



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 Issue: VII Month of publication: July 2025

DOI: <https://doi.org/10.22214/ijraset.2025.73426>

www.ijraset.com

Call:  08813907089

E-mail ID: ijraset@gmail.com

AI Agents and Their Role in Shaping the Future

Agrim Farkiya

Student, Department of Electronics and Telecommunication Engineering, International Institute of Information Technology, Bhubaneswar

Abstract: Artificial intelligence (AI) agents are rapidly developing reaching the level of accomplishing more and more advanced tasks and providing the professionals in a broad spectrum of the areas with conclusive assistance. The current article examines the broad implications of such agents on labour markets and organisational structures, and the society in general. Even though these systems are used to automate repetitive tasks and create impressive results in terms of productivity, widespread use of these technologies also raises many concerns about occupation displacement, security of personal information, algorithmic discrimination, and transparency. The contextual discussion highlights the need of a humanistic integration policy as defined with the aid of morally increased product designs, progressive visionary policies, and contextual initiatives that strive to ensure that AI agents transform the economy towards a level that cannot but have a positive impact on society as it develops into the limelight even further.

Keywords: AI Agents, Autonomous Systems, Large Language Models, LangChain, Agent Frameworks, AI Workflows, Automation, Intelligent Systems

I. INTRODUCTION

Artificial intelligence (AI) has made significant progress in recent years, especially with large language models (LLMs) such as GPT, Claude, and Gemini. These models were created for basic text generation and answering questions. Now, they have transformed into advanced autonomous agents capable of complex reasoning, multi-step planning, and completing tasks on their own.

AI agents using LLMs go beyond the limits of traditional chatbots. They can interact with APIs, perform web searches, run code, compile documents, and make strategic decisions based on the goals set by users. New frameworks such as AutoGPT (2023), BabyAGI, and LangChain-based systems allow agents to break down complex tasks into smaller parts, pull in external tools when needed, and produce detailed outputs beyond simple answers.

The technological advances come from improvements in transformer architecture and memory-augmented processing systems. These breakthroughs have led to extraordinary capabilities in planning, decision-making, and self-correction. This change marks a shift from reactive systems to forward-thinking intelligent agents that operate with a high degree of independence.

This evolution offers exciting opportunities in various fields, including education, healthcare, legal services, governance, and personal productivity. Industry forecasts suggest that the market for AI agents will increase from 5.25 billion in 2024 to 52.62 billion by 2030, with 60 percent of companies expected to use AI agents for 30 percent of their operations by then.

This research explores how AI agents could change social structures and economic systems by 2030. It looks at development frameworks, impacts on different sectors, and challenges in deployment. The goal is to develop a deeper understanding of what AI agents can do, evaluate potential risks, and suggest ways to maximize social benefits while reducing negative effects.

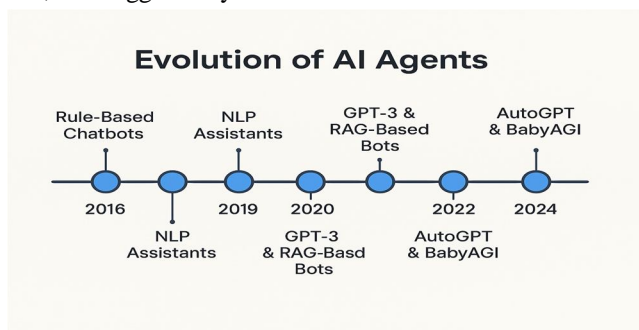


Fig. 1: Evolutionary Timeline of AI Agents

II. LITERATURE REVIEW

Autonomous AI agents have evolved from simple chat systems to advanced entities capable of planning, managing memory, utilizing tools, and self-correcting. This section reviews important developments in transformer-based models, reinforcement learning, memory enhancement, tool integration, and safety issues, focusing on groundbreaking frameworks and applications.

A. Transformer Architectures and LLM Foundations

Transformer models like GPT-4 Turbo and Claude 3 use self-attention to understand long-range dependencies. This enables a deeper grasp of language, which is crucial for complex reasoning and breaking down tasks. Retrieval-augmented generation (RAG) adds to the model's context by using external vector stores, such as Pinecone and FAISS. This allows agents to access documents or past interactions beyond the limits of in-model tokens.

