

BOOTSTRAPPING YOUR AI JOURNEY: A STEP-BY-STEP STRATEGY TO GETTING STARTED WITH GENERATIVE AI

Sreenadh Payyambally*¹

*¹Staff Software Engineer, Automation Anywhere, San Jose, California, USA.

DOI: <https://www.doi.org/10.56726/IRJMETS68484>

ABSTRACT

This article outlines a thorough, step-by-step methodology for learning Artificial Intelligence (AI) from scratch. It begins by clarifying objectives and scope, then reviews fundamental AI literature to establish essential concepts and frameworks. A comparative analysis between conventional programming and AI highlights key distinctions in instructions, data usage, adaptability, and decision-making. Next, the article synthesizes critical terminology, explaining how methods like Supervised and Unsupervised Learning, Deep Learning, and Natural Language Processing fit into the broader AI ecosystem. Recognizing challenges such as data security, transparency, and bias, it outlines specialized techniques—prompt engineering, prompt tuning, retrieval-augmented generation, and fine-tuning large language models—to help students and practitioners address complex tasks efficiently. Additionally, the text explores AI agents, leveraging autonomy and learning capabilities to transform customer service and decision-making in various sectors. Practical best practices and real-world examples guide newcomers in crafting effective prompts, managing computational resources, and aligning AI tools with organizational objectives. Ultimately, readers learn to navigate and implement AI responsibly by considering performance needs, data quality, and ethical constraints. This structured, incremental approach ensures a solid foundation for understanding AI's evolving landscape, positioning learners for future advancements in the field. By following these steps meticulously, learners gain confidence in building AI solutions.

Keywords: Conventional Programming Vs Artificial Intelligence, Key AI Terminologies, Traditional AI Vs Generative AI, Prompt Engineering, Retrieval-Augmented Generation, Fine-Tuning, AI Agents.

I. INTRODUCTION

Artificial Intelligence (AI) has rapidly evolved from a niche field of computer science to a cornerstone of modern innovation, influencing sectors ranging from healthcare to finance. For those looking to build a foundation in AI, understanding its core principles, terminologies, and practical applications is essential. This article offers a structured, step-by-step methodology to guide learners—from absolute beginners to more advanced practitioners—through the fundamentals and intricacies of AI.

The methodology begins by distinguishing AI from traditional programming, underscoring how AI's data-driven adaptability contrasts with fixed, rule-based code. Next, it introduces essential AI concepts, including Supervised and Unsupervised Learning, Artificial Neural Networks (ANNs), and Deep Learning, highlighting their interconnected roles in processing and interpreting data. Recognizing real-world challenges—such as data security, transparency, and computational demands—the article then delves into specialized techniques like prompt engineering, retrieval-augmented generation (RAG), and fine-tuning large language models (LLMs). These methods streamline AI development and customize models for specific tasks, ensuring effectiveness and efficiency.

Finally, the article explores AI agents, showing how they combine learning and autonomous decision-making to drive transformative solutions. By following this roadmap, newcomers gain clarity on both the conceptual underpinnings and the practical steps required for embracing AI—whether in personal projects, business applications, or large-scale enterprise initiatives.

II. METHODOLOGY

A step-by-step plan for developing a comprehensive article on Artificial Intelligence (AI). It starts by clarifying the article's goals and scope, then moves on to a review of key AI literature. The methodology includes a comparative analysis of traditional programming and AI, clarifies core terminology, and tackles challenges like data security. It also highlights specialized AI practices—namely prompt engineering, prompt tuning, retrieval-

augmented generation, and fine-tuning large language models. Finally, it demonstrates how AI agents bring these methods together into cohesive, evolving systems, offering a clear framework for explaining today's AI landscape

1. Comparative Analysis

Conventional Programming vs. AI: Examine the differences in how each approach handles instructions, data, adaptability, and decision-making. Traditional AI vs. Generative AI: Present a side-by-side comparison of their purposes, typical data requirements, and industry applications.

2. Terminology Synthesis

Identify Core AI Terms: Collate definitions for Machine Learning, ANN, Deep Learning, Supervised vs. Unsupervised Learning, etc.

Clarify Relevance: Explain how each term fits into the broader AI ecosystem (e.g., how Deep Learning differs from basic ML, why data labeling matters).

