



e-ISSN: 2278-8875
p-ISSN: 2320-3765



International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

Volume 9, Issue 9, September 2020



ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA

Impact Factor: 7.122

9940 572 462

6381 907 438

ijareeie@gmail.com

www.ijareeie.com



A Conceptual Framework for Understanding Machine Learning and Artificial Intelligence Roles in Technological Advancements

N Rajasekhar Reddy¹, Nanda Kumar A N², Nandish A C³, Vishvakiran R C⁴, Gopikishan J⁵

Professors, Department of Computer Science and Engineering, City Engineering College, Bengaluru,
Karnataka, India^{1,2}

Assistant Professors, Department of Electronics and Communication Engineering, City Engineering College,
Bengaluru, Karnataka, India^{3,4,5}

ABSTRACT: In recent years, the terms "machine learning" and "artificial intelligence" have become increasingly prevalent across both scientific literature and media outlets, often used interchangeably despite their distinct meanings. This study aims to clarify the relationship between these two concepts, with a particular focus on the role machine learning plays in the development of artificial intelligence. By reviewing relevant literature, we identify key distinctions between the two fields and explore how machine learning serves as a subset of artificial intelligence, providing the foundational algorithms that enable intelligent decision-making and autonomous behaviour in artificial agents. Additionally, we introduce a conceptual framework to provide a clearer understanding of how machine learning techniques, such as supervised learning, unsupervised learning, and reinforcement learning, contribute to the broader goals of artificial intelligence. This framework is intended to serve as a foundation for interdisciplinary discussions and guide future research by offering precise definitions and a more structured understanding of the intersection between these fields. Ultimately, this work seeks to foster a better comprehension of the roles that machine learning and artificial intelligence play in modern technological advancements and how they can be distinguished in both academic and practical applications.

KEYWORDS: Machine Learning, Artificial Intelligence, Conceptual Framework, Interdisciplinary Research

I. INTRODUCTION

In his April 2018 US Senate hearing, Mark Zuckerberg emphasized the need for Facebook's "AI tools" to effectively identify hate speech and terrorist propaganda. Typically, such tasks are categorized as classification tasks within the realm of (supervised) machine learning. However, with the growing popularity of artificial intelligence (AI), the term AI is frequently used interchangeably with machine learning. This usage is not only seen in Zuckerberg's statements and interviews but also across various theoretical and application-oriented contributions in recent literature. Carner (2017) even acknowledges using AI as a synonym for machine learning, despite knowing this is not entirely accurate. This ambiguity can lead to significant imprecision in both research and practice when discussing methods, concepts, and results. It is surprising that despite the frequent use of these terms, there is a lack of clear scientific delineation. This paper aims to clarify the relationship between machine learning and artificial intelligence by examining the role of machine learning in the context of intelligent agents. We approach this by focusing on the machine learning perspective of intelligent agent capabilities and their implementation.

Our contribution is threefold. First, we build on the theoretical framework provided by Russel & Norvig (2015) by refining the "thinking" layer of intelligent agents into distinct "learning" and "executing" sublayers. Second, we demonstrate how this distinction allows us to better understand the various contributions of machine learning to intelligent agents. Third, we use the implementation of these sublayers ("backend") to define a continuum between human involvement and agent autonomy. The paper proceeds by reviewing relevant literature on machine learning and artificial intelligence, presenting and elaborating on our conceptual framework that highlights the contribution of machine learning to AI. We then outline an agenda for future research and conclude with a summary, current limitations, and future outlook.



II. RELATED WORK

To establish the foundation for our conceptual work, we first review the various notions, concepts, and definitions of machine learning and artificial intelligence present in existing research. Additionally, we provide a detailed discussion of the theories that underpin our framework.

2.1 Terminology

Machine learning, artificial intelligence, data mining, deep learning, and statistical learning are related terms that frequently appear in similar contexts and are sometimes used interchangeably. Although these terms are commonly used across various communities, their specific usage and meanings can vary significantly.

