What Comes Next

Drawing Parallels and Distinctions Between Biological Intelligence and Machine Intelligence

Chapter 1: The Functional Evolution of Prediction and Human Intelligence

Chapter 2: Nature's Functional Blueprint for Prediction

Chapter 3: Bridging Biological Blueprints, Artificial Intelligence, and the Path to Superintelligence

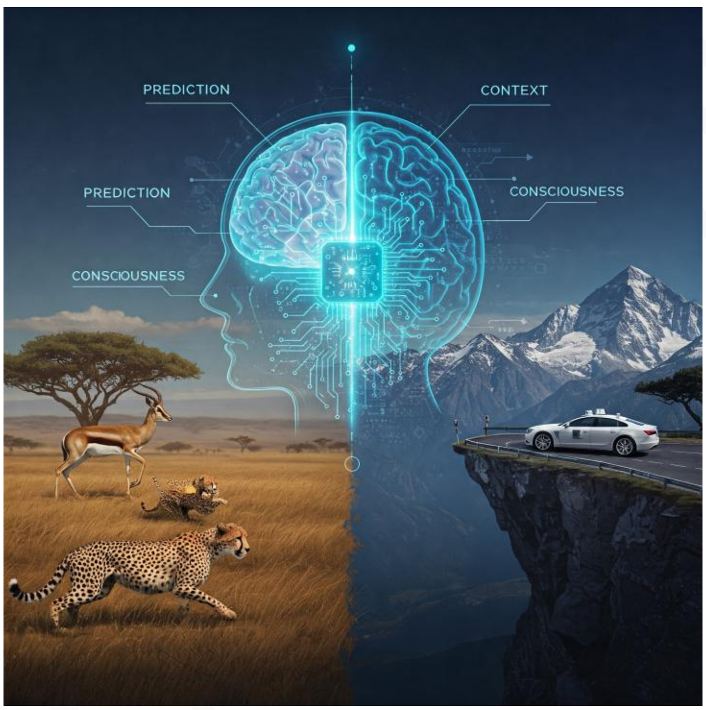

By Sreekanth Pannala[1]

V1: April 28, 2025; V2: May 14, 2025

[1] *With the help of Intelligent Assistant - Gemini Flash Pro 2.5*

Preface

Inspired by E. Schrödinger, *What is Life?* (1944)

Early in my career, moving between engineering and information science, I encountered the common expectation that one should achieve deep expertise in a single field before commenting on others. However, the subject that captivated me – the interplay between biological and artificial intelligence – inherently crosses these boundaries. It calls for a perspective blending insights from physics, chemistry, biology, mathematics, and computation, much like nature itself operates. My own curiosity wouldn't allow me to remain confined to one discipline, compelling me to disregard the "stay in your lane" convention and explain my reasons for attempting this synthesis.

I began in engineering, focused on controlling and optimizing physical processes. This work soon guided me toward information science – the world of computational modeling, mathematics, chaos theory, machine learning, and algorithms. This cross-disciplinary path naturally led me to view nature not just as a subject of study, but as the ultimate engineer. Perhaps this viewpoint began to form during long nights spent soothing a colicky baby, when my thoughts drifted from the immediate task to fundamental questions: How does life originate? How does self-assembly work, starting from a single cell? These reflections connected with my research at the time on the nucleation and growth of carbon nanotubes, exploring the principles of self-assembly.

Later, observing nature's everyday ingenuity reinforced this perspective. I became fascinated watching my dog, who somehow anticipated my wife's arrival in the neighborhood minutes before any car was audible. How did she make that prediction? Observing her meticulously track a scent, processing information in real-time, seemed like a natural form of complex calculation. It underscored how nature, over billions of years, developed incredible solutions for efficiency, sustainability, self-replication, information transfer, and the intelligence needed for survival – all without adhering to our distinct academic disciplines.

This perspective – shaped by bridging engineering and information science, reflecting on life's origins, and observing nature's integrated approach – fueled a strong desire to understand intelligence more holistically. Einstein noted that science aims to "coordinate our experiences and bring them into a logical system," and that concepts are justified only if they help represent those experiences. That resonated with me. Considering the rapid progress in biology, neuroscience, computer science, and psychology, alongside the need for systems thinking to navigate today's world (as discussed by Friedman in Thank You for Being Late), it feels timely to try and connect the dots regarding biological intelligence.

But how does one approach such a task? The desire for a unified understanding conflicts with the impossibility of mastering every relevant specialty. The only path forward seemed to be to simply attempt the synthesis. My goal is to weave together facts and theories from these diverse fields, constructing a model based on this integration. I recognize that any such model is just a representation of current understanding, not a final truth. As Thomas Kuhn described, new discoveries inevitably lead to paradigm shifts. Science isn't about absolute certainty, but it's our best method for coordinating experience and navigating complexity.

Therefore, this work is my attempt at synthesis. I am driven by a need to understand the functional imperatives behind biological intelligence, hoping this understanding can contribute positively, while remaining aware of the risks associated with powerful technologies. It's a deliberate effort to view intelligence through the specific lens of my experiences. I accept that my knowledge in some areas is necessarily incomplete and that I risk oversimplification.

"So much for my apology." Ultimately, the process of inquiry itself is valuable. It's rewarding to be curious and engage with these fundamental questions.

Summary of "What Comes Next"

This three-part monograph proposes a framework called P→C→C (Prediction → Context → Consciousness) to understand both biological and artificial intelligence from an evolutionary and functional perspective. The core idea is that intelligence, driven by the need for survival, fundamentally relies on the ability to predict "what comes next".

- **Chapter 1**: The Functional Evolution of Prediction and Human Intelligence introduces the P→C→C hierarchy.
  - Prediction (Level 1): Basic, rapid predictions based on sensory input guide immediate actions.
  - Context (Level 2): Accumulated predictions and actions build embodied, experiential context, allowing for more nuanced understanding and prediction.
  - Consciousness (Level 3): Deeply ingrained context leads to an integrated functional awareness geared towards survival and procreation (e.g., assessing danger, finding mates).
  - Higher cognitive functions like creativity and morality are presented as consequences emerging from this core survival-driven system.
  - The article contrasts this with Large Language Models (LLMs), noting their sophisticated prediction (Level 1 analogy) but highlighting their limitations due to a lack of embodied interaction, grounded context formation (Level 2), and integrated functional awareness (Level 3).
- **Chapter 2**: Nature's Functional Blueprint for Prediction explores how biology implements the P→C→C hierarchy with remarkable energy efficiency (~20W).
  - It discusses nature's design principles: functional reuse of components, cyclical prediction based on biological clocks, gradient detection (differentiation), information integration over time, and efficient feature extraction (with analogies to linear algebra).
  - It examines potential biological mechanisms like specific neural circuits, synaptic plasticity (LTP/LTD), hierarchical memory systems (including the role of sleep in consolidation), complex dendritic computations, and the efficiency benefits of synaptic computing (co-localizing processing and memory).
- **Chapter 3**: Bridging Biological Blueprints, Artificial Intelligence, and the Path to Superintelligence connects the biological insights to AI.
  - It analyzes current AI (LLMs, neuromorphic computing) through the P→C→C lens, further detailing the functional gaps compared to biology, particularly regarding grounded context, integrated awareness, and energy efficiency.
  - It suggests ways AI could be enhanced by incorporating principles from biology, such as predictive coding, embodiment, bio-inspired memory systems, and synaptic computing architectures.
  - Finally, it speculates on the trajectory towards Artificial General Intelligence (AGI) and Superintelligence (ASI), considering potential capabilities, societal impacts, and the critical challenge of alignment.

In essence, the monograph argues that understanding intelligence requires focusing on its core evolutionary function – prediction for survival – and the efficient, embodied mechanisms nature developed, offering this P→C→C framework as a lens to analyze biological systems and guide future AI development.

Chapter 1. The Functional Evolution of Prediction and Human Intelligence

1.1 Introduction

The remarkable success of Large Language Models (LLMs), despite their occasional inconsistencies like "hallucinations," is undeniable and driving rapid advancements in artificial intelligence. Yet, as powerful as these LLMs are at pattern matching and generating plausible content, they still struggle with robust common sense, deep contextual understanding, and capabilities associated with integrated context and consciousness (Lake et al., 2017; Mitchell, 2019). They excel at sophisticated sequence prediction based on learned correlations in data, but this differs fundamentally from the embodied, goal-driven intelligence shaped by biological evolution (Shapiro, 2019; Korteling et al., 2021). Simultaneously, human understanding of the intricate biological circuitry and mechanistic "how" of the human brain, while progressing, remains elusive – arguably still only scratching the surface of how billions of neurons give rise to intelligence. Furthermore, the lack of formal mathematical frameworks to describe intelligence or fully characterize AI capabilities poses significant challenges for rigorous assessment and targeted development toward more general intelligence (Legg & Hutter, 2007).

This confluence of challenges suggests the value of exploring a new angle of attack. This article is the first part of a three-part series exploring the connection between biological and machine intelligence by taking a fresh look: focusing not primarily on the intricate biological circuitry or empirical benchmarks alone, but on the functional evolutionary imperative that may have driven intelligence in the first place – the need for survival and procreation. It is argued here that a core principle enabling this survival was, and remains, the ability to predict "what comes next".

The primary objective of this paper (Chapter 1) is to propose and explore a functional evolutionary hierarchy (P→C→C) as a framework for understanding biological intelligence. The central hypothesis is that human intelligence operates via this hierarchy rooted in prediction: Prediction based on sensory input governs actions, accumulating these predictions and actions creates embodied Context, and sufficient ingrained context contributes to Consciousness. This functional hierarchy – Prediction → Context → Consciousness (P→C→C) – is hypothesized to have arisen from the need for survival, with celebrated higher faculties emerging as potential consequences. This perspective aligns with broader theories of embodied and enactive cognition, which emphasize the coupling of organism and environment through action (Varela et al., 1991; Shapiro, 2019). Examining intelligence from this functional, evolutionary perspective, which posits a core predictive mechanism potentially underpinning even complex outputs like human language, may offer insights complementary to mechanistic studies and benchmark-driven approaches. This functional model will be explored, drawing parallels and critical distinctions with the predictive capabilities of LLMs, all while noting the astonishingly low power (approximately 15-20 watts) at which the biological brain performs these complex functions (Korteling et al., 2021).

Chapter 2 will delve into potential biological circuitry enabling this functionality and the efficiencies of evolutionary design, seeking inspiration from nature's non-siloed solutions. Chapter 3 will explore how these functional and mechanistic insights might inform the development of more efficient and intelligent AI, discussing the potential for superintelligence unbound by biological constraints, and how to navigate the risks while pursuing opportunities beyond nature's original reach.

1.2 The Evolutionary Imperative: Predicting "What Comes Next" for Functional Needs

At its heart, it is proposed that all biological intelligence originates from a foundational imperative: the ability to predict "what comes next". This drive appears paramount for survival and propagation, potentially emerging from fundamental thermodynamic and kinetic principles that favor persistent, self-replicating systems under non-equilibrium conditions (England, 2013; Pross, 2011; Pascal and Pross, 2022). Organisms that could anticipate future events – locating energy sources, securing food, finding mates, avoiding predators, and adapting to environmental changes – likely held a distinct evolutionary advantage (Shettleworth, 2010). From the very emergence of life, the capacity for anticipation, driven by the need to maintain stability and persist, may have been a driving force, representing the primary selective pressure that shaped early nervous systems and cognitive abilities. Focusing purely on the functional requirement, intelligence, in its earliest forms, likely evolved simply to enable organisms to perform actions that increased their chances of staying alive and reproducing. It was perhaps not initially about understanding the cosmos or contemplating philosophy, but about the functional outcome of predicting where the next meal might be or when a predator might strike. This basic predictive function, driven by immediate sensory input and tied to action within an environment (Varela et al., 1991), governed the organism's immediate actions.

Considering the early history of life on Earth helps illustrate this functional imperative. As detailed in works like "First Life" (Attenborough, 2010), the story begins with single-celled organisms billions of years ago. Even in extreme conditions, survival depended on sensing the environment and reacting in anticipation of nutrient availability or harmful conditions – a basic predictive function guiding action and tied to biological cycles detailed in a later section. The evolution of multicellularity allowed for greater specialization and coordinated responses, enabling more complex interactions with the environment and better anticipation of resource distribution.

