

Accelerating Knowledge Acquisition with Help from Large Language Models: From Digital Documents to Database Models

Carlos Henrique Serrado¹, Geraldo Xexéo¹, Carlos Eduardo Barbosa^{1,2}, Matheus Argôlo¹, Lucas Nóbrega¹, Luiz Felipe Martinez¹, and Jano de Souza¹

¹Systems and Computing Engineering Program of the Alberto Luiz Coimbra Institute for Graduate Studies and Research in Engineering, Federal University of Rio de Janeiro, Brazil

²Naval Systems Analysis Center, Brazilian Navy, Rio de Janeiro, Brazil

chserrado@cos.ufrj.br

xexeo@cos.ufrj.br

eduardo@cos.ufrj.br

matheusargolo@cos.ufrj.br

lucasnobrega@cos.ufrj.br

felipem@cos.ufrj.br

jano@cos.ufrj.br

Abstract: Knowledge Management (KM) processes are essential for organizations, allowing them to effectively capture, store, and use their knowledge to make informed decisions. Modern enterprises use computerized systems and relational databases to manage their operational processes. However, a significant challenge remains in transforming insights found in digital documents into actionable data models without overloading business analysts or necessitating constant updates and modifications. This work introduces a method for modeling dynamic environments using a knowledge base. The approach involves creating a world model within a relational database that can be updated using Structured Query Language (SQL) expressions derived from documents that describe changes in that world. The techniques discussed include using agents and Large Language Models (LLMs) to generate SQL commands to keep the database current. The proposed world model aims to remain sufficiently generic and adaptable to handle a variety of entities and relationships across multiple organizational domains. Representing events, objects, and their interactions in a flexible structure ensures that real-world transformations are accurately mirrored in the database. This versatility allows the model to be implemented in different sectors without significantly modifying the underlying data architecture. Integrating these processes with advanced language models, such as ChatGPT, aims to improve the generation of data models and streamline the KM workflow by automating the interpretation of explicit knowledge. This integration of language models and relational databases is intended to enhance the organization, storage, and retrieval of insights, thereby reducing manual effort and improving the knowledge base's adaptability to changing needs. Overall, the proposed solution seeks to leverage LLMs to assist in modeling data and managing knowledge from explicit sources, providing a practical framework for organizations looking to stay competitive in evolving environments.

Keywords: Knowledge management, Data modeling, Automated knowledge extraction, Large language models, Dynamic world modeling

1. Introduction

In dynamic systems, knowledge is in a constant state of transformation, requiring a deep understanding of how it is represented, shared, and updated across evolving environments. World-building in such systems involves modeling complex relationships between regions, organizations, and individuals, all of which can shift due to internal developments or external events. These transformations must be coherently reflected within the structure of the world, ensuring consistency and narrative continuity. Typically, these changes are first internalized during interaction with the environment and later externalized through revisions that update the world's state. As complexity increases, managing this process manually becomes increasingly challenging, emphasizing the importance of automated methods for maintaining accurate and responsive representations. Fictional universes and simulation platforms often face these challenges—prominent examples include narrative-driven environments such as those found in Role-Playing Games (RPGs) (N.R and Ikhtiyorovna, 2023), where world states evolve through continuous player interaction and storytelling (Cook et al., 2003; Halliwell et al., 1986; N.R and Ikhtiyorovna, 2023).

However, recent advancements in Large Language Models (LLMs) offer a promising solution to this problem. These models, particularly Generative Pre-trained Transformers (GPT), have been increasingly recognized as general-purpose technologies capable of performing various tasks (Brown et al., 2020; Devlin et al., 2018; OpenAI, 2023). Their ability to process and generate human-like text allows them to streamline complex workflows by extracting key information from unstructured text and converting it into structured data (Hadi et

al., 2023; Raiaan et al., 2023). This makes LLMs powerful tools for automating tasks such as managing databases (Li et al., 2024; Zhou et al., 2023) and translating text into Structured Query Language (SQL) (International Organization for Standardization, 2016) commands (Dong et al., 2023; Sun et al., 2024; Tan et al., 2024), allowing for efficient and accurate data management in dynamic systems.

Recent advancements in LLMs have significantly improved SQL generation tasks, known as Text-to-SQL, making database interactions more accessible (Zhu et al., 2024). Key improvements stem from integrating expert knowledge (Hong et al., 2024), leveraging structural information (Zhang et al., 2024), and decomposing complex tasks into manageable subtasks (Pourreza and Rafiei, 2024). Additionally, benchmark evaluations reveal the need for more complex datasets to assess LLM capabilities effectively (Ma et al., 2024).