In the field of statistics, the focus is on statistical learning, which involves methods and algorithms designed to gain knowledge, predict outcomes, and make decisions by constructing models from data. From this perspective, machine learning can be seen as an implementation of statistical learning. In computer science, machine learning centres on designing efficient algorithms to address problems using computational resources. Although machine learning incorporates statistical approaches, it also includes methods that are not entirely derived from traditional statistical work, leading to new and influential contributions to the field. In recent years, deep learning, a subset of machine learning, has garnered significant attention. Deep learning models consist of multiple processing layers that learn data representations with various levels of abstraction, greatly enhancing machine learning capabilities, particularly in speech and image recognition.

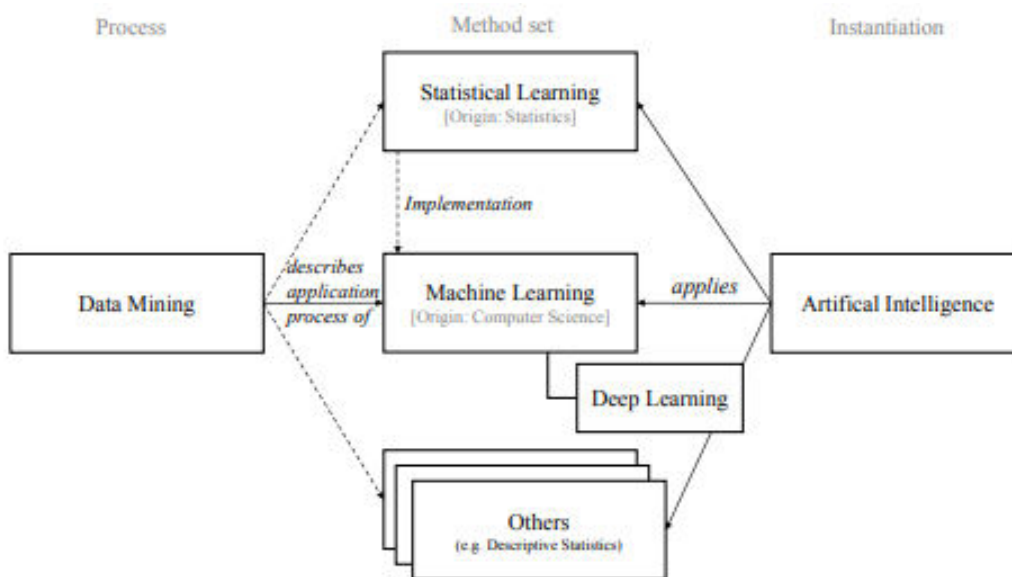


Figure 1. General Terminology

Distinct from these terms, data mining refers to the application of quantitative analytical methods to solve practical problems, such as those encountered in business settings. In the context of machine learning, data mining involves generating meaningful machine learning models with the aim of gaining insights from data, rather than advancing knowledge about the algorithms themselves. Thus, machine learning serves as a foundational element for data mining. Artificial intelligence, on the other hand, employs techniques like machine learning, statistical learning, and other methods such as descriptive statistics to simulate intelligence in machines. Figure 1 and the terms defined in this paragraph provide the basis for the rest of this work. However, the terminology and conceptual relationships remain a topic of debate [22]. This paper aims to provide clarity on these terms, with a specific focus on elucidating the role of machine learning within AI. To achieve a deeper understanding, we examine both machine learning and AI in greater detail.



2.2 Machine Learning

Machine learning encompasses a range of techniques designed to address various real-world problems using computer systems that can learn to tackle these issues without being explicitly programmed [23]. Generally, machine learning is divided into unsupervised and supervised categories. This work concentrates on supervised machine learning, as the most commonly used methods fall into this category. In supervised machine learning, the system learns by using a series of examples ("past experience") to develop knowledge about a specific task. Although statistical methods are employed during the learning process, manual adjustments or programming of rules or strategies to solve the problem are not necessary. Specifically, supervised machine learning techniques aim to construct a model by applying an algorithm to a set of known data points in order to make inferences about unknown data.

The creation of a machine learning model generally involves three main phases: model initiation, performance estimation, and deployment. During model initiation, a human user defines the problem, prepares and processes the dataset, and selects an appropriate machine learning algorithm for the task. In the performance estimation phase, various parameter configurations of the algorithm are tested and validated, and the best-performing setup is chosen based on its effectiveness in solving the task. Finally, the model is deployed and used to tackle the task on new, unseen data.