3. Examination of AI Challenges

Data Security and Governance: Address privacy regulations like GDPR.

Transparency and Bias: Discuss the "black box" nature of some AI models and the risk of bias.

Resource Constraints: Evaluate the high computational costs and expertise required.

4. Technique Exploration

Prompt Engineering: Outline best practices for crafting prompts and why clarity affects AI output

Prompt Tuning: Explain how soft prompts can refine large language models with minimal changes. Retrieval-Augmented Generation (RAG): Show how linking LLMs to external databases or documents provides up-to-date, context-specific information.

5. AI Agents in Practice

Conceptual Overview: Define AI agents, their key features, and workflow (Perception → Decision → Action → Learning).

Real-World Applications: Examine how AI agents are used for customer service, automation, and decision support.

Conventional programming vs Artificial intelligence

Conventional programming	Artificial intelligence
Requires explicit instructions	Requires data
Requires human programmer	Requires to be machine to learn from the data
Terminology is deterministic programming, follows algorithms, patterns and code to accomplish specific purposes which is spec out already.	Autonomous in making decisions
Limited against New use case/ scenarios adaption	Adapt to new information/data

A) Key AI Terminologies

• **Machine Learning:**

A branch of AI that applies mathematics to large datasets to find trends and patterns, mapping inputs to outputs.

• **Model:**

Represents the learned mapping between inputs and outputs. It is trained from data so the system can make intelligent decisions through observation.

• **Artificial Neural Network (ANN):**

Simulates the structure of the human brain with input, hidden, and output layers, requiring vast amounts of data to learn effectively.

• **Deep Learning:**

A subset of ANN with multiple layers (deep neural networks) that can capture complex patterns in data.

• **Supervised Learning:**

Involves labeled datasets where humans (data scientists) guide the model by showing correct outputs for training.

• **Unsupervised Learning:**

Allows the model to learn patterns and relationships from unlabeled data without explicit human guidance.

• **Training Data Set:**

A segment of data used for teaching the model to recognize patterns.

• **Test Data Set:**

Unlabeled data used to evaluate the model’s performance and generalization.

• **Classify Data:**

Organizing data by assigning it to predefined categories.

• **Cluster Data:**

Grouping data based on similarity—often used in unsupervised learning.

• **Binary Classification:**

Classification with only two possible outcomes (e.g., yes/no, true/false).

• **Natural Language Processing (NLP):**

AI that enables the understanding of language context and meaning.

• **IoT Devices:**

Internet-connected devices (e.g., home assistants, wearables) that serve as significant data sources.

• **Big Data:**

Very large datasets that support machine learning by providing abundant examples for training.

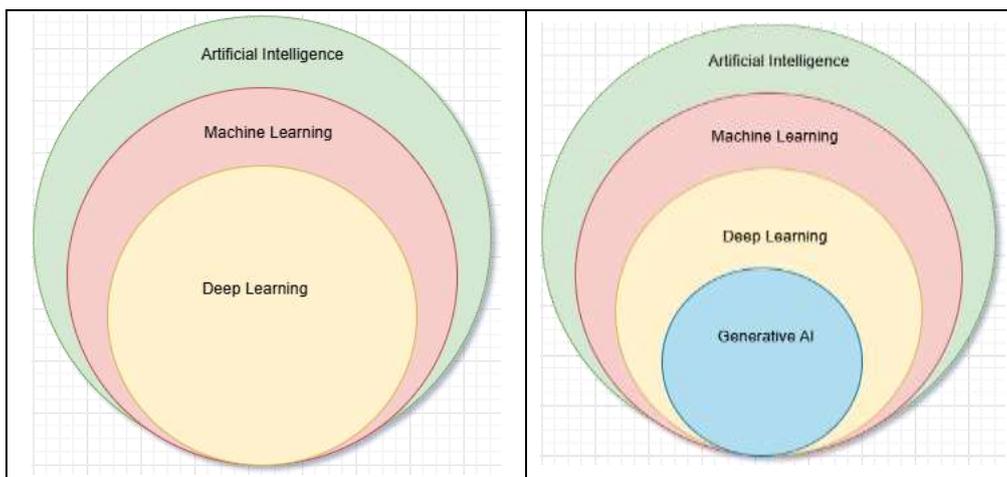
• **Cloud Computing:**

Enables quick access to powerful computing resources (e.g., GPUs) for efficient AI training and deployment.