In this work, we propose an approach for modeling dynamic worlds using a knowledge base. Our methodology involves creating a dynamic world model in a relational database, continuously updated through SQL expressions derived from narrative documents detailing world changes. We employ a combination of agents and LLMs to generate SQL commands to manipulate the database's data, including altering and deleting data. This approach enables automated updates to the state of the world represented in the database, ensuring that it accurately reflects the evolving game narrative and maintains consistency across sessions.

Our work aims to make various contributions, including automating dynamic world updates and expanding the use of LLMs in managing complex systems. By combining LLMs with relational databases, the research introduces a new method for converting unstructured text into structured data, reducing the need for manual intervention in the update process. It has practical applications in real-world scenarios such as smart city simulations, where factors like population, infrastructure, and economics undergo continuous evolution. The proposed framework offers urban planners, policymakers, and professionals in dynamic fields such as project management and financial forecasting a reliable tool for precise, efficient, and scalable data management. Additionally, the research identifies potential avenues for improvement, such as multi-agent systems and advanced knowledge representation techniques, which could significantly enhance the model's scalability, precision, and applicability in more complex environments.

2. Theoretical Background

This section presents the theoretical background, introducing fundamental concepts related to role-playing games, world representation, and knowledge representation. These concepts are essential for understanding the developments discussed in subsequent sections.

2.1 World-building

Fictional world-building involves the creation of intricate environments populated by diverse regions, societies, and systems that interact in complex ways. These constructed worlds often feature detailed representations of cities, defined by characteristics such as population size, security dynamics, and economic conditions. A coherent and dynamic world requires careful modeling of interactions between regions, institutions, and inhabitants, as these elements collectively shape the setting's internal logic and narrative potential.

- **Regions:** These include cities, towns, and other locations. Each region has its own set of attributes, including but not limited to population size, economic resources, and security levels. Accurate modeling of regions is essential as they serve as the primary setting for most adventures and influence the narrative's direction and scope.
- **Organizations:** These encompass guilds, councils, and other groups that operate within the regions. Organizations play significant roles in the storyline, often serving as quest givers, antagonists, or support structures for the inhabitants. Understanding their hierarchies, goals, and relationships with other entities is vital for a nuanced world representation.
- **Inhabitants:** The non-player characters (NPCs) and player characters (PCs) interact within the regions and organizations. Inhabitants have various attributes, such as health, status, and affiliations, which are crucial for driving the narrative forward. Their interactions, statuses, and evolutions must be accurately tracked to maintain consistency and depth in the storytelling.

This work focuses on the relation between these entities, representing any changes during sessions. Modeling these entities and their relationships in SQL allows for a structured and dynamic representation of the world setting. This modeling enables updates and changes to be reflected accurately as the game progresses. SQL facilitates data manipulation, ensuring that the database can handle the complex and evolving nature of the game world. The entity-relationship model, illustrated in Figure 1, captures the relationships between these

entities by defining the tables and relationships necessary to represent regions, organizations, and inhabitants. This model includes schemas for each entity, outlining their names and how they interact.

Accurate and detailed modeling of these entities and their relationships is crucial for the integrity of the game. It provides a framework that can be adapted for other applications, such as intelligent city simulations. Understanding these entities' ontologies and interactions helps design systems that reflect complex, dynamic environments.

