010). *A Death in the Family: Opportunities for Designing Technologies for the Bereaved*. CHI.
5. Moncur, W., et al. (2012). *When Personal Becomes Professional: Exploring the Digital Legacies of Professionals*. CSCW.
6. M. Bolaños, M. Dimiccoli, and P. Radeva, "Toward storytelling from visual lifelogging: An overview," *IEEE Trans. Human-Machine Systems*, vol. 47, no. 1, pp. 77–90, 2016.
7. K. Nowaczyk-Basińska et al., "Safeguards and ethical frameworks for AI 'deadbots' and digital afterlife platforms," Leverhulme Centre for the Future of Intelligence, Univ. of

Cambridge, 2023–2024.

8. A. Jangra, A. Jatowt, et al., "A survey on multimodal summarization," arXiv preprint arXiv:2109.05199, 2021.
9. D. Ribeiro, D. Santos, and A. Rodrigues, "Lifelog retrieval from daily digital data: A narrative review," *JMIR mHealth and uHealth*, vol. 10, no. 5, e30517, 2022.
10. T. Taipalus, M. Niinimäki, and J. Laaksonen, "Vector database management systems: Fundamental concepts, use-cases, and current challenges," *Cognitive Systems Research*, 2024 (in press).
11. J. Vincent, "Reporting on AI 'deadbots,' griefbots, and digital resurrection," *The Verge*, 2022–2024.
12. J. N. Li, Z. Zhang, and J. Ma, "OmniQuery: Contextually Augmenting Captured Multimodal Memory to Enable Personal Question Answering," arXiv (Cornell University), Sep. 2024. DOI: 10.48550/arxiv.2409.08250.
13. D. Dang-Nguyen, M. A. Riegler, L. Zhou, and C. Gurrin, "Challenges and Opportunities within Personal Life Archives," *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval (ICMR '18)*, vol. 176, p. 335, Jun. 2018. DOI: 10.1145/3206025.3206040.
14. D. Cai, S. Wang, P. Chen, Z. Zhang, and M. Xu, "Recall: Empowering Multimodal Embedding for Edge Devices," arXiv (Cornell University), Sep. 2024. DOI: 10.48550/arxiv.2409.15342.
15. N. Alam and Y. Graham, "Memento: a prototype search engine for LSC 2021," *Multimedia Tools and Applications*, vol. 82, no. 24, p. 37807, Apr. 2023. DOI: 10.1007/s11042-023-15067-9.
16. S. Jabeen, X. Li, A. M. Shoib, S. Li, and J. Abdul, "Recent Advances and Trends in Multimodal Deep Learning: A Review," arXiv (Cornell University), Jan. 2021. DOI: 10.48550/arxiv.2105.11087.
17. S. Jabeen, X. Li, M. S. Amin, O. E. F. Bourahla, S. Li, and A. Jabbar, "A Review on Methods and Applications in Multimodal Deep Learning," *ACM Transactions on Multimedia Computing Communications and Applications*, vol. 19, Oct. 2022. DOI: 10.1145/3545572.
18. T. Taipalus, "Vector database management systems: Fundamental concepts, use-cases, and current challenges," *Cognitive Systems Research*, vol. 85, p. 101216, Feb. 2024. DOI: 10.1016/j.cogsys.2024.101216.
19. Y. Han, C. Liu, and P. Wang, "A Comprehensive Survey on Vector Database: Storage

and Retrieval Technique, Challenge," arXiv (Cornell University), Oct. 2023. DOI: 10.48550/arxiv.2310.11703.

20. B. Plaie, S. N. Jyesta, and S. Withana, "Vector embedding of multi-modal texts: a tool for discovery?", arXiv (Cornell University), Sep. 2025. DOI: 10.48550/arxiv.2509.08216.
21. D. Cai et al., "Ubiquitous memory augmentation via mobile multimodal embedding system," *Nature Communications*, vol. 16, no. 1, Jun. 2025. DOI: 10.1038/s41467-025-60802-5.
22. H. Hao, J. Han, C. Li, Y. Li, and X. Yue, "RAP: Retrieval-Augmented Personalization for Multimodal Large Language Models," arXiv (Cornell University), Oct. 2024. DOI: 10.48550/arxiv.2410.13360.
23. H. Hao, J. Han, C. Li, Y.-F. Li, and X. Yue, "RAP: Retrieval-Augmented Personalization for Multimodal Large Language Models," p. 14538, Jun. 2025.
24. L. Patel, S. N. Jha, C. Guestrin, and M. Zaharia, "Semantic Operators: A Declarative Model for Rich, AI-based Data Processing," arXiv (Cornell University), Jul. 2024. DOI: 10.48550/arxiv.2407.11418.
25. S. Kang, Y. Zhang, P. Jiang, D. Lee, J. Han, and H. Yu, "Taxonomy-guided Semantic Indexing for Academic Paper Search," arXiv (Cornell University), Oct. 2024. DOI: 10.48550/arxiv.2410.13918.
26. Y. Tang et al., "Semantic-Enhanced Differentiable Search Index Inspired by Learning Strategies," *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Aug. 2023.
27. S. Hoseini, J. Lipp, and C. Quix, "A survey on semantic data management as intersection of ontology-based data access, semantic modeling and data lakes," *Journal of Web Semantics*, vol. 81, p. 100819, Apr. 2024.
28. W. Joel, "MemoriesDB: A Temporal-Semantic-Relational Database for Long-Term Agent Memory / Modeling Experience as a Graph of Temporal-Semantic Surfaces," arXiv, Oct. 2025. DOI: 10.48550/arxiv.2511.06179.
29. M. Mondal, P. Cudré-Mauroux, H.-Â. Sandlin, and P. Cudré-Mauroux, "A survey of multimodal event detection based on data fusion," *The VLDB Journal*, vol. 34, no. 1, Dec. 2024.
30. V. Kalokyri, A. Borgida, and A. Marian, "Supporting Human Memory by Reconstructing Personal Episodic Narratives from Digital Traces," *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 16, p. 453, May 2022.