• **Hallucination (in AI):**

Occurs when a generative AI model produces incorrect or fabricated information not grounded in real-world facts or training data

B) Traditional AI vs Generative AI



Traditional AI	Generative AI
Predictive AI which is designed for a task	Training systems to take existing data and generate into something new, Autonomous content creation, learning, mimicking patterns
Data intensive	Much more data intensive

Heavily depends on Supervised learning, labeled data to train models to recognize the pattern	Often follows unsupervised or self-supervised learning process.
Common industrial applications are 1) Predictive modeling 2) Classification tasks 3) Recommendation systems	Common retail and industrial applications are 1) LLM, language generations models 2) Code generation models 3) Image and Video Synthesis
Always focused improvement on accuracy, reliability and consistency	Focused wide range of use case-based creativity, adaptability and personalization

C) Challenges in Building Traditional AI Systems

• **Data Security and Governance Compliance:**

Ensuring data privacy and adhering to regulations (e.g., GDPR) can be complex.

• **Transparency Challenges:**

AI models—especially deep learning—can be opaque, making it hard to explain decisions.

• **High Costs:**

Collecting and storing large datasets, as well as computing resources (GPUs, TPUs), can be expensive.

• **Limited Data Assets:**

Some domains lack sufficient data, hampering model accuracy and robustness.

• **High Computing Capacity:**

Training sophisticated models often requires parallel computing on powerful hardware clusters.

• **Bias in AI Decision-Making:**

Biased training data can yield discriminatory results.

• **Technology Expertise:**

Building advanced AI systems requires a specialized skill set, from data scientists to cloud engineers.

D) Prompt Engineering

Prompt engineering involves giving clear, specific instructions to AI to ensure accurate and useful responses. It is similar to asking someone for help -be explicit about what we want, provide context, and set boundaries like length or detail. Techniques include defining tasks clearly, assigning roles (e.g., tutor or coach), breaking complex tasks into steps, and providing examples to clarify expectations. These strategies help the AI understand your needs and deliver relevant, targeted results.

‘Context’ in prompt engineering means giving the AI enough details to understand what we want. Without context, the AI might not give useful answers. For example, instead of saying, "Write an email," you could say, "Write a friendly email inviting a coworker to a team lunch next Friday at 1 PM." Adding details like the purpose, tone, or audience helps the AI give better and more relevant responses.

Language models can take a prompt and create completion by constantly predicting the next token: (Token: A unit easily understood by model)

E) Prompt Engineering Best Practices

Prompt engineering is about providing clear, concise instructions to an AI model to ensure accurate, relevant, and useful responses. By carefully crafting your prompts, you can guide the AI to produce precisely the type of information, style, and level of detail you require. Below are key recommendations for effective prompt engineering:

a) Be Specific and Clear

- Instead of: "Tell me about the economy."
- Use: "What are the major economic trends in the United States for 2025?"
- Why? Giving focused details helps the AI understand your exact requirements and reduces ambiguity.

b) Define the Desired Format

- Specify whether you want a short summary, bullet points, a step-by-step explanation, or a more narrative style.

- For instance: “Explain the stock market in three sentences for beginners.”
- c) Provide Examples
 - Show the AI what style, tone, or structure you are aiming for.
 - For example: “Write a poem like this one, but about winter.”
- d) Iterate and Refine
 - If the initial answer isn’t what you expected, rephrase your question or add more context.
 - Trying alternative approaches can help the AI better grasp your intent.
- e) Break Down Complex Tasks
 - For large or multifaceted requests, ask for one piece at a time.
 - Example: Instead of “Write a report on climate change,” try “First, summarize the causes of climate change.”
- f) Leverage the AI’s Strengths
 - If you need concise facts, ask for them specifically.
 - If you need creative brainstorming, frame your prompt accordingly.
- g) Create Reusable Templates
 - For tasks you do repeatedly, develop a standardized prompt structure or template.
 - This helps ensure consistency and saves time.
- h) Specify Length or Detail
 - Clarify whether you want one sentence, a paragraph, or a more extensive discussion.
 - The AI is more likely to meet your expectations if it knows how detailed or brief the response should be.
- i) Experiment and Provide Feedback
 - Try different phrasing or additional context to see how the AI’s output changes.
 - Offer feedback to further refine future responses.

