

# Semantic Web – A Forgotten Wave of Artificial Intelligence?

**Abstract.** The history of Artificial Intelligence (AI) is a narrative of waves—rising optimism followed by crashing disappointments. AI winters, such as the early 2000s, are often remembered as barren periods of innovation. This paper argues that such a perspective overlooks a crucial wave of AI that seems to be forgotten: the rise of the Semantic Web, which is based on knowledge representation, logic, and reasoning, and its interplay with intelligent Software Agents. Fast forward to today, and ChatGPT has reignited AI enthusiasm, built on deep learning and advanced neural models. However, before Large Language Models (LLMs) dominated the conversation, another ambitious vision emerged—one where AI-driven Software Agents autonomously served Web users based on a structured, machine-interpretable Web. The Semantic Web aimed to transform the World Wide Web into an ecosystem where AI could reason, understand, and act. Between 2000 and 2010, this vision sparked a significant research boom, only to fade into obscurity as AI’s mainstream narrative shifted elsewhere. Today, as LLMs edge toward autonomous execution, we revisit this overlooked wave. By analyzing its academic impact through bibliometric data, we highlight the Semantic Web’s role in AI history and its untapped potential for modern Software Agent development. Recognizing this forgotten chapter not only deepens our understanding of AI’s cyclical evolution but also offers key insights for integrating emerging technologies.

**Keywords.** Artificial Intelligence, Software Agents, Semantic Web, World Wide Web, Hypertext

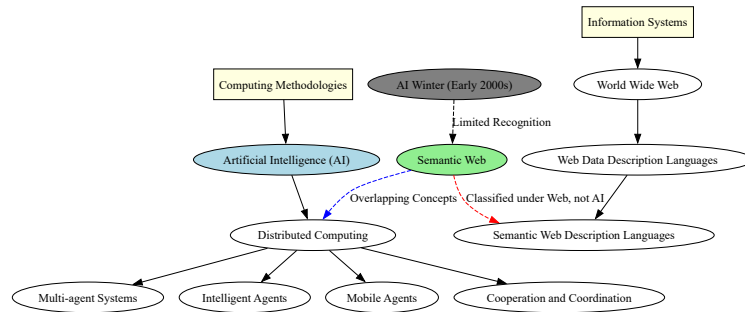
## 1. Introduction

The notion of Artificial Intelligence (AI) as a distinct discipline within Computer Science was introduced in 1956 at a workshop at Dartmouth College by early pioneers such as John McCarthy, Marvin Minsky, and others. Since then, the history of AI has experienced recurring cycles of enthusiasm and subsequent disillusionment, commonly referred to as “AI summers” and “AI winters” [1]. An AI winter refers to a period of reduced funding, interest, and progress in AI research, typically following unmet expectations or technological limitations. In contrast, an AI summer denotes a phase of heightened enthusiasm, investment, and innovation, often driven by breakthroughs in algorithms, computing power, or applications. These cycles are associated with two main approaches to creating AI: symbolic (“white box”) approaches, based on explicit knowledge representation, logic, reasoning, and search, and sub-symbolic (“black box”) approaches, based on neural networks, machine learning, and statistical models.

Following the first AI summer in the 1960s, the first AI winter occurred in the 1970s after it was argued in [2] that sub-symbolic neural networks (then based on single-layer perceptron) would never be useful for solving real-world tasks. As a response, the symbolic approach, including knowledge-based expert systems and logic programming, fu-

elled the next AI summer in the 1980s. However, challenges in, e.g., knowledge representation and acquisition [3] soon led to another AI winter. Research on multi-layer sub-symbolic neural models and learning systems continued despite—or perhaps because of—the shortcomings of the knowledge-based approach. By expanding machine learning models through deep learning, employing new training algorithms such as backpropagation, introducing novel neural architectures and non-linear activation functions, leveraging Big Data from the Web and available databases for training, , and by using new efficient AI chips for parallel neural computations AI has entered an unprecedented summer in the 2020s [4]. A prominent research topic today, aimed at addressing challenges in neural systems such as hallucinations, is the combination of symbolic and sub-symbolic approaches in neuro-symbolic (hybrid) AI systems [5,6].

The early 2000s are often labelled as an AI winter, with little recognition of the relevant AI research conducted during that time. However, this narrative overlooks significant developments from that era, including the emergence of the Semantic Web, which gained momentum in 2001 [7,8]. Although rooted in knowledge representation, logic, reasoning, and ontologies—core subjects of traditional symbolic AI—this branch of research is typically not classified as part of AI. For example, in the authoritative ACM Computing Classification System<sup>1</sup> of Computer Science, AI is a major first-level category under Computing Methodologies while Semantic Web is only mentioned under the minor category under Information Systems. According to the original Semantic Web vision outlined in 2001 [7], the Semantic Web is centered around intelligent agents on the Web. The category of Intelligent Agents can be found as a minor subcategory under Artificial Intelligence. This categorization issue is also illustrated in Figure 1.



**Figure 1.** Categorization issue: Semantic Web research vision had strong artificial intelligence focus but it is categorized under WWW branch

Even review papers on the Semantic Web do not focus on the AI or Software Agent research conducted during that period within the Semantic Web domain. For example, Hitzler [9] provides a comprehensive review of the Semantic Web, detailing its evolution, core principles, and ongoing challenges. The paper highlights how the Semantic Web extends the traditional Web by incorporating structured and linked data to enhance machine interoperability. Key technologies [10] such as RDF (Resource Description Framework), OWL (Web Ontology Language), and SPARQL are discussed, emphasizing their roles in enabling semantic interoperability. In addition, the paper explores real-world ap-

<sup>1</sup>ACM Computing Classification System: <https://dl.acm.org/ccs>

plications in domains such as AI, knowledge graphs, and data integration. Hitzler also addresses challenges such as scalability, reasoning complexity, and data quality, underscoring the need for continued research and development to fully realize the Semantic Web's potential.

The rapid advancement of Generative Artificial Intelligence (GAI) and the widespread adoption of Large Language Models (LLMs) were catalyzed by the introduction of the Transformer architecture, which enabled more efficient training and superior contextual understanding in deep learning models [11,4]. This breakthrough laid the foundation for applications such as ChatGPT, which demonstrated the potential of LLMs in generating human-like text, leading to a surge in research and commercial interest. More recently, the focus has expanded beyond standalone LLMs to the integration of Software Agents and agentic capabilities, where models exhibit autonomy, goal-directed behaviour, and enhanced reasoning abilities [12,13]. These agentic systems, which incorporate planning, memory, and interaction mechanisms, are now being extensively studied to augment the functionality of LLMs, enabling them to operate dynamically in complex environments, solve multi-step problems, and interact more effectively with users and digital ecosystems.

It appears that for some reason the Semantic Web is a forgotten wave of AI although the current AI trend is very much happening on the Web. To study and analyze this argument, we review previous trend analyses related to the Semantic Web and AI and present statistical analyses of academic literature on these topics.

