Regular Issue

# The Application of Large Language Models and Virtual Assistants in Project Management Research: A Review

Jesus Gil Ruiz[1]*, Javier Zayas Gallardo[2], Hernán Díaz Rodríguez[3]*

[1] Department of Computer Science, Universidad Europea (Spain)
[2] Department of Information Technology and Telematic Systems Engineering, University of Extremadura (Spain)
[3] Department of Computer Science, University of Oviedo (Spain)

* Corresponding author: jesus.gil@universidadeuropea.es, diazhernan@uniovi.es

Received 18 February 2025 | Accepted 5 February 2026 | Published 23 February 2026

LA UNIVERSIDAD EN INTERNET

## Abstract

The rapid evolution of generative artificial intelligence (AI) is transforming project management practices by enhancing efficiency, productivity and adaptability in decision-making processes. The integration of large language models (LLMs) into project management research and practice is reviewed, with a particular focus on virtual assistants as decision support tools. State-of-the-art models such as Mistral, Large Language Model Meta AI (LLaMa), Bidirectional Encoder Representations from Transformers (BERT) and T5, are assessed for their potential to automate complex project tasks, extract insights from project datasets, and optimize decision-making across various project management domains and business sectors. Generative AI is shown to surpass traditional project management systems by not only analysing historical project data but also generating new strategies and solutions in real time. Applications include project risk assessment, resource allocation optimisation, stakeholder communication and project performance prediction. The role of fine-tuning and retraining LLMs to adapt them to industry-specific project management challenges is also examined enhancing relevance and performance across diverse business environments.

## Keywords

Business Decisions, Generative AI, LLMs, Virtual Assistants.

DOI: 10.9781/ijimai.2026.2230

## I. Introduction

In recent years, projects have become larger and more complex, evolving into significant undertakings. Alongside this, industrial expansion has created a greater need for specialised skills in project management, which is critical given the typically slim profit margins involved. Although established project management methodologies offer essential frameworks for initiating and advancing projects, enabling issues to be addressed quickly and effectively, traditional methods alone may prove insufficient. Managing complex and evolving challenges requires a more structured and closely monitored approach across all project areas [1].

The objective is to strengthen the project manager's capacity to handle difficult situations during development and to reduce errors arising from inadequate planning or oversight, particularly in areas such as portfolio management [2], [3], [4]. Even with widely used management practices such as the project management professional (PMP) [5] in place, a clearer structure and comprehensive oversight across all project domains, including information systems, remains essential.

Traditional approaches often leave project managers relying heavily on intuition and prior experience when making decisions, especially when confronted with issues involving numerous variables and outcomes. It is nearly impossible for a single project manager to address all the complexities of modern projects alone, which contributes to the high rate of project failures [6].

Generative artificial intelligence can be applied not only to project management [7], but also to business decision-making more broadly. In today's rapidly changing and data-centric landscape, organizations must make fast and well-informed decisions across various domains such as marketing, finance and operations. Generative AI models, including generative pre-trained transformer (GPT)-based systems and advanced machine learning techniques, can process large datasets, predict future trends and recommend effective strategies. By automating repetitive tasks, providing valuable insights, and simulating different scenarios, these models enable decision-makers to focus on higher-level strategies while reducing risks and inefficiencies, ultimately improving overall decision-making and fostering business growth [8], [9], [10], [11].

Please cite this article as:
J. Gil Ruiz, J. Zayas Gallardo, H. Díaz Rodríguez. The Application of Large Language Models and Virtual Assistants in Project Management Research: A Review, International Journal of Interactive Multimedia and Artificial Intelligence, vol. 9, no. 6, pp. 105-115, 2026, http://doi.org/10.9781/ijimai.2026.2230

Previous studies have explored this concept within the specific context of applying machine learning to project management [12], [13]. Gil et al. [12] examined how machine learning techniques can optimise project workflows, enhance resource allocation, and support decision-making processes in project environments. Building on these findings, the present research expands this approach by applying similar principles to novel domains.

In classical AI systems, decision support tools are designed to assist businesses by analysing historical data, identifying patterns, and offering recommendations to guide strategic decisions [14], [15]. These systems typically rely on machine learning algorithms and statistical models to process large datasets, make predictions, and assess risks, helping businesses optimise their decisions in areas such as resource management [16], market strategies and financial forecasting. They provide a structured approach to decision-making by focusing on existing information and using predictive analytics to suggest data-driven solutions, reducing the reliance on human intuition.

These conventional approaches are constrained in managing intricate and ever-changing problems. The role of generative AI, as a development of classical models addresses this limitation. In addition to analysing and forecasting based on existing data, generative AI models such as GPT and other language models can generate new ideas, optimise tactics in real time and provide highly flexible insights. These models expand the capabilities of previous systems by producing additional content, strategies or potential outcomes, offering organisations a more adaptable and powerful toolfor complex decision-making tasks [17], [18], [19].

Various generative AI models exist, each with unique capabilities that can be tailored to different applications, including the development of specialised chatbots through retraining open models [20]. These chatbots can be designed to handle specific tasks, such as customer support [21], [22], [23], team collaboration, or complex decision-making processes in business environments [14], [15]. Large language models (LLMs), can be adapted to support strategic decision-making by analysing data trends, forecasting outcomes, and providing insights in real time, which is particularly valuable in corporate settings. An AI-powered chatbot, for example, could assist managers by streamlining information [24], identifying potential risks [25], [26], and suggesting data-driven strategies [27], [28], thereby enhancing productivity and enabling more informed decision-making in dynamic business scenarios.

The present study examines the various architectures that have contributed to the development of contemporary models used in the creation of virtual assistants, with a particular focus on project management and business decision-making environments. It also highlights practical applications that leverage existing models through retraining techniques, demonstrating their adaptability and relevance in addressing complex organisational challenges.

A structrured literature review was conducted by querying the Scopus and Google Scholar databases. The search strategy employed a combination of keywords representing core project management knowledge areas (e.g., "resource allocation," "risk assessment," "portfolio management") alongside technical terms including "Large Language Models," "LLMs," "Generative AI," and "fine-tuning".

Initially, approximately 100 records were identified. After removing duplicates and screening titles and abstracts, the remaining candidates were evaluated against specific inclusion and exclusion criteria. Works were included if they:

- Demonstrated an explicit link to a recognized project management process or knowledge area.
- Presented empirical or design science research involving the deployment or evaluation of an LLM customized via fine-tuning.

Conversely, studies were excluded if they:

- Focused solely on general-purpose AI without specific domain adaptation for business or project environments.
- Lacked sufficient technical detail regarding the underlying architecture or training methodology.
- Were not published in English.

This process yielded a final set of 29 studies for in-depth analysis. In addition to peer-reviewed publications, the review incorporates preprints and technical reports in order to reflect the fast-paced evolution of the LLM field, where state-of-the-art models, benchmarks, and applications are frequently disseminated through open repositories prior to formal journal publication.

This study has several limitations. Despite efforts to capture recent advances, the rapid evolution of LLMs implies that some emerging models or applications may not be fully represented. Moreover, while the review identifies general trends and insights, the practical effectiveness of LLM-based approaches in project management contexts is likely to vary depending on organizational characteristics, domain-specific requirements, and data availability. Finally, the scope of this review is primarily limited to text-based LLMs and does not extensively address multimodal or hybrid AI systems, which may become increasingly relevant in future project management applications.

Based on the systematic analysis of the literature, the main contributions and key insights of this review can be summarized as follows:

- LLMs are increasingly used as decision support tools in project management, moving beyond descriptive analytics towards generative and adaptive decision-making capabilities.
- Different project management domains tend to favour different LLM architectures. Open-source models such as LLaMa and Mistral are preferred in contexts requiring customisation, data privacy, and regulatory compliance, while proprietary models such as GPT are more commonly adopted for rapid deployment and stakeholder communication tasks.
- Fine-tuning is a critical enabler for effective adoption of LLMs in project management. Across domains such as finance, risk management, logistics, and human resources, fine-tuned models consistently outperform general-purpose LLMs in terms of relevance, accuracy, and contextual alignment.
- Generative capabilities represent a qualitative shift with respect to traditional decision-support systems. Rather than relying solely on historical data analysis, LLMs are increasingly used to simulate scenarios, generate alternative strategies, and support real-time adaptive planning.
- Trust, transparency, and human oversight remain central challenges. While LLMs improve efficiency and decision quality, the literature highlights the risks of over-reliance and the importance of explainability and governance mechanisms in practical deployments.