By following these practices, you can significantly enhance the quality and relevance of the AI-generated content, making your interactions with AI more efficient and productive.

Why It Matters:

- Reduces the risk of ambiguous or irrelevant answers.
- Improves AI’s accuracy and efficiency.
- Saves time by minimizing trial-and-error interactions.

F) Prompt Tuning

Definition: Prompt tuning refines the behavior of large language models (LLMs) by training a small set of “soft prompts” instead of updating the model’s entire set of parameters. The pre-trained model remains largely unchanged, preserving its general language skills.

How It Works:

- Soft Prompt Initialization:

A small trainable vector is created to represent the new prompt tokens.

- Input Concatenation:

The soft prompts are prepended to the input text.

- Forward Pass and Optimization:

Only the prompt embeddings are updated through backpropagation, while the main model weights stay frozen.

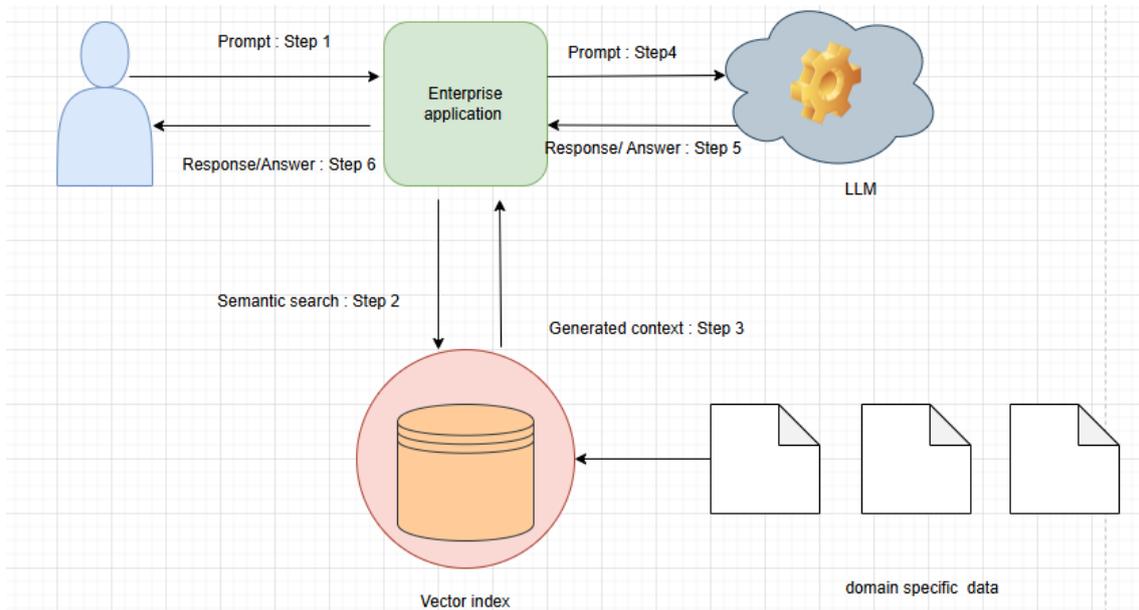
Advantages:

- Efficiency: Far fewer computational resources and training time are needed compared to full fine-tuning.
- Flexibility: A single LLM can be used for multiple tasks by swapping out different soft prompts.
- Preserved Generalization: The base model retains its broad language capabilities.

G) Retrieval-Augmented Generation (RAG): Empowering LLMs with Real-World Knowledge

Retrieval-Augmented Generation (RAG) is a game-changing technique that addresses a fundamental weakness of Large Language Models (LLMs): their dependence on static training data. RAG supercharges LLMs by connecting them to external knowledge sources, enabling them to access and integrate real-time information

directly into their responses. This results in significantly more accurate, relevant, and current outputs, unlocking the true potential of LLMs for real-world applications.



How RAG Works: A Detailed Look

RAG seamlessly blends the generative power of LLMs with the dynamic nature of information retrieval through a well-defined, multi-stage process

Key Steps:

a) Indexing (Knowledge Preparation):

Transform external documents or databases into vector embeddings stored in a specialized database.

b) Retrieval (Finding Relevant Information):

When a query is received, the system searches the database for similar or matching embeddings.

c) Augmentation (Contextual Enrichment):

Relevant documents are added to the model's prompt, providing it with real-time data.

d) Generation (Informed Response):

The model produces output grounded in the retrieved information, reducing hallucinations and improving reliability.