B. From Chatbots to Autonomous Agents

AutoGPT was the first to create self-directed loops. It generates "Thought," "Plan," and "Critique," while using tools continuously. Benchmark studies show that AutoGPT versions with GPT-4 do better than GPT-3.5 on simulated decision tasks when enhanced by light supervised "Additional Opinions" modules. BabyAGI brought in cyclical workflows that include generating tasks, retrieving context, and executing subtasks. This showcases goal-driven behavior similar to that of virtual employees.

C. Reinforcement Learning and Self-Correction

Reinforcement learning models treat agent interactions as Markov decision processes, aiming to improve policies to maximize rewards such as task completion or user satisfaction. Critic networks offer feedback, helping agents refine actions and reduce errors. New methods combine chain-of-thought prompting with RL-style reward models to improve reliability in changing environments.

D. Memory Augmentation

In order to make up with constraints of an L-long context window, modern Artificial Intelligence solutions combine several memory models. Vector databases allow storing interactions, structured knowledge and external resources as embeddings, and therefore, it is easy to retrieve relevant knowledge quickly. Retrieval-Augmented Generation (RAG) then incorporates this stored memory to prompts thus guaranteeing consistency at long and multi-getting procedures. More complex agents also use a hierarchy of memory: working memory is present to take care of the immediate context relevant to current subtasks, episodic memory stores the records of past workflows and choices, whereas long-term knowledge has area-specific information stored in the form of knowledge graphs. All these strategies together allow agents to maintain situational awareness and access important information as well as reason with a decent degree of assurance even when confronted with complex tasks.

E. Modular Tool Ecosystems

Frameworks like LangChain simplify the use of external APIs, such as web search, code execution, and database queries, into "tools" that agents can access. Agents handle complete workflows like data scraping, transformation, and report generation. This modular approach turns LLMs into orchestrators in digital spaces, making way for specialized copilots and teamwork among agents.

F. Multi-Agent Coordination

New architectures allow for coordination between specialized sub-agents like data extractors, analytics units, and report generators. Affordance-based planners, such as AutoGPT+P, combine classical symbolic planners with LLMs to match environmental possibilities with actions. These systems achieve success rates of up to 98 percent on robotics tests with incomplete information.

G. Safety, Governance, and Ethics

Greater autonomy increases safety risks, like model poisoning and adversarial attacks, along with ethical issues such as bias, privacy violations, and lack of accountability. Research on "Safe BabyAGI" focuses on safe training environments, controlled feedback, and human oversight to reduce unexpected behaviors. Clear audit trails and regulatory guidelines are vital before deployment.

H. Applications

Automatic artificial intelligence agents make up many as- is functions in various industries, which further increase their decision-making capabilities and efficiency. As productivity copilots, they automate daily tasks by scheduling meetings, handling calendars, and synthesizing communications, causing users to not take care of tedious administrative duties. In the legal field, autonomous agents help to retrieve pertinent precedents and essential information out of a huge library of case filings and law-related documents, facilitating the process of legal research as a whole. These agents automatically review the code and report defects, and provide suggestions on code refactoring, enhancing product code quality, and reducing the time of deployment in software development. In health care, these agents can summarize the data provided by different sources and produce treatment advice or a tl;dr of the clinical picture, giving access to crucial data to the practitioners. Collectively, these applications demonstrate how autonomous AI agents are becoming an ever more important part of automating complex, multistep processes and providing professionals with reliable, context-sensitive information.

I. Future Directions

Important research areas for 2030 include learning contin- uously without full retraining, creating explainable reasoning paths for audits, improving human–AI collaboration interfaces, and establishing global standards for certifying agents and assigning liability. Cooperation between technologists, ethi- cists, and policymakers is crucial to using agent capabilities responsibly.






Name	LLM Support	Planning	Tool Integration	Use Case
 Auto-GPT	GPT	Yes	Plugins	Automating complex tasks
 BabyAGI	GPT	Yes	Plugins	TASK PRIORITISATION
 LangChain	OpenAI Anthropic	Yes	Extensive	Building versatile
 ReAct	LLM-agnostic	Yes	Tools API	Tool experimentation
 OpenAgents	GPT-3.5 GPT-4	Yes	Tools Framework	Developing agentic workflows

Fig. 2: Comparison Table of Major Frameworks

III. AI AGENT FRAMEWORKS & METHODOLOGY

Modern AI agents operate within structured frameworks that combine language models, planning modules, memory systems, and tool integration. The basic strategy is always the same, despite variations in different agent designs.

