computers

MDPI

Article

Exploring the Potential of the Bicameral Mind Theory in Reinforcement Learning Algorithms

Munavvarkhon Mukhitdinova [1] and Mariana Petrova [2,*]

[1] Tashkent State University of Economics, Islom Karimov 49, Tashkent 100066, Uzbekistan; munavvarkhon7@gmail.com
[2] Department of Information Technologies, St. Cyril and St. Methodius University of Veliko Tarnovo, 5003 Veliko Tarnovo, Bulgaria
* Correspondence: m.petrova@ts.uni-vt.bg

Abstract: This study explores the potential of Julian Jaynes' bicameral mind theory in enhancing reinforcement learning (RL) algorithms and large language models (LLMs) for artificial intelligence (AI) systems. By drawing parallels between the dual-process structure of the bicameral mind, the observation–action cycle in RL, and the "thinking"/"writing" processes in LLMs, we hypothesize that incorporating principles from this theory could lead to more efficient and adaptive AI. Empirical evidence from OpenAI's CoinRun and RainMazes models, together with analysis of Claude, Gemini, and ChatGPT functioning, supports our hypothesis, demonstrating the universality of the dual-component structure across different types of AI systems. We propose a conceptual model for integrating bicameral mind principles into AI architectures capable of guiding the development of systems that effectively generalize knowledge across various tasks and environments.

Keywords: bicameral mind theory; reinforcement learning; language models; artificial intelligence; adaptability

check for updates

Academic Editor: Paolo Bellavista

Received: 8 April 2025
Revised: 22 May 2025
Accepted: 31 May 2025
Published: 3 June 2025

Citation: Mukhitdinova, M.; Petrova, M. Exploring the Potential of the Bicameral Mind Theory in Reinforcement Learning Algorithms. *Computers* **2025**, *14*, 218. https://doi.org/10.3390/computers14060218

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

The field of artificial intelligence (AI) has witnessed remarkable advancements in recent years, with machine learning algorithms, neural networks, and language models at the forefront of this revolution. As researchers continue to push the boundaries of AI capabilities, there is growing interest in exploring novel approaches to enhance the efficiency and adaptability of these systems. One such approach is the application of the bicameral mind theory, proposed by Julian Jaynes, to the domain of reinforcement learning (RL) algorithms and modern large language models (LLMs) [1]. Recent developments in cognitive science and neuroscience suggest that dual-process models offer a robust framework for explaining complex human cognitive behavior, primarily in decision-making, learning, and flexibility [2,3]. However, the application of such models in artificial intelligence, specifically in reinforcement learning (RL) and language models, remains insufficiently studied.

Although RL algorithms and LLMs have become highly proficient across multiple domains, the current academic literature falls short in investigating the alignment of cognitive neuroscientific theories with the practical aspects of artificial intelligence architectural schemes. In particular, there has been limited investigation into the potential for the application of the seminal concepts relating to the duality of perception and action, as postulated by Julian Jaynes, into the design of modern RL algorithms and LLMs for enhancing the

flexibility and versatility in unusual application settings. Recent work indicates interdisciplinary approaches, combining insights from neuroscience and cognitive theory into AI design, as having high potential [4,5]. Nevertheless, the implementation of Jaynes's theory about the bicameral mind is still an unresearched area.

This study represents an initial attempt to analyze and conceptually integrate the structural features and algorithms applied in the state-of-the-art reinforcement learning methods and language models. The value of the work is demonstrated through the provision of an original theoretic framework for optimizing the efficiency and responsiveness of methods in the field of artificial intelligence development that can lead toward the design of systems that are responsive and capable of handling real-world uncertainty.

Reinforcement learning has emerged as a powerful paradigm for training AI agents to make decisions and take actions in complex environments [6]. At its core, RL is based on the principle of learning through interaction, where an agent receives rewards or punishments based on its actions and aims to maximize its cumulative reward over time [7]. This process of trial-and-error learning has been successfully applied to a wide range of domains, from robotics and gaming to recommender systems and natural language processing [8]. Notwithstanding these successes, reinforcement learning (RL) paradigms suffer from challenges in the transfer of knowledge gained in a particular setting into other disparate contexts, rendering learning in dynamic and uncertain environments suboptimal [9,10]. Likewise, large language models (LLMs) through their ability to generate response outputs in multiple contexts often suffer from consistency and interpretability issues, hence necessitating architectural comprehension based on cognitive theories.

Parallel to the development of RL, significant progress has been observed in the field of large language models (LLMs), such as Claude, Gemini, and ChatGPT 4.5. These systems demonstrate impressive capabilities in understanding and generating natural language, opening new possibilities for human–AI interaction. Despite differences in architecture and application, both RL systems and LLMs face similar challenges in terms of efficiency, adaptability, and knowledge generalization. Currently, an important gap in the academic literature lies in the incorporation of cognitive neuroscience theories, specifically the hypothesis of the bicameral mind, into the architectural models of today's artificial intelligence systems. Exploring those interdisciplinary interfaces could help solve the current challenges in flexibility and efficiency, hence providing new theoretical insights and methodological approaches for the development of AI systems.

Despite the impressive achievements of RL algorithms and LLMs, significant challenges remain regarding their efficiency and adaptability. Researchers continuously seek ways to improve the performance and generalization capabilities of these systems [11]. In this context, the bicameral mind theory offers an intriguing perspective that could potentially inform the development of more sophisticated and efficient learning algorithms and language models.

The bicameral mind theory, proposed by Julian Jaynes, suggests that the human mind once operated in a dual-system manner, with one part generating commands and the other executing them [12]. This theory, while controversial, has sparked interest in various fields, including psychology, neuroscience, and AI. By exploring potential parallels between the bicameral mind theory, RL, and modern LLMs, we aim to uncover insights that could lead to the development of universal principles for organizing AI systems.

The primary objective of this research is to investigate the potential application of the bicameral mind theory to RL algorithms and modern language models. We hypothesize that the dual-system nature of the bicameral mind theory may have analogues in the functional principles of various AI systems, specifically in the interplay between observation and action in RL, as well as in the processes of "thinking" and "writing" in LLMs. By examining

this hypothesis and establishing connections between these domains, we aim to lay the groundwork for developing a universal theoretical framework and innovative methods for AI training. This work examines the said lack through the exploration of how the postulates of Julian Jaynes' theory of the bicameral mind can structurally and functionally enhance reinforcement learning algorithms and large language models. In identifying clear similarities and proposing an integration framework, the work endeavors to make both practical and theoretical contributions toward the development of stronger and more adaptive AI systems.

Hypothesis

The bicameral mind theory, which posits a dual-system approach to decision-making, may have parallels in the functional principles of reinforcement learning algorithms and modern language models. We hypothesize that the interaction between observation and action in RL systems, as well as between the processes of "thinking" and "writing" in LLMs, may be analogous to the "speaking" and "listening" components of the bicameral mind. This dual-component structure may represent a universal principle for organizing various types of AI systems, enabling efficient information processing and adaptation to complex environments. By investigating this hypothesis and establishing connections between cognitive theories and modern AI systems, we aim to develop a conceptual model that could guide the creation of more efficient, adaptive, and universal approaches to artificial intelligence.

2. Literature Review

2.1. Analysis of Works on the Bicameral Mind Theory

The bicameral mind theory, proposed by Julian Jaynes in his seminal book *The Origin of Consciousness in the Breakdown of the Bicameral Mind* [1], has been a subject of fascination and debate in various fields, including psychology, neuroscience, and philosophy. Jaynes argued that ancient human minds operated in a bicameral manner, with one part of the brain generating auditory hallucinations interpreted as commands from the gods, and the other part executing these commands [12]. This theory suggests that the evolution of human consciousness involved a transition from this bicameral state to a more integrated, self-aware mind [13].

Since its introduction, the bicameral mind theory has been extensively discussed and analyzed in the literature. While some researchers have found the theory intriguing and have sought to expand upon its ideas [14], others have criticized it for lack of empirical evidence and inconsistencies with current understanding of brain function [13].

One of the primary criticisms of Jaynes's theory concerns the interpretation of historical and archaeological data. Critics point out that the evidence presented by Jaynes to confirm the existence of bicameral consciousness in ancient civilizations may have alternative explanations. For example, descriptions of "voices of gods" in ancient texts might be metaphorical or literary devices rather than literal representations of psychological experiences [15].