This paper is structured into two main sections. The first focuses on generative AI models, exploring their evolution and the key architectural advancements that have shaped their current capabilities. The second section addresses the practical applications of these models within project management, emphasizing their role in optimising decision-making processes and improving efficiency in organisational settings.

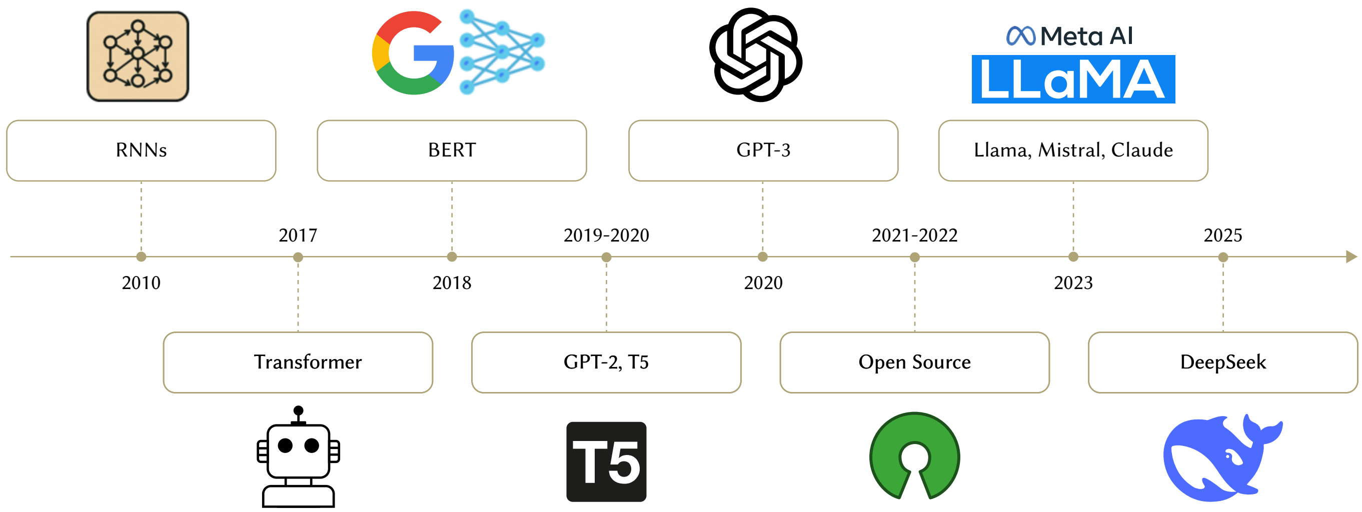


Fig. 1. Timeline of LLM evolution: Key milestones.

## II. Generative AI Models

Generative AI refers to a subset of artificial intelligence focused on creating new and original content [29] such as text, images, audio or code, rather than merely analysing or classifying data. It uses advanced models such as neural networks [30], often based on architectures such as Transformers or generative adversarial networks [31], to learn patterns from existing data and generate outputs that emulate human creativity. Applications include chatbots [32] and realistic voice synthesis enabling automation and innovation across diverse fields.

The development of LLMs has progressed through significant milestones over the past decade, driven by breakthroughs in neural network architectures and increases in computational power (Fig. 1) [33]. In the early 2010s, recurrent neural networks (RNNs) [34] were the dominant models for sequential data processing, providing an initial framework for tasks such as text generation and sentiment analysis. However, their limitations in handling long-range dependencies led to Vaswani et al.'s introduction of the revolutionary Transformer architecture in the seminal 2017 paper Attention Is All You Need [35].

Transformers removed the sequential processing bottleneck of RNNs, enabling parallelisation and scalability, that are critical for modern LLMs. This advance has contributed to the emergence and proliferation of virtual assistants and generative AI systems. The improvement was achieved by replacing the recurrent and convolutional layers [36] of traditional neural networks with a self-attention mechanism, which allows the model to evaluate the importance of each word in a sentence relative to others, regardless of distance. This capability to capture long-range dependencies made Transformers highly effective for understanding and generating human language.

The Transformer neural network architecture (Fig. 2), consists of an encoder (left block) and a decoder (right block), both composed of stacked identical layers. Each encoder layer has two sub-layers: a multi-head self-attention mechanism and a position-wise fully connected feed-forward network. The decoder adds a third sub-layer, which performs multi-head attention over the encoder stack's output. Each sub-layer incorporates a residual connection and layer normalisation. Inputs and outputs are embedded into vectors, and positional encoding is added to provide information about word order in the sequence. The Transformer relies solely on self-attention to compute representations of input and output without the need for sequence-aligned RNNs or convolution. This scalability has enabled the development of advanced LLMs such as GPT, Bidirectional Encoder Representations from Transformers (BERT), Mistral, Large Language Model Meta AI (LLaMa) and Text-To-Text Transfer Transformer (T5) which have demonstrated exceptional abilities in language comprehension, generation, and even reasoning.

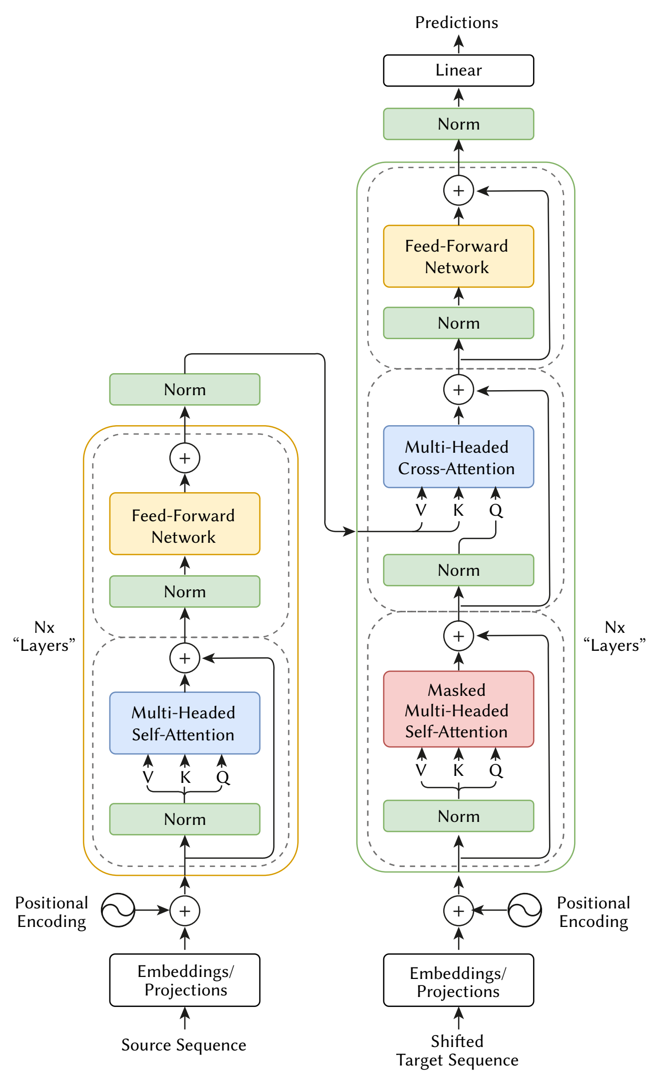


Fig. 2. The transformer - model architecture, reproduced from [37].