This section presents a generalized pipeline of a typical LLM-based AI agent’s operations:

A. Goal Interpretation

The AI agent first translates the user’s high-level request, such as “Schedule a trip to Goa” or “Summarize this 20- page legal contract” into an internal representation suitable for planning. The language model analyzes the prompt’s semantics, identifies intent, constraints, and desired outputs, and then formulates a series of preliminary subtasks. For “Schedule a trip,” it may extract destination, dates, budget, and accommodation preferences; for “Summarize,” it determines document structure, key sections, and length requirements. This interpretation phase ensures that subsequent planning, tool selection, and execution steps align precisely with the user’s goals, laying the foundation for accurate and efficient task completion.

B. Task Planning

Using sophisticated reasoning techniques, the agent breaks down a high-level objective into a structured series of action- able steps. Initially, it produces a series of ideas, delineating rational steps in a straight line. It uses a Tree of Thoughts to explore several possible solution paths, prune low-value branches, and go back as necessary for more complex goals. The ReAct paradigm combines distinct tool-invocation actions (“Action: . . .”), like memory lookups, code execution, web searches, and API calls, with explicit reasoning (“Thought:. . .”). The agent iterates the reasoning or action loop until all subtasks are successfully finished, processes the observation after each action, and updates its plan.

C. Tool Selection & API Calling

Once a step requires external action (like searching the web or accessing a document), the agent calls tools or APIs like:

- Web Browsing
- File Reading
- Python Code Execution
- Databases / Vector Stores

D. Memory Retrieval

In order to maintain consistency throughout lengthy work-flows, contextual memory allows an AI agent to remember and make use of data from previous interactions, retrieved documents, and earlier tool executions. When processing new tasks, the agent can retrieve pertinent context by storing embeddings or structured records of previous conversations, choices, and intermediate outcomes in external memory stores (such as vector databases or knowledge graphs). Because the agent uses prior findings to prevent duplication, preserve consistency, and expand on earlier insights, this long-term coherence facilitates advanced decision-making. Reliable, customized, and context-aware agent behaviors are supported by efficient memory management, which distinguishes between temporary buffers and permanent repositories.

E. Looping and Correction

The iterative feedback loop that modern AI agents work in is similar to how humans self-correct. The agent compares the intermediate result, whether it be a tool response, reasoning "Thought," or API output, to the initial objective and predetermined quality standards. A critique module initiates revision when disparities, mistakes, or less-than-ideal results are found: the agent modifies its strategy, improves prompts, or calls tools again. The agent locks in the outcome and moves on to the next subtask once the output satisfies the requirements. This ongoing cycle of assess, revise, and proceed improves dependability, lowers hallucinations, and guarantees logical multi-step execution across complicated workflows.

F. Final Output Generation

The AI agent combines intermediate outputs, such as reasoning chains, tool results, and memory insights, into a cohesive final product once all subtasks have been completed. This could be a draft email, a filled-out form, a data visualization (such as a table or chart), or a comprehensive report, depending on the user's goal. To ensure smooth handoff and user satisfaction, the agent formats content to meet predetermined requirements (style, length, and structure), applies any necessary post-processing (such as proofreading or data validation), and packages the final product for delivery through the proper channel, email, dashboard, API endpoint, or file export.

G. Use Cases

In order to provide nearly human reasoning and autonomy, a variety of domain-specific AI agents are supported by the generalized methodology previously discussed. These agents all use the same core pipeline, which consists of goal interpretation, planning, tool invocation, memory management, and iterative self-correction.

The agent in AI-powered personal assistants (e.g., ChatGPT with web browsing) interprets conversational prompts, plans multi-step workflows (e.g., researching flight options), retrieves current information using a text-based or visual browser plugin, stores user preferences in vector memory, and iterates until a comprehensive itinerary or synthesized summary is generated. While the agent independently explores the web and incorporates findings, users enjoy smooth back-and-forth communication.