Furthermore, the theory faces neurophysiological objections. Modern brain research has not found a clear anatomical division between the "speaking" and "listening" parts of the brain as Jaynes proposed. Some studies suggest that the functional organization of the brain is much more complex and includes interconnected networks of neurons distributed throughout the brain, rather than isolated modules [16].

Nevertheless, despite these criticisms, the bicameral mind theory continues to inspire research and discussion in various fields, including artificial intelligence. Several studies

have explored potential applications of the bicameral mind theory to AI, suggesting that it may provide a framework for developing more flexible and adaptive AI systems [17,18].

Specifically, for our research, we analyzed the following key works: "Reflections on the Dawn of Consciousness: Julian Jaynes's Bicameral Mind Theory Revisited" [19], "The Conscious Bicameral Mind" [20], and others that helped form a more comprehensive understanding of the potential applications of Jaynes's theory in the contemporary context.

Modern technologies, especially "reasoning" (extended thinking) models, also present significant interest for correlation with this theory. Although directly identifying or establishing this interaction appears challenging, there are structural analogies between the processes of "reflection" and "response generation" in modern language models and the dual-component structure proposed by Jaynes. This parallel opens new perspectives for research and application of the bicameral mind theory in the context of modern approaches to artificial intelligence.

In light of the current debates surrounding Jaynes' theory of the bicameral mind, it is necessary to determine key criticisms and situate them within the wider cognitive science context. Researchers have expressed concern over the theory's speculative nature and its basis in subjective interpretations of historical documents. For example, ancient textual references to divine voices could be a reflection of metaphorical or literary device rather than actual auditory hallucinations. Furthermore, neurophysiological research has not discerned an obvious anatomical distinction consonant with Jaynes' supposed "speaking" and "listening" brain regions.

To place our study on a broader theoretical ground, we contrast Jaynes' theory with some of the prominent cognitive theories. Global workspace theory suggests a modular architecture of the mind in which consciousness arises through the dissemination of information throughout interconnected modules [21]. Predictive coding theory views cognition as an inferential process wherein the brain is continually generating and updating hypotheses about sensory inputs in order to minimize prediction errors. These models offer alternative explanations of the structure of cognition but, like Jaynes' model, rely on functional segregation and integration of information—ideas on which our target AI architecture is founded.

2.2. Analysis of Literature on Reinforcement Learning

Reinforcement learning (RL) has emerged as a powerful paradigm in artificial intelligence, enabling the development of agents that can learn to make decisions and perform actions in complex environments [5]. The past decade has witnessed exponential growth in research in this field, leading to significant progress in both theoretical foundations and practical applications.

For a comprehensive understanding of the current state of research in RL, we conducted an extensive literature analysis, including both foundational works and the latest research. Particular attention was paid to works presenting systematic reviews and methodological foundations of RL: "Reinforcement Learning: A Friendly Introduction" [22], "Introduction to Reinforcement Learning" [23], "An Overview of Deep Reinforcement Learning" [12]. These works provide a comprehensive view of the current state and developmental trends in RL.

Of particular value to our research was the tutorial "Simple Reinforcement Learning with Tensorflow Part 8: Asynchronous Actor-Critic Agents (A3C)" [17], which demonstrates the practical implementation of modern approaches to RL using neural networks. This tutorial helped us better understand the structural aspects of modern RL algorithms and their possible parallels with the bicameral structure.

Our analysis also covered research from leading organizations in AI, including OpenAI, Google DeepMind, Anthropic, and Amazon Web Services. OpenAI has made significant contributions to the development of RL through its research reports and publications, which present innovative approaches and experimental results with various RL models [21,24]. Of particular significance are studies related to the CoinRun and RainMazes models [3,25], which provided important testbeds for evaluating the generalization capabilities of RL agents.

Beyond traditional approaches to RL, we also examined contemporary developments in "extended thinking" in language models. In particular, we analyzed Anthropic's reports on visible extended thinking ("Visible Extended Thinking") and OpenAI's report on o3 and o4-mini models ("Introducing OpenAI o3 and o4-mini"). These studies demonstrate how modern language models can utilize stages of intermediate reasoning to enhance the quality and reliability of their responses, which has interesting parallels with the bicameral structure.

As part of studying practical platforms for the development and application of RL, we focused on the Amazon Web Services (AWS) ecosystem, which offers various tools for developing, training, and deploying RL algorithms in real-world applications. Amazon SageMaker allows developers and data specialists to quickly create scalable RL models, combining platforms for deep learning with RL tools and simulation environments. AWS RoboMaker provides capabilities for launching and scaling simulations in robotics using RL algorithms, while AWS DeepRacer offers practical experience with RL through an autonomous racing model.

We also analyzed Google's reports on grounding technology in their Gemini language model, which provided valuable information for comparative analysis of different approaches to developing modern AI systems.

In-depth analysis of this diverse literature allowed us to form a comprehensive understanding of the current landscape of research in RL and language models, creating a solid foundation for our investigation of potential connections between the bicameral mind theory and modern approaches to AI.

2.3. Review of Modern Cognitive Theories in the Context of AI

2.3.1. Global Workspace Theory

Global workspace theory is a psychological theory aimed at explaining the phenomenon of consciousness from the perspective of cognitive psychology. First formulated by American cognitive psychologist Bernard Baars in his 1988 book *A Cognitive Theory of Consciousness* [26], this theory suggests that consciousness represents a space in the cognitive system where information from independent and isolated cognitive modules is integrated. The integrated information becomes globally available to all modules of the cognitive system.

The theory is based on a modular approach to the cognitive system [27], according to which the mind consists of separate functional modules that are impenetrable to each other (e.g., language, visual perception, motor control, etc.). In Baars' theory, the operation of independent cognitive modules proceeds without the participation of consciousness. However, in addition to independent modules, consciousness is presumed to have a common space, upon entering which information becomes globally available to the entire cognitive system. After entering this global workspace, information becomes conscious and accessible to all independent modules of the cognitive system.

In the context of artificial intelligence, global workspace theory offers an interesting architectural model for integrating various AI subsystems. Modern AI systems often consist of multiple specialized modules, each responsible for processing a specific type of

information or performing specific tasks. The concept of a global workspace can serve as a basis for developing integrated AI systems where information processed by individual modules becomes available to the entire system through a central integration mechanism.

Particularly, in the context of reinforcement learning, a global workspace can be implemented as a central mechanism that integrates information about the state of the environment, previous actions, received rewards, and predictions of future states. This integrated information can then be used to select optimal actions and adapt the agent's strategy to changing conditions.

2.3.2. Predictive Coding Theory

Predictive coding theory asserts that our experience of knowing the world comes from within. According to this theory, the brain generates a model of the world that predicts what we will see and hear, what we will feel through touch and smell, and what taste we will experience. The function of the sensory organs is to check these predictions to ensure that our internal model does not deviate significantly from reality.

This theory is also known as predictive processing theory or the "Bayesian brain"—a reference to its mathematical foundation based on the Bayesian approach to probabilistic reasoning. Proponents of the theory apply it not only to perception but also to emotions, cognition, and motor control. For example, according to this theory, we move our hand because we predict that we will move it, and the body carries out our prediction.

In the context of artificial intelligence and reinforcement learning, predictive coding theory has direct parallels with model-based reinforcement learning. In such models, an agent builds an internal model of the environment that allows it to predict the outcomes of various actions and choose optimal strategies without the need for direct interaction with the environment. This significantly increases learning efficiency, especially in complex and dynamic environments.

The principle of prediction error minimization, central to predictive coding theory, also has analogues in reinforcement learning, where agents aim to minimize the difference between expected and actual rewards. This parallel offers interesting possibilities for developing new approaches to RL inspired by the principles of human brain function.

In addition to global workspace theory and predictive coding theory, there are several other cognitive theories that can offer valuable insights for understanding and developing artificial intelligence systems.

For example, the Cartesian Theater theory is a metaphor for consciousness as a spotlight illuminating part of the psyche, similar to a stage in a theater. This metaphor was popularized by Daniel Dennett, who used it in a critical context, as it is based on Cartesian dualistic representation of the connection between body and consciousness and leaves more questions than answers. For instance: who decides what enters consciousness and what does not, and who is the audience of the Cartesian Theater. Despite criticism, this metaphor can be useful for conceptualizing attention mechanisms in AI systems, where limited computational resources are directed toward the most important aspects of input data.