Building on the Transformer architecture, Google's BERT [38], [39], [40] (2018) introduced bidirectional context understanding to language representation, marking a paradigm shift in natural language understanding tasks. Models such as GPT-2 and T5 [41], [42] (2019–2020) further extended LLM capabilities, excelling in both generative and comprehension tasks. OpenAI's GPT-3 [43], [44] (2020) demonstrated the potential of large-scale pretraining, with 175 billion parameters, establishing itself as a benchmark in generative AI.

The architecture of the GPT-3 language model (Fig. 3), represents an advanced evolution of the Transformer paradigm, incorporating structural modifications to enhance contextual processing and semantic representation. The Text and Position Embed block indicates the embedding of input text and its position within the sequence, which is essential for understanding context. Unlike the original architecture, which used both encoder and decoder stacks, GPT-3 employs a decoder-only structure. This configuration focuses on generating text by predicting the next word in a sequence, using masked multi-head self-attention [45] to focus on relevant preceding words. As with the Transformer, GPT-3 also incorporates residual connections, layer normalisation, and feed-forward networks. The primary distinction lies in GPT-3's specialisation for text generation, with a streamlined architecture compared to the more general-purpose encoder decoder structure of the original Transformer.

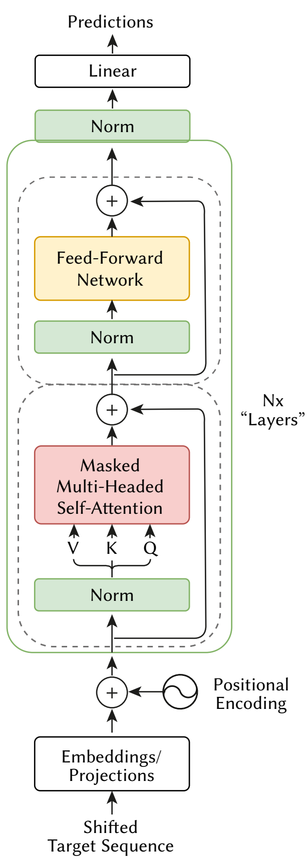


Fig. 3. GPT Transformer decoder, adapted from [37].

More recently, the emergence of open-source models (2021–2022) and the exploration of graph-based LLMs have highlighted a diversification in model architectures and applications. Open-source efforts have democratised AI research, fostering innovation and accessibility. Platforms such as Hugging Face [46] have become central to this democratisation, offering a vast repository of pre-trained models and tools for natural language processing NLP, computer vision and other applications. Hugging Face serves as a hub for both researchers and practitioners, enabling access to state-of-the-art models and supporting practical applications across diverse domains, including healthcare, education and entertainment [47], [48], [49].

Mistral [50], [51] is an emerging open-source AI model launched in 2023 and designed to emphasise efficiency and innovation in generative AI. Mistral 7B, a decoder-only language model, shares its foundation with Transformers but incorporates distinct features. While both utilise stacked decoder layers with self-attention and feed-forward networks, Mistral introduces optimizations such as grouped-query attention [52] and sliding window attention for enhanced efficiency. It also applies root mean square (RMS) normalisation and SiLU activation, differing from the original Transformer. Additional differences include rotary positional embeddings (RoPE) and a focused design for text generation, setting it apart from the Transformer's broader encoder decoder arquitecture.

LLaMA [53], [54] was developed by Meta (formerly Facebook) with a research-focused design and launched in 2023. LLaMa adopts a decoder-only language model that, while inspired by Transformers, includes key differences. Unlike the original Transformer's encoder decoder structure, LLaMa focuses solely on the decoder for text generation. It retains core elements such as self-attention and feed-forward networks within each decoder layer but incorporates specific modifications. LLaMa applies RMS normalisation in place of layer normalisation, and its feed-forward network features an expansion and contraction of dimensionality (H Ñ 2.7H Ñ H) using SiLU activation. Compared to the general-purpose Transformer, LLaMa represents a streamlined architecture specialised for text generation with distinct internal configurations.

Recently, additional models have gained prominence, including Anthropic's Claude [55], Nemotron [56] and DeepSeek [57], [58], [59].

DeepSeek has emerged as a family of advanced language models, distinguished by its emphasis on computational efficiency and knowledge generation. DeepSeek leverages optimised training techniques and enhanced scalability to deliver strong performance in text comprehension and generation tasks. Its architecture introduces innovations in context handling and adaptability, making it particularly suitable for business and research applications. With these advances, DeepSeek positions itself as a competitive alternative within the LLM ecosystem, rivaling open-source models such as LLaMa and Mistral.

These developments underscore the dynamic and competitive nature of the field, with a growing variety of architectures tailored for specific tasks and contexts, further enriching the LLM ecosystem.

## III. Applications of LLMs in Project Management

This section details where these models are being retrained and applied in real-world project management environments. A summary is provided in Table I.

### A. Tenders

In project management, tenders [60] represent a structured process in which interested parties (vendors, contractors, suppliers) submit formal bids to undertake specific work within a project. This competitive process allows project managers to evaluate proposals based on criteria such as cost, expertise, and timeline, ensuring the selection of the most suitable partner. Tenders are crucial for promoting transparency, fairness and value.

Ranjan Parida et al. [61] examine the integration of LLMs into retail customer relationship management (CRM) systems, highlighting their impact on customer support, engagement, and sales strategies. LLMs enhance CRM capabilities by automating customer interactions, providing real-time responses and offering multilingual

TABLE I. Summary of Projects Retraining Models in Each Area of Project Management

| Category | Source | Model |
|---|---|---|
| | Tenders | |
| Financial Analysis | [63] | Finantial analyst AI |
| Sales Prediction | [61]<br>[62] | N/D<br>Bert |
| | Project | |
| Decision Making | [65] | N/D |
| Retail | [66] | RetailGPT |
| Customer lifecycle | [67] | Llama |
| Business Applications | [68] | Bert |
| Business Process | [69] | Bert |
| Product Decisions | [70] | Bert |
| Risk Management | [71]<br>[72] | CMITS3-30E<br>T5 |
| | Human Resources | |
| Human Resources | [73] | GPT-2 |
| Recruitment | [74]<br>[75] | JobReco<br>GPT-2 |
| | Information Technology | |
| Cibersecurity | [76]<br>[77] | Hackphyr<br>Bert |
| Innovation Support | [78] | LLaMa2, GPT-2 |
| | Logistics | |
| Logistic | [79] | LLaMa, Vicuna, Mistral |
| Automation | [80] | LLama3-7B |
| Object Detection | [81]<br>[82] | DINOv2<br>ContextDET |
| | Engineering & Design | |
| Proyect Data Analytics | [83] | N/D |
| Design Automation | [84] | ChatEDA |
| | Operations | |
| Maintenance | [85] | Mistral-7B, LLaMa3-8B, Phi3-mini |
| Finantial Management | [86]<br>[87]<br>[88] | LLaMa 2<br>LLaMa 2<br>LLaMa 2 |
| | Supply Chain | |
| Supply Chain | [89] | T5 |
| | Construction | |
| Portfolio selection | [90] | LLM with RAG |

support, thereby improving customer satisfaction and retention. These models also personalise content and conduct sentiment analysis, to drive customer loyalty. In addition, LLMs support sales strategies by analysing customer data for lead scoring and product recommendations. The study addresses technical challenges, including data privacy, model interpretability, and computational efficiency, and suggests best practices for successful integration.

Varma et al. [62] present a custom BERT architecture for sales and support conversations, demonstrating how incorporating the generative capabilities of GPT with a clean, small dataset improves performance. The enhancements bring the model's output closer to that of fine-tuned LLMs. Future research includes applying the architecture to chat and voice data, expanding datasets to over 100 million conversations, and developing customer-specific models to further optimise prediction performance.

The Financial Analyst AI on Hugging Face [63] assists businesses by analysing financial documents such as earnings calls. It transcribes audio to text, summarises content and evaluates fiscal sentiment using specialised finance models, helping to identify trends and risks. It also extracts key company-related information through named entity recognition and detects forward-looking statements. These features enable faster, more informed decision-making by simplifying complex financial data.