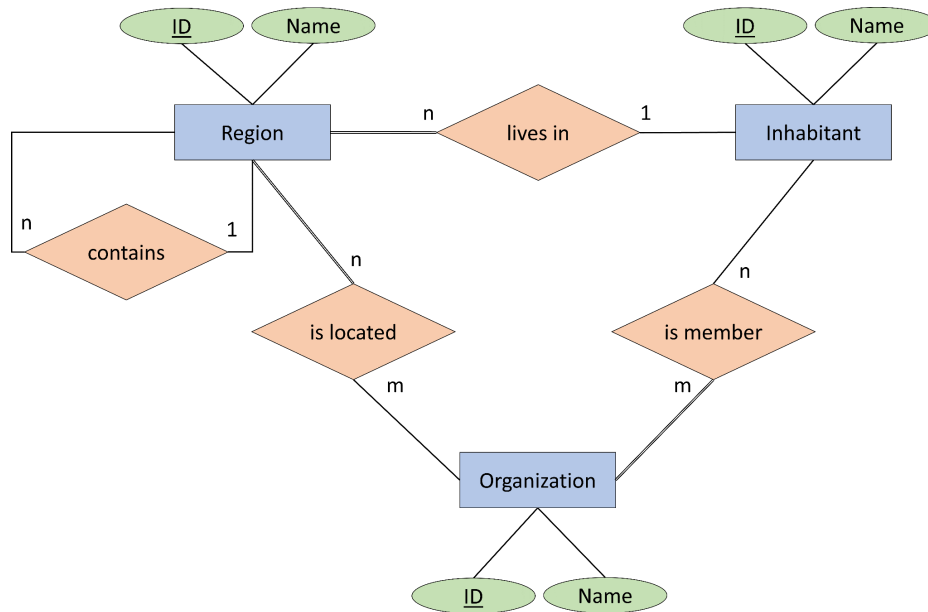


Figure 1: An Entity-Relationship model focused on the Region, Inhabitant, and Organization

2.2 Representing Worlds

Suppose a world W , in a state s , denoted by W^s , and a world model M_W that, when updated to represent W^s , is denoted by M_W^s . Since worlds are dynamic, at some point, there will be a state $s+1$, defining W^{s+1} , and it will be necessary to update M_W , i.e., calculate M_W^{s+1} .

If the changes in W from W^s to W^{s+1} are described by a set of documents $D = \{d_1, d_2, \dots, d_n\}$, it is possible to extract information from D to discover a series of functions $\{f_1, f_2, \dots, f_k\}$ that should be applied to M_W^s to achieve M_W^{s+1} . For example, M_W^s can be the map of a city (W_s), and D can be the set of documents describing civil engineering interventions that changed the city to a new configuration W^{s+1} . Those changes must be represented on the map. Extracting information from D can generate instructions to change the map accordingly.

The world W is a set of entities w_i modeled as objects o_i in M . Each object has, at the world and model state s , its own state o_s . While the world changes, some changes must be tracked so that the model can reflect them. These changes are modeled as events e . From e , one must calculate the necessary function f to make the transition from o_s to o_{s+1} , and from M_s to M_{s+1} .

An event e can create, destroy, or change the state of o from o_s to $o_{(s+1)}$. Events can change the state of an object o permanently, for a specific interval of time, or until another event happens. For example, suppose M is modeling transit in a city. In that case, an event e can represent the creation of a new road, a permanently closed road to pedestrians, or a temporary interdiction of a road. This interdiction, however, can have a specific time frame—such as closing a street for a one-day event—or require another event to return it to its original state, such as an accident that must be cleared, with the clearing time being uncertain. However, changes described in real-world documents are not guaranteed to be detected in the same order as they occur in the world. For example, an accident on a street can be detected on a social network seconds after it happens. Still, a report stating that the road is clear might arrive at a monitoring system hours after the road is actually cleared. This discrepancy would maintain incorrect information about the road being closed even though it is already evident. Moreover, a notice declaring the road clear may arrive before the notice that it was closed. Therefore, the system must understand the sequence of detected events and possibly implement virtual time or a similar algorithm (Jefferson, 1985). In this work, we will create a world model in a relational database and update it

with SQL expressions derived from documents describing the changes. Table 1 explains the terms introduced in this section.

Table 1: Term Descriptions

Term	Description
W	World
s	State
W^s	World W in state s
M_W	Model for world W
M_W^s	Model M_W that represents world W^s
D	Set of documents that describe the changes from W^s to W^{s+1}
d	Document that details a specific change
f	Function used to update the model from M_W^s to M_W^{s+1}
w	Entity in the world W
o	Object that represents an entity w in the model
e	Event affecting the state of an object or the model

2.3 Representing Knowledge

Knowledge Management (KM) is a critical process that formalizes and disseminates the collective experiences, knowledge, and expertise within an organization, transforming them into accessible and actionable resources (Beckman, 1999). A key element in understanding and applying KM is recognizing the dual nature of knowledge, which is fundamentally divided into two types: tacit and explicit (de Castro Peixoto et al., 2022; Prat, 2006; Stollenwerk, 2001). Tacit knowledge represents individual experiences—context-specific and often difficult to formalize and communicate—encompassing skills, experiences, and insights. In contrast, explicit knowledge is codified, systematic, and easily transferable through formal languages, such as documents and databases (de Castro Peixoto et al., 2022; Nonaka, 1994; Prat, 2006; Stollenwerk, 2001).