The remainder of this paper is structured as follows.

- Section 2 provides an overview of the AI hype cycle, detailing the recurring patterns of “AI summers” and “AI winters”. We review and synthesize insights from previous studies on AI history and introduce a normalized “temperature” scale to quantify these fluctuations.
- Section 3 explores the concept of the Semantic Web, its foundational vision, and its relationship with AI-driven Software Agents.
- Section 4 presents an analysis of bibliometric data and academic trends related to Semantic Web and Software Agent research. We examine search term frequencies and notable researcher contributions to illustrate the presence of the “Forgotten AI Wave”.
- Section 5 provides a critical discussion of our findings, highlighting the broader implications of the legacy of the Semantic Web, its challenges, and its potential resurgence in modern AI research.
- Section 6 summarizes the key takeaways from our study and outlines directions for future research at the intersection of the Semantic Web and contemporary AI advancements.


## 2. AI Summers and Winters

In this section, we first discuss the general hype cycle model, a framework used to describe the waves of enthusiasm followed by subsequent disappointment that technological fields often experience. This pattern, commonly referred to as the “hype cycle,” has been observed in various industries and technologies. We then review and synthesize illustrations of AI history, exploring the alternate periods of “AI winters” and “AI

summers” and placing them on a normalized scale. These phases highlight the cycles of overenthusiasm, disillusionment, and eventual productivity, from the early promises of symbolic AI to the breakthroughs of the deep learning revolution.

**The Hype Cycle** The hype cycle model, introduced by Gartner Inc. in 1995, [14] illustrates the evolution of technological expectations over time. As shown in Figure 2 [14], the model is formed by merging two distinct curves: a human-centric expectation curve, represented by the dashed blue line, and a classical technology maturity S-curve, represented by the dotted green line. The first equation describes the hype level, which follows a bell-shaped curve driven by excitement, social contagion, and heuristic decision making. This curve captures the tendency of people to overestimate the potential of a technology, leading to an initial peak of inflated expectations, followed by a sharp decline when real-world limitations become evident.

### Hype Cycle



**Figure 2.** Hype Cycle: Hype Level, Engineering or Business Maturity and combination Hype Cycle [14]

The second equation, the S-curve of technological maturity, represents the gradual advancement of a technology over time. Initially, progress is slow due to limited understanding and minimal performance gains from early investments. However, as knowledge accumulates and improvements accelerate, the technology experiences rapid development until it eventually stabilizes at a plateau defined by its inherent limitations. Together, these two curves form the hype cycle (solid gray line), which models the typical trajectory of emerging technologies, highlighting both the initial overhype and the eventual stabilization at a realistic level of maturity. The stages outlined in the hype cycle depicted in Figure 2 are:

1. **Innovation Trigger:** A new technology or concept emerges, generating early excitement and media interest, although practical applications are not yet fully understood.
2. **Peak of Inflated Expectations:** Media hype and early success stories fuel high expectations, leading to overenthusiasm and unrealistic predictions about the technology’s potential, with many projects launched but most failing to deliver.
3. **Trough of Disillusionment:** Initial excitement fades as challenges and limitations become evident, causing skepticism due to failed projects and unmet expectations, leading to a decline in interest and reduced investments.

4. **Slope of Enlightenment:** A gradual understanding of the technology's real potential and limitations develops, leading to practical applications and increasing adoption as early adopters recognize value in realistic use cases.
5. **Plateau of Productivity:** The technology becomes mainstream and widely adopted, with stable, practical use cases delivering real benefits, achieving sustainable growth and a proven track record.

Although the hype-cycle is typically illustrated as a single expectation curve, it actually overlays the S-curve of technology diffusion; how these two curves interact determines both the tempo and the amplitude of successive boom-and-bust phases, yet this relationship remains only sketchily articulated in the literature. Consequently, existing research offers limited guidance on when firms and investors should initiate, scale, or exit their involvement in order to capture upside while containing downside risk. Closing this gap will require analytical frameworks that fuse expectation dynamics with adoption kinetics and strategic positioning along the hype-cycle.

**Review and Synthesis of AI History Illustrations** In the following we synthesise evidence from peer-reviewed papers, technical blogs, and public presentations to trace the recurring boom-and-bust pattern that has characterised artificial-intelligence research and commercialisation. Importantly, AI is not a single invention but a shifting portfolio of techniques – from symbolic reasoning and expert systems to machine learning, deep learning, and today's large-scale generative models. Each element in this portfolio follows its own trajectory, so the broad history of AI is best understood as the superposition of many overlapping hype cycles rather than one monolithic wave.

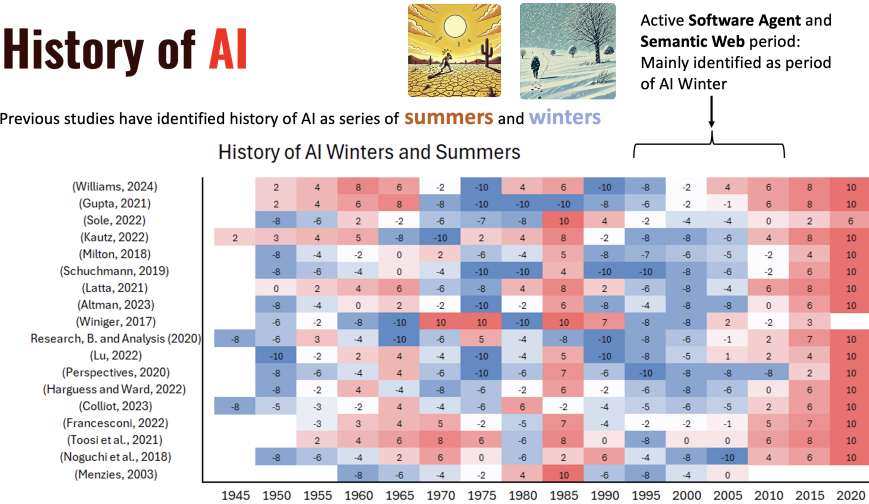
Take *software agents* as an illustrative example. Introduced in the early 1990s, agent technology ascended to widespread media enthusiasm, encountered technical and market barriers, resurfaced with the rise of web services, and is now experiencing renewed interest through autonomous AI assistants. In other words, the same underlying concept has already traversed the Gartner hype-cycle several times.

Viewed in aggregate, these technology-specific curves explain why the field repeatedly climbs to a *Peak of Inflated Expectations*, descends into a *Trough of Disillusionment*, and, where genuine utility is present, climbs the *Slope of Enlightenment* to a durable *Plateau of Productivity*. Recognising AI as a constellation of evolving technologies helps clarify why optimism and disappointment can coexist—and why practical value often emerges only after multiple passes through the cycle.

All available AI history illustrations were analyzed, as listed below: [15,16,17,18,19,20,21,22,23,24,25,26,27,28,29,30].