### B. Project

A project [64] is a temporary endeavour undertaken to create a unique product, service, or result. It has a defined beginning and end, involving coordinated efforts and resources to achieve specific objectives. Projects are characterised by distinct goals, limited resources and a focus on delivering value within the constraints of time, cost and scope. Effective project management involves planning, organising, executing and closing projects to meet stakeholder expectations.

Eigner et al. [65] identify key determinants of LLM-assisted decision-making, emphasising their influence on decision quality. Findings show that trust and reliance on LLMs strongly affect outcomes, with over-reliance leading to errors when users accept inaccurate suggestions without scrutiny. Transparency, including clear explanations of LLM outputs, enhances user trust but may also expose model limitations. Prompt engineering is shown to improve decision quality by guiding LLMs to produce more relevant and accurate outputs. Task-specific factors such as complexity and accountability also shape the effectiveness of LLM support.

Sharee et al. [66] present RetailGPT, a tailored LLM for the retail industry. RetailGPT significantly improves product recommendations, customer sentiment analysis and inventory forecasting. The model's understanding of customer preferences is enhanced by the contextualized item attention mechanism, which increases the accuracy of recommendations and enriches customer experience. Its sentiment analysis provides real-time feedback, allowing retailers to respond quickly to changing consumer demands. In addition, its forecasting capabilities reduce operational costs, highlighting the model's operational impact. Despite this advancements, RetailGPT represents only an early step in the continuing evolution of AI in retail.

Soni [67] explores the use of LLMs in customer lifecycle management, focusing on stages including awareness, acquisition, purchase, and post-purchase engagement. LLMs improve lead identification, targeting and marketing strategy through advanced analytics and A/B testing. By analysing customer data patterns, LLMs help to personalise interactions, optimise sales processes and enhance post-purchase. The study also discusses challenges such as data biases, lack of emotional nuance in personalised content and over-reliance on technology, stressing the need for ongoing human oversight.

Bilal et al. [70] address information overload on online review platforms by fine-tuning a BERT base model to predict the helpfulness of customer reviews. Tests on Yelp shopping reviews show that BERT-based classifiers significantly outperform traditional bag-of-words methods, demostrating superior generalisation and accuracy in classifying helpful versus unhelpful reviews.

Chen et al. [69] propose a multi-task prediction method using BERT and transfer learning for predictive business process monitoring (PBPM) [91]. The approach is tested on 11 real-world event logs and achieves high accuracy, F-score [92], and area under the receiver operating characteristic curve (auc-roc [93], [94]) compared to state-of-the-art methods, particularly for complex datasets such as BPIC2012 and Sepsis. Fine-tuning the pre-trained BERT model significantly reduces training time while improving prediction performance.

Lee et al. [71] demonstrate the effectiveness of fine-tuned LLMs, particularly CMITS3-30E, in construction site management during severe weather. Using domain-specific datasets, the model generates context-specific responses aligned with safety best practices. CMITS3-30E outperformes foundational models in accuracy, specificity and efficiency. The study highlights the importance of comprehensive fine-tuning and domain knowledge for improving LLM performance in safety and operational resilience.

Xia et al. [72] report on the use of the fine-tuned Typhoon-T5 model, to enhance disaster information retrieval related to typhoons. By integrating spatiotemporal features and factual typhoon data, the model provides accurate, context-specific responses. Leveraging retrieval-augmented generation (RAG) technology, it outperforms other models in reasoning and information retrieval. Future improvements include expanding to other disaster domains, incorporating multimodal data and knowledge graphs and addressing limitations in current question-and-answer systems.

### C. Human Resources

Human resources (HR) in project management [95] involves managing the people who contribute to a project. This includes recruitment, hiring, training, motivating and managing the project team. Effective HR management ensures that the right personnel with the necessary skills are available at the right time and fosters a collaborative and productive environment.

Bhatnagar [73] presents the use of transfer learning to enhance response generation in HR peer learning. By fine-tuning the GPT-2 model with employee survey data improvements were observed in domain-specific vocabulary and language comprehension. The study illustrates transfer learning as a promising approach for tailoring LLMs to specialised tasks and improving NLP performance. The paper also adresses the challenges of ethical considerations, bias and handling complex scenarios. These remain critical for practical deployment, with the findings offering insights for further refinement and responsible AI application in resource-constrained settings.

### D. Information Technology

Information technology (IT) [96] plays a critical role in modern project management, providing tools and systems to plan, execute and monitor projects effectively. Project management software, communication platforms, and data analytics tools enable efficient collaboration, information sharing, and decision-making. IT facilitates the tracking of progress, resource allocation and risk management, contributing to improved project outcomes and stakeholder satisfaction.

Rigaki et al. [76] highlight the creation and fine-tuning of a 7-billion-parameter LLM, Hackphyr, tailored for network security tasks. Addressing privacy concerns and hardware limitations, Hackphyr was designed as an on-premises solution and achieved performance comparable to leading commercial models such as GPT-4. Using a custom dataset, the model's weaknesses in complex environments and generating actionable outputs were mitigated. Extensive behavioural analysis provides a foundation for explaining LLM agents' actions, supporting further advancements in explainability and effectiveness.

Amine Ferrag et al. [77] explore the use of the BERT architecture in cybersecurity with the SecurityBERT model, which demonstrated an accuracy of 98.2% in detecting 14 types of cyber attacks. The model outperformed traditional machine learning and deep learning approaches, including CNNs and RNNs. The paper proposes future improvements, through fine-tuning and retraining to handle a wider range of cyber threats and incorporate real-time data updates. The automation of mitigation actions based on the model classifications is also suggested to enhance cybersecurity management.

Nguyen et al. [78] demonstrate the importance of fine-tuning as a method to adapt LLMs to specific domains and tasks, resulting in enhanced performance and contextual understanding. Focusing on Japanese-language tasks, the study shows how retraining models on domain-specific data can address linguistic and cultural nuances, improving accuracy and making LLMs more effective specialized environments.

### E. Logistics

Logistics in project management involves the planning, execution, and control of the flow of goods, services, and information between the point of origin and the point of consumption to meet project requirements. This includes transportation, warehousing, inventory management, and distribution. Efficient logistics ensure timely delivery of materials, equipment and resources, minimising delays and optimising project execution.

Olena [79] presents the importance of fine-tuning LLMs for improving customer service in the logistics industry. The study focuses on integrating AI chatbots using LLMs to enhance external communication with customers and improve user satisfaction. By employing parameter-efficient techniques such as QLoRA, pre-trained models including Llama, Vicuna and Mistral were fine-tuned on a logistics-specific dataset. The research highlights the value of domain-specific fine-tuning in achieving superior chatbot performance and contributing to more effective customer service.

Xia et al. [80] present an approach to integrating LLMs into industrial automation systems, significantly improving task automation and system flexibility. By incorporating LLMs within a multi-agent framework and linking them with digital twins, the study demonstrates their ability to manage production planning and control autonomously. This integration supports dynamic decision-making and reduces the need for human intervention. A fine-tuned open-source model achieved performance comparable to proprietary GPT-4 in this context. Despite challenges such as real-time latency and economic concerns, the study highlights the potential of LLMs to advance intelligent autonomous systems in industrial settings.

### F. Engineering and Design

Engineering and design [97] encompass the technical aspects of project management, involving the application of scientific and engineering principles to plan, design, and develop project deliverables. This includes creating technical specifications, conducting feasibility studies, developing prototypes, and overseeing the construction or implementation of project outputs. Effective engineering and design ensure that project objectives and performance requirements are met.

Bratić et al. [83] address the challenge of centralised storage for educational materials in fragmented databases, aiming to improve retrieval efficiency for educators and students. The study introduces a hybrid model based on the Transformer framework and an API for an existing LLM-based chatbot, to provide precise responses from a comprehensive materials database. The model uses advanced word embedding techniques for accurate text processing, enhancing efficiency and adaptability. The research offers a technical solution and demonstrates the potential for integrating emerging technologies in education to promote a more accessible and efficient learning environment.