Understanding how tacit and explicit knowledge interact is crucial for effective management. In this context, the SECI model outlines the dynamic knowledge conversion process through four modes: Socialization, Externalization, Combination, and Internalization. This model underscores the interplay between tacit and explicit knowledge, facilitating the continuous creation and sharing of knowledge within an organization. Socialization involves sharing tacit knowledge through direct interaction, while Externalization translates tacit knowledge into explicit forms. Combination synthesizes different bodies of explicit knowledge, and Internalization converts explicit knowledge back into tacit understanding through practice and reflection (de Castro Peixoto et al., 2022; Nonaka et al., 2000; Nonaka and Takeuchi, 1995).

To effectively manage and utilize knowledge, choosing an appropriate representation method is essential. There are many ways to represent knowledge, such as using an ontology (Noy and McGuinness, 2001) or Prolog (International Organization for Standardization, 1995). We have opted to represent knowledge using a relational database (Codd, 1970). A relational database is a structured system for managing data that organizes information into tables, also known as relations. Each table is defined by a schema, which specifies the structure of the data, including its attributes (columns) and their data types. The data within the table is stored as tuples, or rows, with each tuple representing a specific instance of the schema. This tabular structure supports complex queries and ensures efficient data management, retrieval, and manipulation (Mannila and Rähjä, 1992; Yang et al., 2009).

- Schema - A schema R for a relation is defined as follows: $R(A_1, A_2, \dots, A_n)$ where A_i (for $1 \leq i \leq n$) are the attributes of the relation. Each attribute A_i has a domain D_i , which specifies the possible values that the attribute can take: $A_i \in D_i$.

- Relation - A relation r on a schema R is a finite set of tuples: $r = \{ t_1, t_2, \dots, t_m \}$ where each tuple t_j (for $1 \leq j \leq m$) is an ordered list of values: $t_j = (v_{j1}, v_{j2}, \dots, v_{jn})$ such that $v_{ji} \in D_i$ for all $1 \leq i \leq n$.
- Database - A database \mathcal{D} is a collection of relations: $\mathcal{D} = \{ r_1, r_2, \dots, r_k \}$ where each r_i is a relation defined on its respective schema R_i .

Although there are theoretical models for database handling, such as relational calculus (Codd, 1972) and relational algebra (Codd, 1970), the standard language used to operate databases is SQL (International Organization for Standardization, 2016). SQL provides a comprehensive framework for managing and manipulating relational databases, encompassing several key components—the Data Definition Language (DDL), Data Manipulation Language (DML), Data Control Language (DCL), and Transaction Control Language (TCL) (Kumar et al., 2012). DDL defines and manages the database’s schema and objects, such as tables, indexes, and views. The primary commands in DDL include CREATE, ALTER, DROP, and TRUNCATE. DCL is used to control access to the database’s data, with primary commands GRANT and REVOKE. TCL manages transactions in the database, ensuring that operations are completed successfully and maintaining data integrity; its primary commands include COMMIT, ROLLBACK, and SAVEPOINT (Kumar et al., 2012). DML is the most relevant component of this article. It is used to retrieve, insert, delete, and update data in database tables (Kumar et al., 2012). The primary DML commands include SELECT (to query and retrieve data), INSERT (to add new rows), UPDATE (to modify existing rows), and DELETE (to remove rows).

3. Method

This research explores an approach to modeling dynamic worlds by leveraging a knowledge base. Our methodology entails constructing a world model using a relational database, which is subsequently updated with SQL expressions derived from documents that describe changes in the world. The approach includes employing agents and LLMs to generate DML commands used to update the state of the world represented within the database.

3.1 Web Scraping

This subsection presents the methodology used to extract structured data from a data source on a fictional universe, such as an encyclopedia. The collected information—covering locations, populations, and organizations—served as the foundation for building the knowledge base of our dynamic world model. Automated web scraping techniques were applied, including URL generation, HTTP requests, and HTML parsing, to retrieve and organize relevant data efficiently.