Learning, in a broader sense, is a fundamental aspect of human cognition that involves the processes by which sensory input is transformed, reduced, elaborated, stored, retrieved, and utilized [28, p. 4]. Humans process a vast amount of information using abstract knowledge to better understand new input. Machine learning models, due to their adaptive nature, can emulate certain cognitive abilities of humans in a limited context. However, machine learning is fundamentally a collection of methods for detecting patterns in existing data, resulting in analytical models that can be integrated into larger IT systems.

2.3 Artificial Intelligence

The topic of artificial intelligence (AI) spans various research disciplines, including computer science, philosophy, and futures studies. In this work, we concentrate primarily on computer science because it is most pertinent for identifying the role of machine learning within AI and for distinguishing between the two terms.

AI research can be categorized into different streams. These streams vary based on the goals of AI applications (thinking vs. acting) and the type of decision-making involved (human-like decisions vs. ideal, rational decisions). This distinction results in four main research currents, which are outlined in Table 1. The "Cognitive Modelling" stream (i.e., thinking humanly) posits that AI should be a machine capable of human-like cognition. This approach involves not only achieving the same output as a human when presented with the same input but also replicating the reasoning processes that lead to that conclusion. In this context, the terms rational and intelligent are often used interchangeably in related work. The "Laws of Thought" stream (i.e., thinking rationally) demands that an AI make decisions that are rational, regardless of what a human might conclude.

Table 1. AI Research Streams

Objective Application to	Humanly	Rationally
Thinking	Cognitive Modeling	"Laws of thought "
Acting	Turing Test	Rational Agent



Therefore, an AI must adhere to the laws of thought by utilizing computational models that reflect logical principles. The "Turing Test" stream (i.e., acting humanly) suggests that an AI should demonstrate intelligent behaviour during interactions with humans. To meet this criterion, an AI must perform tasks as well as or better than humans. The Turing Test can be used to assess these capabilities.

The "Rational Agent" stream views AI as a rational or intelligent agent that not only operates autonomously but also aims to achieve the most rationally ideal outcomes. An alternative approach to defining AI involves establishing a general definition of intelligence and applying these insights to create intelligent machines. Legg and Hutter use intelligence tests, theories of human intelligence, and psychological definitions to develop a measurement of intelligence. Their agent-environment framework, which describes general intelligence and artificial intelligence when the agent is a machine, closely resembles the "acting rationally" stream. In addition to defining AI, classifying AI is another important area of research. Searle proposes a distinction between weak AI, which only simulates thought, and strong AI, which possesses actual mental states. Gubrud, on the other hand, categorizes AI based on the type of task it performs. According to Gubrud, artificial general intelligence (AGI) is an AI that operates at least at the level of a human brain across various domains, though it does not require consciousness.

In contrast, narrow AI refers to systems that excel in specific, limited tasks but do not rival the overall capabilities of the human brain. In the following, we will explore the "Rational Agent" stream in greater detail, as it is crucial for understanding the implementation of machine learning within AI. We will revisit the other three research streams, where we demonstrate their compatibility with our framework for agent-based AI.

According to the "Rational Agent" stream, intelligence is manifested through the actions of agents. These agents are characterized by five key features: they "operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals". An agent's actions are defined not for itself but in interaction with its environment. It perceives the environment through sensors, uses an agent program to process input data, and performs actions through actuators. To be considered a rational agent, the agent must act to achieve the highest expected outcome based on its performance measure, drawing from current and past knowledge of the environment and potential actions.

Russell & Norvig categorize agents into four types based on their programs. A simple reflex agent reacts solely to sensor data, while a model-based reflex agent also considers its internal state. A goal-based agent makes decisions aimed at achieving specific goals, which are evaluated as either fulfilled or not. In contrast, a utility-based agent seeks to maximize a utility function rather than achieving binary goals. By extending its program, an agent can become a learning agent, which includes a performance element for action selection based on sensor data and a learning element that receives feedback from the environment, generates its own problems, and improves the performance element.