### G. Operations

Operations in project management [98] refer to the ongoing activities and processes involved in executing the project plan and delivering the project's outputs. This includes tasks such as managing resources, coordinating tasks, monitoring progress, and ensuring quality control. Efficient operations focus on optimising workflows, minimising waste and maximising productivity to achieve project goals effectively.

Fatem et al. [86] apply instruction fine-tuning to smaller-scale LLMs, including Mistral-7B, LLaMa3-8B and Phi3-mini, to improve performance in financial text classification tasks. The retraining process adapts these models to specialised financial texts, resulting in significant improvements in task-specific performance and better zero-shot capabilities on complex financial tasks.

Agarwal et al. [87] fine-tune the LLaMa 2 7b-hf model for financial sentiment analysis using parameter-efficient fine-tuning methods and the Simple Fine-tuning Trainer. These adaptations enhance the model's ability to handle the complex linguistic characteristics of financial texts, achieving significant improvements in sentiment classification accuracy across negative, neutral, and positive categories.

Robert [88] explores the application of fine-tuned LLaMa 2 in financial sentiment analysis, showcasing its role in interpreting complex financial narratives. By leveraging fine-tuning techniques, the model improves analysis of diverse data sources, such as news articles, social media and financial reports, enabling accurate trend identification, investor sentiment gauging and financial outcome prediction.

Myöhänen et al. [85] examine the potential of LLMs for enhancing predictive maintenance and process optimisation in industrial businesses. They show how LLMs can automate the analysis of complex text data, such as maintenance logs and process documents, to detect patterns, identify anomalies and optimise processes. The study emphasises the importance of fine-tuning with domain-specific data to improve performance and highlights the risks of bias and performance limitations, urging careful integration and monitoring of LLMs in industrial systems.

### *H. Supply Chain*

The supply chain in project management [99] encompasses the network of organisations, people, activities, information and resources involved in moving a product or service from its origin to the end customer. It includes suppliers, manufacturers, distributors and the project team itself. Effective supply chain management ensures timely and cost-effective procurement of materials, equipment and services required for project completion.

Cheng et al. [89] explore the application of fine-tuned LLMs to extract hierarchical schemas and predict disruptions in complex supply chains, highlighting their value in structured schema learning and real-time predictions. By fine-tuning models such as T5 on domain-specific data, the research shows how they can identify patterns and relationships in event sequences, even under complex conditions like geopolitical tensions or legislative changes. The methodology includes annotated datasets, robust evaluation metrics such as precision, recall, and F-score, and real-world case studies, demonstrating how LLMs provide actionable insights. These findings highlight the potential of LLMs to enhance decision-making, offering timely predictions and proactive strategies to mitigate risks.

### *I. Construction*

Construction refers to the process of building or assembling infrastructure or physical structures within a project. It includes activities like site preparation, foundation laying, structural erection, system installation and finishing works. Construction project management focuses on coordinating resources, managing subcontractors, ensuring safety, and delivering the project within budget, schedule, and quality standards.

Zhao [90] introduces an innovative portfolio management framework that integrates advanced NLP and deep reinforcement learning (DRL). Leveraging Transformer-based models and LLMs the framework enhances return predictions and sentiment extraction from financial texts, while DRL agents optimise portfolio performance. It uses a RAG module with dynamic weight adjustments, fine-tuned LLMs and a hierarchical reinforced trader for strategic stock selection and precise trade execution. The approach outperforms benchmarks such as the S&P 500, improving Sharpe ratios in bullish markets and reducing losses in bearish conditions. The study also addresses interpretability challenges using a univariate flagging algorithm extended through a generalized additive model, offering predictive accuracy even with missing or imbalanced data. This work sets a new benchmark for intelligent financial decision-making through advanced AI integration.

## IV. Discussion

In this section, the key points presented earlier, along with the contents of Table I, are analysed to identify broader patterns and emerging trends in decision-making processes within project management influenced by the use of LLMs.

A key observation is the increasing adoption of LLMs to support and automate project management tasks. These include risk assessment, resource allocation, stakeholder communication and performance forecasting, functions that have traditionally required human expertise and manual effort. The ability of LLMs to process unstructured data and generate coherent outputs in natural language provides a level of responsiveness previously unattainable with classical project management tools.

The review reveals that different types of LLMs are employed based on the nature of the task and the operational requirements of specific sectors. In financial project management, for example, open-source models such as LLaMa have gained considerable traction. These models are deployed for budgeting, financial forecasting and regulatory compliance, particularly in settings where data privacy and traceability are essential. Their open architecture and capacity for fine-tuning allow organisations to retain control over sensitive data and adapt model outputs to specific reporting standards.

Risk analysis is another area where LLMs are making a notable impact. Mistral models, with their efficient architecture and competitive performance, are being fine-tuned for industries such as infrastructure, energy and construction. These models support dynamic risk identification and can generate mitigation strategies based on real-time data and project documentation. Their computational efficiency makes them well suited to integration with existing project monitoring systems that have limited processing capacity.

Tasks involving knowledge extraction, document summarisation, and decision support often rely on models such as BERT and T5. These models are particularly effective in sectors where project documentation is dense and structured knowledge is critical, including healthcare, legal services and scientific research. Their ability to classify information and produce concise summaries helps project teams navigate complex requirements and streamline compliance.

In manufacturing and logistics, the use of LLMs for resource allocation and scheduling is becoming increasingly common. Fine-tuned versions of T5 and LLaMa are integrated with enterprise resource planning systems to optimise workforce deployment and material flow. By anticipating resource constraints and proposing alternative allocation strategies, these models help to enable more resilient and adaptive project execution.

A clear distinction also emerges between the adoption of open-source and proprietary models. Proprietary solutions, such as GPT based models, dominate in contexts where rapid deployment and integration with existing cloud-based tools are priorities. Open-source alternatives, by contrast, are favoured where customisation, regulatory compliance, and cost control are paramount. The ability to fine-tune models on domain specific datasets is particularly valuable, allowing organisations to align model outputs with internal processes.

Another significant trend is the shift from purely analytical to generative applications of LLMs. Beyond summarising historical data, these models are increasingly used to simulate alternative project scenarios, formulate strategic responses and assist with adaptive planning. This generative capability marks a fundamental change in the role of AI in project management, transitioning from passive data analysis to active involvement in strategic decision-making. LLMs also provide real-time insights, automate routine tasks and support complex leaders to focus on strategy and accelerate decision-making through rapid data processing.

In conclusion, the integration of LLMs into project management practices represents a major transformation in how organisations address complexity, uncertainty and decision-making. The strategic choice between open-source and proprietary models reflects an awareness of the trade-offs between performance, control and regulatory compliance. LLMs are emerging as essential tools in the digital transformation of project-based work across a range of sectors.

## V. Conclusion

This study examined the integration of LLMs into project management, with particular emphasis on their role as decision support tools. By synthesising current applications and model-specific deployments across sectors, it has been shown how generative AI, particularly LLMs such as GPT, LLaMa, Mistral, BERT and T5, is reshaping project workflows. Unlike traditional decision-support systems, which rely heavily on static historical data and predefined rules, generative AI enhances adaptability by generating new strategies, offering creative solutions, and optimising real-time decision-making.

Section IV highlighted clear patterns linking specific models to particular project domains. For example, the widespread use of LLaMa 2 in financial management is driven by its open-source nature and adaptability for secure environments. GPT's language generation capabilities make it particularly suited for stakeholder communication, while Mistral models are being fine-tuned for dynamic risk assessment in infrastructure and energy projects. This differentiated adoption underscores the importance of aligning model selection with both technical requirements and sector priorities.

Furthermore, it was found that that fine-tuning models for specialised tasks such as financial sentiment analysis, market trend forecasting and PBPM allows organisations to operate more efficiently. These models can analyse complex data and generate insights that closely align with organisational objectives, supporting decision-making based on a more thorough understanding of the business environment. The capability of AI to interpret nuanced data, including financial language or customer sentiment, further enhances its value as a strategic decision-making aid.