- URL Construction: URLs were dynamically constructed for each category (locations, inhabitants, organizations) and each city of interest.
- HTTP Requests: Requests were sent to these URLs to retrieve the HTML content of the web pages.
- HTML Parsing: The BeautifulSoup (Richardson, 2007) library was used to parse the HTML content and extract relevant data.
- Data Filtering: Specific HTML elements containing the needed information were filtered and processed. For example, the tags were examined to extract lists of names while ensuring that irrelevant sections (e.g., appendices) were excluded.
- Regular Expression Matching: Regular expressions were used to accurately extract the names of individuals and sub-categories from the text.
- Database Insertion: Extracted data was organized and inserted into an SQLite database, which involved creating tables and establishing relationships between entities.

3.2 Model Update

This subsection discusses the methodology used to update the model, using agents, tasks, and tools. Agents are units that execute tasks using available tools (Wang et al., 2024; Weng, 2023; Xi et al., 2023). By analogy, imagine hiring a Database Engineer (agent) who first reads the text (task 1) and then writes the SQL commands (task 2) while checking current data to verify the status (tool).

LLMs have emerged as transformative technologies in Natural Language Processing. These neural networks, with billions of parameters, are trained on vast amounts of text data to understand and generate human-like text. By employing deep learning techniques, LLMs can identify intricate patterns and relationships within data, enabling them to perform a range of tasks with proficiency. Their capabilities include text generation, translation, summarization, and few-shot question-answering. The sophisticated pattern recognition of LLMs allows them to produce responses that are both coherent and contextually appropriate (Brown et al., 2020; Devlin et al.,

2018; OpenAI, 2023). By employing an LLM as the agent, we leverage its advanced language understanding to interpret textual information and execute the required database operations. In this research, the model employed is llama3-70b-8192 (AI@Meta, 2024). This model's architecture and training enable it to handle complex text parsing tasks and perform precise data manipulations (Huang et al., 2024; Touvron et al., 2023).

To facilitate collaborative efforts within a multiagent system, we use CrewAI (crewAI, 2024), a framework designed to simplify the processes of generating thought, using tools, and communicating among agents. This framework creates agents with specific roles, backstories, and goals while defining tasks with detailed descriptions and expected outputs.

4. Discussion

As demonstrated in our results, LLMs' ability to manage, maintain, and optimize a relational database system aligns with the research of Zhou et al. (2023)—which explores how LLMs can enhance database maintenance through real-time diagnosis and optimization—and Li et al. (2024), which utilizes LLMs to optimize data management systems. Our findings suggest that LLMs have significant potential as database administrators.

Regarding the translation of natural language into SQL queries, our results demonstrate that LLMs can effectively generate SQL instructions corresponding to session summaries, accurately reflecting the described events and changes. This aligns with investigations by Sun et al. (2024), Tan et al. (2024), and Dong et al. (2023), who have used various prompting techniques to improve text-to-SQL performance. These findings highlight the growing versatility of LLMs not only as data management tools but also as efficient translators between natural language and structured database commands.

The convergence of our results with findings from various researchers underscores the growing capacity and applicability of LLMs as database administrators and text-to-SQL translators. The techniques employed—initially applied within the context of dynamic world modeling in games—have broader potential in areas such as project management, supply chain optimization, and financial forecasting, where constant data updates and dynamic decision-making are essential. For example, by modeling a smart city's dynamic nature—including inhabitants, organizations, and regions—it becomes possible to accurately represent and manage changes such as population growth, economic fluctuations, and infrastructure development. Consequently, urban planners and policymakers can leverage this approach to make informed decisions and optimize resource allocation. Moreover, the potential applications extend to other domains where managing complex and dynamic environments is crucial.

Limitations include the dependency of generated SQL commands on clear and precise input documents, with ambiguities potentially causing inaccuracies. Performance and scalability challenges may also arise in more complex scenarios, highlighting the need for further system improvements.

5. Conclusion

LLM models are general-purpose technologies capable of transforming and automating complex workflows by converting unstructured text into structured data. Dynamic systems require constant updates to reflect ongoing changes, necessitating methods that ensure accurate and efficient updates. This work proposes an approach for modeling dynamic worlds through a relational database continuously updated by SQL expressions derived from textual descriptions of changes. By employing agents and LLMs to generate SQL commands for inserting, altering, and deleting data, the system ensures minimal manual intervention.

Our research demonstrates that the proposed approach effectively automates dynamic updates, providing a robust framework for managing complex, evolving environments. Future work could explore advanced knowledge representation methods, integrate multiagent systems to handle complex updates efficiently, develop visualization tools for dynamic environments, and evaluate scalability in real-time, large-scale contexts.