Use Cases:

- Intelligent customer support
- Research assistance
- Dynamic content creation
- Personalized recommendations
- Real-time data-driven answers

The Future of Intelligent Systems

RAG is a revolutionary approach that significantly amplifies the capabilities of LLMs by bridging the gap between generative AI and the vast world of real-world knowledge. Its ability to access and incorporate external information makes LLMs more reliable, accurate, adaptable, and ultimately, more useful. As research in this exciting field continues, RAG is poised to play a pivotal role in shaping the future of intelligent systems and transforming how we interact with information.

When to Choose Retrieval-Augmented Generation (RAG)

RAG is the preferred approach in several scenarios:

- Need for Up-to-Date Information: RAG excels at retrieving the latest data in real-time, making it ideal for rapidly changing fields.

- Detailed and Accurate Responses: Particularly valuable in domains like medical research, legal analysis, and technical documentation.
- Handling Complex Queries: Effective for answering intricate questions that require extensive knowledge and fact-checking.
- Ensuring Information Consistency: Essential in fields such as law and healthcare, where accuracy and coherence are critical.
- Frequent Data Updates: Ideal for applications that rely on constantly evolving information sources.
- Cost-Effective Knowledge Access: Reduces infrastructure costs while maintaining deep and diverse knowledge retrieval.
- Scalability and Security: Provides a reliable, scalable, and secure framework for integrating trustworthy information.
- Niche and Specific Queries: Useful when dealing with domain-specific, recent, or less commonly available information.
- Explainability and Transparency: Enables traceability by providing sources for retrieved information.
- Real-Time Knowledge Tasks: Crucial for applications that depend on immediate and dynamic data access.

RAG is especially useful for customer support systems, research tools, and knowledge retrieval platforms, where accuracy, scalability, and real-time adaptability are essential.

H) Fine-Tuning Large Language Models (LLMs)

Fine-tuning Large Language Models (LLMs) is a crucial process for adapting pre-trained models to specific domains and tasks, bridging the gap between general-purpose capabilities and specialized applications. This process involves training the pre-trained model on a targeted dataset, enhancing its performance and relevance for particular use cases.

Key Objectives of Fine-Tuning

Fine-tuning primarily aims to achieve two interconnected goals:

1. Knowledge Injection: Infusing the model with new, domain-specific information that wasn't present in the original training data. This allows the model to answer questions, generate content, and perform tasks within that specialized domain.
2. Alignment: Refining the model's output to adhere to specific formats, styles, guidelines, or desired behaviors. This ensures the model's responses are appropriate, consistent, and aligned with user expectations.

Data Requirements and Importance

The success of fine-tuning hinges on the quality and quantity of the training data. While larger datasets generally provide more information, data quality is paramount. Clean, relevant, and well-structured data is essential for effective learning.

Research, such as the LIMA paper, has demonstrated that high-quality alignment can be achieved even with relatively small datasets, provided they are carefully curated and representative of the target task.

Fine-Tuning Techniques

Several techniques are employed for fine-tuning LLMs, each with its strengths and weaknesses:

- a) Continued Pre-training: Further training the pre-trained model on a large corpus of domain-specific text.
- b) Instruction Tuning: Training the model on a dataset of task-specific instructions and corresponding outputs.
- c) Supervised Fine-Tuning (SFT): Training the model on a dataset of labeled examples, where each example consists of an input and the desired output.
- d) Reinforcement Learning from Human Feedback (RLHF) and Direct Preference Optimization (DPO): These techniques leverage human feedback to optimize the model's responses.

Best Practices for Effective Fine-Tuning

To maximize the effectiveness of fine-tuning, consider the following best practices:

- a) Prioritize Data Quality: Ensure the training data is clean, relevant, representative of the target task, and sufficiently large. Invest time in data cleaning and preprocessing.

b) Careful Hyperparameter Tuning: Experiment with different learning rates, batch sizes, and training epochs to find the optimal configuration for the specific task and dataset. Consider using learning rate schedules and early stopping.

c) Regular Evaluation and Monitoring: Continuously monitor the model's performance on a separate validation dataset during the fine-tuning process.