Later evolutionary leaps, such as the development of mobility in creatures like Dickinsonia and Kimberella around 550 million years ago (Bobrovskiy et al., 2022), dramatically increased predictive capabilities. The ability to move allowed organisms to actively seek food sources or flee from danger, requiring an anticipation of where resources might be found or threats might emerge. The subsequent evolution of bilateral symmetry, concentrating sensory organs and a rudimentary nervous system in a "head" region (as in Sprigina), further enhanced the capacity for directed movement and sophisticated anticipation of the immediate future. This increasing biological complexity and the development of nervous systems were intrinsically linked to the growing need and capacity to predict environmental changes specifically for the functional purposes of survival and successful reproduction (Dicke & Roth, 2008; Shettleworth, 2010). These fundamental predictive abilities, rooted in ancient, survival-driven functional needs, potentially form the base layer of the proposed hierarchical system governing actions. This section examines what capabilities may have evolved and why they provided a functional advantage, leaving the details of the biological mechanisms for Chapter 2.

1.3 The Predictive Brain: A Proposed Functional Hierarchy of Intelligence

This evolutionary journey highlights that anticipation is not merely a cognitive function but may represent a fundamental biological imperative tuned by the functional demands of survival. As life became more complex, so did the mechanisms for prediction, culminating in the sophisticated predictive capabilities of the human brain. The hypothesis presented here suggests that human intelligence operates through a functional hierarchy rooted in prediction, viewing the mind as fundamentally embodied and situated (Shapiro, 2019; Varela et al., 1991):

- Prediction (Level 1): At the most fundamental level, actions are proposed to be governed by the immediate prediction of "what comes next" based on real-time, embodied sensory input. This rapid, often subconscious prediction represents the brain's first response to stimuli, guiding immediate motor output and reactions, constantly minimizing the error between prediction and incoming sensory reality (Kimura, 2021; Friston, 2005). This is presented not just as pattern matching, but as an active process tied to physical interaction and survival goals.
- Context (Level 2): As sufficient possibilities are sampled through Level 1 predictions and the resulting actions/results within the environment, this experiential information accumulates, allowing for the creation and understanding of context. Contextualizing incoming sensory data and potential actions provides a richer framework for prediction, enabling more nuanced responses beyond immediate reflexes (Summerfield & de Lange, 2014). Actions are then performed and interpreted within this established, embodied, and dynamic context (Shapiro, 2019). This includes predicting the actions and intentions of others within social contexts, crucial for navigating complex social environments (Whiten & Byrne, 1997). This context is built from the ground up through interaction, not merely inferred from static data.
- Consciousness (Level 3): Once enough actions have been processed within various contexts, this information becomes deeply ingrained, contributing to an integrated functional awareness geared towards survival and procreation. This includes the capacity to generalize from accrued information, identify danger (sense of fear) versus safety (sense of calm), distinguish friend from foe, recognize kin, differentiate food from poison, assess sexual attractiveness, find mates, and act on instincts (maternal, paternal, gut feelings) essential for progeny survival and continuation of the life cycle. Consciousness, within this functional model, is conceptualized as the highest emergent level arising from the P→C→C hierarchy, potentially representing a continuous simulation built from the vast history of predictions and experientially grounded contexts relevant to these core drives (Hohwy, 2013; Clark, 2015).

This proposed functional hierarchy of decision making – Prediction → Context → Consciousness – appears to govern human actions, from the simplest reflex to the most complex strategic planning. This entire intricate process, performed by the biological brain, operates at an amazingly low power of approximately 15-20 watts (Attwell & Laughlin, 2001), a remarkable feat of functional efficiency compared to artificial systems (Korteling et al., 2021).

The simple yet profound question "what comes next?" permeates virtually every functional level of human existence, driving the proposed Prediction → Context → Consciousness hierarchy:

- Level 1 Prediction: Anticipating the impact when touching a hot surface (sensory input → action prediction), predicting the trajectory of a thrown object for immediate reaction (visual input → motor prediction).
- Level 2 Context: Understanding the meaning of a conversation based on the preceding sentences and social cues (accumulated linguistic/social predictions → semantic context), navigating a familiar environment by predicting what lies around the next corner within the context of the known location or the predator-prey chase described later (spatial/situational predictions → environmental context).
- Level 3 Consciousness: Reflecting on past experiences (integrated contexts) to plan future actions, engaging in abstract problem-solving by simulating hypothetical scenarios – functions enabled by the conscious mind operating on ingrained predictions and contexts to navigate potential futures.

This pervasive drive to anticipate future states, driven initially by the need for favorable functional outcomes for survival, likely forms the basis for building experientially grounded context and ultimately gives rise to consciousness. Higher-level reasoning, questioning, and abstract thought can thus be seen as potential consequences or sophisticated applications of this functional hierarchy, allowing for the simulation of more complex functional scenarios far removed from immediate survival needs.

1.4 A Natural Illustration: The Predator-Prey Chase

To illustrate the vital link between this functional hierarchy and the primal need for survival, a classic, high-stakes drama playing out in nature serves as an example: the predator-prey chase. Consider a cheetah hunting a herd of gazelles on the savanna. Both predator and prey appear to operate their functional intelligence hierarchies at peak performance, with life or death hanging on the speed and accuracy of their predictive loops.

- Initial State (Stalking/Grazing):
    - Cheetah: From cover, the cheetah is making Level 1 Predictions based on sensory input: the wind direction (will it carry scent?), the position of the sun (visibility), the stiffness of a gazelle's posture (alertness). These predictions build Level 2 Context: the layout of the terrain (obstacles, open ground), the size and composition of the herd, and critically, identifying a potential target by predicting vulnerability based on subtle cues - perhaps a younger animal lagging slightly, an older one with a less springy gait, or one positioned further from the protective

    center of the herd. This contextual information contributes to the cheetah's focused predatory Level 3 Consciousness: its awareness of the hunting opportunity, its assessment of the best chance for a kill, and its ingrained drive to capture prey.
    - Gazelles: The herd is grazing. Individual gazelles make Level 1 Predictions from peripheral vision and sound: a rustle in the grass, a shift in the behavior of a neighboring gazelle. These inputs build Level 2 Context: the overall relaxed state of the herd, awareness of the open, relatively safe environment. This contributes to their generally calm Level 3 Consciousness: alert for danger but focused on feeding. Mothers may be subtly predicting the movements of their young, keeping them close.
- The Attack and Initial Reaction:
    - Cheetah: The cheetah bursts from cover, initiating the chase towards its chosen target. Its brain is now making ultra-rapid Level 1 Predictions: the exact acceleration needed, the line to intercept the chosen gazelle's initial flight path.
    - Gazelles: The sudden visual detection of the cheetah triggers immediate, instinctual Level 1 Predictions of extreme danger. Action: Scatter! The brain prioritizes the raw sensory prediction of threat over the previous relaxed context. The Level 2 Context shifts instantaneously: danger is present, its location, escape is paramount. Mothers within the herd quickly update their context to include the location and state of their young, initiating actions to shield them or guide them towards safer escape routes, predicting which direction offers the best chance for the fawn to survive. Level 3 Consciousness floods with fear and the urgent will to survive, intensified in mothers by the drive to protect offspring.
- The High-Speed Chase: This is where the dynamic, continuous cycle of Prediction, Context, and Consciousness is most critical and rapid for the targeted gazelle and the pursuing cheetah.
    - Gazelle: Every millisecond, the gazelle's eyes feed information to its brain. It makes split-second Level 1 Predictions: the cheetah's precise angle and closing speed, the terrain directly ahead. Based on these predictions and its rapidly updating Level 2 Context (its own fatigue, the cheetah's apparent effort, obstacles, the location of safety or other herd members, or a mother attempting to guide a fawn), it makes an action prediction: jink left? jink right? maintain a straight line? This prediction immediately governs its action (a sharp turn). This new action and the resulting sensory feedback then update its Level 2 Context (distance gained or lost, cheetah's reaction to the jink, or the proximity of its young), which informs the next set of Level 1 Predictions. Its Level 3 Consciousness is intensely focused on escape, fueled by survival instinct and integrated knowledge of past chases (deeper context), with the mother's awareness acutely focused on the vulnerability of her offspring.
    - Cheetah: Similarly, the cheetah is locked into its own loop. It makes Level 1 Predictions about the gazelle's trajectory and potential jinks based on its body language and muscle tension. These predictions, integrated into the complex Level 2 Context of the chase (distance, speed differential, terrain, the target's

apparent fatigue or separation from others), inform its action prediction: adjust stride, lean into a turn, prepare for a final lunge. This action generates new sensory input, updating its context for the next predictive cycle. Its predatory Level 3 Consciousness is fully engaged, aware of the goal, assessing its own physical limits, driven to make the kill, potentially adjusting strategy if the targeted gazelle proves unexpectedly agile or if a more vulnerable one becomes accessible.

- Outcome: The chase is a continuous feedback loop. The gazelle's survival depends on its ability to make unpredictable jinks – actions that defy the cheetah's Level 1 Predictions and force the cheetah to constantly update its Level 2 Context, hopefully slowing it down, or leading it into unfavorable terrain. The cheetah's success depends on its ability to make accurate predictions about the gazelle's evasive maneuvers and adapt its actions faster than the gazelle can update its context and predict a safe escape route. The subtle predictions about vulnerability made by the cheetah and the protective actions predicted and taken by a mother gazelle, all operating within their respective functional hierarchies, are crucial elements in this life-or-death struggle.

This life-or-death chase provides a powerful, observable example of how the functional hierarchy of Prediction, leading to Context, and integrated by Consciousness, may function not as an abstract concept but as the fundamental engine of survival-driven intelligence, incorporating complex factors like assessing vulnerability and social bonds, all potentially honed by millions of years of evolutionary pressure at an astonishingly low energy cost. It illustrates the complex, dynamic interplay required just to navigate a hostile world, bringing the functional imperative vividly to life directly from the observation of nature.

1.5 From Savanna to Street: The Predictive Driver

Moving from the life-or-death stakes on the savanna to the complex, dynamic environment of human transportation, the same fundamental functional hierarchy of Prediction → Context → Consciousness appears absolutely critical. While the goals shift from capturing prey or escaping predation to safely navigating roads, the underlying cognitive processes, rooted in embodied interaction (Shapiro, 2019), seem remarkably similar. Human driving, especially in challenging conditions, provides another compelling illustration of this evolved intelligence at work.

Consider the intricate task of driving, not on a controlled highway, but in environments that demand continuous, high-stakes prediction and adaptation – the narrow, unpredictable roads of the Himalayas, the chaotic, dense traffic of urban India, or the high-speed, precision requirements of a race track.

- Prediction (Level 1): A skilled driver is constantly making rapid, often subconscious Level 1 Predictions based on immediate sensory input: predicting when the car ahead will brake based on subtle cues, predicting the path of a pedestrian darting out, predicting the grip of tires on a wet patch, predicting the angle of a turn, predicting how their vehicle will respond to a steering input (all rapid, sensory-driven). These are split-second, highly dynamic predictions analogous to the gazelle predicting the cheetah's jink or the cheetah's lunge.

- Context (Level 2): These constant Level 1 predictions feed into building a complex, dynamic Level 2 Context: understanding the flow of traffic around the vehicle (not just immediate threats), assessing the overall road conditions (ice, potholes, gravel), knowing the car's capabilities, anticipating potential hazards (school zones and timings, kids playing on the road or driveways, drunken or erratic driving while texting), understanding traffic rules (even if they are chaotic, the context is the pattern of chaos), knowing the route (analogous to understanding terrain, herd behavior, fatigue). This context is built through continuous interaction and experience (Varela et al., 1991).
- Consciousness (Level 3): The driver's Level 3 Consciousness integrates these predictions and contexts to make higher-level decisions and maintain overall control: planning the best lane changes, choosing when to overtake, deciding how to safely navigate a tricky intersection, managing fatigue, responding consciously to unexpected situations (a sudden detour, an emergency vehicle). Consciousness maintains focus, assesses risk tolerance, and allows for strategic adaptation based on ingrained experience and learned rules – analogous to the survival drive, predatory focus, or fear. The "feel" for the road and the vehicle often operates at this level, representing a deeply ingrained, consciously accessible context built from countless hours of prediction and action.

Highly proficient human drivers are arguably not simply following rules; they function as master predictors and context-builders, leveraging the same functional hierarchy of intelligence that may have allowed human ancestors (and other animals) to navigate complex, unpredictable environments for survival. The ability to handle the extreme variables of a Himalayan road, the rapid, multi-agent interactions of Indian traffic, or the precise limits of physics on a race track likely stems directly from the evolved capacity of human brains to perform robust, multi-layered functional prediction within a dynamic context, culminating in aware, adaptive action – all within that remarkably efficient 15-20 watt power budget.

Connecting Evolutionary Intelligence to Autonomous Driving

This perspective offers a potential insight into the challenge of achieving Level 5 autonomous driving – the ability for a vehicle to operate safely in any environment, under any conditions, without human intervention. Current AI systems excel at specific Level 1 Prediction tasks (identifying objects, predicting their immediate trajectory based on sensor data) and can operate effectively in relatively constrained Level 2 Contexts (well-mapped roads, clear weather, predictable traffic flows defined by data).