Ethics Declaration: Ethical clearance was not required for the development of this research.

AI Declaration: AI tools such as Grammarly and ChatGPT were used solely for language revision. The authors' analysis and interpretations are their own.

References

- AI@Meta, 2024. Llama 3 Model Card.
- Arjoranta, J., 2011. Defining Role-Playing Games as Language-Games. *Int. J. Role-Play*. 3–17.
<https://doi.org/10.33063/ijrp.vi2.190>

- Beckman, T.J., 1999. The Current State of Knowledge Management, in: Liebowitz (Ed.), Knowledge Management Handbook. CRC Press.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., others, 2020. Language models are few-shot learners. *Adv. Neural Inf. Process. Syst.* 33, 1877–1901.
- Codd, E.F., 1972. Relational Completeness of Data Base Sublanguages, in: Rustin, R. (Ed.), Database Systems. Prentice Hall, Englewood Cliffs, NJ, USA, pp. 65–98.
- Codd, E.F., 1970. A Relational Model of Data for Large Shared Data Banks. *Commun. ACM* 13, 377–387.
- Cook, M., Tweet, J., Williams, S., 2003. Dungeons & Dragons Player’s Handbook. Wizards of the Coast, Renton, WA.
- crewAI, 2024. CrewAI [WWW Document]. URL <https://www.crewai.com/> (accessed 6.25.24).
- Daniau, S., 2016. The Transformative Potential of Role-Playing Games—: From Play Skills to Human Skills. *Simul. Gaming* 47, 423–444. <https://doi.org/10.1177/1046878116650765>
- de Castro Peixoto, L., Barbosa, R., Ferreira de Faria, A., 2022. Management of Regional Knowledge: Knowledge Flows Among University, Industry, and Government. *J. Knowl. Econ.* 13. <https://doi.org/10.1007/s13132-020-00702-9>
- Devlin, J., Chang, M.-W., Lee, K., Toutanova, K., 2018. BERT: Pre-training of deep bidirectional transformers for language understanding. *ArXiv Prepr. ArXiv181004805*.
- Dong, X., Zhang, C., Ge, Y., Mao, Y., Gao, Y., Chen, lu, Lin, J., Lou, D., 2023. C3: Zero-shot Text-to-SQL with ChatGPT.
- Hadi, M.U., Al-Tashi, Q., Qureshi, R., Shah, A., Muneer, A., Irfan, M., Zafar, A., Shaikh, M., Akhtar, N., Wu, J., Mirjalili, S., 2023. Large Language Models: A Comprehensive Survey of its Applications, Challenges, Limitations, and Future Prospects. <https://doi.org/10.36227/techrxiv.23589741.v1>
- Halliwel, R., Priestley, R., Davis, G., Bamba, J., Gallagher, P., 1986. Warhammer Fantasy Roleplay. Games Workshop, Nottingham.
- Hong, Z., Yuan, Z., Chen, H., Zhang, Q., Huang, F., Huang, X., 2024. Knowledge-to-SQL: Enhancing SQL Generation with Data Expert LLM.
- Huang, W., Zheng, X., Ma, X., Qin, H., Lv, C., Chen, H., Luo, J., Qi, X., Liu, X., Magno, M., 2024. An Empirical Study of LLaMA3 Quantization: From LLMs to MLLMs.
- International Organization for Standardization, 2016. ISO/IEC 9075-1:2016 Information technology – Database languages – SQL – Part 1: Framework (SQL/Framework). ISO/IEC, Geneva, Switzerland.
- International Organization for Standardization, 1995. ISO/IEC 13211-1:1995 Information technology — Programming languages — Prolog — Part 1: General core.
- Jefferson, D.R., 1985. Virtual Time. *ACM Trans. Program. Lang.* 7, 405–425.
- Kumar, V., Raheja, G., Sachdeva, S., 2012. DATABASE MANAGEMENT. *Int. J. Comput. Technol.*
- Li, G., Zhou, X., Zhao, X., 2024. LLM for Data Management. *Proc. VLDB Endow.* 17, 4213–4216.
- Ma, L., Pu, K., Zhu, Y., 2024. Evaluating LLMs for Text-to-SQL Generation With Complex SQL Workload.
- Mannila, H., Rähkä, K.-J., 1992. The design of relational databases. Addison-Wesley Longman Publishing Co., Inc., USA.