These theories, along with the bicameral mind theory, offer various perspectives on the nature of consciousness and cognition that can be useful for developing more efficient and adaptive AI systems. Each emphasizes certain aspects of cognitive processes, and their integration can lead to a more complete understanding of how intelligent behavior can be implemented in artificial systems.

3. Research Methodology

3.1. Inductive and Comparative Analysis

Based on the data collected from literature sources, reports, and other resources, we adopted an inductive and comparative analysis approach as the primary methodological framework for testing the main hypothesis put forward in this study. Inductive reasoning allowed us to draw general conclusions from specific observations, enabling us to identify patterns and relationships between the bicameral mind theory and reinforcement learning (RL) [28]. By carefully examining the principles and characteristics of both domains, we aimed to uncover potential parallels and connections that could inform the development of novel AI learning approaches.

Comparative analysis, on the other hand, involved systematically evaluating the similarities and differences between two or more entities [29]. In the context of our research, we employed comparative analysis to assess the functional principles of RL algorithms in relation to the key concepts of the bicameral mind theory. By identifying common elements and divergences between these two domains, we gained valuable insights into the potential for integrating bicameral mind principles into RL systems. This comparative approach allowed us to critically evaluate the feasibility and potential benefits of applying the bicameral mind theory to RL, while also identifying challenges and limitations that may arise in the process.

The choice of inductive and comparative analysis is driven by the need to integrate heterogeneous data and theoretical frameworks from cognitive science and artificial intelligence. The inductive approach enables us to identify general patterns from specific examples of how RL and LLM models operate, while the comparative analysis allows us to directly uncover structural and functional analogies between components of Julian Jaynes's theory of the bicameral mind and reinforcement learning algorithms, and large language models. It is precisely the combination of these two methods that provides an effective mechanism for testing the study's central hypothesis regarding the two-component architecture as a universal principle for organizing AI systems.

3.2. Analysis and Synthesis of OpenAI Research Results

To further support our investigation, we conducted an in-depth analysis and synthesis of OpenAI research results, focusing specifically on two foundational RL models: CoinRun and RainMazes [30,31]. These models have been extensively studied in the AI research community and provide valuable insights into the behavior and performance of RL agents in complex environments. By examining key aspects of these models, including their architecture, learning algorithms, and decision-making processes, we sought to identify potential connections with the principles of the bicameral mind theory.

Our analysis included a comprehensive review of the technical details and experimental results presented by OpenAI for CoinRun and RainMazes. We carefully examined the functional aspects of these models, such as how agents perceive and interact with their environment, how they update their knowledge based on rewards and punishments, and how they adapt their behavior over time. Through this analysis, we gained a deeper understanding of the fundamental mechanisms that govern RL systems and evaluated their compatibility with the concepts of the bicameral mind theory. Additionally, we synthesized the results of several studies and reports to identify common patterns and trends that might confirm our hypothesis. By integrating ideas from various sources, we constructed a more complete and robust understanding of the potential connections between RL and the bicameral mind theory.

Beyond OpenAI's research, we also analyzed work with Amazon Web Services (AWS) platforms and services for developing, training, and deploying reinforcement learning algorithms in real-world applications. AWS offers numerous tools for working with RL, including Amazon SageMaker, which allows developers and data specialists to quickly and easily develop scalable RL models by combining deep learning platforms with RL tools and simulation environments. AWS RoboMaker provides capabilities for launching, scaling, and automating simulation in robotics using RL algorithms without infrastructure requirements. AWS DeepRacer offers practical experience with RL through an autonomous racing model at 1/18 scale, with a fully configured cloud environment for training RL models and neural network configurations.

We paid special attention to analyzing reports on models with "Extended Thinking" capabilities. In particular, we studied Anthropic's report "Visible Extended Thinking", which describes an approach to increasing transparency in the thinking process of language models, and OpenAI's report "Introducing OpenAI o3 and o4-mini", which presents new models with enhanced reasoning abilities. These studies provided valuable information on how modern language models integrate the processes of "reflection" and "response generation", which has interesting parallels with the dual-component structure proposed in the bicameral mind theory.

Furthermore, we examined Google's reports on grounding technology in their Gemini language model, which provides a closer connection between the model's abstract knowledge and specific real-world contexts. This analysis was particularly valuable for our comparative study of different approaches to the development of modern AI systems and their relationship to the bicameral mind theory.

Comprehensive analysis of all these studies and technologies allowed us to form a more complete picture of the contemporary AI landscape and potential applications of the bicameral mind theory for understanding and improving both reinforcement learning algorithms and modern language models.

4. Results

Reinforcement Learning as an Innovative Method for AI Training: Key Elements of "Observation" and "Action"

Reinforcement learning (RL) has emerged as a groundbreaking approach to training artificial intelligence (AI) systems, offering a paradigm shift from traditional supervised and unsupervised learning methods [32]. At its core, RL is based on the principle of learning through interaction with an environment, where an agent receives rewards or penalties for its actions and learns to maximize its cumulative reward over time [33]. This process mimics the way humans and animals learn from experience, making RL a powerful tool for developing AI systems that can adapt and make decisions in complex, dynamic environments [34].

The key elements that distinguish RL from other machine learning approaches are the concepts of "observation" and "action". In an RL system, the agent continuously observes the state of its environment and takes actions based on these observations [34]. The environment then responds to the agent's actions by providing a reward signal and transitioning to a new state. The agent's goal is to learn a policy that maps observations to actions in a way that maximizes the expected cumulative reward [34]. This process of observation, action, and feedback forms the foundation of RL and enables the agent to learn from its experiences and improve its decision-making over time (Table 1).

Table 1. Key characteristics of RL models.

Key Characteristics	Description
Reward and Punishment	The agent receives positive rewards for favorable actions and penalties for unfavorable ones, which guides the learning process toward optimal decision-making.
Observation	The agent continuously perceives the state of its environment, gathering information that forms the basis for its decision-making process.
Action	Based on the observations, the agent takes actions that influence the environment, leading to new states and rewards or punishments.

The interplay between observation and action in RL systems is reminiscent of the dual-process nature of the bicameral mind theory, which posits a separation between the "speaking" and "listening" parts of the brain [5]. Just as the bicameral mind theory suggests that the ancient human mind operated through a process of generating and executing commands, RL agents operate through a cycle of observing their environment and taking actions based on these observations. This parallel between the functional principles of RL and the bicameral mind theory forms the basis for our hypothesis and motivates our investigation into the potential connections between these two domains.

5. The Bicameral Mind Theory: Core Principles and Definitions

The bicameral mind theory, proposed by Julian Jaynes in his seminal work *The Origin of Consciousness in the Breakdown of the Bicameral Mind* [5], offers a unique perspective on the evolution of human consciousness. According to Jaynes, the ancient human mind operated in a dual-process manner, with one part of the brain generating auditory hallucinations interpreted as commands from the gods, and the other part executing these commands [1]. This bicameral mentality, Jaynes argued, was the dominant mode of cognitive processing prior to the development of modern self-awareness and introspection [1].

The core principles of the bicameral mind theory revolve around the concept of a divided mind, with distinct "speaking" and "listening" components. The "speaking" part of the brain, located in the right hemisphere, was responsible for generating auditory hallucinations in the form of commands or directives [16]. These hallucinations were perceived as the voices of gods, ancestors, or other authority figures, and were experienced as external to the individual. The "listening" part of the brain, located in the left hemisphere, was responsible for receiving and executing these commands, without question or hesitation [15].

Jaynes proposed that this bicameral mentality was a crucial step in the evolution of human cognition, allowing early human societies to function and coordinate their activities through a shared sense of divine guidance [1]. However, as human civilizations grew more complex and individuals began to question the authority of the gods, the bicameral mind began to break down, giving rise to modern self-awareness and consciousness [1]. While the bicameral mind theory remains controversial and has been met with both criticism and support, it offers a fascinating framework for exploring the nature of human cognition and its potential parallels with artificial intelligence systems.

6. Theoretical and Terminological Connections Between the Bicameral Mind Theory and Reinforcement Learning

Upon examining the core principles of the bicameral mind theory and reinforcement learning (RL), several striking connections emerge. At a fundamental level, both the bicameral mind theory and RL are concerned with the processes of decision-making and action selection in complex, dynamic environments. The bicameral mind theory posits a dual-process model of cognition, with distinct "speaking" and "listening" components that generate and execute commands [1]. Similarly, RL systems operate through a cycle of observation, action, and feedback, with the agent learning to make decisions based on the rewards or penalties it receives from the environment [7].