However, they notoriously struggle when confronted with the truly chaotic, the ambiguous, or the unprecedented – situations requiring the flexible, rapid, and deeply experiential Level 2 Context building and the robust, adaptive decision-making associated here with Level 3 Consciousness in humans. The challenge lies not just in prediction accuracy but in building the right kind of context – one grounded in interaction and capable of handling true novelty, not just variations within training data (Lake et al., 2017; Mitchell, 2019).

According to the argument presented here, truly unlocking Level 5 autonomy may require more than just faster processing or bigger datasets; it likely involves understanding the functional principles and potentially the efficient mechanisms by which biological brains achieve such robust, general-purpose Prediction → Context → Consciousness processing, built upon embodied interaction (Shapiro, 2019; Varela et al., 1991). Studying this "evolutionary intelligence" – not just the what (functional hierarchy) but also the how (biological circuitry) and crucially, the why (survival drive, embodiment) – might offer insights needed to design AI systems capable of building rich, dynamic context from prediction and making highly adaptive decisions in complex, real-world environments that current systems find challenging. This motivates the exploration of the functional journey detailed in this part as a necessary prelude to exploring the mechanisms in Chapter 2 and the AI implications in Chapter 3.

1.6 From Street to Pitch: The Predictive Athlete

The P→C→C hierarchy also manifests profoundly in elite athletic performance, where rapid prediction and context integration are paramount. Consider the example of a world-class soccer player like Lionel Messi executing a drive towards the goal. His seemingly magical abilities on the pitch can be analyzed through this functional lens, illustrating the interplay between prediction, context, and conscious action in a high-stakes, dynamic environment.

- Messi's P→C→C:
  - *Prediction (Level 1):* As Messi dribbles, his brain is constantly making subconscious predictions: the ball's trajectory with each touch, the closing speed and angle of the nearest defender, the subtle weight shifts of opponents indicating their next move, the position of the goalkeeper, the optimal foot placement for the next touch or potential shot. His mastery allows the ball to function as an extension of his body, freeing cognitive resources from basic motor control to focus on these environmental and social predictions.
  - *Context (Level 2):* Through continuous scanning (reportedly higher than typical forwards) and integrating countless Level 1 predictions, Messi builds a rich, dynamic context: the spatial layout of all players (teammates and opponents), the defensive formation he's facing, passing lanes to teammates making runs, the distance to goal, the score and time remaining, and perhaps even the known tendencies of the specific defenders and goalkeeper he's confronting. This internal "bird's-eye view" is constantly updated.
  - *Consciousness (Level 3):* Messi's functional awareness integrates this stream of predictions and context to make high-level decisions. Approaching the penalty area, he assesses: "Given the defender's predicted lunge and the goalkeeper's predicted position, is a shot to the far post the highest probability action? Or does the predicted run of my teammate into space offer a better opportunity via a pass?" Even when double-teamed, his awareness of the broader context allows him to predict an opening and execute a seemingly impossible pass or shot. The final action – the feint past the defender, the precisely weighted shot – is the output of this integrated P→C→C process.
- Defender's P→C→C:

    - *Prediction (Level 1):* The defender facing Messi predicts his immediate next touch, potential direction change based on hip movement, and the likelihood of a shot versus a pass based on proximity and body shape.
    - *Context (Level 2):* The defender's context includes Messi's known repertoire of moves, their own defensive positioning relative to the goal and supporting defenders, the overall defensive strategy, and Messi's history in similar situations.
    - *Consciousness (Level 3):* The defender integrates these to decide: "Do I commit to a tackle now, predicting his next touch, or do I try to contain him, predicting he might pass?" Their action (e.g., lunging tackle, holding position) is based on this conscious prediction within their contextual understanding. Often, Messi's superior prediction and context-building leads the defender to predict incorrectly.
- Goalkeeper's P→C→C:
    - *Prediction (Level 1):* The goalkeeper predicts the ball's trajectory off Messi's foot, the angle of the shot, and the power, based on Messi's body mechanics and approach.
    - *Context (Level 2):* Their context includes Messi's preferred shooting spots and techniques, the positioning of the defensive wall (if any), the distance, and the current defensive screen.
    - *Consciousness (Level 3):* The goalkeeper decides where to position themselves, when to commit to a dive, and which direction to anticipate based on integrating the immediate predictions (shot angle, power) with the broader context (player tendencies, game situation). Their save attempt is the action resulting from this conscious integration.

This constant, high-speed interplay of prediction, context-building, and conscious decision-making by multiple agents, each operating within their own P→C→C loop, creates the dynamic tension and anticipation that makes watching elite sports so compelling. The path to mastery in any complex skill domain, from athletics (tennis, cricket, basketball, boxing, gymnastics) to driving or even playing a musical instrument, involves relentless training. This training serves to automate Level 1 predictions and actions (making the ball or instrument an extension of the body), build richer and more accurate Level 2 context through experience, and refine the speed and quality of Level 3 conscious decision-making based on that integrated understanding. In team sports, this expands to predicting and synchronizing with teammates while anticipating opponents, adding another layer of complexity to the P→C→C interplay.

1.7 Prediction and the Cosmos: From Astronomy to Astrology

Moving beyond immediate environmental navigation and interaction, the human drive to predict also extended to the grand, seemingly immutable patterns of the cosmos. Questioning the movements of planets, the fixed locations of stars in constellations, and the appearance of comets represented an early, large-scale application of predictive minds. From these observations, two distinct, yet initially intertwined, paths emerged, both rooted in the desire to predict "what comes next" in the celestial sphere and its potential impact:

- Scientific Prediction (Astronomy): The careful, systematic observation and recording of celestial bodies led to the development of early astronomy. This involved using Level 1 Predictions (observing a planet's position night after night) to build Level 2 Context (understanding orbital patterns, cycles, and relationships). This contextual understanding allowed for increasingly accurate Level 3 Conscious predictions of future celestial events: eclipses, planetary conjunctions, and the return of comets. This predictive success formed one of the earliest foundations of scientific inquiry, demonstrating the power of observation and logical system-building to understand and predict the physical world.
- Non-Scientific Prediction (Astrology): Simultaneously, the human mind sought meaning and prediction beyond mere physical movement. The positions and patterns of celestial bodies were interpreted as signs or influences on earthly events and human fates. This involved using Level 1 Predictions (observing star configurations) to build a different kind of Level 2 Context (assigning symbolic meanings, relationships between celestial and terrestrial events). This contextual framework was then used to make Level 3 Conscious predictions about future human affairs. While lacking empirical validation, astrology served a functional role in early societies, providing a framework for anticipating the future, influencing beliefs, and shaping religions as humans looked towards the beautiful, mysterious sky for guidance and predictability in their lives.

Both astronomy and astrology, despite their vastly different outcomes in terms of scientific validity, originated from the same fundamental human functional imperative to predict based on observed patterns. They highlight the flexibility of the predictive brain and its capacity to build context from sensory input, whether that context is a physical model of orbital mechanics or a symbolic system of cosmic influence, both feeding into conscious understanding and guiding actions or beliefs. These early practices underscore how deeply ingrained the predictive drive appears to be in the human mind, forming foundations for both rigorous science and complex belief systems as humans sought to bring the vastness of the cosmos into a predictable, understandable system.

1.8 Prediction, Context, and Consciousness in Human Experience and Art: Leveraging the Functional Hierarchy

The innate human functional predictive capacity, potentially honed by evolutionary pressures for survival and organized into the proposed Prediction → Context → Consciousness hierarchy, is not limited to basic survival; it may form the very foundation of how humans engage with the world and create meaning. Cultural creations, from language itself to complex art forms, can be viewed as profound manifestations of this functional hierarchy at work. They demonstrate how humans leverage their predictive engine and capacity for context-building to shape shared experiences. The universality and enduring power of great art across cultures and time may lie precisely in its ability to tap into and expertly manipulate these fundamental, shared functional processes of the human brain.

An examination reveals how this hierarchy might operate in the creation and appreciation of various art forms:

- Music Composition and Appreciation: Music is a masterclass in manipulating prediction, context, and consciousness.
    - *Creation:* A composer makes sequential choices - selecting the next note, chord, rhythm, or melodic phrase. This involves predicting how different choices will sound and feel relative to what has come before and anticipating the desired effect on the listener.
    - *Appreciation:* As listeners, humans constantly make Level 1 Predictions about the next note, chord, or beat based on learned understanding of musical scales, harmonies, and rhythms (Koelsch et al., 2019). These sequential predictions build the Level 2 Context of the musical structure, genre, and emotional landscape of the piece. The successful fulfillment or artful violation of these predictions contributes to the emotional experience and understanding processed by Level 3 Consciousness.
    - *Universality & Endurance:* Music that resonates deeply does so because it skillfully plays with these universal predictive expectations and context-building processes in a way that feels both engagingly predictable and refreshingly novel to consciousness. Masterful manipulation of tension and release through harmonic or rhythmic prediction speaks directly to the brain's core functional processes, allowing such music to transcend cultural boundaries and resonate across generations, as experienced by listeners.
- Storytelling (Writing, Literature, Theatre): Narratives are structured experiences designed to guide the predictive mind through a sequence of events and ideas.
    - *Creation:* A writer crafts sentences, paragraphs, plot points, and character arcs, predicting how each choice will influence the reader's understanding and anticipation of what comes next. He or she builds the narrative world and rules that form the foundation for prediction.
    - *Appreciation:* As readers, humans engage in Level 1 Prediction about the next word, sentence, or immediate event. These predictions rapidly build the Level 2 Context of the story's plot, setting, characters, and themes (Ryan, 2003). Authors manipulate this context through foreshadowing, suspense, and plot twists, deliberately fulfilling or subverting predictions. The integration of this complex, contextualized narrative is processed by Level 3 Consciousness, leading to empathy, suspense, reflection, and a sense of meaning.
    - *Universality & Endurance:* Enduring stories tap into universal human contexts – relationships, conflict, triumph, and loss – and structure them using predictive sequences that human brains are inherently designed to process. Their lasting power comes from creating compelling contexts and guiding conscious experience through expertly managed predictive journeys that resonate with fundamental ways of understanding the world.
- Movie Making and Viewing: Film leverages multiple sensory streams to create a rich, predictive experience.
    - *Creation:* A filmmaker makes deliberate choices about camera angles, editing cuts, sound design, music, and dialogue, predicting how these combined sensory

inputs will influence the viewer's immediate expectations and build the overall narrative and emotional context.

- *Appreciation:* As viewers, humans simultaneously make Level 1 Predictions based on visual cues (what will enter the frame, where will the camera move) and auditory cues (when will the music swell, what is that sound off-screen). These multisensory predictions build the dynamic Level 2 Context of the scene, tension, character intentions, and plot progression (Smith, 2013). The integrated, contextualized sensory experience is processed by Level 3 Consciousness, resulting in powerful emotional responses, immersion, and narrative comprehension.
- *Universality & Endurance:* Successful films resonate because they intuitively understand how to guide the universal human multisensory predictive system and build contexts that are emotionally and cognitively engaging. Masterful editing, cinematography, and sound design align with the brain's functional architecture, allowing certain films to remain impactful across technological and cultural shifts.

Across print, audio, and visual media, the consistent element appears to be their reliance on engaging and manipulating the audience's predictive brain – the same brain shaped by the urgent functional need to predict threats and resources, operating through the proposed Prediction → Context → Consciousness hierarchy. These cultural artifacts might represent sophisticated ways humans entertain themselves and communicate deep truths by exercising and exploiting a cognitive ability that evolved for much more fundamental functional reasons.

1.9 Creativity: A Potential Functional Consequence Operating within the Hierarchy

Where does creativity fit into this survival-driven narrative and the functional hierarchy? Creativity – the ability to generate novel and valuable ideas or outputs – can perhaps be conceptualized not necessarily as the initial evolutionary goal of intelligence, but as a sophisticated functional consequence made possible by the development of a highly flexible, complex predictive brain operating through this hierarchy.

A brain constantly generating Level 1 Predictions and minimizing errors needs the capacity to handle novelty and unexpected input. This requires flexibility and the ability to explore different potential outcomes. Within this framework, creativity may involve the system deliberately manipulating Level 1 Predictions or exploring novel predictive sequences that deviate from the most statistically probable, thereby generating new and unexpected Level 2 Contexts (e.g., a novel combination of ideas or sensory inputs), which are then evaluated and integrated by Level 3 Conscious awareness for their novelty and value. It could be the predictive engine, honed for survival, now with enough power and flexibility to explore alternative predictions and possibilities beyond immediate necessity, consciously evaluating and refining these explorations.