31. H. Xu, P. Tong, M. Li, and M. Srivastava, “AutoLife: Automatic Life Journaling with Smartphones and LLMs,” arXiv (Cornell University), Dec. 2024. DOI: 10.48550/arxiv.2412.15714.
32. P.-W. Kao, A.-Z. Yen, H. Huang, and H. Chen, “ConvLogMiner: A Real-Time Conversational Lifelog Miner,” *IJCAI 2021 Proceedings*, p. 4992, Aug. 2021. DOI: 10.24963/ijcai.2021/710.
33. Y.-P. Chen, A.-Z. Yen, H. Huang, H. Nakayama, and H. Chen, “LED: A Dataset for Life Event Extraction from Dialogs,” 2023, p. 384, Jan. 2023.
34. Q. Xu et al., “Lifelog Image Retrieval Based on Semantic Relevance Mapping,” *ACM Transactions on Multimedia Computing, Communications and Applications*, vol. 17, no. 3, Jul. 2021. DOI: 10.1145/3446209.
35. G. Savcinsen et al., “Using sequences of life-events to predict human lives,” *Nature Computational Science*, vol. 4, no. 1, p. 43, Dec. 2023.
36. B. S. Santana, R. Campos, E. Amorim, A. Jorge, P. Silvano, and S. Nunes, “A survey on narrative extraction from textual data,” *Artificial Intelligence Review*, vol. 56, no. 8, p. 8939, Jan. 2023. DOI: 10.1007/s10462-022-10383-7.
37. S. Talaei, M. Li, K. Grover, J. K. Hippler, D. Yang, and A. Saberi, “StorySage: Conversational Autobiography Writing Powered by a Multi-Agent Framework,” arXiv (Cornell University), Jun. 2025. DOI: 10.48550/arxiv.2506.11540.
38. N. Gutherlé, A. Doucet, and A. Jatowt, “Archive TimeLine Summarization (ATLS): Conceptual Framework for Timeline Generation over Historical Document Collections,” arXiv (Cornell University), Jan. 2023.
39. N. Gutherlé, A. Doucet, and A. Jatowt, “Archive TimeLine Summarization (ATLS): Conceptual Framework for Timeline Generation over Historical Document Collections,” HAL (Le Centre pour la Communication Scientifique Directe), Oct. 2022.
40. J. L. Jeung and Y. Huang, “Unlocking Memories with AI: Exploring the Role of AI-Generated Cues in Personal Reminiscing,” May 2024.
41. Y. Jin, W. Cai, L. Chen, Y. Zhang, G. Doherty, and T. Jiang, “Exploring the Design of Generative AI in Supporting Music-based Reminiscence for Older Adults,” May 2024.
42. P. Pataranutaporn et al., “Living Memories: AI-Generated Characters as Digital Mementos,” *IUI ’23 Proceedings*, p. 889, Mar. 2023. DOI: 10.1145/3581641.3584065.
43. J. J. Pan, J. Wang, and G. Li, “Vector Database Management Techniques and Systems,” tutorial, May 2024.
44. K. Hatalis et al., “Memory Matters: The Need to Improve Long-Term Memory in LLM-

Agents," *Proceedings of the AAAI Symposium Series*, vol. 2, no. 1, p. 277, Jan. 2024.