To quantify “summer” and “winter”, a grading scale from +10 to -10 was used, normalizing each illustration to this same scale. This means that all figures receive a maximum value of +10 when they visually reach their peak and -10 when they visually reach their lowest point. This grading serves as an indication of “temperature,” verbally expressed as “AI summer” (hot) or “AI winter” (cold). The aggregated illustration is presented in Figure 3. This figure also indicates that the period of active Semantic Web research is generally considered an AI winter. Furthermore, none of the studies mention either the topic *Semantic Web* or *Software Agents* as representing any form of AI boom in historical AI illustrations.

AI history illustrations known to the authors of this study depict the period from 1995 to 2010 as an AI winter rather than a time of substantial research advancements. These accounts consistently emphasize stagnation and minimal innovation, portraying



**Figure 3.** History of Artificial Intelligence as illustrated by previous studies. Period of active Semantic Web development is identified typically as AI Winter period

this period as one characterized by diminished funding, waning public interest, and a general lack of transformative breakthroughs in AI. While some research activities continued during this time, the prevailing perception is that AI did not experience any major surge or revolution.

Furthermore, none of the AI history illustrations reviewed in this study explicitly mention Software Agents or the Semantic Web. This absence suggests that these topics were not widely recognized or highlighted as pivotal elements within the broader narrative of AI development during this period. Their exclusion reinforces the notion that this era was marked more by stagnation than by the emergence of significant new paradigms or technologies.

3. The Semantic Web and Its Vision of AI Software Agents

**Hype Cycles: Software Agent Technology** Software agent technology has gone through a hype cycle a couple of times during its history. The following quote is from 1996:

”Agents are here to stay, not least because of their diversity, their wide range of applicability and the broad spectrum of companies investing in them. As we move further and further into the information age, any information-based organization which does not invest in agent technology may be committing commercial hara-kiri.” [31]

Authors of the above quote wrote another article a few years later [32] stating that the assumptions were too high. If anything, companies that invested in agent technology in the 1990s were perhaps the ones to commit a form of hara-kiri.

Recently, in December 2024, Microsoft CEO Satya Nadella stated that AI-driven software agents are fully revolutionizing business-to-business (B2B) applications. According to Nadella, AI-driven agents will take over business logic and enable new AI-

powered applications, placing software agents at the heart of the current AI hype. The full quote from an episode of the BG2 Pod series, where Nadella stated in an interview video [33], e.g., is:

“I think that the notion that business applications (B2B Apps) exist, that is probably where they all collapse, in the (Software) Agent-era.” – “Because if you think about it, they (B2B Apps) are essentially CRUD-databases with a bunch of business logic. The business logic will go to these Agents. These Agents are going to be multi-repo-CRUD so they are not going to discriminate between what the backend is. They are going to be updating multiple databases. All the (business) logic will be in the AI-tier (Agent-layer) so to speak.” “Once the AI-tier becomes the place where all the logic is, then people will start replacing the backends. – As we speak, I think we are seeing quite high rates of wins in our Dynamics backends and the Agent use. We will go very aggressive trying to collapse it all. Whether it is customer service, be it finance and operations. . .” – “People want more AI-native business apps. It means the logic tier can be orchestrated by AI and AI Agents in a very seamless way.”

Obviously, noting that agent technology has had inflated expectations in the past does not prove anything about the future. This time, it might be different, as the overall technology infrastructure has evolved significantly since the 1990s. We also have to remember that, for example, neural networks have experienced several hype cycles before the major breakthrough in the 2010s.

**Software Agents Before the Semantic Web** Long before the Web era, the conceptual foundations for software agents was established by research in *Distributed Artificial Intelligence* (DAI) and the *Actor model* [34]. DAI cast intelligence as an *emergent* phenomenon arising from multiple, loosely-coupled problem solvers that cooperate across a network, thereby elevating autonomy and negotiation to primary design concerns [35,36]. In parallel, the Actor model supplied a concrete computational metaphor of independent, concurrently executing entities that communicate solely through asynchronous message passing. Together these two traditions contributed the twin conceptual pillars of *distribution* and *concurrency*, which later researchers distilled into the now-standard agent properties of social ability and proactive–reactive behaviour [37].

The 1990s witnessed the emergence of software agents as a concept in computer science, driven by the need for autonomous and intelligent systems in increasingly complex environments. This decade saw the consolidation of core agent characteristics, including autonomy, proactivity, reactivity, and social ability [31,38]. The rapid expansion of the World Wide Web around 1994 underscored the potential for agents to operate within this new digital space [39]. Foundational research explored diverse aspects of agent technology, from theoretical underpinnings and formal methods to practical applications for information management and personalized assistance [38,39].

The development of various agent architectures, such as reactive, deliberative, and hybrid models, further illustrated the field’s commitment to creating sophisticated autonomous entities [40]. Researchers also focused on defining the boundaries of agency, distinguishing agents from conventional programs through taxonomies and formal definitions [41]. The intellectual roots of software agents were traced back to Distributed Artificial Intelligence (DAI) and Carl Hewitt’s Actor model from 1977, highlighting the long-standing interest in distributed and interacting intelligent systems [31,34,37].

**From Definitions to Standards.** The explicit taxonomies and formal definitions produced during this period played a pivotal role in *standardising* the concept of agency. By establishing a common vocabulary, they enabled the Foundation for Intelligent Physical Agents (FIPA) to publish its initial specifications—most notably the *FIPA Agent Communication Language* (FIPA ACL) and a reference architecture that has since underpinned almost every industrial-grade multi-agent platform [42]. These efforts, together with subsequent W3C initiatives (e.g. the *Agent Markup Language* submission), ensured that later Semantic-Web agents could interoperate across organisational and domain boundaries on the basis of well-defined semantics rather than ad-hoc conventions.

**The Vision of the Semantic Web and Intelligent Agents** Tim Berners-Lee, James Hendler, and Ora Lassila envisioned in their seminal 2001 paper “The Semantic Web” [7] an evolution of the World Wide Web into a more intelligent and interconnected system. This Semantic Web would be an infrastructure where information has explicitly defined meaning, facilitating human-computer cooperation and enabling machines to understand and process web data through a universal framework of metadata and ontologies. This involved enriching web resources with semantic markup using languages like RDF and OWL, making content machine-processable [43]. The goal was a transition from unstructured documents to interconnected data, enhancing knowledge sharing, reuse, and integration [44], aligning directly with AI’s goals of machine understanding and reasoning [7]. This highly influential paper has garnered over 30 000 citations by today.

Intelligent agents were conceived as crucial actors within this vision, able to automatically access, interpret, and reason about the semantically marked-up information [45]. These agents would act on behalf of users, performing tasks such as information retrieval, service discovery, automated decision-making (e.g., scheduling meetings or making purchases), and seamless interaction across different platforms [46]. Personal assistants, guided by rules, ontologies, and user profiles, were highlighted as prime examples of potential applications [47].