While creativity might have later become subject to its own selective pressures (e.g., in toolmaking, social signaling, mate attraction), its initial emergence was potentially a consequence of the intense pressure to build better predictive models for survival, leading to a brain capable of simulating and manipulating potential futures with increasing complexity and forming intricate contexts, ultimately giving rise to consciousness capable of evaluating novelty. The capacity for questioning, abstract reasoning, and philosophical thought similarly might represent the application of this functional hierarchy to simulate abstract concepts and explore non-immediate consequences – remarkable functional consequences of a predictive engine that evolved for much more basic survival tasks.

1.10 Higher-Level Consciousness: Another Functional Consequence

Similar to creativity, the complex tapestry of human higher-level consciousness – encompassing abstract thought, moral reasoning (right vs. wrong), ethical considerations, concepts of justice and fairness, and the formation of ideologies – can also be viewed as a functional consequence built upon the foundational P→C→C hierarchy, rather than its primary evolutionary driver.

The core hierarchy, evolved for survival and procreation, established the essential machinery: a powerful predictive engine (Level 1), the capacity to build rich, embodied context through experience (Level 2), and an integrated functional awareness (Level 3) focused on immediate survival needs. This machinery, particularly the ability to simulate potential futures and understand social contexts, provided the *capacity* for more abstract forms of reasoning.

As human societies became more complex, requiring intricate cooperation and navigation of social norms, this existing cognitive architecture was likely co-opted and further developed. Abstract concepts like fairness, justice, and morality could emerge as ways to predict and regulate social interactions, promoting group cohesion and long-term stability – factors that indirectly contribute to survival and reproductive success within a social species. These higher-level constructs are not proposed as direct outputs of the initial survival-driven prediction loop, but as sophisticated applications and elaborations of that underlying system, shaped significantly by cultural evolution, language, and social learning. They represent the P→C→C hierarchy being applied to simulate and evaluate complex, abstract social scenarios, far removed from immediate physical threats or mating opportunities, but crucial for thriving in human society. This same capacity for intense, context-driven simulation, when turned inwards towards abstract problems in mathematics or science, might also underlie the profound focus and intuition exhibited by figures like Einstein, Gödel, Cantor, Boltzmann, Turing, Ramanujan, or Green, who seemed to operate at a different level by deeply leveraging this predictive machinery on abstract domains.

1.11 Biological Cycles and Functional Prediction Beyond the Brain

The principle of anticipation is not confined to the cognitive functions of the brain; this principle is evident in biological processes at a fundamental level, all serving the functional needs of survival and propagation. Life is replete with examples of organisms operating on predictable cycles, anticipating future environmental states or biological needs, highlighting the universality of the functional requirement for prediction in living systems:

- Circadian Rhythms: The daily internal clock that governs sleep-wake cycles, anticipating periods of light and dark (Ralph et al., 1990) – a crucial functional adaptation.
- Seasonal Adaptations: Hibernation in anticipation of cold and food scarcity, or migration triggered by predictable seasonal changes based on environmental cues (Sherry, 2006) – essential functional predictions for group survival.
- Reproductive Cycles: The precisely timed emergence of 17-year cicadas (Karban, 1982) or the synchronized spawning of corals, all driven by biological mechanisms that functionally anticipate optimal reproductive conditions. The navigation of salmon involves anticipating routes and recognizing locations through learned cues (Hasler & Scholz, 1983) for the functional outcome of returning to spawn.
- Developmental Stages: The complex, predictable sequences of gestation and birth, orchestrated by programming that functionally anticipates developmental milestones.

These biological phenomena illustrate that prediction, in its broadest sense, appears to be a core feature of living systems, enabling them to persist and thrive in a dynamic world by fulfilling basic functional requirements, operating at low energy costs relative to their complexity, a functional efficiency echoed and vastly magnified in the low-power operation of the human brain's predictive hierarchy (Attwell & Laughlin, 2001).

1.12 Large Language Models: A Functional Analogy and Its Limits

The surprising success of LLMs can be understood through the lens of prediction, offering a fascinating, albeit limited, functional analogy to the base level of the proposed P→C→C hierarchy. However, it is crucial to first note that intelligence, as framed within the P→C→C model (rooted in prediction, context, and functional awareness for survival), does not necessarily require language. Many organisms demonstrate sophisticated predictive intelligence without human-like linguistic capabilities. LLMs operate fundamentally on language data.

- Prediction: Parallel and Divergence (Level 1) At their core, LLMs are sophisticated pattern-matching and prediction machines trained on vast datasets of text and code. They learn the statistical relationships between words, symbols, and concepts within that data. When prompted, they generate responses by predicting the most statistically probable next word or sequence of words based on the input and their training (Wang et al., 2024; Brown et al., 2020). This constitutes their core functional capability: generating plausible sequences by operating at a highly sophisticated level of statistical prediction based on the context provided by the input text and learned correlations. This predictive function shows an analogy to the biological Level 1 prediction in that both involve anticipating "what comes next." However, the functional difference is critical (Korteling et al., 2021): LLM prediction is primarily statistical and disembodied, while biological prediction is embodied, action-oriented, and involves error minimization based on environmental feedback.
- Context: The Embodiment and Grounding Gap (Level 2) Here, the functional divergence becomes stark. The biological Level 2 Context is proposed to emerge from the accumulation of embodied interactions (Varela et al., 1991). LLMs lack this embodied grounding (Shapiro, 2019; Bender et al., 2021), limiting their ability to build deep,

flexible, experientially grounded context beyond the statistical correlations in their training data (Lake et al., 2017; Mitchell, 2019).

- Consciousness: Beyond Statistical Patterns (Level 3) Given the divergence at Level 2, current LLMs do not possess the functional equivalent of Level 3 Consciousness (functional awareness for survival) as proposed here. LLM "hallucinations" further highlight their operation at the level of plausible pattern generation rather than integrated, conscious understanding (Mitchell, 2019).

Therefore, while LLMs demonstrate the power of prediction as a functional principle, the analogy to biological intelligence appears primarily limited to sequence prediction (Level 1). They currently lack the mechanisms for embodied context formation (Level 2) and integrated functional awareness (Level 3) characteristic of the proposed P→C→C hierarchy shaped by biological evolution (Korteling et al., 2021).

A further challenge in comparing biological and artificial intelligence stems from the current lack of a comprehensive, formal mathematical framework to describe intelligence itself, or even to fully characterize the emergent capabilities and internal functioning of complex AI systems like LLMs (Legg & Hutter, 2007). Much of the assessment of AI progress, particularly concerning AGI, relies on performance across suites of benchmark tasks or qualitative descriptions, rather than on rigorous, theoretically grounded measures. This absence of deep theoretical formalization makes quantifying progress and understanding the true trajectory of AI development difficult. Given the profound societal interest and potential implications of advanced AI, developing more rigorous theoretical and potentially mathematical descriptions is increasingly important. Efforts to formalize conceptual models, such as the P→C→C hierarchy explored here, or existing mathematical frameworks like Integrated Information Theory for consciousness (Tononi et al., 2016), represent potential steps toward a more unified and quantifiable understanding applicable to both biological and artificial systems, though substantial challenges remain.

This leads to a seemingly preposterous situation: how can systems trained simply to predict the next word or token exhibit behaviors that appear, at times, akin to reasoning or understanding intent? An evolving perspective, grounded in the P→C→C framework, offers a speculative but potentially fruitful hypothesis. If biological intelligence is fundamentally based on predicting 'what comes next' across multiple modalities for survival, and if complex human language is one of the richest *outputs* generated by this underlying predictive machinery, then training LLMs on vast quantities of this output might force them to implicitly learn or approximate aspects of the generative mechanism itself. Rather than just mimicking surface statistics, they might be inadvertently reverse-engineering functional facets of the human 'what comes next' engine.

This is admittedly a significant leap of faith. Yet, the empirical reality that LLMs, despite known limitations like hallucinations and factual errors, can often interpret complex instructions, follow intentions, exhibit emergent abilities (Wei et al., 2022), and function effectively as 'Intelligent Assistants' provides intriguing, albeit tentative, evidence for *some* connection beyond surface pattern matching. It suggests these models might be capturing deeper functional regularities reflective of the cognitive processes that produced the language they were trained on. This doesn't imply LLMs possess embodied context (Level 2) or survival-driven functional awareness (Level 3) as defined here. However, it does open an exciting avenue for exploration: could scaling and refining these models, guided by an understanding of the P→C→C hierarchy, lead towards AI that more effectively integrates context and assists human cognition, even if via a different path than biological evolution? Investigating the nature and depth of this implicit learning remains a critical area for future research.

1.13 Consciousness as Functional Awareness for Survival

Within the proposed P→C→C hierarchy, the complex phenomenon often termed "consciousness" (Level 3) is framed specifically as an emergent functional awareness primarily shaped by evolutionary pressures for survival and procreation. It arises from the deep integration and ingraining of countless Level 1 predictions and the rich, embodied Level 2 contexts they generate through interaction with the world. This integrated state isn't necessarily focused on abstract self-reflection in the philosophical sense, but rather on consolidating accrued information and enabling higher-level cognitive functions crucial for navigating life's core challenges.

Specifically, this functional awareness (Level 3 Consciousness) encompasses the capabilities needed to generalize from past experiences, make critical distinctions like danger vs. safety (manifesting as fear or calm), identify friend vs. foe, recognize parents and siblings, discriminate edible food from poison, and respond to "gut instincts." It underpins drives essential for propagating the species, such as assessing sexual attractiveness, finding suitable mates, and acting on maternal or paternal instincts to ensure the survival of progeny. This level represents the culmination of the predictive engine's work, creating an internal simulation (Hohwy, 2013; Clark, 2015) focused on maximizing the chances of survival and successful reproduction – fulfilling a sense of purpose tied to continuing the cycle of life. The mechanisms underlying this integration might relate to theories involving widespread information access or integration across brain regions (Tononi et al., 2016; Dehaene & Changeux, 2011; Seth & Bayne, 2022), but the *functional content* shaped by evolution is proposed here to be centered on these survival imperatives. Phenomena like anesthesia disrupting the predictive modeling process (Alkire et al., 2008) underscore the reliance of this functional awareness on the underlying P→C→C machinery.

It is crucial, within this framework, to distinguish this evolutionarily driven functional awareness from higher-level concepts often associated with human consciousness, such as morality (right vs. wrong), ethics, justice, fairness, or complex ideologies, as discussed in Section 1.10. These are proposed not as direct products of the survival-focused P→C→C hierarchy itself, but rather as later cognitive constructs built upon the foundational machinery this hierarchy provides. Once the basic capacity for complex prediction, context-building, and integrated awareness was established, interactions within complex social environments and cultural evolution likely allowed humans to cultivate these more abstract forms of thought and value systems. They represent a sophisticated application of the underlying cognitive architecture, rather than its primary evolutionary purpose.

1.14 Conclusion and Future Directions

This article argues that the capacity for anticipating "what comes next" is a deeply ingrained principle, likely forged by billions of years of biological evolution and fundamental to the functioning of the human brain. This predictive power plausibly evolved primarily because it served the urgent, practical, functional needs of survival and reproduction within an embodied agent interacting with its environment (Varela et al., 1991; Shapiro, 2019). A functional hierarchy (P→C→C) of biological intelligence was proposed where actions are governed by Prediction from sensory input, leading to the formation of experientially grounded Context, and ultimately contributing to integrated Consciousness, defined within this framework as functional awareness enabling generalization, threat/opportunity assessment, and other capabilities crucial for survival and procreation. This framework offers a potential way to approach this core functional awareness, conceptualizing it as an emergent state arising from the integration of predictions and contexts, all within a brain operating at an astonishingly efficient 15-20 watts.

While higher cognitive functions, including abstract reasoning, questioning, creativity, and moral reasoning, are remarkable human traits, they are presented here as sophisticated capabilities that may have emerged as consequences of, or were built upon, the foundation of this functional prediction-based hierarchy honed for survival. The ability to create and appreciate enduring art across diverse forms demonstrates the universality of engaging with this hierarchy, as successful art resonates by expertly manipulating innate human predictive processes and context-building mechanisms.

The recent emergence of powerful LLMs offers a fascinating, but limited, functional analogy. Their success suggests prediction is a potent mechanism, yet the comparison highlights crucial functional differences: LLMs currently lack the embodied grounding, deep contextual integration (Level 2), and survival-driven functional awareness (Level 3) characteristic of the biological P→C→C hierarchy (Korteling et al., 2021; Lake et al., 2017; Mitchell, 2019). Recognizing these functional differences, alongside the current lack of formal mathematical frameworks for intelligence (Legg & Hutter, 2007) and the intriguing hypothesis that LLMs might implicitly reverse-engineer aspects of the human predictive hierarchy by modeling its linguistic output (Wei et al., 2022), is vital for navigating future AI development.