- Nonaka, I., 1994. A dynamic theory of organizational knowledge creation. *Organ. Sci.* 5, 14–37.
- Nonaka, I., Takeuchi, H., 1995. The Knowledge-creating Company: How Japanese Companies Create the Dynamics of Innovation, Everyman’s library. Oxford University Press.
- Nonaka, I., Toyama, R., Konno, N., 2000. SECI, Ba and Leadership: a Unified Model of Dynamic Knowledge Creation. *Long Range Plann.* 33, 5–34. [https://doi.org/10.1016/S0024-6301\(99\)00115-6](https://doi.org/10.1016/S0024-6301(99)00115-6)
- Noy, N.F., McGuinness, D.L., 2001. Ontology Development 101: A Guide to Creating Your First Ontology (No. KSL-01-05). Stanford Knowledge Systems Laboratory.
- N.R, Q., Ikhtiyorovna, K.G., 2023. DEVELOPMENT OF FANTASY GENRE IN 20TH CENTURY. *Intent Res. Sci. J.* 2, 1–5.
- OpenAI, 2023. GPT-4 Technical Report. OpenAI.
- Pourreza, M., Rafiei, D., 2024. DTS-SQL: Decomposed Text-to-SQL with Small Large Language Models.
- Prat, N., 2006. A Hierarchical Model for Knowledge Management. *Encycl. Knowl. Manag.* 1. <https://doi.org/10.4018/978-1-59904-931-1.ch036>
- Raiaan, M., Hossain, Md.S., Fatema, K., Fahad, N., Sakib, S., Mim, Most.M.J., Ahmad, J., Ali, M.E., Azam, S., 2023. A Review on Large Language Models: Architectures, Applications, Taxonomies, Open Issues and Challenges. <https://doi.org/10.36227/techrxiv.24171183>
- Richardson, L., 2007. Beautiful soup documentation. April.
- Stollenwerk, M.F.L., 2001. Gestão do conhecimento: conceitos e modelos, in: Tarapanoff, K. (Ed.), Inteligência Organizacional e Competitiva. UNB, Brasília, pp. 143–163.
- Sun, R., Arik, S.Ö., Muzio, A., Miculicich, L., Gundabathula, S., Yin, P., Dai, H., Nakhost, H., Sinha, R., Wang, Z., Pfister, T., 2024. SQL-PaLM: Improved Large Language Model Adaptation for Text-to-SQL (extended).
- Tan, Z., Liu, X., Shu, Q., Li, X., Wan, C., Liu, D., Wan, Q., Liao, G., 2024. Enhancing Text-to-SQL Capabilities of Large Language Models through Tailored Promptings, in: Calzolari, N., Kan, M.-Y., Hoste, V., Lenci, A., Sakti, S., Xue, N. (Eds.), Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024). ELRA and ICCL, Torino, Italia, pp. 6091–6109.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., Rodriguez, A., Joulin, A., Grave, E., Lample, G., 2023. LLaMA: Open and Efficient Foundation Language Models.
- Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W.X., Wei, Z., Wen, J., 2024. A survey on large language model based autonomous agents. *Front. Comput. Sci.* 18. <https://doi.org/10.1007/s11704-024-40231-1>

- Weng, L., 2023. LLM-powered Autonomous Agents. lilianweng.github.io.
- Xi, Z., Chen, W., Guo, X., He, W., Ding, Y., Hong, B., Zhang, M., Wang, J., Jin, S., Zhou, E., Zheng, R., Fan, X., Wang, X., Xiong, L., Zhou, Y., Wang, W., Jiang, C., Zou, Y., Liu, X., Yin, Z., Dou, S., Weng, R., Cheng, W., Zhang, Q., Qin, W., Zheng, Y., Qiu, X., Huang, X., Gui, T., 2023. The Rise and Potential of Large Language Model Based Agents: A Survey.
- Yang, X., Procopiuc, C.M., Srivastava, D., 2009. Summarizing relational databases. *Proc VLDB Endow* 2, 634–645. <https://doi.org/10.14778/1687627.1687699>
- Zhang, Q., Dong, J., Chen, H., Li, W., Huang, F., Huang, X., 2024. Structure Guided Large Language Model for SQL Generation.
- Zhou, X., Li, G., Liu, Z., 2023. LLM As DBA.
- Zhu, X., Li, Q., Cui, L., Liu, Y., 2024. Large Language Model Enhanced Text-to-SQL Generation: A Survey.