Financial report summarizers use file-reading tools to read large documents (PDF, XLSX), use prompt-driven decomposition to extract important metrics, and run Python routines to calculate ratios or format tabular data. Agents use vector databases to retrieve industry standards, evaluate generated summaries based on predetermined accuracy thresholds, and produce succinct executive summaries or slide decks.

Academic research copilots use Copilot-style scholarly search modules to build search strategies, parse research questions, retrieve articles from online repositories, and synthesize literature reviews. They can track citations and conduct coherent, multi-paper analyses because they preserve episodic memory of previous queries. Summaries are improved by the self-reflection loop to conform to formatting requirements and publication standards.

Agents accept high-level objectives (e.g., “Draft a nondis- closure agreement”), break down tasks into clause-level sub- tasks, use contract-law databases and document-assembly APIs, and iterate over critique modules to ensure compliance and reduce ambiguity when performing legal analysis and document preparation. Consistent drafting across engagements is supported by a persistent memory of client preferences and jurisdictional regulations. Lastly, completely formatted legal briefs, redlined contracts, or brief explanatory memos prepared for attorney review are examples of final outputs.

In all of these areas, the unified pipeline makes it possible for AI agents to operate as dependable “virtual employees,” combining memory, tool usage, and reasoning to produce results of a high caliber for both individual users and business clients.

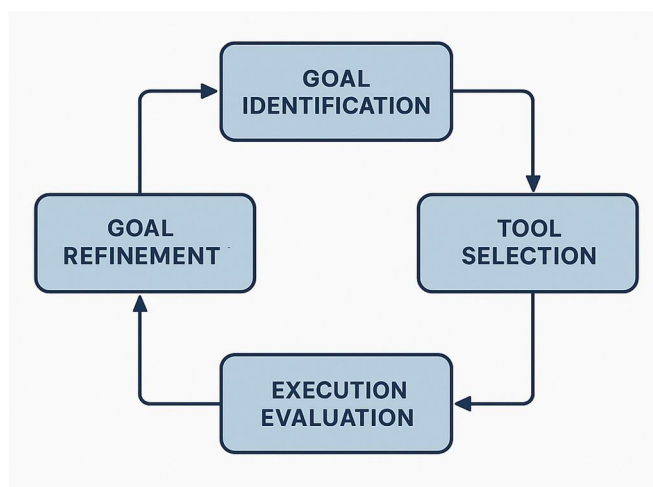


Fig. 3: Flowchart of Methodology Loop

IV. PREDICTED IMPACT ON FUTURE SECTORS

By 2030, the quick development of autonomous AI agents has the potential to completely transform a number of in- dustries. Large Language Models (LLMs) are driving these intelligent systems, which are already being used in a variety of fields, including business, governance, healthcare, and edu- cation. Their anticipated impact on these sectors is examined in this section. AI agents are evolving into individualized digital tutors in education that can adjust to a student’s performance, learning preferences, and pace. Real-time progress tracking, assignment customization, research support, and even natural language feedback evaluation of written work are all possible features of future systems. The result is a personalized, highly scal- able educational system that lessens reliance on conventional resources.

AI agents are being incorporated into the healthcare industry to help with mental health support, patient monitoring, and diagnostics. By compiling medical records, they can reduce medical professionals’ workload by deciphering sensor data from wearable technology and providing round-the-clock con- versational support. These agents are anticipated to support human practitioners rather than replace them, particularly in telemedicine and primary healthcare, even though ethical and regulatory issues still pose difficulties.

AI agents are simplifying daily tasks in corporate settings. They are used to create meeting agendas, draft emails, ana- lyze customer data, and provide real-time insights to support decision-making. They are already being used by a large number of startups and businesses for market analysis, market- ing automation, and content creation. This leads to increased operational efficiency and productivity, particularly in settings with limited resources.

The legal and governance sectors are also changing. The workload for legal teams can be lessened by using legal agents to review contracts, evaluate case law, and help draft formal documents. Artificial intelligence (AI) agents can be used in government to improve responsiveness and accountability by simulating policy outcomes, assessing public sentiment, and producing real-time analytics from citizen feedback.