For practitioners and organisations seeking to integrate LLMs into project management processes, the following roadmap is proposed: First, identify project tasks where language understanding or generation plays a central role, including reporting, planning or documentation. Second, assess whether data sensitivity or regulatory constraints require the use of open-source models. Third, conduct pilot deployments using models such as LLaMa or T5 for targeted tasks, and iteratively fine-tune them with internal data. Finally, clear evaluation metrics, such as time savings, decision accuracy or stakeholder satisfaction should be established to measure impact and justify scaling efforts.

However, the implementation of generative AI in business requires attention to trust and transparency. Although LLMs offer impressive capabilities, their effectiveness depends on how well they are integrated into business processes and on the level of trust users have in their outputs. Over-reliance on AI-generated suggestions without human scrutiny can result in errors, making it essential to balance AI recommendations with human oversight. Additionally, providing clear explanations of model outputs is critical to fostering trust and understanding, especially in complex decision-making scenarios.

In conclusion, the integration of generative AI into business decision-making offers a substantial opportunity to optimise operations, enhance strategic planning and support long-term growth. By leveraging advanced AI models organisations can gain valuable insights, improve decision quality and adapt to rapidly changing market conditions. The ability of these models to process vast amounts of data, generate new strategies and provide real-time insights empowers businesses to make more informed and effective decisions. As AI technologies become more widely adopted, the potential for innovation and competitive advantage will continue to grow, establishing generative AI as a vital tool for the future of business.

## Credit Authorship Contribution Statement

Jesús Gil Ruíz: Conceptualization, Methodology, Investigation, Project administration, Supervision.

Javier Zayas Gallardo: Data curation, Investigation, Visualization, Writing – original draft, Writing – review & editing.

Hernán Díaz Rodríguez: Formal analysis, Investigation, Validation, Visualization, Writing – review & editing, Supervision.

## Data Statement

No primary data were generated for this systematic literature review. The study is based on published articles retrieved from Google Scholar and Scopus. The search strategy and inclusion/exclusion criteria are reported in the manuscript to enable replication.

## Declaration of Conflicts of Interest

No conflict of interest exists.

## Acknowledgment

The authors have no acknowledgements to declare. Funding: This research did not receive funding.

## References

[1] J. Chan, Y. Li, "Enhancing team diversity with generative ai: A novel project management framework," in *2024 IEEE 48th Annual Computers, Software, and Applications Conference (COMPSAC)*, 2024, pp. 1648–1652, IEEE.

[2] F. Costantino, G. Di Gravio, F. Nonino, "Project selection in project portfolio management: An artificial neural network model based on critical success factors," *International Journal of Project Management*, vol. 33, no. 8, pp. 1744–1754, 2015.

[3] A. Barcaui, A. Monat, "Who is better in project planning? generative artificial intelligence or project managers?," *Project Leadership and Society*, vol. 4, p. 100101, 2023.

[4] B. Haidabrus, *Generative AI in Agile, Project, and Delivery Management*, pp. 100–110. Springer, 2024.

[5] H. Singh, P. S. Williams, "A guide to the project management body of knowledge: Pmbok (®) guide," in *Project management institute*, 2021, pp. 1–8.

[6] J. Holmström, N. Carroll, "How organizations can innovate with generative ai," *Business Horizons*, 2024.

[7] V. Kanabar, *The AI Revolution in Project Management: Elevating Productivity with Generative AI*. Sams Publishing, 2023.

[8] E. L. Chuma, A. M. Alves, G. G. de Oliveira, "Evolution of generative ai for business decision-making: A case of chatgpt," *Management Science and Business Decisions*, vol. 4, no. 1, pp. 5–14, 2024.

[9] E. L. Chuma, G. G. De Oliveira, "Generative ai for business decision-making: A case of chatgpt," *Management Science and Business Decisions*, vol. 3, no. 1, pp. 5–11, 2023.

[10] E. Kromidha, R. M. Davison, "Generative ai-augmented decision-making for business information systems," in *IFIP International Conference on Human Choice and Computers*, 2024, pp. 46–55, Springer.

[11] P. R. Kiran, S. Khaiyum, A. R. Palandye, S. D. Apoorva, R. Arjun, S. A. Akshaya, S. Annapoorna, "Leveraging llama3 and langchain for rapid ai application development," *Journal of Electrical Systems*, vol. 20, pp. 2146–2153, 2024.

[12] J. G. Ruiz, J. M. Torres, R. G. Crespo, "The application of artificial intelligence in project management research: A review," *International Journal of Interactive Multimedia and Artificial Intelligence*, vol. 6, no. 6, pp. 54–66, 2021, doi: https://doi.org/10.9781/ijimai.2020.12.003.

[13] J. G. Ruiz, H. Díaz, R. G. Crespo, "The application of artificial intelligence planning and scheduling in photovoltaic plant construction projects," *Expert Systems*, vol. 42, no. 2, p. e13798, 2025, doi: https://doi.org/10.1111/exsy.13798.

[14] J. Merkert, M. Mueller, M. Hubl, "A survey of the application of machine learning in decision support systems," in *Proceedings of the 23rd European Conference on Information Systems (ECIS 2015)*, 2015, Association for Information Systems (AIS). Paper 133.

[15] A. Awaysheh, J. Wilcke, F. Elvinger, L. Rees, W. Fan, K. L. Zimmerman, "Review of medical decision support and machine-learning methods," *Veterinary pathology*, vol. 56, no. 4, pp. 512–525, 2019.

[16] F. Hussain, S. A. Hassan, R. Hussain, E. Hossain, "Machine learning for resource management in cellular and iot networks: Potentials, current solutions, and open challenges," *IEEE communications surveys & tutorials*, vol. 22, no. 2, pp. 1251–1275, 2020.

[17] D. Patel, "Revolutionizing project management with generative ai," *International Scientific Journal of Engineering and Management*, vol. 2, pp. 1–8, 2023.

[18] V. Kanabar, *The AI Revolution in Project Management: Elevating Productivity with Generative AI*. Sams Publishing, 2023.

[19] J. C. Weng, "Putting intellectual robots to work: Implementing generative ai tools in project management," NYU SPS Applied Analytics Laboratory, July 2023. [Online]. Available: http://hdl.handle.net/2451/69531.

[20] J. Kietzmann, A. Park, "Written by ChatGPT: AI, large language models, conversational chatbots, and their place in society and business" *Business Horizons*, vol. 67, no. 5, p. 453 – 459, 2024, doi: https://doi.org/10.1016/j.bushor.2024.06.002.

[21] R. K. Verma, N. Kumari, "Generative ai as a tool for enhancing customer relationship management automation and personalization techniques," *International Journal of Responsible Artificial Intelligence*, vol. 13, no. 9, pp. 1–8, 2023.

[22] S. Bamberger, N. Clark, S. Ramachandran, V. Sokolova, "How generative ai is already transforming customer service," *Boston Consulting Group*, 2023.

[23] N. Sadhotra, N. Gupta, "Generative ai in customer service," *Digital Paradigm Shift: Unravelling Technological Disruption in Business*, vol. 31, 2023.

[24] N. Rane, "Role and challenges of chatgpt and similar generative artificial intelligence in human resource management," *Available at SSRN 4603230*, 2023.

[25] Y. Wang, "Generative ai in operational risk management: Harnessing the future of finance," *Operational Risk Management: Harnessing the Future of Finance (May 17, 2023)*, 2023.

[26] S. Bidah, K. Akdim, M. Zahid, "Application of generative ai (chatgpt as example) in risk management," Available at *SSRN 4841318*, 2024.

[27] T. J. Marion, M. Srour, F. Piller, "When generative ai meets product development," *MIT Sloan Management Review*, vol. 66, no. 1, pp. 14–15, 2024.

[28] R. Kasar, T. Kumar, "Digital twin and generative ai for product development," *Procedia CIRP*, vol. 128, pp. 905–910, 2024.

[29] E. Brynjolfsson, D. Li, L. R. Raymond, "Generative ai at work," National Bureau of Economic Research, 2023.