Realising this vision depends on the *ability of autonomous agents to converse*. The FIPA ACL specification meets this requirement by supplying a performative-based message envelope whose semantics are defined in a modal logic of communicative acts [48]. When the content field of such messages is expressed in RDF or OWL, the resulting utterances are simultaneously interpretable at the *data* and the *speech-act* levels, thereby enabling richly contextualised dialogues among heterogeneous Web agents [49]. Accordingly, ACLs form the critical bridge between the Semantic Web’s data layer and the multi-agent community’s interaction layer.

The paper’s emphasis on AI Software Agents, and the backgrounds of Hendler and Lassila in intelligent agents and knowledge representation, underscore the Semantic Web’s connection to advancing autonomous systems capable of reasoning, decision-making, and communication [7]. By leveraging standardized semantics, agents could interpret and act upon information with greater sophistication. This vision linked the future of the Semantic Web to developments in AI and positioned Software Agents as a critical component in realizing the Web’s full potential. While the full vision remains unrealized, it remains a cornerstone in the intersection of Semantic Web technologies and AI Software Agent development.

**Key Enabling Technologies and the Semantic Web Stack** Several technologies were crucial for realizing the vision of the Semantic Web and its integration with intelligent agents. The Resource Description Framework (RDF) provided a standard model



for data interchange on the web, enabling the expression of relationships between resources [50]. The Web Ontology Language (OWL), built upon RDF, was designed to represent rich and complex knowledge about entities, their classifications, and relationships [43,51]. Ontologies, formal and explicit specifications of shared conceptualizations, provided a common vocabulary and understanding of specific domains, which was crucial for enabling agents to interpret information consistently [52].

Agent Communication Languages (ACLs), such as the Foundation for Intelligent Physical Agents (FIPA) ACL [53], were recognized as essential for facilitating interaction and information exchange between autonomous agents. Furthermore, the *DAML-S* ontology—renamed *OWL-S* after the release of OWL [54]—was created to add a machine-interpretable semantic layer to conventional web service descriptions. It exposes three mutually supportive components: a *Service Profile* that advertises what the service does, a *Process Model* that specifies how the service can be executed, and a *Grounding* that binds abstract inputs and outputs to concrete message exchanges. Together these components allow an intelligent agent to match its goals to advertised profiles and thus discover relevant services, to use the process model and grounding information to construct and send the correct messages for direct invocation, and to chain multiple process models into coherent workflows for automatic composition of complex tasks. The development of DAML-S/OWL-S therefore represents a focused effort to supply the semantic infrastructure—rooted in ontologies, RDF, and OWL—along with the inter-agent communication mechanisms required for truly autonomous agents on the Semantic Web.

**Software Agents in the LLM Age 2017 onwards** The advent of large language models (LLMs) after the advent of the Transformer [11] has dramatically reshaped the landscape of software agent development, opening up new possibilities and challenges. Pre-LLM, agents often relied on meticulously crafted rules, ontologies, and knowledge bases, limiting their adaptability and generalizability [37,55]. While these approaches provided a degree of control and explainability, they struggled to handle the complexity and ambiguity of real-world scenarios. LLMs, trained on massive datasets, offer a fundamentally different approach. They possess an unprecedented ability to understand and generate natural language, reason, and even exhibit some degree of common sense, albeit imperfectly [56,57,58]. This inherent capability allows LLM-powered agents to interact with users and environments in a much more natural and flexible way, reducing the need for extensive manual programming and domain-specific knowledge engineering.

The integration of LLMs with agent frameworks has led to the emergence of “LLM-powered agents” or “LLM agents,” capable of performing a wider range of tasks than their predecessors. The simple definition for an LLM Agent is an artificial entity with prompt specifications (initial state), conversation trace (state), and ability to interact with the environments such as tool usage (action) [59,60]. These agents can leverage LLMs for various functions, including natural language understanding (NLU), natural language generation (NLG), planning, decision-making, and even tool use [61,62,63]. For instance, an LLM agent can be prompted to break down a complex task into smaller, manageable sub-tasks, search the web for relevant information, interact with APIs to execute actions, and summarize the results for the user [13,64]. Frameworks like LangChain [65] and AutoGen [66] provide tools and abstractions to facilitate the development and deployment of these agents, enabling developers to chain together LLM calls, external tools, and custom logic. This marks a significant shift from symbolic AI to a more data-driven, emergent approach to agent development.

Despite the significant advancements, LLM agents also present unique challenges. One major concern is the potential for "hallucinations," where the LLM generates plausible-sounding but factually incorrect or nonsensical information [67]. This can lead to unreliable agent behavior and requires careful mitigation strategies, such as grounding the LLM's responses in external knowledge sources or employing techniques like retrieval-augmented generation (RAG) [68]. Other challenges include ensuring the safety and ethical behavior of agents, managing the computational cost of running large models, and addressing biases inherent in the training data [69,70]. Furthermore, the "black box" nature of LLMs can make it difficult to understand the reasoning behind an agent's actions, hindering debugging and limiting trust in critical applications [71].

The ongoing research in LLM-powered agents focuses on addressing these limitations and exploring new capabilities with the emerging multi-agent systems (MAS). This includes developing methods for improving the factual accuracy and reliability of LLMs, designing more robust and explainable agent architectures, and creating effective evaluation metrics to assess agent performance across diverse tasks [72,73]. The convergence of LLMs and software agents represents a significant step towards building more general-purpose, adaptable, and intelligent systems, with the potential to transform how we interact with computers and automate complex tasks. However, responsible development and deployment, with a keen awareness of the inherent challenges, are crucial to realizing the full potential of this technology.

**Could Semantic Web Principles Complement LLMs?** Unlike the high-dimensional token spaces in which large language models (LLMs) operate, an ontology offers a *canonical, machine-interpretable vocabulary* of entities, relations and axioms. Mapping generated spans to URI-identified concepts therefore gives each symbol an unambiguous referent, providing the first step toward faithful explanations [74]. OWL and RDF guarantee sound, monotonic inference but require careful manual modelling, limit expressivity to stay decidable, adopt an Open World Assumption and incur worst-case exponential reasoning cost. LLMs sit at the opposite end of the spectrum: they learn broad coverage automatically, answer in closed-world style with constant-time decoding, and tolerate noisy data—yet offer no built-in notion of truth, identity or provenance. Combining the two thus pairs the symbolic rigour of OWL/RDF with the statistical reach of LLMs.