Potential Pitfalls and Mitigation Strategies

Fine-tuning can present several challenges:

a) Overfitting: The model learns the training data too well, leading to poor generalization on unseen data. Mitigation: Use more data, regularize the model, or use techniques like dropout.

b) Underfitting: The model fails to learn the underlying patterns in the data. Mitigation: Increase the model's capacity, train for longer, or adjust the learning rate.

c) Catastrophic Forgetting: The model loses previously acquired knowledge while specializing in the new domain. Mitigation: Use techniques like elastic weight consolidation or continual learning strategies.

d) Data Leakage: Information from the validation or test set inadvertently leaks into the training data. Mitigation: Ensure strict separation between datasets and careful data preprocessing.

Recent Advances and Future Directions

a) Parameter-Efficient Fine-Tuning (PEFT): Reduces computational costs and memory requirements by adjusting only a small subset of the model's parameters. Techniques like adapter modules and low-rank adaptation (LoRA) fall into this category.

b) Synthetic Data Generation: Leveraging high-quality synthetic data, especially in conjunction with techniques like instruction tuning, can be effective for fine-tuning, particularly when real-world data is scarce.

Fine-tuning LLMs is a powerful technique for adapting pre-trained models to specific domains and tasks. By carefully considering data quality, training strategies, and evaluation methods, practitioners can unlock the full potential of these models and achieve significant performance improvements. Ongoing research in areas like PEFT and synthetic data generation promises to further enhance the efficiency and effectiveness of fine-tuning in the future

When to Fine-Tune a Large Language Model (LLM)

Fine-tuning a large language model (LLM) can be transformative for tailoring general-purpose AI capabilities to specific tasks or domains. However, this process is resource-intensive, so it's critical to identify when fine-tuning truly adds value. Consider fine-tuning under the following circumstances:

a) Prompt Engineering Limitations

- Indicator: You have already tried various prompt engineering strategies, but the model consistently fails to deliver the desired accuracy or format.
- Implication: The LLM lacks domain-specific knowledge or struggles with nuanced requirements, indicating a need for deeper customization.

b) Domain-Specific Expertise is Crucial

- Indicator: Your case requires specialized vocabulary or knowledge not covered by a generic model (e.g., legal, medical, financial).
- Implication: Fine-tuning enables the LLM to learn and apply task-specific terminology or protocols more reliably.

c) Availability of High-Quality Training Data

- Indicator: You possess a substantial dataset that accurately reflects the types of inputs and outputs the model will handle.
- Implication: Clean, relevant data is essential; even smaller datasets can be effective if their quality is high.

d) Performance Optimization is Paramount

- Indicator: Your production environment demands reduced latency and minimized computational costs.
- Implication: A fine-tuned, smaller model may outperform a larger, general model for targeted tasks, leading to more efficient inference.

e) Proprietary or Sensitive Information

- Indicator: Your application deals with confidential or internally unique data.
- Implication: Fine-tuning avoids sharing sensitive data with external services, preserving privacy and security.

Key Considerations Before Fine-Tuning

- Prioritize Data Quality and Quantity

A robust, well-structured training set is critical. Depending on task complexity, aim for thousands of clean, representative examples.

- Rigorous Evaluation

Define clear metrics and use a held-out validation set to measure progress. Continuous evaluation helps in pinpointing improvement areas and in avoiding overfitting.

- Iterative Refinement

Fine-tuning is rarely final on the first try. Plan for ongoing monitoring and updates as your dataset, requirements, or the underlying LLM evolves.

- Exhaust Prompt Engineering First

In many cases, prompt engineering can achieve the desired results with less time and effort. Only proceed to fine-tuning once you have fully explored prompt-based solutions.