Economically speaking, the widespread use of AI agents could result in the automation of standard white-collar tasks like customer service, data entry, and basic accounting. Never- theless, it is anticipated that this disruption will also lead to the creation of new jobs in fields like system integration, prompt engineering, AI ethics compliance, and agent supervision, which will cause employment to shift rather than decline.

AI agents still raise big moral and social problems, even with these changes. To deal with privacy issues, data misuse, and possible bias in autonomous decision-making, we need strong policy frameworks. The digital divide could also get worse if only certain groups of people can use advanced systems. So, for growth to be fair and long-lasting, it is important that agents are deployed responsibly, that there is regulatory oversight, and that their operations are open to the public.

V. LIMITATIONS AND FUTURE SCOPE

This review is a critical theoretical perspective on AI agents based on large language models (LLMs) and evaluates their transformative capacity in various industries; however, a few limitations should be mentioned. In particular, the argument is confined to tertiary sources, i.e., peer-reviewed papers, paper-substitute blogs, and industry white papers, so that the analysis of agent frameworks cannot rely on primary experimentation or direct comparison. This way, measurements of agent performance, reliability, and scalability are based on reported results, instead of proper empirical tests, including limiting the generality of the review and narrowing applicability to real-world environments of unpredictable requirements, heterogeneous information, and system complexity.

Another limitation arises due to the highly dynamic nature of the agent ecosystem; important frameworks like AutoGPT, LangChain, ReAct, and BabyAGI are under active development and undergo frequent changes. Even their APIs, toolchains, and performance characteristics may change significantly in the course of a relatively short period, making the data provided in this review obsolete very quickly after publication, which is a common issue with modern artificial intelligence work.

As much as the said issues are saliently discussed therein, the review however, fails to perform a thorough investigation on the possible legal, ethical and socio-technical implications. Dilemmas of liability in independent decision-making, transparency in rationale, human checks, the potential of ill intentioned use and systemic biases all require nuanced, cross disciplinary scrutiny to support efficient regulation and use policy.

To study the problem further, a number of research vectors are of special importance. Above all, there is empirical benchmarking that is carried out within a wide array of real-life situations. Formal testing of performance based on accuracy, scalability, robustness, and user experience would provide a more critical, comparative understanding of the strengths and weaknesses of emerging agent platforms. Just as significant is a direct analysis of agent orchestration over heterogeneous toolchains and multi-agent collaboration settings and which may reveal empirical bottlenecks and optimisation strategies. Second, the development of explainable, auditable, and regulation-friendly agent schemas should also be enhanced. The existential gap between technology innovation and policy, especially on issues relating to safety, accountability, and fair results, is what matters in transforming hands-off yet accountable and sustainable integration of AI agents in education, healthcare, and governance, among other applications of critical needs. It will take an interdisciplinary, concerted empirical approach to safely work through such intricate challenges and harness all the positive potential of next-generation AI agents.

VI. CONCLUSION

The rapid spread of AI agents, popularized by Large Language Models, is an important change in the direction that the field of computer augmentation of human abilities is taking. In 2030, it is expected that autonomous workers will be ubiquitous, whose task will be to act as continuous virtual employees to detect complex and multi-step work processes, data-driven insights, and to provide personalized services to a huge number of individuals. With the ability to break down large-scale goals into a series of small steps, they are able to bring targeted instruments whenever the situation requires, and they maintain situational understanding at all times. Therefore, they can perform up to the kind of tasks that were only performed by human practitioners, such as the writing of legal briefs and medical stories, supply chains, and tailoring education curricula.

In a school, these agents will serve as adaptive tutors and administrative assistants, personalizing learning paths and automating grading processes and the process of distributing resources. They will also diagnose conditions, supervise the management of the patients, and partake in the administrative tasks, thus relieving the clinicians of the practice and increasing continuity. In the governmental areas, AI agents have the potential to improve policy analyses, control the regulatory compliance, and develop citizen engagement platforms, but there will be a need to have a vigorous observant role in order to ensure the inclusion of fairness and transparency. Most business processes, including CRM, finance, and logistics, will be running across companies: all these processes will be administered by the agents to increase efficiency, rapidity, and innovation.