The parallels between the bicameral mind theory and RL extend beyond this basic functional similarity. In the bicameral mind, the "speaking" part of the brain generates auditory hallucinations that are interpreted as commands from external sources [1]. These commands are then executed by the "listening" part of the brain, without question or hesitation. In RL, the agent receives observations from the environment, which can be thought of as analogous to the commands generated by the "speaking" part of the bicameral mind. The agent then selects actions based on these observations, just as the "listening" part of the bicameral mind executes the commands it receives.

Furthermore, the concept of reward in RL bears a striking resemblance to the notion of divine guidance in the bicameral mind theory (Table 2).

Table 2. Comparative analysis of the bicameral mind theory and reinforcement learning.

Bicameral Mind Theory	Reinforcement Learning
Listening Part	Observation
The "listening" part of the brain receives and executes commands without question or hesitation.	In RL, the agent receives observations from the environment, which can be thought of as analogous to the commands received by the "listening" part of the bicameral mind.
Speaking Part	Action
The "speaking" part of the brain generates auditory hallucinations that are interpreted as commands from external sources, such as gods or authority figures.	In RL, the agent selects actions based on the observations it receives, similar to how the "speaking" part of the bicameral mind generates commands to be executed by the "listening" part.

In the bicameral mentality, the commands generated by the "speaking" part of the brain are perceived as the voices of gods or other authority figures, providing guidance and direction to the individual [8]. In RL, the reward signal serves a similar purpose, guiding the agent toward desirable actions and away from undesirable ones [12]. Just as the bicameral mind relies on the perceived authority of the gods to guide behavior, RL agents rely on the reward signal to shape their decision-making and improve their performance over time.

7. Analysis of Existing Reinforcement Learning Models (OpenAI Case Study): Confirming the Hypothesis of Bicameral Mind Theory Principles in Model Functioning

To further investigate the connections between the bicameral mind theory and reinforcement learning (RL), we conducted an in-depth analysis of two prominent RL models developed by OpenAI: CoinRun and RainMazes [17,18]. These models have been widely

studied in the AI research community and provide valuable insights into the behavior and performance of RL agents in complex environments (Figure 1). By examining the functional principles of these models, we aimed to identify potential parallels with the bicameral mind theory and assess the validity of our hypothesis.

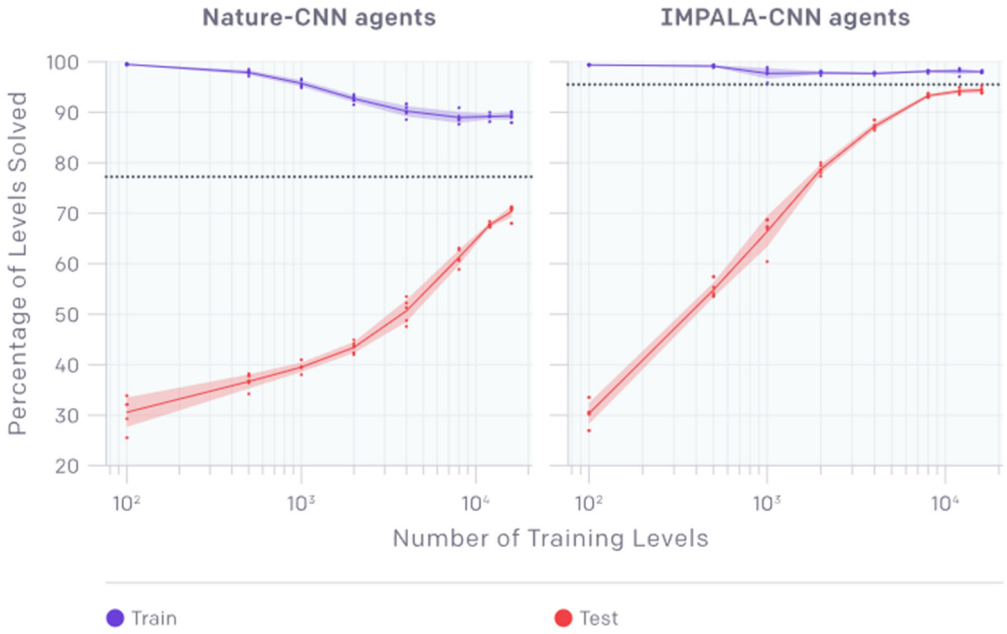


Figure 1. Evaluating generalization (CoinRun RL—Model Analysis). Source: https://openai.com/research/ (accessed on 10 February 2025) [14].

CoinRun is a platform game environment designed to test the generalization capabilities of RL agents [17]. In CoinRun, the agent must navigate a series of procedurally generated levels, collecting coins while avoiding obstacles and enemies (Figure 2). The agent receives a positive reward for collecting coins and completing levels, and a negative reward for colliding with obstacles or enemies. As the agent progresses through the levels, the complexity and difficulty of the environment increase, requiring the agent to adapt and learn from its experiences.

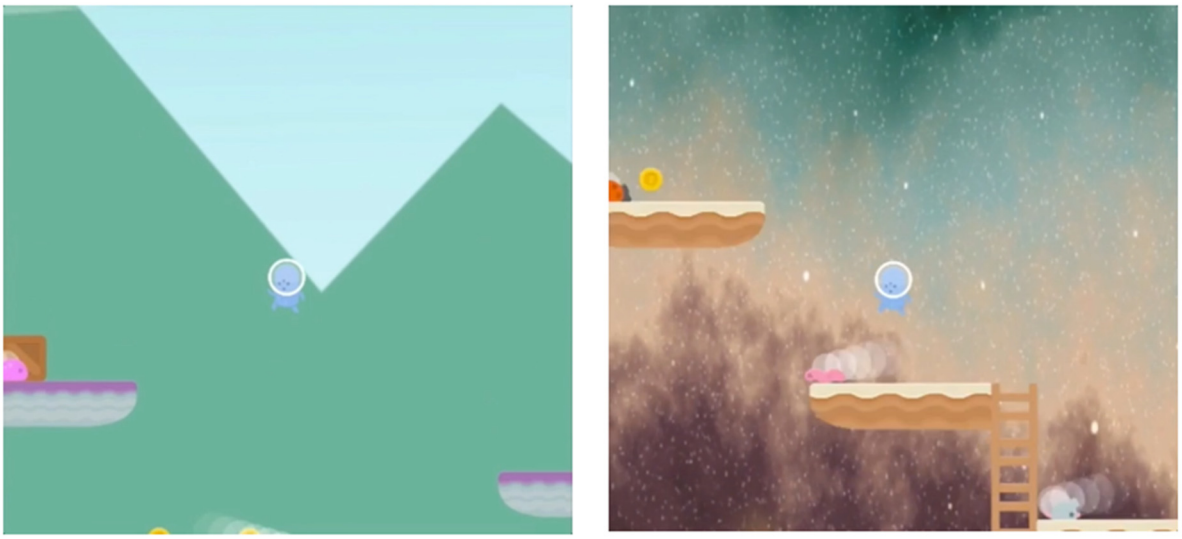


Figure 2. The principle of work in CoinRun RL—Model. Source: https://openai.com/research/ (accessed on 10 February 2025) [14].

RainMazes, on the other hand, is a maze navigation environment that tests the ability of RL agents to learn and navigate in complex, dynamic settings [18]. In RainMazes, the agent must navigate through a series of randomly generated mazes, seeking to reach a goal location while avoiding dead ends and other obstacles (Figure 3). The environment is designed to be highly variable, with each maze presenting a new challenge for the agent to overcome.