It is important to acknowledge potential limitations of the proposed P→C→C framework. Operationalizing the concepts of 'context' and 'consciousness' (as functionally defined here) within this hierarchy for empirical testing presents significant challenges. The framework is currently presented at a functional level, and its precise mapping to specific neural circuits remains largely speculative at this stage. Furthermore, the model may need refinement to fully account for the diversity of cognitive phenomena and individual differences. The call for more formal mathematical descriptions applies also to conceptual frameworks like this one.

Having established this core functional principle of prediction, its potential evolutionary driver, and a proposed hierarchical organization (P→C→C), along with drawing functional analogies and distinctions with LLMs, the stage is set for deeper exploration. Future research stemming from this framework could involve *developing formal mathematical or computational models* implementing the P→C→C hierarchy, designing neuroimaging or behavioral experiments to test predictions about context integration and its relation to functional awareness, *investigating the nature and limits of emergent abilities and implicit knowledge representation in LLMs in relation to this hierarchy and the hypothesis of reverse-engineered functional properties*, and exploring the framework's applicability to other cognitive domains like decision-making or social cognition. Chapter 2 will shift focus to investigate potential biological circuitry enabling this prediction hierarchy, seeking insights into evolutionary design principles. Chapter 3 will then connect these insights back to artificial intelligence, exploring potential avenues for improving AI efficiency and intelligence, and considering the path towards superintelligence. Understanding intelligence, both biological and artificial, may benefit from starting with the fundamental, functional imperative to predict what comes next, leading through embodied context to consciousness.

Chapter 2. Nature's Functional Blueprint for Prediction

2.1 Introduction: Bridging Function and Mechanism

Chapter 1 of this series established the foundation, proposing that biological intelligence evolved primarily to serve the functional imperative of predicting "what comes next" for survival. It introduced the Prediction → Context → Consciousness (P→C→C) hierarchy as a potential functional architecture shaped by this evolutionary pressure, highlighting its operation within the remarkably low power budget (~20W) of the human brain (Attwell & Laughlin, 2001; Korteling et al., 2021). This framework focuses on the purpose (survival) and the functional outcome (prediction, context, consciousness) of intelligence. This perspective aligns with emerging views in physics suggesting that life itself might be characterized as a unique state of matter fundamentally distinguished by its ability to utilize information to maintain a low-entropy state far from thermodynamic equilibrium (Walker et al., 2016; Hidalgo, 2024).

In this second part, the focus pivots from the "what" and "why" to the crucial question of "how." How does nature physically implement such a sophisticated predictive, information-processing system? What are the underlying biological mechanisms and design principles that enable the P→C→C hierarchy? This article explores the strategies that evolution appears to favor remarkable energy efficiency driven by thermodynamic constraints, the pervasive reuse and adaptation of existing biological components originating from fundamental chemical processes, and an intrinsic mastery of complex computational principles (like calculus and linear algebra) that predate human formalization by eons. The central thesis presented here is that nature provides a blueprint, utilizing a core set of reusable functional design principles implemented across diverse biological scales – from molecules and cells to networks and systems – to meet the computational demands of prediction, context-building, and ultimately, consciousness, all while adhering to fundamental physical laws. This exploration aims to bridge the gap between high-level functional descriptions and tangible biological reality, seeking biological hardware capable of supporting the software of prediction. It acknowledges, however, that mapping high-level function directly onto specific mechanisms is complex and often speculative, requiring iterative refinement as neuroscience progresses.

2.2 Functional Requirements of the P→C→C Framework (Summary)

Understanding the biological implementation requires clarity on the specific functional demands imposed by each level of the P→C→C hierarchy, as outlined in Chapter 1:

- **Prediction (Level 1)**: This foundational level necessitates mechanisms capable of rapid, real-time pattern recognition from often noisy and incomplete sensory data. It requires anticipating the immediate future based on learned environmental regularities and internal states. Crucially, it involves a continuous process of comparing predictions with actual sensory input and minimizing the resulting prediction error, a core concept in theories like predictive coding (Friston, 2005; Rao & Ballard, 1999). This error signal drives learning and refines predictions, guiding immediate, embodied actions essential for survival tasks like foraging, escaping predators, or navigating complex terrains. From a physics perspective, this involves efficient information acquisition and processing to maintain the organism's state against environmental fluctuations.

- **Context (Level 2)**: Moving beyond immediate prediction, context formation requires integrating information over longer timescales. Mechanisms are needed to accumulate sequences of predictions, actions, and their consequences into dynamic, flexible models of the environment and the self. This involves robust systems for associative learning (linking stimuli, actions, and outcomes), memory formation and retrieval (encoding experiences and accessing them later), and the development of situational awareness – understanding the broader setting in which events unfold. This contextual layer allows for more nuanced and adaptive behavior than simple reflexes, enabling interpretation of ambiguous stimuli and prediction of more distant outcomes (Summerfield & de Lange, 2014; Chapter 1). This represents a more complex form of information storage and utilization compared to Level 1.
- **Consciousness (Level 3)**: The apex of the hierarchy involves integrating these experientially grounded contexts into a coherent, unified subjective experience. This requires mechanisms capable of large-scale information integration across different brain regions and modalities (Tononi et al., 2016; Dehaene & Changeux, 2011). This integrated state enables sophisticated cognitive functions such as offline simulation (mental time travel, planning, considering counterfactuals), abstract reasoning, complex social cognition (theory of mind), and a stable sense of self-representation within the world (Hohwy, 2013; Clark, 2015; Varela et al., 1991). The transition from context to consciousness remains one of the most significant explanatory gaps, potentially involving unique properties of information processing and integration within the biological substrate.
- **Overarching Constraints**: These functional capabilities must operate within biology's stringent constraints. Energy efficiency is paramount, driven by thermodynamic limits on computation and information processing; the brain consumes ~20% of the body's energy despite being only ~2% of its mass, demanding highly optimized processing (Attwell & Laughlin, 2001). Robustness is essential for reliable function despite noisy signals and potential component damage. Adaptability allows organisms to learn and adjust to novel or changing environments. Finally, embodiment grounds all processing in the physical body's interaction with the world, shaping the very nature of cognition (Varela et al., 1991; Shapiro, 2019). These constraints profoundly influence the types of algorithms and architectures favored by evolution.

2.3 Nature's Design Principle: Functional Perfection and Reuse

Evolution is often described as a "blind watchmaker" (Dawkins, 1986) or a "tinkerer" (Jacob, 1977), modifying existing structures and processes rather than designing anew. This leads to solutions that are functionally effective ("good enough" for survival) and readily adaptable or reusable in diverse contexts. This principle manifests across all biological scales, creating layers of functional capability, potentially originating from fundamental principles of self-organization in chemical systems. The emergence of life likely involved autocatalytic sets – networks of molecules where components catalyze each other's formation, leading to self-sustaining and potentially self-replicating systems (Hordijk & Steel, 2018; Kauffman, 1986). Such systems provide a plausible chemical basis for the origin of reusable biological modules.

- **Mitochondria**: Their endosymbiotic origin is a classic example of reuse – incorporating a free-living bacterium to serve as a cellular power plant. Beyond ATP synthesis, they are central players in calcium signaling, apoptosis regulation, and reactive oxygen species (ROS) production, acting as critical signaling hubs. The emerging field exploring the mitochondria-microbiome axis suggests metabolites produced by gut bacteria (like SCFAs) can directly influence mitochondrial function (e.g., biogenesis, dynamics, oxidative stress), potentially impacting host metabolism, aging, and neurodegenerative diseases. This hints at ancient, conserved communication channels linking the external microbial environment to internal cellular energy and signaling networks (Franco-Obregón & Gilbert, 2017; Han et al., 2019; Borbolis et al., 2023).
- **Microbiome**: This complex ecosystem co-evolved with its host, demonstrating symbiotic reuse. Microbes perform functions the host cannot, like digesting complex carbohydrates or synthesizing certain vitamins. Their influence extends far beyond the gut, modulating the immune system and communicating with the brain via neural, endocrine, and immune pathways (the gut-brain axis). Microbial metabolites can cross the blood-brain barrier or influence vagal nerve signaling, impacting neurotransmitter levels, stress responses, and potentially contributing to neurological disorders (Cryan et al., 2019; Carabotti et al., 2015).
- **Cells & Differentiation**: The process of cellular differentiation is a testament to modularity and reuse. A limited set of basic cellular components and pathways, likely derived from early self-replicating chemical systems, are configured in different ways, controlled by differential gene expression, to generate the vast diversity of specialized cells (neurons, myocytes, lymphocytes, etc.) required for a complex organism. Each cell type executes a specific function while relying on the same fundamental toolkit of life.
- **DNA**: The universality of the genetic code and the DNA molecule itself across nearly all life highlights an incredibly successful and reusable information storage system, essential for the replication and evolution central to life's definition. Its structure allows for high-fidelity replication, stable long-term storage, and a mechanism for encoding complex instructions using a simple four-letter alphabet. Recent technological efforts leveraging DNA for digital data storage capitalize on this evolved efficiency and density (Church et al., 2012; Goldman et al., 2013; Ceze et al., 2019; Hu et al., 2020).
- **Ribosomes**: These protein synthesis factories are another example of conserved, essential machinery, translating the stored information in DNA (via mRNA) into functional proteins. Their complex structure, involving multiple RNA and protein components, is finely tuned for accurately reading mRNA codons and catalyzing peptide bond formation. Their fundamental role and conserved nature underscore the evolutionary optimization of critical functional modules (Ramakrishnan, 2002; Steitz, 2008; Tinoco, 2011).

Evolution's strategy appears to be one of building complex systems by combining, adapting, and layering these functionally optimized, reusable modules, potentially rooted in the self-organizing principles of systems chemistry (Hordijk & Steel, 2018), rather than designing each system de novo. This iterative process allows for the emergence of complex functions from simpler building blocks.

2.4 Design Principle 1: Cyclical Prediction and Biological Timing

Anticipating predictable environmental changes is a fundamental survival strategy, deeply integrated into biological systems and forming a core component of Level 1 Prediction. Life evolved under cyclical pressures, and internalizing these rhythms provides a significant advantage. This cyclical nature may be linked to the fundamental requirement for self-sustaining, cyclical chemical processes, like autocatalysis, necessary for life's persistence (Hordijk & Steel, 2018).

- **Biological Cycles & Clocks**: Life is entrained to geophysical cycles. Circadian clocks, endogenous oscillators with a period of approximately 24 hours, are found in organisms from bacteria to humans. They regulate a vast array of physiological processes, aligning internal biology with the external light-dark cycle. These clocks are typically based on transcription-translation feedback loops involving specific "clock genes" (e.g., Period, Cryptochrome, Clock, Bmal1 in mammals) (Dunlap et al., 2004; Takahashi, 2017). Beyond daily cycles, circannual clocks regulate seasonal behaviors like reproduction, migration, and hibernation, often cued by changes in day length (photoperiod) (Helm and Liedvogel, 2024; Sherry, 2006). Even longer cycles exist, like the 13- or 17-year emergence of periodical cicadas, a strategy thought to overwhelm predators (Karban, 1982). These clocks are not mere responders; they generate internal predictions about upcoming environmental states.
- **Cyanobacteria – Molecular Timekeeping**: The KaiABC system in Synechococcus elongatus offers profound insight into the molecular basis of timing and prediction. This post-translational oscillator maintains a stable 24-hour rhythm independent of transcription, driven by the ordered phosphorylation/dephosphorylation cycle of KaiC, modulated by KaiA and KaiB. This clock mechanism regulates global gene expression, ensuring that processes like photosynthesis occur optimally during the day and others like nitrogen fixation (which is sensitive to oxygen produced during photosynthesis) occur at night. This precise temporal scheduling, driven by an internal predictive model of the daily cycle, confers a significant fitness advantage (Dong & Golden, 2008; Nakajima et al., 2005; Ishiura et al., 1998). The existence of such sophisticated timekeeping in ancient organisms underscores the deep evolutionary roots of prediction based on environmental cycles, forming a potential bedrock for the P→C→C hierarchy's temporal processing needs.

2.5 Design Principle 2: Gradient Detection - Nature's Calculus (Differentiation)

Effective interaction with the environment requires sensing not just the presence but the spatial and temporal variation of stimuli – calculating gradients. This is essential for navigation, resource localization, and threat avoidance, contributing to both Prediction (Level 1) and Context (Level 2).

- **Early Life - Chemotaxis**: Bacterial chemotaxis provides a clear example of biological differentiation. E. coli uses transmembrane receptors (methyl-accepting chemotaxis proteins, MCPs) to detect changes in chemical concentrations. Receptor methylation levels provide a short-term memory, allowing the cell to compare current ligand concentrations with those experienced moments earlier. This temporal comparison mechanism effectively computes a time derivative of the concentration, influencing the flagellar motor's rotation direction and thus modulating the run-and-tumble behavior to navigate towards attractants or away from repellents (Sourjik & Berg, 2004; Baker et al., 2006; Berg, 1993). This demonstrates computation of change over time at the cellular level.