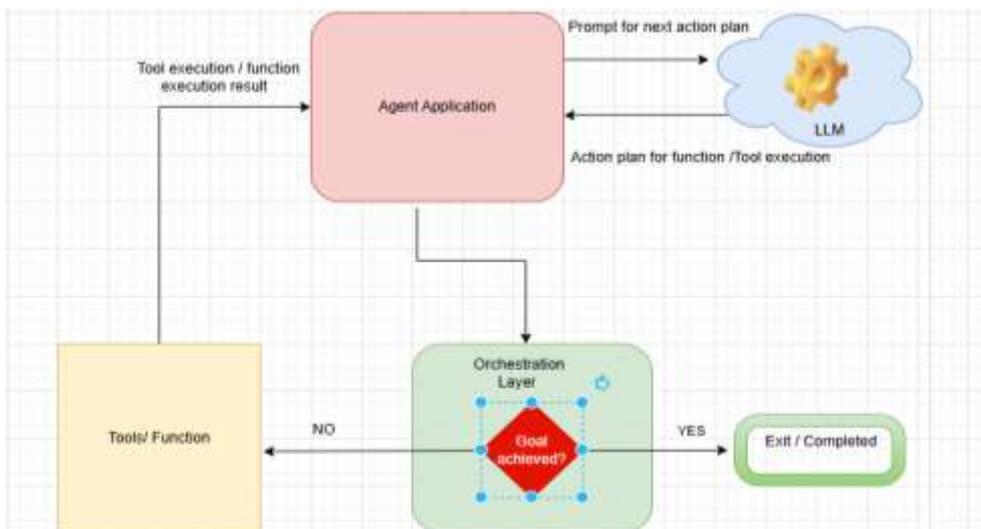
By carefully weighing these factors, particularly data availability, performance needs, and security considerations, you can determine whether fine-tuning is a worthwhile investment for your specific application.

I) AI agents

AI agents are sophisticated programs designed to perform tasks, make decisions, and interact with their environment autonomously, much like a dedicated assistant working independently. They leverage artificial intelligence (AI) technologies, including natural language processing (NLP) and machine learning, to understand, learn, and continuously improve their performance.

What is an AI Agent?

An AI agent is a system capable of perceiving its surroundings, reasoning, and acting to achieve specific objectives without constant human intervention. These agents power a range of applications, from chatbots and virtual assistants to complex automation tools that streamline processes in both personal and professional settings.



Key Features of AI Agents

- **Autonomy:** AI agents operate independently, making decisions and completing tasks without requiring direct human control.
- **Learning and Adaptation:** Many AI agents possess the ability to learn from experience, refining their performance and becoming more effective over time.

• **User Interaction:** They can often understand and respond to human language, facilitating seamless interaction and intuitive use.

The AI Agent Workflow

AI agents typically operate through a four-stage process:

- a) **Perception:** The agent gathers information from various sources, such as user input, databases, sensors, or other relevant data streams.
- b) **Decision-Making:** Utilizing AI algorithms, the agent analyzes the collected data to determine the optimal course of action.
- c) **Action:** The agent executes the chosen action, which might involve answering a question, automating a process, or performing a physical task.
- d) **Learning and Adaptation:** The agent learns from feedback and new data, continuously improving its decision-making and overall performance.

Applications of AI Agents

AI agents are transforming numerous sectors, including:

- **Customer Service:** Chatbots provide instant support and answer common customer inquiries.
- **Process Automation:** AI-powered tools automate repetitive tasks in offices, factories, and other industries, increasing efficiency and productivity.
- **Personal Assistance:** Virtual assistants like Siri and Google Assistant manage schedules, provide information, and assist with various tasks.
- **Healthcare:** AI agents can assist with diagnosis, personalized medicine, and patient monitoring.
- **Finance:** They can be used for fraud detection, risk assessment, and algorithmic trading.

Democratizing AI for All: No-code platforms are making AI accessible to businesses of all sizes, regardless of technical expertise. Companies now have the power to implement AI agents easily, unlocking new levels of automation and optimization across functions.

Prepare for the Future of Decision-Making: According to Gartner, AI agents will influence 15% of all organizational decisions by 2028. Embracing AI today ensures your business is ready for this inevitable transformation, positioning you for sustained success and adaptability.