- **Pattern 1 – ontology-constrained decoding** During generation the model is prompted or fine-tuned so that every entity mention must be resolvable to an ontology concept; illegal continuations are pruned. This drastically reduces hallucinations and yields provenance for every token. A practical example is the plant-disease classifier of [75], where OWL reasoning is used online to check that the captioned visual attributes entail a unique disease class.
- **Pattern 2 – knowledge graph RAG.** Retrieval-augmented generation in which the retriever indexes a knowledge graph rather than free text grounds answers in triples whose semantics are explicit. Recent systems such as ITEMIZE show that even small models achieve state-of-the-art accuracy *and* produce subgraph-level rationales that users can inspect [76].
- **Pattern 3 – neuro-symbolic post-hoc explanation.** After an LLM produces an answer, a reasoning engine aligns the answer with the ontology, derives minimal entailment paths, and verbalises them as "because-clauses." Confalonieri and Guizzardi identify reference modelling, commonsense reasoning and knowledge-

May 2025


refinement as the three main roles ontologies play in such explanatory pipelines [74].

#### 4. Academic Significance and Statistics

In this section, we use Google Scholar and other bibliometric tools to evaluate the significance of the research wave surrounding the Semantic Web. By analyzing citation counts, publication trends, and related metrics, we assess the impact and reach of Semantic Web research within the academic community. This evaluation provides insight into how the Semantic Web concept influenced subsequent studies, interdisciplinary collaborations, and its adoption in both AI and broader technological domains. Such an analysis helps contextualize the Semantic Web’s role in shaping research agendas and its interplay with advancements in AI, particularly in the development of intelligent systems and software agents.

##### Was the Semantic Web a Significant Research Wave? - Google Scholar Analysis


To assess the significance of the Semantic Web and its associated AI trends compared to other common AI research themes, we conducted Google Scholar searches <sup>2</sup>. The Semantic Web was undoubtedly a notable research wave, as evidenced by 3.41 million Google Scholar search results (March 2025). While broader topics like “Artificial Intelligence” (7.25M) and “Machine Learning” (5.47M) have higher volumes, the Semantic Web’s specific focus on structured, meaningful data integration, along with its connection to related areas such as Artificial Intelligence (1.66M) and Agents (1.19M), underscores its significant impact. Despite its narrower scope, the Semantic Web established foundational principles that continue to influence modern AI and data technologies, solidifying its importance in the research landscape.



**Figure 4.** Google Books: number of hits using “Semantic Web” as key words (Data not available after 2022)

<sup>2</sup>Google Scholar-based metrics have well-known caveats: (i) uneven disciplinary and language coverage (e.g., conference proceedings and non-English venues indexed less consistently than major English journals); (ii) inclusion of non-peer-reviewed or duplicate records that can inflate counts; (iii) continuous re-indexing that causes citation numbers to fluctuate over time; and (iv) keyword matching, where the search string may appear only tangentially in a document, so the retrieved set is not guaranteed to be semantically aligned with the intended topic. Accordingly, all figures reported here should be read as indicative rather than definitive.

**Semantic Web in Google Books** Figure 4 presents a screenshot from the Google Books Ngram Viewer, illustrating the historical prominence of the term Semantic Web in English-language books from 1990 to 2022. The graph shows a significant rise in usage beginning in the early 2000s, coinciding with the publication of the seminal 2001 paper by Tim Berners-Lee, James Hendler, and Ora Lassila, which introduced the concept. Usage peaks between 2005 and 2010, reflecting the increased academic and industrial interest in its potential to transform data structuring and access on the Web. However, the decline after this peak suggests a waning emphasis on the Semantic Web in scholarly and public discourse, possibly due to challenges in realizing its full vision or a shift toward emerging technologies such as machine learning and deep learning. This trend underscores the cyclical nature of technological interest, aligning with broader discussions of the hype cycle.



**Figure 5.** Search results with “Software Agents” with five year time intervals Google Scholar (hits) and Google Books (frequencies)


**When Did Software Agent Research Emerge?** One of the most promising future approaches is integrating software agents with large language models (LLMs) to enable not only discussions but also the execution of actions and commands by applications. To analyze the origins and activity trends of software agent research, we compiled statistics from Google Scholar and Google Books Ngram Viewer, with the results presented in Figure 5.

Google Scholar searches indicate that the term “Software Agents” appeared with minimal activity in the 1970s and 1980s, with only a handful of mentions. Significant growth occurred in the 1990s, reaching 463 results between 1990 and 1994. The field experienced exponential growth after 2000, peaking between 2005 and 2015, with over 16,000 search results in each five-year period. Activity has slightly declined in recent years, with 15,400 results from 2015 to 2020 and a projected 13,700 from 2020 to 2024, indicating sustained but slightly reduced interest in the field.

Figure 5 highlights the historical prominence of the term “Software Agents” in English-language books from 1990 to 2019, showing a significant rise in usage between 2000 and 2010. This trend suggests that Software Agent research received substantial attention during this period, coinciding with the heightened focus on the Semantic Web. The overlap is noteworthy, as the original Semantic Web vision in 2001 envisioned a machine-readable web where Software Agents could perform complex tasks on behalf of users. The simultaneous growth in interest in Software Agents and the Semantic Web underscores their interconnectedness, with Software Agents positioned as a key applica-

tion of Semantic Web technologies. This period represents a peak in academic and industrial enthusiasm for intelligent, autonomous systems capable of leveraging structured web data, aligning with broader trends in artificial intelligence research.

The illustration in Figure 6 depicts the number of Google Scholar search hits for the general terms Software Agent and specific combination "Software Agents" from 1990 to 2010, highlighting their academic prevalence over time. The data indicates a steady increase in publications for both terms, with Software Agent experiencing continuous growth and peaking around 2010 at approximately 413,000 hits. In contrast, the more specific term "Software Agents" saw rapid growth until about 2005, reaching a peak of approximately 5,230 hits before experiencing a slight decline. This suggests that while interest in the broader concept of software agents continued to grow, research using the specific phrase "Software Agents" peaked mid-decade before stabilizing or declining.



**Figure 6.** Citation Trends in Software Agent Research: Google Scholar Cites Per Year. Data extracted from Google Scholar.

## 5. Discussion

This paper has provided a review and synthesis of previous studies on AI history, particularly the cycles of AI summers and AI winters in Section 2. We introduced a normalized temperature scale to quantify these fluctuations. Section 3 explored the Semantic Web's foundational vision and its relationship with AI-driven Software Agents, while Section 4 analyzed bibliometric data and academic trends, highlighting the Forgotten AI Wave.

**Software Agents: A Hot Topic—But Where Did They Come From?** As discussed in Section 3, Software Agents are currently a hot topic, with Microsoft's CEO recently predicting that AI-driven Software Agents will disrupt B2B software applications. Where did agents originate? AI history is marked by cycles of optimism and setbacks. In Section 2, we reviewed prior studies and illustrations of AI history, concluding that the Semantic Web and Software Agents represent a neglected phase in AI's evolution. This section explores why this research wave did not gain widespread recognition, the challenges that limited its impact, and how modern AI developments may revive its core principles.