45. G. He et al., "Zero-Indexing Internet Search Augmented Generation for Large Language Models," arXiv preprint arXiv:2411.19478, 2024.
46. J. Shen, J. J. Dudley, and P. O. Kristensson, "Encode-Store-Retrieve: Augmenting Human Memory through Language-Encoded Egocentric Perception," arXiv, 2023.
47. R. Wu, J. Meng, J. Xu, H. Wang, and K. Rong, "Rethinking Similarity Search: Embracing Smarter Mechanisms over Smarter Data," arXiv (Cornell University), Jan. 2023. DOI: 10.48550/arxiv.2308.00909.
48. L. Patel, P. Kraft, C. Guestrin, and M. Zaharia, "ACORN: Performant and Predicate-Agnostic Search Over Vector Embeddings and Structured Data," arXiv (Cornell University), Mar. 2024.
49. Y. Su et al., "Modernizing Facebook Scoped Search: Keyword and Embedding Hybrid Retrieval with LLM Evaluation," arXiv preprint arXiv:2509.13603, 2025.
50. S. B. Prajapati, "Data Privacy and Security in AI: Strategies for protecting user data while maintaining the functionality and scalability of AI solutions," *World Journal of Advanced Research and Reviews*, vol. 25, no. 1, p. 2142, Jan. 2025.
51. H. D. Bell, J. Kwesi, H. Laabdali, and P. Emami-Naeini, "'Is it always watching? Is it always listening?' Exploring Contextual Privacy and Security Concerns Toward Domestic Social Robots," arXiv (Cornell University), Jul. 2025.
52. D. Korobenko, A. Nikiforova, and R. Sharma, "Towards a Privacy and Security-Aware Framework for Ethical AI: Guiding the Development and Assessment of AI Systems," p. 740, Jun. 2024.
53. C. Sandeepa, B. Siniarski, N. Kourtellis, S. Wang, and M. Liyanage, "A Survey on Privacy of Personal and Non-Personal Data in B5G/6G Networks," *ACM Computing Surveys*, vol. 58, no. 10, 2024.
54. S. Yunusov, H. Sidat, and A. Emami, "MirrorStories: Reflecting Diversity through Personalized Narrative Generation with Large Language Models," arXiv (Cornell University), Sep. 2024. DOI: 10.48550/arxiv.2409.14845.
55. V. L. Lemieux et al., "Clio-X: A Web3 Solution for Privacy-Preserving AI Access to Digital Archives," arXiv (Cornell University), Jul. 2025.
56. S. Duan, M. Khona, A. Iyer, R. Schaeffer, and I. Fiete, "Uncovering Latent Memories: Assessing Data Leakage and Memorization Patterns in Frontier AI Models," arXiv (Cornell University), Jun. 2024. DOI: 10.48550/arxiv.2406.14549.
57. V. Smith, A. S. Shamsabadi, C. Ashurst, and A. Weller, "Identifying and Mitigating

Privacy Risks Stemming from Language Models: A Survey," arXiv (Cornell University), Sep. 2023. arxiv:2310.01424.

58. S. Zhang et al., "'Ghost of the past': identifying and resolving privacy leakage from LLM's memory through proactive user interaction," arXiv (Cornell University), Oct. 2024. arxiv:2410.14931.
59. E. Lee, "Towards Ethical Personal AI Applications: Practical Considerations for AI Assistants with Long-Term Memory," 2024.
60. E. Lee, "Towards Ethical Personal AI Applications: Practical Considerations for AI Assistants with Long-Term Memory," arXiv (Cornell University), Sep. 2024. arxiv:2409.11192.
61. N. Maleki, B. Padmanabhan, and K. Dutta, "AI Hallucinations: A Misnomer Worth Clarifying," 2024.
62. A. L. Smith, F. Greaves, and T. Panch, "Hallucination or Confabulation? Neuroanatomy as metaphor in Large Language Models," *PLOS Digital Health*, vol. 2, no. 11, Nov. 2023.
63. S. Barros, "I Think, Therefore I Hallucinate: Minds, Machines, and the Art of Being Wrong," 2025.
64. P. Pataranutaporn, C. Archiwaranguprok, S. Chan, E. Loftus, and P. Maes, "Synthetic Human Memories: AI-Edited Images and Videos Can Implant False Memories and Distort Recollection," 2024.
65. R. McNamara, "Virtual Representations and Their Ethical Implications," *Baltic Screen Media Review*, vol. 11, no. 1, p. 198, Dec. 2023. DOI: 10.2478/bsmr-2023-0011.
66. R. A. Gonzalez and S. DiPaola, "Exploring Augmentation and Cognitive Strategies for AI-based Synthetic Personae," arXiv (Cornell University), Apr. 2024. arxiv:2404.10890.
67. P. N. Venkit, J. Li, Y. Zhou, S. Rajtmajer, and S. Wilson, "A Tale of Two Identities: An Ethical Audit of Human and AI-Crafted Personas," arXiv, 2025. arxiv:2505.07850.
68. A. Reeves, A. Shaghaghi, S. Krebs, and D. Ashenden, "Data After Death: Australian User Preferences and Future Solutions to Protect Posthumous User Data," in *IFIP Advances in Information and Communication Technology*, Springer Science+Business Media, 2024.
69. J. Holt, J. D. Smeddinck, J. Nicholson, V. Vlachokyriakos, and A. Durrant, "Post-mortem information management: exploring contextual factors in appropriate personal data access after death," *Human–Computer Interaction*, Jan. 2024.
70. A. Reeves, A. Shaghaghi, S. Krebs, and D. Ashenden, "Data After Death: Australian

User Preferences and Future Solutions to Protect Posthumous User Data," arXiv (Cornell University), Jul. 2024.