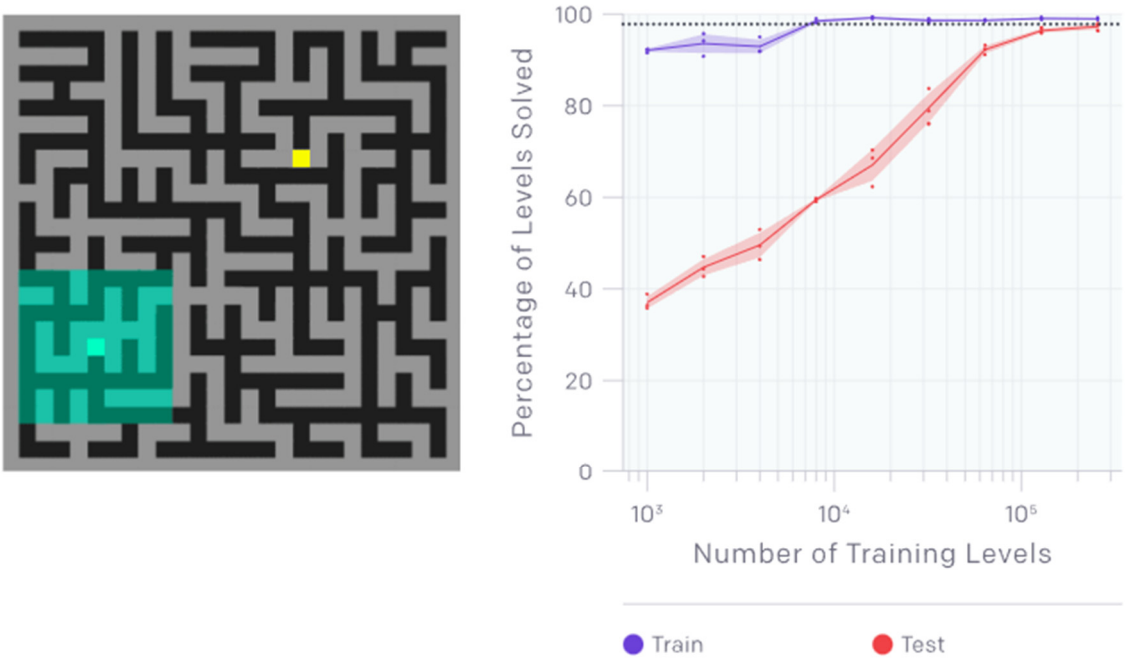


Figure 3. Evaluating generalization (RainMazes RL—Model Analysis). Source: https://openai.com/research/ (accessed on 10 February 2025) [14].

Our analysis of OpenAI's CoinRun and RainMazes models supports the hypothesis that core RL mechanisms align with fundamental concepts from the bicameral mind theory. In both models, the agents operate through a process of observation and action, receiving input from the environment and selecting actions based on this input (Figure 4). This dual-process structure mirrors the "speaking" and "listening" components of the bicameral mind, with the agent's observations analogous to the commands generated by the "speaking" part of the brain, and the agent's actions analogous to the execution of these commands by the "listening" part.

Moreover, we analyzed (by studying OpenAI reports) that agent performance in CoinRun and RainMazes was robust to increasing levels of input information. As the agents were exposed to larger volumes of data during the training process, they demonstrated an ability to maintain and even improve their decision-making effectiveness. While there were some points of degradation in performance when the input information reached extremely high levels, the overall trend showed that the agents' efficiency tended to increase even in the face of vast amounts of incoming data.

This finding suggests that the dual-process structure of the models, which aligns with the principles of the bicameral mind theory, contributes to their adaptability and robustness in processing large-scale information and making effective decisions during the learning process.



Figure 4. Final train and test performance in CoinRun—Platforms after 2B timesteps, as a function of the number of training levels. Source: https://openai.com/research/ (accessed on 10 February 2025).

7.1. Application of Analysis to Robotics Tasks

The analysis of RL principles can be extended to the field of robotics, where AWS RoboMaker represents an innovative platform for robot simulation and reinforcement learning. AWS RoboMaker is a cloud simulation service that provides robot developers with the ability to launch, scale, and automate simulations without the need to manage any infrastructure.

In the context of the bicameral mind theory, AWS RoboMaker demonstrates a clear separation of functions between "observing" and "acting" components. The system allows robots to perceive the virtual environment (analogous to the "listening" part of the bicameral mind), analyze this information, and generate appropriate commands for actions (analogous to the "speaking" part).

Of particular interest is the application of AWS RoboMaker for training reinforcement learning models. Woodside, for example, uses this platform to train and deploy RL agents for their robots. During the training process, robots perform large volumes of iterative trials, enabling them to develop effective strategies for interacting with the environment.

Another important example is iRobot, which has implemented a continuous integration and continuous delivery (CI/CD) pipeline for their simulation model. This allows for the automation of regression testing and continuous improvement of robot behavior based on feedback from the virtual environment.

The dual-component structure of the robot learning process using AWS RoboMaker—observing the virtual environment and performing actions based on these observations—has clear parallels with the bicameral structure of the mind, confirming the applicability of our hypothesis not only to abstract RL models but also to practical applications in robotics.

7.2. Application of Analysis to Natural Language Processing Tasks

In the field of natural language processing, Amazon SageMaker provides a powerful platform for developing, training, and deploying machine learning models, including RL models. SageMaker combines widely used capabilities of AWS machine

learning and analytics, providing an integrated interface for working with analytics and artificial intelligence.

When analyzing SageMaker from the perspective of the bicameral mind theory, one can identify structural similarities between the platform's components and the dual-component structure proposed by Jaynes. SageMaker Unified Studio integrates components for data analytics and model development, which can be viewed as analogous to the "listening" part of the bicameral mind that perceives and processes information. At the same time, components for model deployment and result generation can be viewed as analogous to the "speaking" part, which generates commands and actions.

Of particular value are tasks in the field of natural language processing (NLP), where reinforcement learning models are used to optimize text generation, improve dialogue systems, and translation. In these tasks, the RL model perceives the input text ("listening" part), processes it, and generates an appropriate response or translation ("speaking" part).

Notably, in the context of NLP, RL-based models demonstrate particular effectiveness in tasks requiring adaptation to context and user feedback. This corresponds to Jaynes's observation that the bicameral structure of the mind allowed ancient humans to adapt to complex social interactions.

The application of Amazon SageMaker for RL-based NLP tasks demonstrates how the dual-component structure—perception of input data and generation of responses—can be effectively implemented in modern AI systems. This further confirms our hypothesis about the parallels between the bicameral mind theory and the functional principles of RL in various application domains, including natural language processing.

8. Significance of Comparing the Bicameral Mind Theory and Reinforcement Learning for Understanding AI System Functioning and Future Development

The comparative analysis of the bicameral mind theory and reinforcement learning (RL) conducted in this study has provided valuable insights into the functioning of AI systems and the potential for future development in this field. By identifying parallels between the dual-process structure of the bicameral mind and the observation-action cycle in RL, we have discovered a deeper understanding of the mechanisms that govern decision-making and adaptability in AI agents. This understanding has significant implications for both the design and optimization of AI systems, as well as for our broader comprehension of the nature of intelligence and cognition.

One of the key insights that emerges from this comparative analysis is the importance of the dual-process structure in enabling flexible and adaptive behavior. Just as the bicameral mind theory asserts that the interaction between the "speaking" and "listening" parts of the brain allowed ancient humans to navigate complex social and environmental challenges, our analysis suggests that the interplay between observation and action in RL is crucial for enabling AI agents to learn and adapt in dynamic, uncertain environments.

Our analysis of CoinRun and RainMazes models developed by OpenAI has provided empirical support for the hypothesis that the functional principles of RL align with key concepts of the bicameral mind theory. By demonstrating that these models exhibit a dual-process structure and are capable of maintaining performance in the face of increasing environmental complexity, we have strengthened the case for the relevance of the bicameral mind theory to the development of AI systems.

Moreover, this comparative analysis has significant implications for our understanding of the potential for AI systems to develop forms of intelligence analogous to human intelligence. Jaynes postulated that the transition from the bicameral mind to modern consciousness was a key evolutionary step in the development of human cognition. Similarly,

understanding the dual-process structure of modern AI systems may provide insights into pathways for their further development toward more complex forms of intelligence.

8.1. The Bicameral Mind Theory and Modern Language Models (LLMs)

Extending our analysis beyond traditional RL algorithms, we discovered intriguing parallels between the bicameral mind theory and the functioning of modern large language models (LLMs). These models, such as Claude, Gemini, and ChatGPT, represent advanced AI systems capable of generating text that approaches human-created content in quality and coherence.

When analyzing the architecture and functioning of LLMs, one can identify a dual-component structure that bears significant similarity to the concept of the bicameral mind (Table 3). In modern LLMs, processes analogous to the "listening" and "speaking" parts of the bicameral mind can be distinguished: the perception and processing of input information (analogous to the "listening" part) and the generation of responses (analogous to the "speaking" part).