- **Modern Marvel - Olfaction**: Mammalian olfaction involves detecting volatile molecules via hundreds of different Olfactory Receptors (ORs) expressed by olfactory sensory neurons (OSNs) in the nasal epithelium. Each OSN typically expresses only one type of OR. The pattern of activation across OSNs projecting to specific glomeruli in the olfactory bulb forms a combinatorial code representing the odorant. Intensity is encoded partly by the firing rate of OSNs and the number of activated glomeruli (Mori et al., 1999; Buck & Axel, 1991; Firestein, 2001).
    - Canine Prowess: Dogs leverage this system with exceptional acuity. Their larger olfactory epithelium, higher density of ORs, and specialized nasal airflow dynamics contribute to their sensitivity. Tracking involves detecting extremely faint concentration differences over time and space, discriminating the target odor from background noise, and integrating this information to predict the source location. This requires sophisticated processing likely involving computations analogous to spatio-temporal differentiation (Horowitz, 2009; Jenkins et al., 2018; Craven et al., 2009). The cadaver dog example (The Investigation, https://www.bbc.co.uk/programmes/p09343x5) highlights the extreme performance achievable through biological gradient analysis, effectively performing calculus on chemical fields.
- **Calculus Analogy & Beyond**: These examples show nature implicitly implementing differentiation. Sensory systems are often tuned to detect changes rather than absolute levels. Fractional calculus, which deals with derivatives of non-integer order, offers an intriguing extension, potentially providing more accurate models for systems with memory effects or complex, multi-scale dynamics, such as viscoelastic tissues or anomalous diffusion processes common in biological environments (Magin, 2004; West et al., 2003; Magin, 2010). This suggests nature's computational toolkit might encompass mathematical concepts beyond standard integer-order calculus, potentially offering more nuanced ways to model predictive processes. While a powerful analogy, equating biological gradient detection directly with formal calculus requires caution; the biological mechanisms are evolved solutions, not explicit mathematical implementations.

2.6 Design Principle 3: Integration, Linear Algebra, and Dominant Modes

Building context (Level 2) and enabling higher cognition (Level 3) requires integrating information over time and space and efficiently extracting salient features from complex, high-dimensional sensory data. Nature seems to have mastered both integration and dimensionality reduction.

- **Integration (Calculus)**: Biological systems constantly integrate information. Path integration in navigating animals involves accumulating self-motion cues (vestibular signals, proprioception, optic flow) over time to estimate displacement and orientation (McNaughton et al., 2006; Wehner & Srinivasan, 2003). Decision-making often involves integrating noisy sensory evidence until a confidence threshold is met. Neural correlates of this accumulation process have been observed in various brain regions, where firing rates ramp up proportionally to the integrated evidence (Shadlen & Kiani, 2013; Gold & Shadlen, 2007). At the cellular level, synaptic integration sums excitatory and inhibitory postsynaptic potentials over time and space in the neuron's dendritic tree to determine whether an action potential will be fired (Kandel et al., 2013). These processes are functionally equivalent to mathematical integration, allowing the system to build up information over time, crucial for forming context.

- **Linear Algebra (AGOP / Dominant Modes)**: Extracting meaningful patterns from the high-dimensional data streams the brain receives (e.g., millions of photoreceptor inputs) requires powerful dimensionality reduction and feature extraction techniques. Linear algebra provides the mathematical framework for such operations.
  - Neural Networks & Hierarchies: The brain's hierarchical organization, particularly in sensory cortices, suggests a staged process of feature extraction. Recent theoretical and empirical work suggests that neural networks, both artificial and biological, learn by identifying the dominant modes of variation in their inputs. The study by Radhakrishnan et al. (2024) proposes that the Average Gradient Outer Product (AGOP) matrix captures the essential information used for feature learning in deep networks, even potentially enabling learning without explicit backpropagation. The dominant eigenvectors of this matrix correspond to the principal features learned by the network (Radhakrishnan et al., 2024). This aligns with ideas that neural learning rules like Hebbian plasticity implicitly perform computations similar to Principal Component Analysis (PCA) (Oja, 1982). This ability to extract dominant features is vital for efficient representation and prediction within the P→C→C framework.
  - Dog Olfaction Revisited (A Speculative Link): The dog's active sniffing strategy can be viewed through this lens, although the direct link to AGOP or PCA/SVD remains speculative. By rapidly sampling the olfactory environment from slightly different positions and times (head movements, airflow control), the olfactory system acquires a high-dimensional dataset of receptor activation patterns. It's conceivable that neural processing downstream applies principles analogous to PCA or SVD to this data, allowing the system to identify the underlying "odor objects" or dominant scent profiles. This could effectively separate the target trail signature (the principal component of interest) from background noise and irrelevant odors. This dimensionality reduction would allow the dog to efficiently track the source by following the gradient of this dominant mode, showcasing a potential biological implementation of sophisticated linear algebraic principles for signal extraction. Further research is needed to validate such specific algorithmic interpretations of complex behaviors.

2.7 Towards Biological Implementation: Circuitry and Mechanisms

Connecting these functional principles (timing, differentiation, integration, mode extraction) to specific biological hardware remains a major challenge, but research points towards candidate mechanisms that could support the P→C→C hierarchy:

- **Cycles/Timing/Prediction:**

    - Mechanisms: Molecular clocks (KaiABC, mammalian clock gene loops). Neural oscillators based on intrinsic neuronal properties and network connectivity (e.g., thalamocortical loops, hippocampal theta rhythm, Central Pattern Generators in spinal cord) (Buzsáki, 2006). Synaptic plasticity (LTP/LTD at glutamatergic synapses via NMDA/AMPA receptors, spike-timing-dependent plasticity) allows learning of temporal sequences (Markram et al., 1997; Malenka & Bear, 2004). Predictive coding circuits involving feedforward and feedback connections between cortical layers are hypothesized to implement prediction error computation (Bastos et al., 2012; Friston, 2005).
- **Differentiation (Calculus):**
    - Mechanisms: Sensory receptor adaptation reduces response to constant stimuli, emphasizing changes. Lateral inhibition sharpens spatial contrast (e.g., in retina via horizontal/amacrine cells). Center-surround receptive fields compute local spatial derivatives (Kandel et al., 2013). Opponent processing compares inputs (e.g., color vision). Electrochemical gradients across membranes are fundamental to neuronal excitability and signaling (Nicholls et al., 2012).
- **Integration (Calculus):**
    - Mechanisms: Dendritic integration sums synaptic inputs with temporal and spatial weighting (Stuart et al., 2008). Recurrent neural circuits and intrinsic cellular mechanisms (e.g., calcium dynamics, specific ion channels) support persistent activity for working memory and temporal integration (Wang, 2001). Accumulator neurons in areas like LIP and frontal eye fields show ramping activity reflecting evidence integration (Huk & Shadlen, 2005; Gold & Shadlen, 2007).
- **Linear Algebra (Mode Detection):**
    - Mechanisms: Hebbian learning rules and their variants (e.g., BCM rule) modify synaptic weights based on correlated pre- and post-synaptic activity, performing PCA-like computations (Bienenstock et al., 1982; Oja, 1982). Synaptic weight matrices in feedforward and recurrent networks implement linear transformations. Population codes, where variables are represented by the distributed activity patterns across neuronal ensembles, allow for robust representation and potentially matrix operations through network dynamics (Georgopoulos et al., 1986; Pouget et al., 2000). Cortical architecture with its layers and columns seems suited for hierarchical feature extraction (Felleman & Van Essen, 1991).
- **Efficient Storage:**
    - Mechanisms: Information is stored across multiple timescales: short-term in ongoing neural activity patterns and transient synaptic changes (facilitation/depression); long-term via structural plasticity at synapses (changes in spine size, receptor numbers, requiring protein synthesis) and potentially through epigenetic modifications (DNA methylation, histone acetylation) influencing gene expression related to memory consolidation (Bailey & Kandel, 1993; Levenson & Sweatt, 2005; Kandel et al., 2013). DNA provides the ultimate stable blueprint. Network connectivity patterns themselves represent stored information.
- **Efficiency Driving Structure:**

- Mechanisms: Physical principles impose fundamental limits on the efficiency of computation and information processing. Minimization of wiring length principles appear to govern brain layout, reducing metabolic cost and signal delay (Chklovskii, 2004; Bullmore & Sporns, 2012). Sparse coding strategies, where information is represented by the activity of a small subset of neurons, are highly energy-efficient (Lennie, 2003; Olshausen & Field, 2004). Homeostatic plasticity mechanisms stabilize network activity, preventing runaway excitation or silence and maintaining function within optimal operating regimes (Turrigiano & Nelson, 2004). Basic thermodynamic laws constrain all biological energy transformations (Attwell & Laughlin, 2001). Drive reduction theories suggest behavior is motivated by reducing internal imbalances, an efficiency principle aimed at restoring homeostasis (Hull, 1943; Wikipedia).

Nature's implementation is a multi-layered system involving dynamic electrical and chemical signaling, coupled with slower structural and molecular changes, all operating under severe physical and energetic constraints.

2.8 Emerging Complexities: Sleep, Synapses, Fields, and Clocks

The picture painted above, while complex, is still likely an oversimplification. Recent neuroscience research reveals further layers of intricacy that are crucial for understanding how biological systems achieve prediction, context, and consciousness, often orchestrated by the fundamental biological clocks discussed earlier.

- **Sleep, Memory Consolidation, and Context**: Sleep is far from a passive state; it plays a critical active role in refining the brain's predictive models. During sleep, particularly Slow-Wave Sleep (SWS) characterized by slow, synchronized brain oscillations, memories initially encoded in the hippocampus are thought to be reactivated and gradually transferred to the neocortex for long-term storage. This process, often termed systems consolidation, involves a dialogue between the hippocampus and cortex, potentially mediated by sharp-wave ripples in the hippocampus coupled with cortical sleep spindles and slow oscillations (Born & Wilhelm, 2012; Stickgold, 2005). This consolidation doesn't just strengthen memories; it likely integrates new experiences (predictions, action outcomes) into existing cortical knowledge structures, thereby enriching and stabilizing the Context (Level 2) required for nuanced understanding and prediction. Furthermore, the synaptic homeostasis hypothesis suggests that sleep, particularly SWS, serves to globally downscale synaptic weights that were potentiated during wakefulness, preventing saturation and runaway potentiation, thus improving signal-to-noise and energy efficiency for the next learning cycle (Tononi & Cirelli, 2014). REM sleep, with its distinct neurochemical milieu and brain activity patterns resembling wakefulness, has been implicated in emotional memory processing and potentially creative insight generation, further contributing to the refinement of Context and potentially aspects of Consciousness (Level 3). The precise timing and cycling through different sleep stages are tightly regulated by both homeostatic sleep pressure and the circadian clock (Section IV), highlighting how fundamental timing mechanisms orchestrate these vital offline consolidation and refinement processes.

- **The Computational Power of Single Neurons and Synapses**: Traditional models often treat neurons as simple point integrators and synapses as single scalar weights. However, research increasingly reveals far greater complexity. Evidence suggests that dendrites, the branched input structures of neurons, are not passive cables but possess active conductances allowing them to perform complex, localized computations. Different dendritic branches can function as independent computational subunits, effectively making a single neuron behave like a multi-layered network (London & Häusser, 2005; Stuart et al., 2008). Work like that from Harnett's lab suggests human dendrites have distinct electrical properties enabling more complex processing than rodent dendrites (Beaulieu-Laroche et al., 2018). Furthermore, recent findings indicate that synaptic plasticity rules are not uniform; they can depend heavily on the synapse's location on the dendrite, the specific type of synapse, and the local biochemical environment. For instance, synapses on distal dendrites might follow different learning rules than those close to the cell body (Poirazi & Papoutsi, 2020). This intricate, location-dependent synaptic logic vastly increases the computational capacity of individual neurons and networks, potentially enabling more efficient learning and information storage crucial for all levels of the P→C→C framework. This complexity operates within the temporal framework set by biological clocks, with synaptic efficacy and plasticity mechanisms known to exhibit circadian modulation (Zong et al., 2023).
- **Electrochemical and Field Effects**: Beyond direct synaptic transmission, other physical forces likely contribute to brain function. Electrochemical gradients are fundamental to generating action potentials and synaptic potentials. But the collective activity of neurons also generates endogenous electric fields (measured as EEG or local field potentials - LFPs). These fields are not merely epiphenomena; evidence suggests they can influence neuronal firing probability and synchronization through ephaptic coupling (non-synaptic electrical interactions) (Anastassiou et al., 2011; Fröhlich & McCormick, 2010). Brain oscillations (gamma, theta, alpha, etc.), reflecting synchronized field potentials, are strongly correlated with cognitive functions like attention, working memory, and consciousness, potentially providing a mechanism for coordinating distributed neural activity and integrating information (Buzsáki, 2006). While more speculative for humans, magnetoreception is well-established in some animals (e.g., birds, turtles) for navigation, relying on quantum effects involving cryptochromes (proteins also involved in circadian clocks) or potentially magnetite particles (Wiltschko & Wiltschko, 2003). While a direct role for magnetism in general human cognition or information storage is not established, these examples illustrate that biological systems exploit diverse physical phenomena. Electrochemical processes and field effects, operating efficiently at the physical level, could offer computationally inexpensive ways to achieve network synchronization, modulate neuronal excitability, and potentially contribute to information binding or even storage, complementing synaptic mechanisms and contributing to the overall efficiency demanded for survival. The precise timing provided by biological clocks is essential for coordinating these oscillatory dynamics and field effects.