[30] H. Abdi, D. Valentin, B. Edelman, *Neural Networks*. No. 124 in Quantitative Applications in the Social Sciences, Thousand Oaks, CA: Sage Publications, 1999.

[31] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, "Generative adversarial networks," *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.

[32] E. Adamopoulou, L. Moussiades, "Chatbots: History, technology, and applications," *Machine Learning with applications*, vol. 2, p. 100006, 2020.

[33] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, *et al.*, "A survey of large language models," *arXiv preprint arXiv:2303.18223*, 2023.

[34] S. Grossberg, "Recurrent neural networks," *Scholarpedia*, vol. 8, no. 2, p. 1888, 2013.

[35] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[36] Z. Li, F. Liu, W. Yang, S. Peng, J. Zhou, "A survey of convolutional neural networks: analysis, applications, and prospects," *IEEE transactions on neural networks and learning systems*, vol. 33, no. 12, pp. 6999–7019, 2021.

[37] H. Díaz Rodríguez, *Deep Learning with Python: Algorithms, Architectures and Practical Applications.* Amazon KDP, 2026.

[38] S. Ravichandiran, *Getting Started with Google BERT: Build and train state-of-the-art natural language processing models using BERT.* Packt Publishing Ltd, 2021.

[39] J. D. M.-W. C. Kenton, L. K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of naacL-HLT*, vol. 1, 2019, p. 2, Minneapolis, Minnesota.

[40] S. Yadav, "What is bert? How it is trained? A high level overview," 2023. [Online]. Available: https://medium.com/@Suraj_Yadav/what-is-bert-how-it-is-trained-a-high-level-overview-120

[41] J. Ni, G. H. Abrego, N. Constant, J. Ma, K. B. Hall, D. Cer, Y. Yang, "Sentence-t5: Scalable sentence encoders from pre-trained text-to-text models," *arXiv preprint arXiv:2108.08877*, 2021.

[42] A. Payong, "A unified text-to-text framework for nlp tasks: An overview of t5 model," 2023. [Online]. Available: https://blog.paperspace.com/flan-t5-architecture/#%3A~%3Atext%3DThere%20are%20two%20stages%20to%2Cto%20learn%20abstract%20linguistic%20representations

[43] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, *et al.*, "GPT-4 technical report," *arXiv preprint arXiv:2303.08774*, 2023.

[44] G. Roffo, "Exploring advanced large language models with llmsuite," *arXiv preprint arXiv:2407.12036*, 2024.

[45] N. Ilinykh, S. Dobnik, "How vision affects language: Comparing masked self-attention in uni-modal and multi-modal transformer," in *Proceedings of the 1st Workshop on Multimodal Semantic Representations (MMSR)*, 2021, pp. 45–55.

[46] H. Face, "Hugging face," 2024. [Online]. Available: https://huggingface.co.

[47] Y. Zhao, H. Yin, B. Zeng, H. Wang, T. Shi, C. Lyu, L. Wang, W. Luo, K. Zhang, "Marco-o1: Towards open reasoning models for open-ended solutions," 2024. [Online]. Available: https://arxiv.org/abs/2411.14405.

[48] M. A. Team, "Pixtral-large," 2024. [Online]. Available: https://huggingface.co/mistralai/Pixtral-Large-Instruct-2411.

[49] B. Hui, J. Yang, Z. Cui, J. Yang, D. Liu, L. Zhang, T. Liu, J. Zhang, B. Yu, K. Dang, *et al.*, "Qwen2. 5-coder technical report," *arXiv preprint arXiv:2409.12186*, 2024.

[50] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. de las Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, *et al.*, "Mistral 7b," *arXiv preprint arXiv:2310.06825*, 2023.

[51] A. Torfi, "Mixtral 8x7b: Pioneering the next generation of large language models," 2023. [Online]. Available: https://medium.com/@amirsina.torfi/mixtral-8x7b-pioneering-the-next-generation-of-large 4753.

[52] S. S. Chinnakonduru, A. Mohapatra, “Weighted grouped query attention in transformers,” *arXiv preprint arXiv:2407.10855*, 2024.

[53] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, *et al.*, “The llama 3 herd of models,” *arXiv preprint arXiv:2407.21783*, 2024.

[54] Pratik, “Llama 2 : Explained in simple step by step process,” 2023. [Online]. Available: https://medium.com/@shahip2016/llama-2-explained-in-simple-step-by-step-process-5076e072cb69.

[55] A. J. Adetayo, M. O. Aborisade, B. A. Sanni, “Microsoft copilot and anthropic claude ai in education and library service,” *Library Hi Tech News*, 2024.

[56] B. Adler, N. Agarwal, A. Aithal, D. H. Anh, P. Bhattacharya, A. Brundyn, J. Casper, B. Catanzaro, S. Clay, J. Cohen, *et al.*, “Nemotron-4 340b technical report,” *arXiv preprint arXiv:2406.11704*, 2024.

[57] H. Lu, W. Liu, B. Zhang, B. Wang, K. Dong, B. Liu, J. Sun, T. Ren, Z. Li, H. Yang, *et al.*, “Deepseek-vl: towards real-world vision-language understanding,” *arXiv preprint arXiv:2403.05525*, 2024.

[58] A. Liu, B. Feng, B. Wang, B. Wang, B. Liu, C. Zhao, C. Dengr, C. Ruan, D. Dai, D. Guo, *et al.*, “Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model,” *arXiv preprint arXiv:2405.04434*, 2024.

[59] A. Liu, B. Feng, B. Xue, B. Wang, B. Wu, C. Lu, C. Zhao, C. Deng, C. Zhang, C. Ruan, *et al.*, “Deepseek-v3 technical report,” *arXiv preprint arXiv:2412.19437*, 2024.

[60] D. Watt, B. Kayis, K. Willey, “Identifying key factors in the evaluation of tenders for projects and services,” *International journal of project management*, vol. 27, no. 3, pp. 250–260, 2009.

[61] P. R. Parida, S. Ramalingam, J. Perumalsamy, “Large language models in retail crm systems: A technical evaluation of improving customer support, engagement, and sales strategies,” *Journal of Artificial Intelligence Research and Applications*, vol. 4, no. 2, pp. 85–130, 2024.

[62] A. Varma, C. Bhat, “Sas-bert: Bert for sales and support conversation classification using a novel multi-objective pre-training framework,” *Computer Science & Information Technology (CS & IT), ISSN*, pp. 2231–5403, 2023.

[63] R. Shah, “Financial analyst ai,” 2024. [Online]. Available: https://huggingface.co/spaces/rajistics/Financial_Analyst_AI.

[64] D. L. Cleland, L. R. Ireland, *Project management*. McGraw-Hill Professional, 2006.

[65] E. Eigner, T. Händler, “Determinants of llm-assisted decision-making,” *arXiv preprint arXiv:2402.17385*, 2024.

[66] F. Shareef, R. Ajith, P. Kaushal, K. Sengupta, “Retailgpt: A fine-tuned llm architecture for customer experience and sales optimization,” in *2024 2nd International Conference on Self Sustainable Artificial Intelligence Systems (ICSSAS)*, 2024, pp. 1390–1394, IEEE.

[67] V. Soni, “Large language models for enhancing customer lifecycle management,” *Journal of Empirical Social Science Studies*, vol. 7, no. 1, pp. 67–89, 2023.

[68] E. Haaralahti, “Utilization of local large language models for business applications,” master’s thesis, Aalto University, School of Science, Espoo, Finland, 2024.

[69] H. Chen, X. Fang, H. Fang, “Multi-task prediction method of business process based on bert and transfer learning,” *Knowledge-Based Systems*, vol. 254, p. 109603, 2022.

[70] M. Bilal, A. A. Almazroi, “Effectiveness of fine-tuned bert model in classification of helpful and unhelpful online customer reviews,” *Electronic Commerce Research*, vol. 23, no. 4, pp. 2737–2757, 2023.

[71] J. Lee, K. Jang, A. E. Sparkling, K. Kang, “Large language model–based construction site management for severe weather preparedness,” *Journal of Construction Engineering and Management*, vol. 152, no. 1, p. 04025220, 2026.