The agent-environment framework consists of three components: the agent, the environment, and the goal. Intelligence is defined as the "agent's ability to achieve goals in a wide range of environments". The agent receives input from perceptions generated by the environment, which can be observations or reward signals indicating how well the agent's goals are met. Based on these inputs, the agent performs actions, which are then communicated back to the environment as signals.

2.4 A framework for understanding the role of machine learning in artificial intelligence

To understand the interplay between machine learning and AI, we build our concept on the framework provided by Russell & Norvig. Their distinction between the two objectives of AI acting and thinking provides a crucial foundation for this discussion.

2.5 Layers of agents

To understand the role of machine learning within AI, we need to adopt a perspective focused on the implementation of intelligent agents. This approach helps us map the various tasks and components of machine learning to the capabilities of intelligent agents. By considering the thinking and acting capabilities of an intelligent agent in terms of software design, we can view the acting capabilities as the frontend and the thinking capabilities as the backend.

Software engineers often separate form from function to enhance flexibility and independence and to allow for parallel development. The frontend serves as the interface with the environment and can take many forms, such as a machine-readable web interface, a human-readable application, or even a humanoid with advanced expression capabilities.

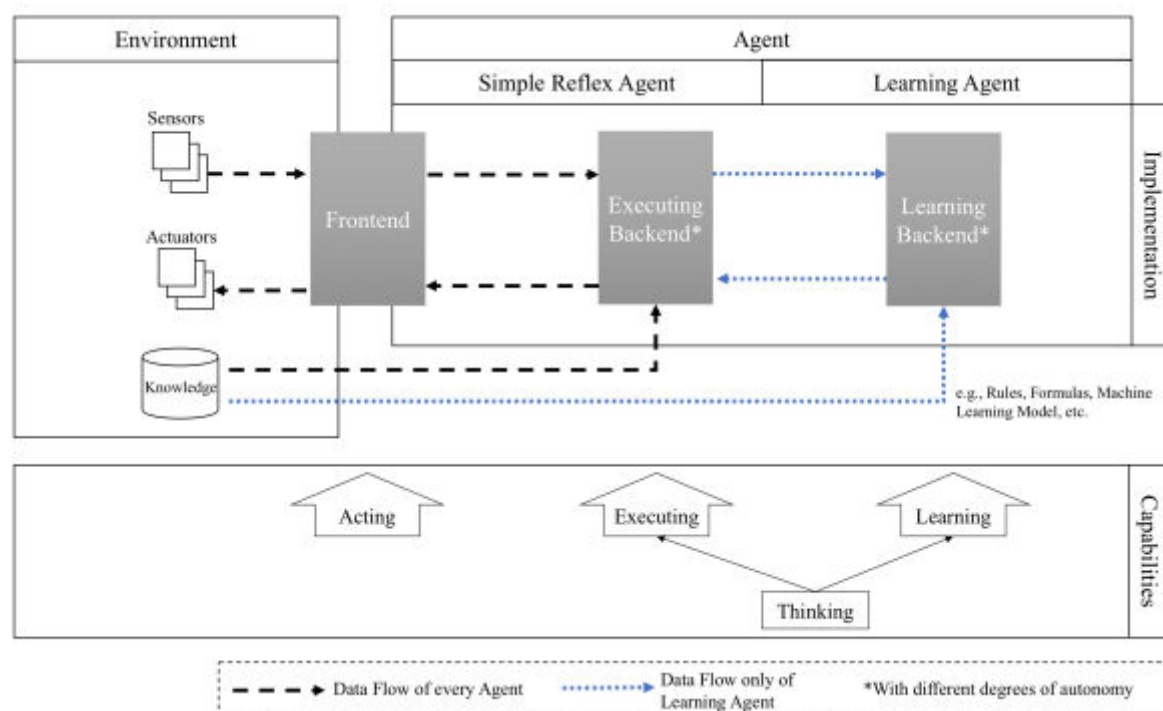


Figure 2. Conceptual Framework

For the frontend to interact with the environment, it requires two technical components: sensors and actuators. Sensors detect events or changes in the environment and relay this information through the frontend to the backend. For example, sensors might measure the temperature of an industrial machine or capture visuals of human interactions. Actuators, on the other hand, are responsible for moving and controlling mechanisms. While sensors process information, actuators perform actions, such as automatically buying stocks or altering the facial expressions of a humanoid. The Turing test can be seen as evaluating the interaction between the environment and the frontend, specifically the combination of sensors and actuators, when assessing an agent's ability to act humanly. Although the specific appearance of the frontend is not crucial to our work, it is important to note that a backend-independent, encapsulated frontend exists.