71. G. Rolan and A. Lewis, "The perpetual twilight of records: consentful recordkeeping as moral defence," *Archival Science*, vol. 24, no. 2, p. 257, Apr. 2024.

72. E. Harbinja, T. Morse, and L. Edwards, "Digital Remains and Post-mortem Privacy in the UK: What do users want?," *SSRN Electronic Journal*, Jan. 2024.

73. J. Danaher and S. Nyholm, "The ethics of personalised digital duplicates: a minimally viable permissibility principle," *AI and Ethics*, vol. 5, no. 2, p. 1703, Jul. 2024.

74. T. Hollanek and K. Nowaczyk-Basińska, "Griefbots, Deadbots, Postmortem Avatars: on Responsible Applications of Generative AI in the Digital Afterlife Industry," *Philosophy & Technology*, vol. 37, no. 2, May 2024.

75. X. Chen, A. Shaghaghi, J. Laeuchli, and S. S. Kanhere, "Beyond Life: A Digital Will Solution for Posthumous Data Management," arXiv (Cornell University), Jan. 2025. arxiv:2501.04900.

76. A. Vaswani et al., "Attention Is All You Need," arXiv (Cornell University), Jan. 2017. arxiv:1706.03762.

77. R. E. Turner, "An Introduction to Transformers," arXiv (Cornell University), Apr. 2023. arxiv:2304.10557.

78. R. M. Aliguliyev, "A new sentence similarity measure and sentence based extractive technique for automatic text summarization," *Expert Systems with Applications*, vol. 36, no. 4, p. 7764, Dec. 2008.

79. F. Tessari, K. Yao, and N. Hogan, "Surpassing Cosine Similarity for Multidimensional Comparisons: Dimension Insensitive Euclidean Metric," 2024.

80. M. A. Salam, M. Gamal, H. F. A. Hamed, and S. Sweidan, "Abstractive text summarization using deep learning models: a survey," *International Journal of Data Science and Analytics*, vol. 20, no. 5, p. 4209, Mar. 2025.

81. Y. Sunusi, N. Omar, and L. Q. Zakaria, "Exploring Abstractive Text Summarization: Methods, Dataset, Evaluation, and Emerging Challenges," *International Journal of Advanced Computer Science and Applications*, vol. 15, no. 7, Jan. 2024.

82. H. Shakil, A. Farooq, and J. Kalita, "Abstractive text summarization: State of the art, challenges, and improvements," *Neurocomputing*, vol. 603, p. 128255, Jul. 2024.

83. R. Nallapati, B. Zhou, C. dos Santos, Ç. Gülçehre, and B. Xiang, "Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond," 2016.

84. Y. M. A. Algani, "A Novel Deep Learning Attention Based Sequence to Sequence

Model for Automatic Abstractive Text Summarization," *International Journal of Information Technology*, vol. 16, no. 6, p. 3597, May 2024.

85. F. Zhao, C. Zhang, and B. Geng, "Deep Multimodal Data Fusion," *ACM Computing Surveys*, vol. 56, no. 9, Sep. 2024.
86. Q. Zhang et al., "Provable Dynamic Fusion for Low-Quality Multimodal Data," *Proceedings of the 40th International Conference on Machine Learning (ICML 2023)*, 2023.
87. N. Sadras, B. Pesaran, and M. M. Shanechi, "Event detection and classification from multimodal time series with application to neural data," *Journal of Neural Engineering*, 2024.
88. S. Ji, S. Pan, E. Cambria, P. Marttinen, and P. S. Yu, "A Survey on Knowledge Graphs: Representation, Acquisition, and Applications," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 2, p. 494, Apr. 2021.
89. T. Gebhart, J. Hansen, and P. Schrater, "Knowledge Sheaves: A Sheaf-Theoretic Framework for Knowledge Graph Embedding," arXiv (Cornell University), Feb. 2022. arxiv:2110.03789.
90. J. Cao, J. Fang, Z. Meng, and S. Liang, "Knowledge Graph Embedding: A Survey from the Perspective of Representation Spaces," *ACM Computing Surveys*, vol. 56, no. 6, 2024.