Agents, as they move toward an infrastructural system, will require strong ethical and functional governance. Hallucinations and biased results jeopardize confidence and fairness, whereas data manipulation and autonomous decision-making evoke some serious concerns with responsibility.

The prospect of displacement of the workforce, in turn, compounds the argument that employees should be reskilled, and new jobs in AI governance created, as well as social safety nets strengthened. Well-built safeguards like retrieval-augmented verification loops, trails of explainable reasoning, and fail-safe procedures would have to go along with more exhaustive regulatory systems whose directives ensure due ethical usage, auditing, and rigorous data protection.

In order to avoid the expansion of current digital gaps and to guarantee more even distribution of benefits, coming studies should focus on the improvement of interpretability via transparent reasoning architectures, democratization of access to new technologies, and increasing agent reliability via adversarial resilience and lifelong learning. Best practices and policies that achieve a balance between innovation and the welfare of society will have to be designed through interdisciplinary collaboration amongst the technologists, ethicists, legislators, and end-users.

The destiny of AI agents will eventually depend on the ability of society to instill human values in the development, use, and monitoring of AI. Ethically and fairly nourished, these systems will deliver potential to reform human-machine teamwork to a then-never-before championship degree of productivity, imagination, and issue-solving energies, driven by AI agents and driving the following (important) step in the history of the digital age.

REFERENCES

- [1] W. Wei, M. Zhang, and C. Zhao, "Payload ChatGPT: LLM-Powered AI Agents in the Real World," *IEEE Access*, vol. 11, pp. 189–200, 2023.
- [2] J. Lee and T. Kwon, "Autonomous LLM Agents for Workflow Automation," in *Proc. AAAI Conf. Artificial Intelligence*, 2024, pp. 122–130.
- [3] J. Goldstein, "Emergent Capabilities in LLM Agents," *ACM Computing Surveys*, vol. 56, no. 1, pp. 1–25, 2025.
- [4] OpenAI, "Superalignment: Ensuring LLM Agent Safety," 2023. [Online]. Available: <https://openai.com/research/superalignment>
- [5] Anthropic, "Claude 3 and Beyond: Safe LLM Agents," 2024. [Online].
- [6] Available: <https://www.anthropic.com/research>
- [7] S. Mehta, "Prompt Chaining for Adaptive AI Agents," *arXiv preprint arXiv:2403.09765*, 2024.
- [8] NVIDIA, "Inference Optimization for LLM Agents on Edge Devices," 2023. [Online]. Available: <https://developer.nvidia.com>
- [9] A. Banerjee and M. Rahman, "Reinforcement Learning in LLM-Based Agents," *Int. J. Neural Syst.*, vol. 35, no. 2, pp. 201–215, 2024.
- [10] A. Zhang, "Next-Gen LLM Agents: Trends and Challenges," *AI Magazine*, vol. 45, no. 1, pp. 66–77, 2025.
- [11] T. Brown et al., "AutoGPT Framework for Autonomous AI Agents,"
- [12] *NeurIPS Workshops*, 2023.
- [13] A. Patel, "GPT Engineer: A Framework for Code-Generating Agents,"
- [14] *arXiv preprint arXiv:2401.01234*, 2024.
- [15] Microsoft Research, "Toolformer: LLMs with Built-in Tool Use," 2023. [Online]. Available: <https://www.microsoft.com/en-us/research/>
- [16] H. Liu, "BabyAGI: Building Memory-Driven Agents Using LLMs,"
- [17] *arXiv preprint arXiv:2404.00981*, 2024.
- [18] LangChain, "LangChain Framework Documentation," 2023. [Online].
- [19] Available: <https://docs.langchain.com>
- [20] S. Yao et al., "ReAct: Reasoning and Acting with Language Models," in *ICLR*, 2023.
- [21] T. Shen, M. Du, and K. Zhang, "HuggingGPT: Large Language Models Meet Machine Learning Services," *arXiv preprint arXiv:2303.17580*, 2023.
- [22] B. Choi and Y. Kim, "AutoGen Framework: Multi-Agent Coordination via LLMs," *arXiv preprint arXiv:2402.07892*, 2024.
- [23] X. Liu, "Ethical Challenges in Deploying Autonomous AI Agents," *AI & Society*, vol. 40, no. 2, pp. 199–212, 2025.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)