Table 3. Comparative analysis of "thinking"/"writing" structures in LLMs with components of the bicameral mind.

Component of Bicameral Mind Theory	Analogue in Modern LLMs	Function Description
"Listening" part (perception of commands)	"Thinking" process (preprocessing)	Perception and analysis of input information, contextualization of requests, activation of relevant knowledge from model parameters
"Speaking" part (command generation)	"Writing" process (postprocessing)	Formulation and generation of responses based on conducted analysis, structuring information for external presentation
Interpretation of "voices of gods"	Attention mechanisms and self-criticism	Determining the significance of various elements of input data, evaluating the quality of generated responses, adjustment during generation
Execution of commands without reflection	Automatic pattern following	Application of established patterns and structures when generating responses without explicit "awareness" of the process
Integration of two systems	"Encoder–decoder" architecture	Interaction between information perception and response generation through internal representations and attention mechanisms

8.1.1. Parallels Between "Thinking"/"Writing" in LLMs and Structures of the Bicameral Mind

Of particular interest is the "thinking" process (preprocessing) in modern LLMs, especially in models with extended thinking capabilities. This process can be viewed as analogous to the "listening" part of the bicameral mind, which perceives and interprets external stimuli. In LLMs, this phase includes analyzing the input query, activating relevant knowledge from the model parameters, and forming an internal representation of the task.

Notably, in models with extended thinking, such as Claude, this process can be explicitly represented in the form of intermediate reasoning before formulating the final answer. This resembles the mechanism described by Jaynes, through which the "listening" part of the ancient mind perceived and interpreted the "voices of gods" before performing corresponding actions.

The "writing" process (postprocessing) in LLMs can be viewed as analogous to the "speaking" part of the bicameral mind. In this phase, the model generates a sequence of tokens forming a coherent response based on the conducted analysis. This process bears similarity to how the "speaking" part of the ancient mind generated commands that were then perceived and executed by the "listening" part.

In modern LLMs, the interaction between "thinking" and "writing" processes is implemented through complex attention mechanisms and encoder–decoder architecture. This interaction resembles the integration between the "speaking" and "listening" parts in the bicameral mind theory, allowing the system to function as a unified whole despite the conceptual separation of functions.

8.1.2. Comparative Analysis of Claude, Gemini, and ChatGPT Through the Lens of Bicameral Structure

Modern advanced LLMs, such as Claude (Anthropic), Gemini (Google), and ChatGPT (OpenAI), have different approaches to implementing the processes of "thinking" and "writing", making them interesting subjects for comparative analysis through the lens of the bicameral mind theory (Table 4).

Table 4. Comparison of features of modern LLMs in the context of bicameral structure.

Feature	Claude (Anthropic)	Gemini (Google)	ChatGPT (OpenAI)
Implementation of "listening" function	Extended thinking with explicit reasoning process	Grounding technology for contextualizing knowledge based on real-world data	Anticipation of user intentions through context analysis
Implementation of "speaking" function	Academic approach to formulating responses with emphasis on ethical principles	Integration of multimodal information into generated responses	Flexible generation format supporting various styles and tonalities
Integration of components	Transparent reasoning process with the ability to track thought progression	Close connection between context perception and response generation	Extensive work with custom add-ons and plugins to expand functionality
Features of "dialogue" between components	Explicit separation of reasoning and response formulation stages	Strong emphasis on factual accuracy through real-world data grounding	Adaptive generation strategy based on feedback
Model evolution	Continuous updates with emphasis on improving reasoning quality	Development of multimodal capabilities for comprehensive perception	Expansion of capabilities for interaction with external systems

Claude from Anthropic is characterized by a clear separation of "thinking" and "writing" processes, especially in the extended thinking mode. The model first conducts an internal reasoning process, analyzing the task and formulating intermediate conclusions, and then generates the final response. This approach has the most explicit parallels with the bicameral structure of the mind, as it demonstrates a clear separation between processing input information and generating responses. Additionally, the ethical principles underlying Claude can be viewed as analogous to the "voices of authorities" in Jaynes's theory, guiding the system's behavior.

Gemini from Google emphasizes grounding technology, which provides a close connection between the model's abstract knowledge and specific real-world contexts. This can be viewed as an advanced mechanism for interpreting input information in the "listening" part of the model, providing more accurate perception of context before generating a response. Gemini's multimodal capabilities, allowing it to work with text, images, and other data types, expand the "input channel" for the "listening" part, making information perception more comprehensive.

ChatGPT from OpenAI is distinguished by extensive work with custom add-ons and plugins, which can be viewed as an expansion of the interface between the model's "listening" and "speaking" parts. Plugins allow the model to access external information sources and tools, expanding its capabilities for perception and action. This resembles the evolution of the bicameral mind described by Jaynes—from a simple bicameral structure to a more integrated system capable of more complex interaction with the environment.

The analysis of these three leading LLMs through the lens of the bicameral mind theory shows that, despite differences in implementation, they all demonstrate a dual-component structure similar to that proposed by Jaynes. This confirms the universality of the dual-process paradigm for AI systems of various types and highlights the potential of the bicameral mind theory as a conceptual framework for understanding and developing future AI systems.

Moreover, the evolution of these models, with continuous improvement in reasoning quality, expansion of perception modalities, and capabilities for interaction with external systems, can be viewed as parallel to the evolution of the human mind described by Jaynes—from a simple bicameral structure to a more integrated and adaptive system capable of complex interaction with the environment.

9. Discussion

9.1. Significance of the Results, Limitations of the Study, and Future Strategies for Overcoming Them

The results of our study provide compelling evidence for the existence of a meaningful connection between the principles of the bicameral mind theory and the functional mechanisms of reinforcement learning (RL) in AI systems. By demonstrating the parallels between the dual-process structure of the bicameral mind and the observation–action cycle of RL, as well as the robustness of RL agents to increasing levels of input information, we have shed light on the potential for the bicameral mind theory to inform the development of more adaptable and efficient AI systems. These findings have significant implications for the field of AI, as they suggest that drawing inspiration from the principles of the bicameral mind theory could lead to the creation of AI agents that are better equipped to handle complex, real-world environments and to generalize their learning across a wide range of tasks.

The main limitation of our study lies in the heavy reliance on established benchmark datasets and traditional performance measures that could fail to capture the complex cognitive dualities implied by bicameral theory. Additionally, our empirical evaluation

is limited to a set of influential AI systems and settings. As a result, other experimental verification by means of special models based on bicameral theory will be required to justify stronger generalizability claims.

However, it is important to acknowledge the limitations of our study and the need for further research to fully realize the potential of the bicameral mind theory in AI. One key limitation is the focus on a specific set of RL models, namely CoinRun and RainMazes, developed by OpenAI. While these models provide valuable insights into the functioning of RL agents, they represent a relatively narrow slice of the diverse landscape of AI systems. To strengthen the generalizability of our findings, future research should examine a broader range of RL models and environments, as well as other types of AI systems, such as those based on unsupervised or supervised learning. Additionally, our study relied primarily on qualitative comparisons between the principles of the bicameral mind theory and the functional mechanisms of RL. To provide a more rigorous and quantitative assessment of the relationship between these two domains, future research could develop formal mathematical models that capture the key features of the bicameral mind theory and use these models to analyze the behavior of RL agents in a more precise and systematic way.

Another significant limitation that we identified during our research is the insufficient attention to the integration of the bicameral mind theory with modern language models (LLMs). Although we expanded our analysis to include parallels between the processes of "thinking" and "writing" in LLMs and components of the bicameral mind, further research is needed to formalize these connections and develop specific architectural solutions based on these parallels.

To overcome these limitations, we propose a conceptual model for integrating bicameral mind principles into the architectures of modern AI systems that combines elements of both RL and LLMs. This model suggests an explicit separation of "perception" and "action" components in the system architecture, with intermediate mechanisms for information integration and decision-making.

9.2. Conceptual Model for Integrating Bicameral Mind Principles into AI Architectures

Our proposed conceptual model is based on the idea of structural separation of "perception" and "generation" components in AI system architectures, with emphasis on mechanisms for their interaction and integration. In the context of RL systems, this model includes the following features:

Perception Module (analogous to the "listening" part of the bicameral mind)—responsible for processing sensory input data, their interpretation, and contextualization. This module includes mechanisms for feature extraction, state representation, and value estimation.

Action Module (analogous to the "speaking" part)—responsible for generating and selecting actions based on processed information. This module includes mechanisms for policy selection, planning, and action execution.