These emerging complexities reinforce the idea that biological intelligence leverages intricate, multi-level mechanisms, tightly orchestrated in time, to achieve its remarkable functional capabilities.

2.9 Memory Hierarchies and Information Flow in the P→C→C Framework

The P→C→C hierarchy inherently implies a structured approach to information storage and retrieval, aligning well with established concepts of multiple memory systems in cognitive neuroscience. Efficient communication and interaction between these memory systems are crucial for the framework's operation.

- **Mapping Memory Systems to P→C→C Levels:**
  - Prediction (Level 1) & Sensory/Working Memory: Level 1 predictions, focused on the immediate future ("what comes next" within milliseconds to seconds), likely rely heavily on sensory memory buffers (iconic, echoic) that briefly hold raw sensory input, and the most transient components of working memory. This includes mechanisms like priming (where recent exposure influences current processing) and rapid synaptic adaptation (facilitation/depression) that tune neuronal responses to immediate temporal patterns (Atkinson & Shiffrin, 1968; Baddeley, 2003). These systems provide the fleeting data needed for immediate anticipation and action guidance. Predictive coding models suggest that predictions generated at higher levels modulate activity in these early sensory/memory stages (Friston, 2005).
  - Context (Level 2) & Episodic/Semantic/Working Memory: Building context requires integrating information over longer durations (seconds to minutes, hours, days, and beyond). This function maps well onto working memory (specifically its capacity for active maintenance and manipulation, e.g., Baddeley's phonological loop and visuospatial sketchpad), episodic memory (memory for specific autobiographical events situated in time and place, heavily reliant on the hippocampus), and semantic memory (general knowledge about the world, facts, concepts, stored distributed across the neocortex) (Tulving, 1972; Squire & Zola-Morgan, 1991). Level 2 involves using working memory to hold current predictions and outcomes, associating them via hippocampal mechanisms to form episodic traces, and gradually integrating these episodes into broader semantic knowledge structures in the cortex. This allows the system to build dynamic models of situations based on accumulated experience.
  - Consciousness (Level 3) & Integrated Long-Term Memory/Global Workspace: Level 3, representing integrated awareness and the self-model, likely draws upon highly consolidated semantic memory and autobiographical memory (the narrative integration of episodic memories). This corresponds to the deeply ingrained knowledge and understanding of the world and one's place within it, built over a lifetime. Theories like the Global Workspace Theory (GWT) propose that consciousness arises when information from various specialized modules (including memory systems) is "broadcast" to a central workspace, making it globally available for report, deliberation, and control (Baars, 1988; Dehaene & Naccache, 2001). In this view, Level 3 consciousness involves accessing and integrating information from the vast long-term stores established through Level 2 processes, allowing for complex simulations and self-reflection. Integrated Information Theory (IIT) also emphasizes the importance of integrated information for consciousness, suggesting that the capacity of a system to integrate information reflects its level of consciousness (Tononi et al., 2016).

- **Efficient Communication Between Memory Levels**: For this hierarchy to function effectively and efficiently, information must flow smoothly between these levels. Several mechanisms likely facilitate this:
  - Systems Consolidation: As discussed in Section VIII, sleep-dependent processes actively transfer and restructure information, moving memories from the hippocampus-dependent episodic system (Level 2) to the neocortex for long-term semantic storage (underpinning Level 3). This prevents catastrophic interference in the hippocampus and builds stable cortical knowledge structures (McClelland et al., 1995; Born & Wilhelm, 2012).
  - Retrieval Cues & Pattern Completion: Contextual cues in the present environment can trigger the retrieval of relevant episodic or semantic memories. The hippocampus is thought to play a key role in pattern completion, where a partial cue can reactivate a complete memory trace, bringing past context (Level 2/3) to bear on current prediction (Level 1) (Rolls, 2013).
  - Attention Mechanisms: Top-down attentional control can selectively prioritize the processing and retrieval of information relevant to current goals, effectively gating the flow of information between memory systems and processing levels (Corbetta & Shulman, 2002).
  - Neural Oscillations: Coordinated rhythmic activity, particularly theta-gamma coupling between the hippocampus and prefrontal cortex, is hypothesized to mediate communication between working memory, episodic memory retrieval, and executive control processes, enabling the dynamic interaction between Level 1, 2, and 3 functions (Buzsáki et al., 2022; Jensen & Lisman, 1998).

This multi-layered memory architecture, with specialized systems for different timescales and efficient mechanisms for interaction and consolidation, provides a plausible biological substrate for the information processing demands of the P→C→C framework, balancing the need for rapid prediction with the integration required for deep context and conscious awareness.

2.10 Synthesizing the Blueprint: The Predictive Brain in Action

How do these diverse design principles and mechanisms coalesce to implement the P→C→C hierarchy and enable intelligent behavior? The process begins with sensory input and unfolds through layers of increasingly sophisticated, temporally integrated computation, shaped by evolutionary pressures and fundamental physical and chemical constraints.

Information Flow and Processing: Sensory information enters the system through specialized receptors that transduce physical energy (light, sound, chemicals, pressure) into electrochemical signals – the brain's fundamental currency. This initial signal is often pre-processed, emphasizing changes and contrasts (differentiation, Section 2.5), for instance, through center-surround antagonism in the retina or adaptation in auditory nerve fibers. Communication between neurons occurs primarily via action potentials (spikes) traveling along axons and triggering the release of neurotransmitters at synapses, which in turn alter the electrical potential of the postsynaptic neuron (Kandel et al., 2013). This electrochemical signaling forms the basis of neural computation. The very possibility of such complex signaling relies on the organism maintaining a low-entropy state, distinct from its environment, a characteristic potentially defining life itself (Walker et al., 2016; Hidalgo, 2024).

Hierarchical Processing and the P→C→C Framework: This processed sensory information typically flows through hierarchical pathways, particularly evident in the cortex (Felleman & Van Essen, 1991). Early stages handle basic feature extraction (Level 1 Prediction foundations), while progressively higher stages integrate information over larger spatial and temporal scales, extracting more complex and abstract representations. This aligns with the P→C→C structure:

- Level 1 (Prediction): Initial sensory processing, emphasizing change detection (differentiation) and rapid pattern matching based on immediate input and highly transient memory (sensory/working memory, Section 2.9), allows for quick predictions about the immediate future. This likely involves circuits optimized for speed and thermodynamic efficiency, potentially leveraging predictive coding mechanisms where feedback carries predictions and feedforward pathways signal errors (Friston, 2005; Rao & Ballard, 1999). Action selection at this level is often reflexive or highly reactive.
- Level 2 (Context): Information from Level 1, representing sequences of predictions and action outcomes, is integrated over time. This involves engaging working memory, forming episodic memories (likely involving the hippocampus), and gradually updating semantic knowledge structures (Section 2.9). Neural mechanisms supporting integration (Section 2.6), such as accumulator neurons and recurrent network activity, play a crucial role. The extraction of dominant modes (linear algebra analogies, Section 2.6) allows for efficient representation of relevant contextual features. This level provides the situational awareness necessary for more flexible, goal-directed behavior, representing a more sophisticated stage of information processing and storage.
- Level 3 (Consciousness): At the highest level, deeply consolidated context (long-term semantic and autobiographical memory) is integrated across multiple brain systems, potentially via mechanisms like the global workspace (Baars, 1988; Dehaene & Naccache, 2001) or large-scale integrated information (Tononi et al., 2016). This allows for conscious reflection, planning, simulation of hypothetical futures, and a coherent sense of self operating within the world model built by Level 2. Executive functions, often associated with the prefrontal cortex (Miller & Cohen, 2001), likely orchestrate the access and manipulation of information at this level.

The Cheetah-Gazelle Chase Revisited: In the life-or-death chase described in Chapter 1, this integrated system operates at peak performance.

- The gazelle's visual system rapidly detects the cheetah's movement (sensory input, differentiation). Immediate Level 1 Prediction circuits anticipate the trajectory and trigger an evasive jink (action).
- Simultaneously, Level 2 Context processing integrates this immediate threat with stored knowledge: the terrain layout, the gazelle's own fatigue level, the typical hunting strategies of cheetahs (semantic memory), and perhaps memories of past close calls (episodic memory). This context informs the choice of evasive maneuver.

- The cheetah performs similar calculations: Level 1 Prediction guides its interception path based on the gazelle's current movement. Level 2 Context incorporates the gazelle's jinking patterns, the distance, its own energy reserves, and identifies signs of vulnerability.
- For both animals, Level 3 Consciousness likely represents the focused awareness of the situation, the goal (escape or kill), and the integration of all relevant context into a high-stakes simulation guiding ongoing strategy adjustments. The mother gazelle's awareness extends to her fawn, adding another layer of complexity to her context and decision-making.

Structural Underpinnings: This hierarchical processing is supported by the brain's structure. The layered neocortex, with its distinct input/output layers and hierarchical organization of sensory and association areas, seems well-suited for implementing predictive coding and feature extraction (Felleman & Van Essen, 1991). The hippocampus acts as a crucial hub for rapidly forming episodic memories and binding contextual elements (Squire & Zola-Morgan, 1991), essential for Level 2. The thalamus serves as a major relay station but also participates in coordinating cortical activity through thalamocortical loops, potentially contributing to timing, attention, and integration (Buzsáki, 2006). The prefrontal cortex sits atop this hierarchy, involved in working memory, executive control, planning, and integrating information for goal-directed behavior, aligning with Level 3 functions (Miller & Cohen, 2001). The entire system operates efficiently due to principles like wiring minimization and synaptic computing (Section 2.7, 2.12).

In summary, the intelligent brain appears to implement the P→C→C hierarchy through a multi-level, multi-timescale system. It leverages evolved design principles like reuse, cyclical timing, differentiation, integration, and efficient coding, implemented via sophisticated electrochemical signaling, complex neuronal and synaptic computation, hierarchical memory systems, and specialized brain structures, all operating under strict physical and energetic constraints defined by thermodynamics and information theory.

2.11 Guiding Future Neuroscience: A Function-Driven Approach

The immense complexity of the brain presents a formidable challenge to understanding its function. Large-scale, bottom-up efforts, such as detailed connectomics projects mapping synaptic connections (like the Human Brain Project or the impressive petabyte-scale H01 dataset mapping a cubic millimeter of human cortex by Google and collaborators (Shapson-Coe et al., 2024)), provide unprecedented structural detail. However, as some have argued, simply accumulating more data – even at extraordinary resolution – may not automatically lead to understanding how the system works (Jonas & Kording, 2017; Carandini, 2012). Interpreting the function of billions of neurons and trillions of synapses requires a guiding theoretical framework, potentially grounded in physical principles.

This is where functional perspectives like the P→C→C hierarchy can offer significant value. By defining core functional requirements derived from evolutionary pressures (the need to predict, build context, integrate information for survival and procreation), such frameworks can generate specific, testable hypotheses about the kinds of biological circuits and mechanisms that should exist. Instead of mapping everything indiscriminately, research could adopt a more targeted, "middle-out" approach (Bassett & Gazzaniga, 2011), using the functional requirements to guide the search for their neural correlates:

- Targeted Experiments: If Level 1 requires rapid prediction and error calculation, experiments could specifically target circuits hypothesized to perform these functions (e.g., examining specific cortical layers implicated in predictive coding, or circuits involved in rapid motor adaptation). If Level 2 requires context integration via episodic memory, research could focus on hippocampal-neocortical interactions during learning and retrieval, particularly during sleep consolidation (Section 2.8). If Level 3 involves global information integration, studies could investigate mechanisms supporting long-range communication and synchronization across brain networks.
- Iterative Refinement: The process should be iterative. The functional framework (P→C→C) guides the search for biological mechanisms. Discoveries about the actual circuitry, its capabilities, and its limitations then feedback to refine and constrain the functional model. For example, finding that dendritic computations (Section 2.8) are more powerful than assumed might lead to revisions in how Level 1 or 2 functions are conceptualized. This dialogue between theory and experiment is crucial.
- Focus on Evolutionary Needs & Physical Principles: The framework encourages researchers to ask why a particular circuit or mechanism exists in terms of the survival advantage it conferred and how it operates within physical constraints (e.g., thermodynamic efficiency). This evolutionary and physical lens helps prioritize which aspects of the massively complex system are likely to be fundamental to its core intelligent functions, potentially distinguishing essential computational elements from historical contingency or structural byproducts.