[72] Y. Xia, Y. Huang, Q. Qiu, X. Zhang, L. Miao, Y. Chen, “A question and answering service of typhoon disasters based on the t5 large language model,” *ISPRS International Journal of Geo-Information*, vol. 13, no. 5, p. 165, 2024.

[73] D. Bhatnagar, *Fine-Tuning Large Language Models for Domain-Specific Response Generation: A Case Study on Enhancing Peer Learning in Human Resource*. PhD dissertation, Dublin, National College of Ireland, 2023.

[74] C. Gan, Q. Zhang, T. Mori, “Application of llm agents in recruitment: A novel framework for resume screening,” *arXiv preprint arXiv:2401.08315*, 2024.

[75] S. H. R. Malik, P. Redij, S. Kulkarni, “Fairhire: A fair and automated candidate screening system,” in *Machine Intelligence, Tools, and Applications: Proceedings of the International Conference on Machine Intelligence, Tools, and Applications–ICMITA 2024*, vol. 40, 2024, pp. 372–382, Springer Nature.

[76] M. Rigaki, C. Catania, S. Garcia, “Hackphyr: A local fine-tuned llm agent for network security environments,” *arXiv preprint arXiv:2409.11276*, 2024.

[77] M. A. Ferrag, M. Ndhlovu, N. Tihanyi, L. C. Cordeiro, M. Debbah, T. Lestable, N. S. Thandi, “Revolutionizing cyber threat detection with large language models: A privacy-preserving bert-based lightweight model for iot/iiot devices,” *IEEE Access*, vol. 12, pp. 23733–23750, 2024, doi: https://doi.org/10.1109/ACCESS.2024.3363469.

[78] C. Nguyen, T. Tran, S. T. Luu, N.-K. Le, D.-T. Do, S. Kawamura, S. Naito, Y. Mogi, L.-M. Nguyen, “Fostering business innovation with ai: Performance of fine-tuned japanese language models under resource restrictions,” in *New Frontiers in Artificial Intelligence*, vol. 14741 of *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence)*, 2024, Springer Nature. Presented at JSAI-isAI 2024 (International Workshop on AI of and for Business - AI-Biz2024).

[79] O. Karaim, “Application of LLMs for a chatbot system in the logistics industry,” bachelor’s thesis, Ukrainian Catholic University, Faculty of Applied Sciences, Lviv, Ukraine, 2024. Department of Computer Sciences.

[80] Y. Xia, J. Zhang, N. Jazdi, M. Weyrich, “Incorporating large language models into production systems for enhanced task automation and flexibility,” *arXiv preprint arXiv:2407.08550*, 2024.

[81] G. Han, S.-N. Lim, “Few-shot object detection with foundation models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 28608–28618.

[82] Y. Zang, W. Li, J. Han, K. Zhou, C. C. Loy, “Contextual object detection with multimodal large language models,” *International Journal of Computer Vision*, vol. 133, pp. 825–843, 2025, doi: https://doi.org/10.1007/s11263-024-02214-4.

[83] D. Bratić, M. Šapina, D. Jurečić, J. Žiljak Gršić, “Centralized database access: transformer framework and llm/chatbot integration-based hybrid model,” *Applied System Innovation*, vol. 7, no. 1, p. 17, 2024.

[84] K. Xu, R. Qiu, Z. Zhao, G. L. Zhang, U. Schlichtmann, B. Li, “Llm-aided efficient hardware design automation,” *arXiv preprint arXiv:2410.18582*, 2024.

[85] J. Myöhänen, “Improving industrial performance with language models: a review of predictive maintenance and process optimization,” bachelor’s thesis in industrial engineering and management, Lappeenranta-Lahti University of Technology LUT, Lappeenranta, Finland, 2023. Examiner: Postdoctoral researcher Kirsi Kokkonen.

[86] S. Fatemi, Y. Hu, M. Mousavi, “A comparative analysis of instruction fine-tuning llms for financial text classification,” *arXiv preprint arXiv:2411.02476*, 2024.

[87] P. Agarwal, A. Gupta, “Strategic business insights through enhanced financial sentiment analysis: A fine-tuned llama 2 approach,” in *2024 International Conference on Inventive Computation Technologies (ICICT)*, 2024, pp. 1446–1453, IEEE.

[88] A. Robert, “AI-Driven Financial Sentiment Analysis for Strategic Business Insights: the Roleof Fine-Tuned Llama 2.” EasyChair Preprint 14073, EasyChair, 2024.

[89] Z. -Q. Cheng, Y. Dong, A. Shi, W. Liu, Y. Hu, J. O’Connor, A. G. Hauptmann, K. S. Whitefoot, “Shield: Llm-driven schema induction for predictive analytics in ev battery supply chain disruptions,” *arXiv preprint arXiv:2408.05357*, 2024.

[90] Z. Zhao, *Next-Generation Intelligent Portfolio Management*. PhD dissertation, Massachusetts Institute of Technology, 2024.

[91] A. E. Márquez-Chamorro, M. Resinas, A. Ruiz-Cortés, “Predictive monitoring of business processes: a survey,” *IEEE Transactions on Services Computing*, vol. 11, no. 6, pp. 962–977, 2017.

[92] M. Sokolova, G. Lapalme, “A systematic analysis of performance measures for classification tasks,” *Information processing & management*, vol. 45, no. 4, pp. 427–437, 2009.

[93] C. E. Metz, “Basic principles of ROC analysis,” *Seminars in Nuclear Medicine*, vol. 8, no. 4, pp. 283–298, 1978, doi: https://doi.org/10.1016/s0001-2998(78)80014-2.

[94] S. Narkhede, “Understanding auc-roc curve,” *Towards data science*, vol. 26, no. 1, pp. 220–227, 2018.

[95] I. Clark, T. Colling, "The management of human resources in project management-led organizations," *Personnel Review*, vol. 34, no. 2, pp. 178–191, 2005.

[96] K. Schwalbe, *Information technology project management*. Cengage Learning, 2016.

[97] H. Eisner, *Essentials of project and systems engineering management*. John Wiley & Sons, 2008.

[98] A. Gunasekaran, E. W. Ngai, "The future of operations management: an outlook and analysis," *International Journal of Production Economics*, vol. 135, no. 2, pp. 687–701, 2012.

[99] J. B. Ayers, *Supply chain project management: a structured collaborative and measurable approach*. CRC press, 2003.


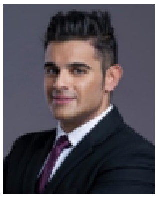

### Jesús Gil Ruiz

Jesús Gil Ruiz holds an International PhD in Computer Science and multiple engineering degrees, including Industrial, Civil, and Industrial Organization Engineering, along with a PMP® certification. He also earned an Executive MBA, a Master in Financial Management and Cost Control, and a Master in Railway Infrastructure Projects from the University of Barcelona. He completed executive programs in AI at MIT and Business Analytics at Wharton. Currently, he is CEO of PROJENER.AI, leading projects in renewable energy and oil and gas, and teaches in the Master of Data Visualization and Project Management at Universidad Europea.

### Javier Zayas Gallardo

Javier Zayas Gallardo is a postgraduate student at the University of Malaga. Software Engineer by the University of Malaga. He has a master's degree in Artificial Intelligence from the Universidad Europea an another in Artificial Intelligence and Software Engineering at the University of Malaga. He is currently doing a PhD in Computer Science from the University of Extremadura. He focuses on the areas of Artificial Intelligence and Quantum Computing, in which he is doing some research and his PhD.

### Hernán Díaz Rodríguez

Hernán Díaz Rodríguez holds a PhD in Artificial Intelligence from the University of Oviedo and an MBA from the Open University (UK). With over 20 years of experience in technology, he has led cloud-based projects for public administration bodies. He also spent five years as a researcher at CERN. He is the author of several publications in Artificial Intelligence, particularly in the areas of Optimization and Machine Learning.