The backend provides the essential functionalities that represent the thinking capabilities of an intelligent agent. Consequently, the agent must be able to learn and apply this learned knowledge. This makes machine learning crucial for this implementation layer. In the context of supervised machine learning, it is important to distinguish between two tasks: building (training) machine learning models and executing the deployed models. To better understand the role of machine learning within intelligent agents, we refine the thinking layer into two sublayers: the learning sublayer (model building) and the executing sublayer (model execution). Thus, the implementation required for the learning sublayer is referred to as the learning backend, while the executing sublayer is denoted as the executing backend.

III. TYPES OF LEARNING

The learning backend determines, first, whether the intelligent agent is capable of learning and, second, how the agent learns. This includes specifying the algorithms used, the data processing methods applied, and how concepts are developed. Russell & Norvig similarly distinguish between learning elements and performance elements to describe this relationship. The learning backend determines not only if the intelligent agent can learn but also how it learns, including the choice of algorithms, data processing methods, and handling of concept drift. We adopt Russell & Norvig's terminology by distinguishing between two types of intelligent agents: simple-reflex agents and learning agents. This distinction is particularly relevant from a machine learning perspective, as it addresses whether the models in the thinking layer are trained once and left unchanged (simple-reflex) or continuously updated and adaptive (learning).



Recent literature provides examples for both types. For instance, Oroszi and Ruhland developed an early warning system for pneumonia in hospitals. While their model showed promising results during development and testing, the system's ability to adapt after deployment could be problematic. Other examples of single-trained models include applications in anaphora resolution, pedestrian prediction, and object annotation. Conversely, examples of learning agents are also present in recent studies. Mitchell et al. introduce the concept of "never-ending learning" agents, which focus on continuously updating and building models. An example of such an agent is demonstrated by Liebman et al., who created a self-learning agent for music playlist recommendations. Other examples of learning agents include the regulation of heat pump thermostats, agents designed to acquire collective knowledge across various tasks, and agents that learn word meanings. The choice between a simple-reflex agent and a learning agent impacts both the overall design of the agent and the role of machine learning. Our framework, illustrated in Figure 2, summarizes this distinction. For a simple-reflex agent, machine learning is applied as a one-time trained model in the execution sublayer. In contrast, for a learning agent, machine learning is integral to the learning sublayer, where it continuously refines the model in the execution sublayer. This ongoing improvement relies on knowledge and feedback obtained from the environment through the execution layer.

IV. CONTINUUM BETWEEN HUMAN INVOLVEMENT AND MACHINE INVOLVEMENT

When considering the executing backend and the learning backend, it is crucial not only to assess whether and how machine learning models are updated but also to evaluate the level of automation in the processes involved. Every machine learning task encompasses various steps, such as selecting data sources, collecting data, pre-processing, building models, evaluating, deploying, executing, and improving. Although a detailed discussion of each step is beyond the scope of this paper, the degree of autonomy and automation in these tasks, as implemented within the agent, is particularly important for the machine learning lifecycle.

For example, while automating the execution of a pre-built model is relatively straightforward, automating tasks such as identifying an appropriate data source for a new problem, retraining, or self-initiating model building is more complex. Consequently, we must consider human involvement in the machine learning tasks of an intelligent agent, as illustrated in Figure 3. Rather than a clear-cut distinction, human involvement in these tasks exists on a continuum. This range spans from minimal agent autonomy with full human involvement to complete agent autonomy with minimal human intervention.