**Were the Terms Artificial Intelligence and Semantic Web a Curse?** A key question arises: Why wasn't the Semantic Web recognized as a distinct AI wave? As revis-

ited in Section 3, the original Semantic Web vision in 2001 sought to create a machine-readable web enabling intelligent reasoning and automation. Software Agents played a central role, designed to operate autonomously, interact with structured data, and execute user tasks. During the early 2000s, Semantic Web research saw a surge in academic publications and funding, as evidenced by our citation analysis. One possible explanation, as suggested by Ora Lassila in his keynote talk [77], is that Artificial Intelligence was perceived negatively at the time, leading researchers to distance themselves from the term. Instead, they used alternative terms, such as Software Agents and Knowledge Representation. A similar trend occurred in the early 1980s when Geoffrey Hinton and other researchers avoided using the term Neural Networks due to its controversial status within the AI community [78].

Also terminology and research focus regarding the Semantic Web has changed during the years [9]. A good deal of research in the early 2000's focused first on ontologies, the *silver bullet* [79] of the Semantic Web and formal description logics for defining them. This line of research encountered the classical challenges of symbolic AI and expert systems. Many people thought that the Semantic Web initiative is doomed to fail and the term, just like AI, got some negative flavor. However, the focus then shifted to semantically simpler (linked) data [80], the new *oil of the digital era* that could be used more easily as a basis for practical applications. Jim Hendler had already in the late 90's in the SHOE project<sup>3</sup> stated: *A little semantics goes a long way*. The next step ahead was introduction of the notion of knowledge graphs (KG) [81] that essentially combines ontologies and linked data in practical application domain settings. This idea was adopted successfully, for example, by major web technology companies, such as Google (Google KG), Microsoft (Satori KG), and Facebook (Open Graph). However, the idea of the semantic web is not necessarily mentioned when talking about linked data and KGs.

**Why Didn't Software Agents on the Semantic Web Succeed?** Despite initial momentum, the Semantic Web failed to become a mainstream AI paradigm. One primary reason was the dominance of symbolic reasoning and logic-based AI at the time, which contrasted with the statistical and data-driven approaches that later fueled AI's success. Software Agents lacked access to the powerful natural language processing (NLP) and deep learning techniques that characterize modern AI applications. Although the Semantic Web did not achieve its grand vision of intelligent software agents, its principles are resurfacing in contemporary AI, particularly through Large Language Models (LLMs) and AI-driven Software Agents.

**Revisiting the Semantic Web in the Age of LLMs** Recent advancements highlight the renewed relevance of Semantic Web concepts such as structured knowledge representation, reasoning, and autonomous agents. Knowledge graphs, an evolution of Semantic Web technologies, are widely used by companies like Google and Microsoft to enhance AI-driven search and recommendation systems. Moreover, modern Software Agents are now integrated with LLMs, enabling them to interpret unstructured text, execute commands, and autonomously interact with various data sources—effectively realizing the Semantic Web's vision through alternative means.

A growing emphasis on AI explainability, data provenance, and structured knowledge integration suggests that Semantic Web technologies could find new applications in the evolving AI landscape. Hybrid neuro-symbolic AI approaches that combine statisti-

---

<sup>3</sup><https://www.cs.rpi.edu/~hendler/LittleSemanticsWeb.html>

cal learning with symbolic reasoning may leverage the Semantic Web's structured data to enhance AI reasoning and interpretability.

**Lessons from the Forgotten AI Wave** The Semantic Web's history offers valuable lessons on AI's cyclical evolution. The shift from logic-based AI to deep learning was not a rejection of earlier methods but rather an adaptation to technological constraints and industry priorities. As AI progresses, integrating structured knowledge with modern techniques could yield more powerful, interpretable, and autonomous systems. Recognizing the Semantic Web as part of AI's history provides a more comprehensive view of technological evolution. Instead of seeing it as a failed wave, we can understand it as a crucial step in AI's development—one whose principles may yet play a key role in the next generation of intelligent systems.

## 6. Conclusion: Reintegrating the Semantic Web into AI's History

The key contributions in this paper are as follows:

- We reviewed previous illustrations of AI history and synthesized them into a uniform scale in Section 2.
- We identified the Semantic Web and Software Agents as a forgotten wave of Artificial Intelligence by connecting them to the current AI boom and analyzing academic statistics, as discussed in Section 1 and Section 4.

We argued that the Semantic Web represents a significant yet often underappreciated chapter in the evolution of artificial intelligence. Emerging in the early 2000s, its vision of a machine-readable web, capable of empowering AI-driven Software Agents, promised transformative applications ranging from intelligent automation to advanced data integration. While the grand vision articulated by Tim Berners-Lee, James Hendler, and Ora Lassila remains unrealized in its entirety, its influence on structuring, integrating, and interpreting web data continues to resonate in modern technologies, including knowledge graphs, ontologies, and standards for Findable, Accessible, Interoperable and Re-usable FAIR data<sup>4</sup> for scientific data management and stewardship.

In this paper, the academic and industrial impact of the Semantic Web was revisited using bibliometric analysis and historical context to argue for its rightful place within AI's broader narrative of hype cycles, AI winters, and eventual breakthroughs. In particular, the overlooked connection between the Semantic Web and AI Software Agent research from 2000 to 2010 highlights its contributions to advancing intelligent systems. Despite challenges in achieving widespread adoption and the subsequent shift in focus toward machine learning and deep learning paradigms, the Semantic Web established foundational principles that remain relevant to contemporary AI development.

Recognizing the contributions of the Semantic Web enriches our understanding of AI's history and underscores the importance of examining past technological waves to draw lessons for the future. By reintegrating this forgotten wave into AI's historical discourse, we gain a more comprehensive perspective on the cyclical nature of innovation and the interplay between vision and practicality in shaping technological progress.