Integration Layer—facilitates interaction between the perception and action modules, managing information flow and decision-making. This layer may include mechanisms for attention, memory, and metacognitive control.

In the context of language models, this structure transforms into the following:

Comprehension Module (extended "thinking")—responsible for analyzing input queries, activating relevant knowledge, and forming an internal representation of the task.

Generation Module (the "writing" process)—responsible for formulating and structuring responses based on the conducted analysis.

Self-reflection Mechanisms—provide evaluation of the quality of generated responses, identification of possible errors or inaccuracies, and adjustment of the generation process.

A key feature of this model is not merely the separation of functions but the establishment of explicit communication mechanisms between components, allowing the system to adapt to complex and changing conditions.

To graphically consolidate the conceptual model introduced in this research, we provide Figure 5, which portrays an integrated bicameral AI design that combines elements of both large language models (LLMs) and reinforcement learning (RL).

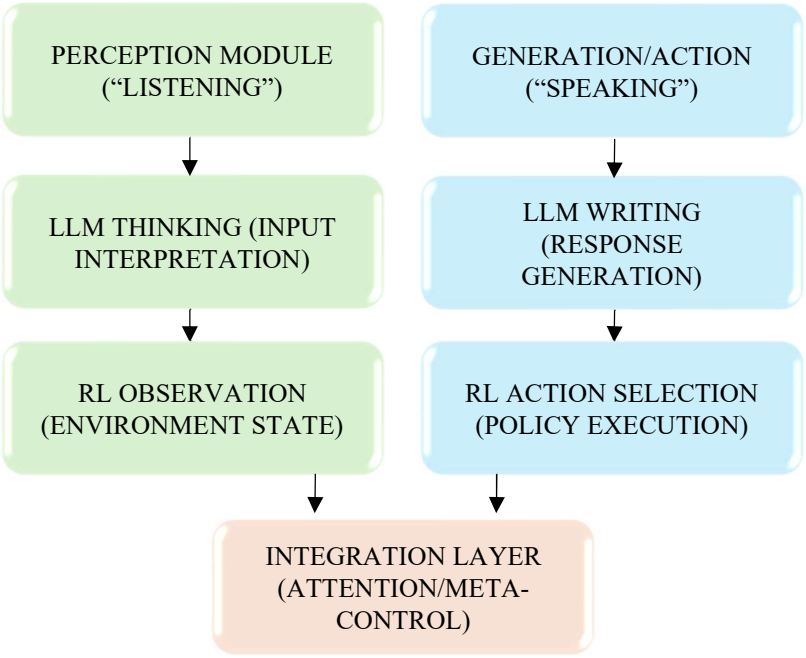


Figure 5. Conceptual diagram of bicameral AI architecture integrating reinforcement learning and language model components. Source: Author's illustration.

The left side of the diagram represents the Perception Module, which is analogous to the "listening" hemisphere under Jaynes' theory. It has two submodules: (i) the RL observation unit, which is responsible for observing states of the world, and (ii) the LLM thinking unit, which processes input queries and context information.

The right part illustrates the Generation/Action Module (the "speaking" counterpart), comprising (i) LLM writing functions, which generate textual responses, and (ii) RL action selection mechanisms, which choose optimal behaviors in response to perceived states.

At the core of the system lies the Integration Layer, the intermediary component that exercises meta-control, manages attention, and coordinates perception and generation processes. This design mirrors Jaynes' bicameral mind-based dual-process model and enables adaptive and scalable behavior in artificial intelligence across modalities.

The diagram also contains a Self-reflection Mechanism, which is shown as a support module involved in evaluating and refining produced outputs, specifically in language model processing.

9.3. Future Research Directions: Prospects for Applying the Bicameral Mind Theory and Other Similar Theories in the Development of AI and RL Technologies

The findings of our study open up a wide range of exciting possibilities for future research at the intersection of the bicameral mind theory, AI, and RL. One promising direction is to explore the potential for incorporating principles of the bicameral mind theory directly into the design of RL algorithms and architectures. For example, researchers could investigate the effects of explicitly modeling the dual-process structure of the bicameral mind within RL agents, with separate modules or pathways corresponding to the "speak-

ing" and "listening" components. Such architectures could potentially allow for more efficient and flexible processing of input information, as well as improved coordination between the perception and action systems of the agent. Additionally, future research could examine the potential for using the bicameral mind theory as a source of inspiration for developing novel RL algorithms that are specifically designed to promote generalization and adaptability, such as by incorporating mechanisms for selectively attending to relevant features of the environment or for dynamically adjusting the balance between exploration and exploitation.

Another important direction for future research is to investigate the potential connections between the bicameral mind theory and other theories of cognition and consciousness, such as the global workspace theory, the integrated information theory, or the predictive processing framework. Each of these theories offers a unique perspective on the nature of the mind and the mechanisms underlying intelligent behavior, and exploring the commonalities and differences between these theories and the bicameral mind theory could lead to a more comprehensive and unified understanding of cognition and its implementation in artificial systems. For example, researchers could examine how the principles of the bicameral mind theory relate to the idea of a global workspace that integrates information from multiple specialized modules, or to the notion of predictive processing as a fundamental mechanism for perception and action. By bringing together insights from these diverse theoretical frameworks, researchers can work toward developing a more complete and coherent account of the mind and its relation to AI, which can inform the design of more sophisticated and human-like artificial agents.

Of particular interest is further investigation of the parallels we identified between the "thinking"/"writing" processes in modern LLMs and components of the bicameral mind. Future research could focus on developing specific architectural solutions for language models that explicitly model the bicameral structure, with separate components for analyzing input data and generating responses, and mechanisms for their efficient interaction.

Moreover, a promising direction is exploring the possibilities of integrating RL and LLM approaches through the lens of the bicameral mind theory. For instance, researchers could develop hybrid architectures where the RL component is responsible for the decision-making process (analogous to the "listening" part), while the LLM component handles explanation generation and communication (analogous to the "speaking" part). Such hybrid systems could combine the strengths of both approaches, providing both efficient decision-making and the ability to communicate and explain their actions.

Another promising direction is the development of a universal theoretical framework for various AI systems based on the dual-process structure. Such a framework could unify RL models, LLMs, and other approaches to AI into a single conceptual structure based on bicameral mind principles, which would allow for more effective comparison and integration of different approaches to artificial intelligence development.

In addition, it would be useful to explore the potential integration of alternative cognitive theories—namely, global workspace theory and predictive processing—into the proposed bicameral-inspired designs. Global workspace theory would be a good starting point for governing distributed subsystems alongside attentional processes, and predictive processing would be able to make the agent better at predicting changes in the local context and adjust accordingly. Exploring the ways in which the two models might interact or complement the dual-process architecture of the bicameral mind might facilitate the construction of more resilient, adaptive, and cognitively plausible AI systems.

10. Conclusions

Summarizing Key Findings and the Potential of the Bicameral Mind Theory in Reinforcement Learning Algorithms for Enhancing AI System Efficiency and Adaptability.

In this study, we have explored the intriguing parallels between the bicameral mind theory and the functional principles of reinforcement learning (RL) in artificial intelligence (AI) systems. By drawing on insights from cognitive science, neuroscience, and AI research, we have developed a novel hypothesis that the dual-process structure of the bicameral mind, with its "speaking" and "listening" components, may have significant implications for the design and optimization of RL algorithms. Our analysis of the CoinRun and RainMazes models, developed by OpenAI, has provided empirical support for this hypothesis, demonstrating that the performance of RL agents is robust to increasing levels of input information, and that the dual-process structure of these models aligns with the principles of the bicameral mind theory. These findings suggest that the bicameral mind theory could serve as a valuable framework for understanding the mechanisms of intelligent behavior in both biological and artificial systems, and for guiding the development of more efficient and adaptable AI technologies.

By extending our analysis beyond traditional RL algorithms, we also investigated parallels between the bicameral mind theory and modern language models (LLMs), such as Claude, Gemini, and ChatGPT. We identified structural similarities between the "thinking" and "writing" processes in LLMs and the "listening" and "speaking" parts of the bicameral mind. This analysis demonstrated that the dual-component structure may serve as a universal principle for organizing various types of AI systems, from RL algorithms to language models.