For instance, an intelligent agent designed to drive a car autonomously and interpret traffic signs demonstrates significant agent autonomy. However, if the agent encounters a new type of traffic sign, it may still require human input to learn about this new situation, as it might not be able to fully adapt on its own [71]. Therefore, understanding the level of human involvement, especially in the thinking layer (i.e., the executing and learning backends), is crucial for describing AI and its underlying machine learning models. Analyzing the degree of autonomy for each machine learning step can provide insights into the agent's overall autonomy concerning its machine learning tasks.

V. RESEARCH PRIORITIES FOR MACHINE LEARNING-ENABLED ARTIFICIAL INTELLIGENCE

The framework presented here on machine learning and its role within intelligent agents is still conceptual. However, given the existing misunderstandings and ambiguities surrounding these terms, there is significant potential for further research to clarify terminology and explore new areas for machine-learning-enabled artificial intelligence. Firstly, empirical validation and ongoing iterative development of the framework are essential. We need to examine various intelligent agent cases across different fields to assess how well the framework applies. It would be valuable to see how both practical and academic projects involving machine-learning-enabled artificial intelligence align with the framework. Additionally, quantifying the proportion of projects that use learning agents versus non-learning agents would provide insights into the necessary human involvement in advanced intelligent agents and help determine their overall autonomy in terms of acting, executing, and learning.

Secondly, reducing human involvement is an important area of interest. As previously noted, this spectrum ranges from full human involvement to complete agent autonomy. One promising approach is transfer learning, which focuses on transferring knowledge (i.e., models) from one environment to another. This could potentially minimize human involvement by automating the adaptation of models to new or modified tasks. Furthermore, addressing changes in the environment with already deployed models is crucial. The subfield of concept drift explores how to detect and adapt to




changes in the environment. Although there are many possibilities, successful applications remain limited. Continued research in these areas could advance the automation and adaptability of intelligent agents.

VI. CONCLUSION

In this paper, we clarify the role of machine learning within artificial intelligence, with a focus on intelligent agents. We present a framework that distinguishes between simple-reflex and learning agents and outlines the role of machine learning in each case. Essentially, machine learning models can be implemented as static, once-trained models within an intelligent agent without the capacity to acquire additional insights from the environment (simple-reflex agent). In this context, we refer to this aspect as the executing backend. Here, the agent can use previously built machine learning models but cannot modify or update them. Conversely, if the agent can learn from its environment and update its machine learning models accordingly, it is classified as a learning agent. Learning agents include an additional layer, the learning backend, which supports model building and training. The implementation of these two sublayers must consider the degree of autonomy required for the machine learning processes within the agent. This involves examining the level of human involvement in tasks such as data collection and algorithm selection. Our research is conceptual and has limitations. Empirical studies are needed to assess how well existing machine-learning-enabled AI applications fit into this framework. Expert interviews with AI designers could validate and refine the model, providing more detail. Additionally, quantifying human involvement in machine-learning tasks is crucial for understanding the autonomy of modern agents. Although preliminary, our framework aims to help scientists and practitioners use more precise terminology when discussing machine learning and AI. It emphasizes the importance of clearly defining the role of machine learning in specific agent implementations rather than using the terms interchangeably.

REFERENCES

1. Russell, S., & Norvig, P. (2016). *Artificial Intelligence: A Modern Approach* (3rd ed.). Pearson Education.
2. Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press.
3. Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science*, 349(6245), 255-260.
4. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436-444.
5. Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Springer.
6. Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529-533.
7. Samuel, A. L. (1959). Some studies in machine learning using the game of checkers. *IBM Journal of Research and Development*, 3(3), 210-229.
8. Hutter, M. (2005). *Universal Artificial Intelligence: Sequential Decisions Based on Algorithmic Probability*. Springer.
9. Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61, 85-117.
10. Domingos, P. (2015). *The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World*. Basic Books.
11. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction* (2nd ed.). MIT Press.
12. Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798-1828.



INNO  **SPACE**
SJIF Scientific Journal Impact Factor

Impact Factor:
7.122

ISSN INTERNATIONAL
STANDARD
SERIAL
NUMBER
INDIA



International Journal of Advanced Research

in Electrical, Electronics and Instrumentation Engineering

 **9940 572 462**  **6381 907 438**  **ijareeie@gmail.com**



www.ijareeie.com

Scan to save the contact details