---

<sup>4</sup>FAIR data: <https://www.go-fair.org/fair-principles/>

## References

- [1] McCorduck P. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. AK Peters Ltd; 2004.
- [2] Minsky M, Papert S. *Perceptrons: An Introduction to Computational Geometry*. MIT Press; 1969.
- [3] Mettrey W. An Assessment of Tools for Building Large Knowledge-Based Systems. *AI Magazine*. 1987;8(4).
- [4] Kelleher JD. *Deep learning*. The MIT Press essential knowledge series. Cambridge, Mass. London: The MIT Press; 2019.
- [5] Hitzler P, Sarker MK, editors. *Neuro-Symbolic Artificial Intelligence: The State of the Art*. IOS Press; 2022.
- [6] Bhuyan A Bikram Pratim nsd Ramdane-Cherif, Tomar R. Neuro-symbolic artificial intelligence: a survey. *Neural Computing and Applications*. 2024;36:12809-44.
- [7] Berners-Lee T, Hendler J, Lassila O. The Semantic Web. *The Scientific American*. 2001 May. URL: <http://www.scientificamerican.com/>.
- [8] Berners-Lee T. Semantic Web Wave. In: Presentation at SIIA; 2003. Accessed: 2024-01-26.
- [9] Hitzler P. A review of the semantic web field. *Communications of the ACM*. 2021;64(2):76-83.
- [10] Hitzler P, Krötzsch M, Rudolph S. *Foundations of Semantic Web technologies*; 2010.
- [11] Vaswani A, Shazeer N, Parmar N, Uszkoreit J, Jones L, Gomez AN, et al. Attention is all you need. *Advances in neural information processing systems*. 2017;30.
- [12] Reed S, Zolna K, et al. A Generalist Agent. *DeepMind*. 2022.
- [13] Yao S, Zhao J, Yu D, Du N, Shafran I, Narasimhan K, et al.. ReAct: Synergizing Reasoning and Acting in Language Models; 2023. Available from: <https://arxiv.org/abs/2210.03629>.
- [14] Steinert M, Leifer L. Scrutinizing Gartner's hype cycle approach. In: *Picmet 2010 technology management for global economic growth*. IEEE; 2010. p. 1-13.
- [15] Menzies T. 21st-century AI: proud, not smug. *IEEE Intelligent Systems*. 2003;18(3):18-24.
- [16] Kautz H. The third ai summer: Aai robert s. engelmore memorial lecture. *Ai magazine*. 2022;43(1):105-25.
- [17] Noguchi T, Hashizume Y, Moriyama H, Gauthier L, Ishikawa Y, Matsuno T, et al. A practical use of expert system" AI-Q" focused on creating training data. In: *2018 5th international conference on business and industrial research (ICBIR)*. IEEE; 2018. p. 73-6.
- [18] Toosi A, Bottino AG, Saboury B, Siegel E, Rahmim A. A Brief History of AI: How to Prevent Another Winter (A Critical Review). *PET Clinics*. 2021 Oct;16(4):449-469. Available from: <http://dx.doi.org/10.1016/j.cpet.2021.07.001>.
- [19] Francesconi E. The winter, the summer and the summer dream of artificial intelligence in law: Presidential address to the 18th International Conference on Artificial Intelligence and Law. *Artificial intelligence and law*. 2022;30(2):147-61.
- [20] Colliot O. A non-technical introduction to machine learning. *Machine Learning for Brain Disorders*. 2023;3-23.
- [21] Harguess J, Ward CM. Is the next winter coming for AI? Elements of making secure and robust AI. In: *2022 IEEE Applied Imagery Pattern Recognition Workshop (AIPR)*. IEEE; 2022. p. 1-7.
- [22] Perspectives A. *AI Perspectives*; 2020. Accessed: 2025-01-12. <https://www.aiperspectives.com/introduction/>.
- [23] Lu A. AI Summer and Winter: The Ups and Downs of AI; 2022. Accessed: 2025-01-12. <https://www.label.ai/blog/four-seasons-of-ai-the-brief-history-of-artificial-intelligence-i>.
- [24] Research B, Analysis. The Rise of AI; 2020. Accessed: 2025-01-12. <https://www.brightworkresearch.com/why-did-ai-rise-for-a-third-time-after-the-1st-and-2nd-ai-winters/>.
- [25] Winiger S. AI Winters = HCL Summers; 2017. Accessed: 2025-01-12. <https://samim.io/p/2018-02-19-ai-winters-hci-summers/>.
- [26] Altmann Y. Small History of AI; 2023. Accessed: December 28, 2024, <https://omq.ai/blog/history-of-ai/>. Available from: <https://omq.ai/blog/history-of-ai/>.
- [27] Latta J. History of AI – from Winter to Winter; 2021. Accessed: 2024-12-29, <https://jaylatta.net/history-of-ai-from-winter-to-winter/>. Available from: <https://jaylatta.net/history-of-ai-from-winter-to-winter/>.
- [28] Schuchmann S. History of the First AI Winter; 2019. Accessed: 2024-12-29, <https://towardsdatascience.com/history-of-the-first-ai-winter-6f8c2186f80b>. Available



May 2025

- from: <https://towardsdatascience.com/history-of-the-first-ai-winter-6f8c2186f80b>.
- [29] Lim M. History of AI Winters; 2018.
  - [30] Solé R. Tweet on Artificial Intelligence History; 2022. Includes reference to Google Books Ngram Viewer: [https://books.google.com/ngrams/graph?content=perceptron&year\\_start=1930&year\\_end=2019&corpus=26&smoothing=3](https://books.google.com/ngrams/graph?content=perceptron&year_start=1930&year_end=2019&corpus=26&smoothing=3), Accessed: 2024-12-29. Available from: [https://twitter.com/ricard\\_sole/status/1497306702181548039](https://twitter.com/ricard_sole/status/1497306702181548039).
  - [31] Nwana HS. Software agents: An overview. *Knowledge Engineering Review*. 1996;11(3):205-44.
  - [32] Nwana HS, Ndumu DT. A Perspective on Software Agents Research. 1999;14(2):1-18. Available from: <http://agents.umbc.edu/introduction/hn-dn-ker99.pdf>.
  - [33] Nadella S. BG2 Pod Series: December 2024 Episode; 2024. Available at: [https://youtu.be/9NtsnzRFJ\\_o?t=2811](https://youtu.be/9NtsnzRFJ_o?t=2811), Accessed: March 2025. YouTube.
  - [34] Hewitt C. Viewing control structures as patterns of passing messages. *Artificial intelligence*. 1977;8(3):323-64.
  - [35] Bond AH, Gasser L. Readings in distributed artificial intelligence. Morgan Kaufmann; 2014.
  - [36] Gasser R, Huhns MN. Distributed Artificial Intelligence: Volume II. vol. 2. Morgan Kaufmann; 2014.
  - [37] Wooldridge M. An introduction to multiagent systems. John Wiley & Sons. 2009.
  - [38] Wooldridge M, Jennings NR. Intelligent agents: Theory and practice. *Knowledge Engineering Review*. 1995;10(2):115-52.
  - [39] Maes P. Agents that reduce work and information overload. *Communications of the ACM*. 1994;37(7):30-40.
  - [40] Weiss G. Multiagent systems: A modern approach to distributed artificial intelligence. MIT press; 1999.
  - [41] Franklin S, Graesser A. Is it an agent, or just a program?: A taxonomy for autonomous agents. In: Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96). Springer; 1996. p. 21-35.
  - [42] Foundation for Intelligent Physical Agents (FIPA). FIPA '97 Specification, Version 1.0 — Part 2: Agent Communication Language. Geneva, Switzerland: FIPA; 1997.
  - [43] McGuinness DL, Van Harmelen F. OWL web ontology language overview. *W3C recommendation*. 2004;10(10):2004.
  - [44] Shadbolt N, Hall W, Berners-Lee T. The semantic web revisited. *IEEE intelligent systems*. 2006;21(3):96-101.
  - [45] Hendler J. Agents and the Semantic Web. *IEEE Intelligent Systems*. 2001;16(2):30-7.
  - [46] McIlraith SA, Son TC, Zeng H. Semantic web services. *IEEE Intelligent Systems*. 2001;16(2):46-53.
  - [47] Lassila O, Swick RR. Enabling the Web of Data with RDF. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*. 2001;24(4):6-10.
  - [48] Foundation for Intelligent Physical Agents (FIPA). FIPA ACL Message Structure Specification. Geneva, Switzerland: FIPA; 2002. SC00061G.
  - [49] Vreeswijk GA. Liberalizing protocols for argumentation in multi-agent systems. In: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems; 2005. p. 1259-60.
  - [50] Lassila O, Swick RR. Resource description framework (RDF) model and syntax specification. *W3C Proposed Recommendation PR-rdf-syntax-19981008*. 1998.
  - [51] Horrocks I, Patel-Schneider PF, Van Harmelen F. From SHIQ and RDF to OWL: The making of a web ontology language. In: *International Semantic Web Conference*. Springer; 2003. p. 7-26.
  - [52] Gruber TR. Toward principles for the design of ontologies used for knowledge sharing. *International journal of human-computer studies*. 1993;43(5-6):907-28.
  - [53] FIPA. FIPA Communicative Act Library Specification; 2002. Accessed: 2024-01-26. Foundation for Intelligent Physical Agents. <http://www.fipa.org/specs/fipa00037/>.
  - [54] Martin D, Burstein M, Hobbs J, Lassila O, McDermott D, McIlraith S, et al. OWL-S: Semantic markup for web services. In: *W3C member submission*. vol. 22; 2004. p. 2004.
  - [55] Russell SJ, Norvig P. Artificial intelligence: a modern approach. 3rd ed. Prentice hall; 2010.
  - [56] Brown T, Mann B, Ryder N, Subbiah M, Kaplan JD, Dhariwal P, et al. Language models are few-shot learners. In: *Advances in neural information processing systems*. vol. 33; 2020. p. 1877-901.
  - [57] Ouyang L, Wu J, Jiang X, Almeida D, Wainwright C, Mishkin P, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*. 2022;35:27730-44.
  - [58] Wei J, Wang X, Schuurmans D, Bosma M, Chi E, Le Q, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*. 2022;35:24824-