Based on our analysis, we proposed a conceptual model for integrating bicameral mind principles into the architectures of modern AI systems, which suggests an explicit separation of "perception" and "action" components, with intermediate mechanisms for information integration and decision-making. This model can serve as a foundation for developing more efficient and adaptive AI systems capable of complex interaction with the environment.

The potential implications of our research are far-reaching and exciting. By leveraging the principles of the bicameral mind theory, researchers and developers in the field of AI may be able to create RL algorithms and architectures that are better equipped to handle the complexities and uncertainties of real-world environments. This could lead to the development of AI systems that are more flexible, robust, and capable of generalizing their knowledge across a wide range of tasks and domains. Furthermore, our findings highlight the importance of interdisciplinary collaboration and cross-pollination between the fields of cognitive science, neuroscience, and AI. By combining insights and methods from these diverse disciplines, we can work toward a more comprehensive and unified understanding of the nature of intelligence, and develop artificial systems that more closely resemble the remarkable adaptability and efficiency of the human mind. While there is still much work to perform to fully realize the potential of the bicameral mind theory in AI, our study represents an important step forward in this exciting and rapidly evolving field.

In conclusion, the universal theoretical framework we have proposed for understanding various types of AI systems through the lens of bicameral structure opens new perspectives for integrating different approaches to AI and developing more efficient and universal artificial intelligence systems. We believe that further research and development in this direction can significantly advance our understanding of artificial intelligence and lead to the creation of systems that combine the best qualities of various approaches to AI.

Author Contributions: Conceptualization, M.M. and M.P.; methodology, M.M.; software, M.M.; validation, M.M. and M.P.; formal analysis, M.M.; investigation, M.M.; resources, M.M.; data curation, M.M.; writing—original draft preparation, M.M.; writing—review and editing, M.M. and M.P.; visualization, M.M.; supervision, M.P.; project administration, M.P.; funding acquisition, M.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Jaynes, J. *The Origin of Consciousness in the Breakdown of the Bicameral Mind*, 1st ed.; Houghton Mifflin: Boston, MA, USA, 1976; ISBN 978-0395329320.
2. Botvinick, M.; Ritter, S.; Wang, J.X.; Kurth-Nelson, Z.; Blundell, C.; Hassabis, D. Reinforcement learning, fast and slow. *Trends Cogn. Sci.* **2019**, *23*, 408–422. [CrossRef] [PubMed]
3. Gazzaniga, M. Forty-five years of split-brain research and still going strong. *Nat. Rev. Neurosci.* **2005**, *6*, 653–659. [CrossRef] [PubMed]
4. Gershman, S.J.; Uchida, N. The computational architecture of value-based decision making. *Nat. Neurosci.* **2021**, *24*, 458–466.
5. Sutton, R.S.; Barto, A.G. *Reinforcement Learning: An Introduction*; MIT Press: Cambridge, MA, USA, 2018.
6. Kaelbling, L.P.; Littman, M.L.; Moore, A.W. Reinforcement learning: A survey. *J. Artif. Intell. Res.* **1996**, *4*, 237–285. [CrossRef]
7. Arulkumaran, K.; Deisenroth, M.P.; Brundage, M.; Bharath, A.A. Deep reinforcement learning: A brief survey. *IEEE Signal Process. Mag.* **2017**, *34*, 26–38. [CrossRef]
8. Cobbe, K.; Klimov, O.; Hesse, C.; Kim, T.; Schulman, J. Quantifying generalization in reinforcement learning. *arXiv* **2019**, arXiv:1812.02341.
9. Cobbe, K.; Hesse, C.; Hilton, J.; Schulman, J. Leveraging procedural generation to benchmark reinforcement learning. *arXiv* **2019**, arXiv:1912.01588.
10. Li, Y. Deep reinforcement learning: An overview. *arXiv* **2017**, arXiv:1701.07274.
11. Jaynes, J. Consciousness and the voices of the mind. *Can. Psychol./Psychol. Can.* **1986**, *27*, 128. [CrossRef]
12. Kuijsten, M. *Reflections on the Dawn of Consciousness: Julian Jaynes's Bicameral Mind Theory Revisited*; Julian Jaynes Society: Harlan, IA, USA, 2006.
13. Cavanna, A.E.; Trimble, M.; Cinti, F.; Monaco, F. The "bicameral mind" 30 years on: A critical reappraisal of Julian Jaynes' hypothesis. *Funct. Neurol.* **2007**, *22*, 11–15. [PubMed]
14. Block, N. Review of Julian Jaynes's Origins of Consciousness in the Breakdown of the Bicameral Mind. *Cogn. Brain Theory* **1978**, *1*, 295–306.
15. Jarrold, W.; Yeh, P.Z. The social-emotional basis of creative production in children with and without autism. *Creat. Res. J.* **2016**, *28*, 206–215.
16. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. [CrossRef] [PubMed]
17. Emilio, G. The Conscious Bicameral Mind. In Proceedings of the Conscious Bicameral Mind, Granada, Spain, 7 November 2019. [CrossRef]
18. Baars, B.J. *A Cognitive Theory of Consciousness*; Cambridge University Press: Cambridge, UK, 1988.
19. François-Lavet, V.; Henderson, P.; Islam, R.; Bellemare, M.G.; Pineau, J. An introduction to deep reinforcement learning. *arXiv* **2018**, arXiv:1811.12560.
20. Mill, J.S. *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*; Longmans, Green, and Company: London, UK, 1884; Volume 1.
21. Dema, D.; Fabiha, I.; Zulfikar, A.; Zeyar, A.; Mohammad, A. Reinforcement Learning: A Friendly Introduction. In Proceedings of the International Conference on Deep Learning, Big Data and Blockchain (Deep-BDB 2021), Virtual Event, 23–25 August 2021. [CrossRef]
22. Neftci, E.O.; Averbeck, B.B. Reinforcement learning in artificial and biological systems. *Nat. Mach. Intell.* **2019**, *1*, 133–143. [CrossRef]
23. Rodionov, A.A.; Fayziev, R.A.; Gulyamov, S.S. Experience in using big data technology for digitalization of information. In Proceedings of the 6th International Conference on Future Networks & Distributed Systems (ICFNDS '22), Tashkent, Uzbekistan, 15 December 2022; Association for Computing Machinery: New York, NY, USA, 2023; pp. 412–415. [CrossRef]

24. Majumder, A. Introduction to Reinforcement Learning. In Proceedings of the 2022 2nd International Conference on Technology Enhanced Learning in Higher Education (TELE), Lipetsk, Russia, 26–27 May 2022. [CrossRef]
25. Cao, L.; Min, Z. An Overview of Deep Reinforcement Learning. In Proceedings of the CACRE2019: Proceedings of the 2019 4th International Conference on Automation, Control and Robotics Engineering, Shenzhen, China, 19–21 July 2019; pp. 1–9. [CrossRef]
26. Juliani, A. *Simple Reinforcement Learning with Tensorflow Part 8: Asynchronous Actor-Critic Agents (A3C)*; Medium: San Francisco, CA, USA, 2016.
27. Saidakhror, G.; Rabim, F.; Andrey, R.; Munavvarkhon, M. The Introduction of Artificial Intelligence in the Study of Economic Disciplines in Higher Educational Institutions. In Proceedings of the 2022 2nd International Conference on Technology Enhanced Learning in Higher Education (TELE), Lipetsk, Russia, 26–27 May 2022; pp. 6–8. [CrossRef]
28. Fodor, J.A. *The Modularity of Mind: An Essay on Faculty Psychology*; MIT Press: Cambridge, MA, USA, 1983.
29. OpenAI. GPT-4 Technical Report. 2023. Available online: https://openai.com/index/gpt-4-research (accessed on 10 February 2025).
30. OpenAI. Hello GPT-4o. 2024. Available online: https://openai.com/index/hello-gpt-4o (accessed on 10 February 2025).
31. Anthropic. Introducing the Next Generation of Claude. 2024. Available online: https://www.anthropic.com/news/claude-3-family (accessed on 10 February 2025).
32. Anthropic. Claude 3.7 Sonnet. 2024. Available online: https://www.anthropic.com/claude (accessed on 10 February 2025).
33. Google DeepMind. Gemini 2.5: Our Most Intelligent AI Model. 2025. Available online: https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025 (accessed on 12 February 2025).
34. Google DeepMind. Gemini 2.5 Pro. 2025. Available online: https://deepmind.google/models/gemini/pro (accessed on 15 February 2025).

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.