A Strategic Roadmap to AI Agent Implementation

To successfully integrate AI agents, you'll need a clear strategy. Here's a practical guide to getting started:

- a. **Identify High-Impact Opportunities:** Focus on areas where AI agents can add the most value—whether it's automating repetitive tasks, improving decision-making, or enhancing customer experience.
- b. **Choose the Right Development Path:** Decide whether to build AI agents in-house, partner with AI vendors, or utilize existing platforms. Consider factors such as cost, technical expertise, and time-to-market.
- c. **Select the Right Tools and Frameworks:** Ensure you choose AI frameworks and tools that match your organization's technical capabilities and specific needs. Invest in scalable solutions that can grow with your business.
- d. **Invest in Quality Data:** AI agents rely on high-quality data to perform effectively. Prioritize data governance, clean data practices, and ensure your data is structured and easily accessible.
- e. **Start Small and Scale:** Begin with pilot projects to refine your approach based on real-world data. Once you have a proven model, expand your AI implementation incrementally, ensuring continuous optimization.

III. CONCLUSION

Artificial intelligence stands at the forefront of transformative innovation, offering the ability to automate tasks, derive insights from massive datasets, and even generate new content. Understanding the distinctions between conventional programming and AI, along with key terminologies and techniques—such as prompt engineering, fine-tuning, retrieval-augmented generation, and the utilization of AI agents—equips organizations to navigate the complexities of AI adoption successfully.

From addressing traditional challenges like data security and bias to embracing advanced approaches such as generative modeling and AI agents, organizations now have the tools to create robust, efficient, and ethically sound AI solutions. As AI technology and methodologies continue to evolve, staying informed and proactive will

enable businesses and institutions to harness AI's full potential, maintaining a competitive edge in a rapidly changing digital landscape.

IV. REFERENCES

- [1] Russell, S., & Norvig, P. (2010). Artificial Intelligence: A Modern Approach (3rd ed.). Pearson.
- [2] Overview: This seminal textbook provides a comprehensive introduction to AI fundamentals, including search, reasoning, learning, and the concept of intelligent agents.
- [3] Citation Example: Russell, S., & Norvig, P. (2010). Artificial Intelligence: A Modern Approach (3rd ed.). Pearson.
- [4] Domingos, P. (2015). The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. Basic Books.
- [5] Overview: This book explores the fundamental differences between rule-based traditional programming and data-driven AI approaches, making it a great resource to understand the evolution of computational methods.
- [6] Citation Example: Domingos, P. (2015). The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World. Basic Books.
- [7] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press.
- [8] Overview: This comprehensive resource covers deep learning techniques, including generative models such as GANs, which underpin many modern generative AI applications.
- [9] Citation Example: Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press
- [10] Reynolds, L., & McDonell, K. (2021). Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm. arXiv preprint arXiv:2102.07350.
- [11] Overview: This paper introduces strategies and best practices for designing effective prompts to guide large language models, a key skill in leveraging modern generative AI systems.
- [12] Citation Example: Reynolds, L., & McDonell, K. (2021). Prompt Programming for Large Language Models: Beyond the Few-Shot Paradigm. arXiv:2102.07350. Retrieved from <https://arxiv.org/abs/2102.07350>
- [13] Hugging Face Documentation. (n.d.). Fine-tuning Transformers.
- [14] Overview: Hugging Face's comprehensive documentation provides practical guidelines and tutorials on fine-tuning pre-trained language models for specific tasks, an essential aspect of adapting LLMs for real-world applications.
- [15] Citation Example: Hugging Face. (n.d.). Fine-tuning Transformers. Retrieved from <https://huggingface.co/docs/transformers/training>
- [16] Lewis, M., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv preprint arXiv:2005.11401.
- [17] Overview: This paper presents the RAG framework, which integrates retrieval systems with generative models to improve output quality on knowledge-intensive tasks, highlighting an important trend in modern NLP.
- [18] Citation Example: Lewis, M., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. arXiv:2005.11401. Retrieved from <https://arxiv.org/abs/2005.11401>
- [19] Russell, S., & Norvig, P. (2010). Artificial Intelligence: A Modern Approach (3rd ed.). Pearson.
- [20] Overview: Beyond foundational AI concepts, this textbook offers detailed insights into AI agents—their design, decision-making processes, and applications—making it a valuable reference for understanding autonomous systems.
- [21] Citation Example: Russell, S., & Norvig, P. (2010). Artificial Intelligence: A Modern Approach (3rd ed.). Pearson.
- [22] The state of AI in early 2024: Gen AI adoption spikes and starts to generate value <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>,
- [23] Superagency: The transformative potential of AI <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/ai-in-action/superagency-the-transformative-potential-of-ai>