- [59] Cemri M, Pan MZ, Yang S, Agrawal LA, Chopra B, Tiwari R, et al.. Why Do Multi-Agent LLM Systems Fail?; 2025. Available from: <https://arxiv.org/abs/2503.13657>.
- [60] Naveed H, Khan AU, Qiu S, Saqib M, Anwar S, Usman M, et al.. A Comprehensive Overview of Large Language Models; 2024. Available from: <https://arxiv.org/abs/2307.06435>.
- [61] Xi Z, Chen W, Guo X, He W, Ding Y, Hong B, et al.. The Rise and Potential of Large Language Model Based Agents: A Survey; 2023. Available from: <https://arxiv.org/abs/2309.07864>.
- [62] Wang L, Ma C, Feng X, Zhang Z, Yang H, Zhang J, et al. A survey on large language model based autonomous agents. *Frontiers of Computer Science*. 2024 Mar;18(6). Available from: <http://dx.doi.org/10.1007/s11704-024-40231-1>.
- [63] Yang S, Nachum O, Du Y, Wei J, Abbeel P, Schuurmans D. Foundation Models for Decision Making: Problems, Methods, and Opportunities; 2023. Available from: <https://arxiv.org/abs/2303.04129>.
- [64] Shinn N, Cassano F, Berman E, Gopinath A, Narasimhan K, Yao S. Reflexion: Language Agents with Verbal Reinforcement Learning; 2023. Available from: <https://arxiv.org/abs/2303.11366>.
- [65] Chase H. Langchain; 2022. Accessed: 2024-01-26. <https://github.com/langchain-ai/langchain>.
- [66] Wu Q, Bansal G, Zhang J, Wu Y, Li B, Zhu E, et al.. AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation; 2023. Available from: <https://arxiv.org/abs/2308.08155>.
- [67] Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, et al. Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*. 2023 Mar;55(12):1–38. Available from: <http://dx.doi.org/10.1145/3571730>.
- [68] Lewis P, Perez E, Piktus A, Petroni F, Karpukhin V, Goyal N, et al.. Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks; 2020. Available from: <https://arxiv.org/abs/2005.11401>.
- [69] Bender EM, Gebru T, McMillan-Major A, Shmitchell S. On the dangers of stochastic parrots: Can language models be too big? In: *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*; 2021. p. 610-23.
- [70] Weidinger L, Mellor J, Rauh M, Griffin C, Uesato J, Huang PS, et al.. Ethical and social risks of harm from Language Models; 2021. Available from: <https://arxiv.org/abs/2112.04359>.
- [71] Doshi-Velez F, Kim B. Towards A Rigorous Science of Interpretable Machine Learning; 2017. Available from: <https://arxiv.org/abs/1702.08608>.
- [72] Nakano R, Hilton J, Balaji S, Wu J, Ouyang L, Kim C, et al.. WebGPT: Browser-assisted question-answering with human feedback; 2022. Available from: <https://arxiv.org/abs/2112.09332>.
- [73] Laskin M, Wang L, Oh J, Parisotto E, Spencer S, Steigerwald R, et al.. In-context Reinforcement Learning with Algorithm Distillation; 2022. Available from: <https://arxiv.org/abs/2210.14215>.
- [74] Confalonieri R, Guizzardi G. On the Multiple Roles of Ontologies in Explainable AI; 2023. Available from: <https://arxiv.org/abs/2311.04778>.
- [75] Amara J, König-Ries B, Samuel S. Enhancing Explainability in Multimodal Large Language Models Using Ontological Context; 2024. Available from: <https://arxiv.org/abs/2409.18753>.
- [76] Li M, Miao S, Li P. Simple Is Effective: The Roles of Graphs and Large Language Models in Knowledge-Graph-Based Retrieval-Augmented Generation; 2025. Available from: <https://arxiv.org/abs/2410.20724>.
- [77] Lassila O. Semantic Web and AI: Can we finally realize the vision?; 2024. Slides: <https://www.lassila.org/publications/2024/lassila-iswc2024-keynote.pdf>, Video: [https://videlectures.net/videos/iswc2024\\_lassila\\_web\\_and\\_ai](https://videlectures.net/videos/iswc2024_lassila_web_and_ai). Keynote Address at the 23rd International Semantic Web Conference (ISWC 2024).
- [78] Ford M. *Architects of Intelligence: The truth about AI from the people building it*. Birmingham, UK: Packt Publishing; 2018.
- [79] Fensel D. *Ontologies: Silver bullet for knowledge management and electronic commerce (2nd Edition)*. Springer-Verlag; 2004.
- [80] Heath T, Bizer C. *Linked Data: Evolving the Web into a Global Data Space (1st edition)*; 2011. Available from: <http://linkeddatabook.com/editions/1.0/>.
- [81] Gutierrez C, Sequeda JF. Knowledge graphs. *Communications of the ACM*. 2021 March;64(3):96-104.