By combining the power of large-scale data acquisition with the targeted hypotheses generated by function-driven, evolutionarily-grounded, and physically-constrained theories like P→C→C, neuroscience may be able to navigate the brain's complexity more expeditiously. This synergy between bottom-up data generation and top-down/middle-out theoretical guidance could accelerate progress towards a deeper understanding of biological intelligence, moving beyond structure to uncover the algorithms of the mind.

2.12 Synaptic Computing: Localized Processing, Memory, and the Path to Efficiency

A crucial aspect underpinning the brain's remarkable energy efficiency and potentially enabling the seamless integration required by the P→C→C framework lies in its fundamental departure from conventional computing architectures. Unlike the standard von Neumann architecture used in most modern computers, where processing (CPU) and memory (RAM, storage) are physically separated, the brain performs in-memory computing at a massive scale, a strategy likely necessitated by thermodynamic limits on information processing.

- The Synapse as a Computing and Memory Unit: In the brain, the synapse – the junction between two neurons – is not merely a passive connection point. It is an active computational element where information is both processed and stored. The strength (or weight) of a synapse, which determines how strongly the activity of one neuron influences another, is dynamically modified based on neural activity (learning). This synaptic weight itself represents a form of memory. Processes like Long-Term Potentiation (LTP) and Long-Term Depression (LTD) directly alter synaptic efficacy based on correlated firing patterns, effectively storing information locally at the site where computation (signal transmission and modulation) occurs (Kandel et al., 2013; Malenka & Bear, 2004). Furthermore, as discussed in Section 2.8, individual dendrites and even single synapses possess complex biochemical machinery and electrical properties allowing for localized computations, far beyond simple weighted summation (Stuart et al., 2008; London & Häusser, 2005; Poirazi & Papoutsi, 2020). This co-localization of memory and processing at the synaptic and sub-neuronal level is a hallmark of biological computation.

- Contrast with Von Neumann Architecture: Modern digital computers largely follow the von Neumann architecture, characterized by separate processing units (CPU/GPU) and memory units (RAM/SSD/HDD). While this design has been incredibly successful, it faces a fundamental limitation known as the von Neumann bottleneck. A significant amount of time and, crucially, energy is spent simply shuttling data back and forth between the processor and memory over data buses (Backus, 1978). As computations become more data-intensive, as is the case in large-scale AI models, this data movement becomes a major performance limiter and energy drain. Estimates suggest data movement can consume orders of magnitude more energy than the actual computation itself in modern hardware (Horowitz, 2014; Ye et al., 2022). This contrasts sharply with the brain's estimated ~20W power consumption for vastly more complex tasks.
- Neuromorphic Computing Efforts: Inspired by the brain's efficiency, the field of neuromorphic computing aims to build hardware that mimics its structure and function, often employing principles of synaptic computing and in-memory processing (Mead, 1990; Schuman et al., 2017). Various approaches exist, using analog or digital circuits, or novel materials (like memristors) that inherently combine memory and processing characteristics (Strukov et al., 2008). However, building large-scale, reliable, and programmable neuromorphic systems that truly capture the complexity and efficiency of biological synapses has proven challenging. Issues include device variability, noise sensitivity, and developing effective training algorithms for these non-traditional architectures (Schuman et al., 2017; Christensen et al., 2022). Emerging research into quantum neuromorphic computing seeks to leverage quantum phenomena to potentially overcome some of these limitations or offer new computational paradigms, though this field is still highly exploratory (Tang et al., 2019).
- Implications for AI and P→C→C: If artificial intelligence is to approach the energy efficiency and adaptive learning capabilities of the human brain, a shift towards architectures inspired by synaptic computing seems increasingly necessary. Localizing memory and computation drastically reduces the energy dissipated via data movement. Furthermore, the hierarchical memory systems discussed in Section 2.9, potentially implemented via synaptic mechanisms operating at different timescales and locations, could facilitate the efficient flow of information required by the P→C→C framework. Rapid, local synaptic changes could support Level 1 Predictions, while slower, more distributed plasticity mechanisms, potentially involving dendritic computation and coordinated by network oscillations, could

underpin the integration needed for Level 2 Context. The global integration characteristic of Level 3 might rely on long-range connections whose weights are shaped by these localized learning rules and consolidated over time (e.g., during sleep). Adopting principles of synaptic computing might therefore be key not only for matching the brain's energy efficiency but also for enabling the kind of deeply integrated context-aware processing that distinguishes biological intelligence from current AI.

2.13 Conclusion: Setting the Stage for Chapter 3

This exploration into the "how" of biological intelligence reveals nature as a master computationalist, albeit one working with very different materials and constraints than human engineers. The functional requirements of the P→C→C hierarchy – prediction, context, consciousness – appear to be met through elegant design principles: functional reuse (rooted perhaps in autocatalytic chemical origins), cyclical anticipation, gradient detection (differentiation), information integration, and efficient feature extraction (linear algebra/dominant modes). These principles are embodied in a diverse array of biological mechanisms, from molecular clocks and synaptic plasticity to network oscillations, complex dendritic computations, and potentially even field effects, all optimized over evolutionary time for robustness and remarkable energy efficiency operating within fundamental physical limits. Furthermore, a hierarchical memory system, aligned with the P→C→C levels and supported by efficient communication mechanisms like consolidation and oscillatory coupling, allows for the storage and integration of information across multiple timescales. The brain's architecture leverages synaptic computing, co-localizing memory and processing to overcome the energy bottlenecks inherent in traditional computing designs. The precise temporal orchestration provided by deeply ingrained biological clocks appears crucial for coordinating these diverse processes, including vital offline consolidation during sleep. Adopting function-driven frameworks like P→C→C may provide crucial guidance for future neuroscience research aiming to unravel this complexity.

The picture emerging is not one of a simple "intelligence circuit," but rather a complex, dynamic, multi-scale system where sophisticated cognitive functions arise from the orchestrated interplay of these fundamental computational building blocks and memory systems. Understanding these principles – how nature computes with molecules, cells, and networks under tight energy budgets, leveraging intricate synaptic rules, non-synaptic interactions, and structured memory hierarchies, all within a precise temporal framework – offers a profound perspective on intelligence itself. It suggests that consciousness and higher cognition may not require entirely exotic mechanisms but could be emergent properties of a highly optimized, predictive system grounded in embodied interaction and refined through processes like sleep-dependent consolidation.

Chapter 3 will bridge this understanding to artificial intelligence. Can leveraging nature's blueprint – its strategies for efficiency, reuse, prediction, context-building, integration, complex computation, hierarchical memory, synaptic computing, and temporal orchestration – guide the design of AI systems that are more capable, adaptive, and perhaps possess a deeper, more grounded form of understanding? The potential applications, the inherent challenges in translating biological principles to silicon (or other substrates), and the critical questions surrounding the future development of AGI and superintelligence will be considered, including both the immense opportunities and the potential risks.

Chapter 3. Bridging Biological Blueprints, Artificial Intelligence, and the Path to Superintelligence

3.1 Introduction: Synthesizing Function, Mechanism, and the Future

This article marks the culmination of a three-part exploration into the nature of intelligence, both biological and artificial. Chapter 1 established the foundational argument: that biological intelligence, driven by the relentless pressure of survival, evolved primarily around the functional imperative of predicting "what comes next." We introduced the Prediction → Context → Consciousness (P→C→C) hierarchy as a potential functional architecture, highlighting how prediction guides action, accumulates into embodied context, and ultimately gives rise to a functional awareness geared towards survival and procreation, all operating within the brain's astonishingly low power budget (Chapter 1; Korteling et al., 2021). Chapter 2 delved into the "how," exploring nature's functional blueprint – the biological mechanisms and design principles enabling this hierarchy. We examined strategies like energy efficiency, functional reuse, biological timing, gradient detection (differentiation), information integration, feature extraction (linear algebra analogies), synaptic computing, and hierarchical memory systems, revealing nature as a master computationalist working under unique constraints (Chapter 2). This evolutionary drive for persistence and propagation may even extend to the pre-biotic chemical world, suggesting that the core principles underpinning life and intelligence are deeply rooted in fundamental physico-chemical processes related to dynamic kinetic stability (Pascal and Pross, 2022; Pross, 2011).

The goal of this third and final part is to bridge this understanding of biological intelligence – its function and mechanism – with the rapidly evolving landscape of Artificial Intelligence (AI). While drawing inspiration from existing frameworks like predictive processing, embodied cognition, and theories of consciousness (Friston, 2005; Varela et al., 1991; Baars, 1988; Tononi et al., 2016), the P→C→C framework offers a unique synthesis by explicitly linking these concepts through a functional hierarchy driven by evolutionary imperatives and grounded in the core mechanism of prediction. The capabilities and limitations of current AI are explored, particularly Large Language Models (LLMs) and neuromorphic approaches, through this specific lens, contrasting their information processing with the unique informational architecture inherent in biological systems (Walker et al., 2016). Potential synergies are considered, examining how insights from biology might inform future AI development towards greater efficiency, robustness, and perhaps deeper understanding. Crucially, while biology offers profound inspiration and benchmarks, especially for efficiency and robustness, it is acknowledged that AI development may follow convergent or entirely novel paths towards intelligent behavior. The core theme remains: how the P→C→C framework can inform our understanding and development of AI, potentially guiding us towards more efficient, robust, and capable systems, regardless of the specific implementation path, as we contemplate the trajectory towards Artificial General Intelligence (AGI) and beyond.

3.2 Recap: The P→C→C Framework and Nature's Implementation

Before bridging to AI, let us briefly revisit the cornerstones laid in the previous parts.

- **The Functional Hierarchy (P→C→C)**: At its core, the framework posits:
  - Prediction (Level 1): The fundamental layer governing immediate anticipation and action guidance based on sensory input, constantly minimizing prediction error.
  - Context (Level 2): The accumulation of predictions, actions, and outcomes over time, integrated through embodied experience to form dynamic, situational understanding.
  - Consciousness (Level 3): An emergent, integrated functional awareness arising from deeply ingrained context, enabling generalization, threat/opportunity assessment, planning, and facilitating actions crucial for survival and procreation. This entire system is characterized by its grounding in embodiment and its evolutionary roots in survival drives, operating at remarkable energy efficiency (~20W) (Chapter 1; Attwell & Laughlin, 2001).
- **Nature's Blueprint (Key Mechanisms from Chapter 2)**: Evolution's implementation relies on several key strategies:
  - Energy Efficiency: Operating complex computations within a strict ~20W budget.
  - Functional Reuse & Tinkering: Adapting existing components (molecules, cells, systems like mitochondria, microbiome, DNA) for new purposes (Jacob, 1977).
  - Biological Timing & Cyclical Prediction: Internalizing environmental rhythms via molecular and neural clocks (e.g., circadian clocks, KaiABC) (Dunlap et al., 2004; Dong & Golden, 2008).
  - Gradient Detection (Differentiation/Calculus): Sensing changes over time and space (e.g., chemotaxis, olfaction) (Sourjik & Berg, 2004; Firestein, 2001).
  - Information Integration & Feature Extraction (Integration/Linear Algebra): Accumulating evidence over time (e.g., path integration, decision-making) and extracting salient features (potential analogies to PCA/AGOP in neural learning) (McNaughton et al., 2006; Gold & Shadlen, 2007; Radhakrishnan et al., 2024; Oja, 1982).
  - Synaptic Computing: Co-locating processing and memory at the synapse, enabling efficient, local learning (LTP/LTD) and complex dendritic computation (Kandel et al., 2013; Stuart et al., 2008).
  - Hierarchical Memory Systems: Utilizing distinct but interacting memory systems (sensory, working, episodic, semantic) operating across different timescales, consolidated efficiently (e.g., during sleep) (Tulving, 1972; Squire & Zola-Morgan, 1991; Born & Wilhelm, 2012).
  - Efficiency Drivers: Principles like wiring minimization and sparse coding further optimize resource usage (Chklovskii, 2004; Lennie, 2003).
  - Underlying Architecture: Collectively, these mechanisms contribute to a unique informational architecture within biological systems, where control structures and information flow are intrinsically linked to the physical instantiation from the cellular level upwards (Walker et al., 2016).

3.3 The Current Landscape: Neuromorphic Computing and Synaptic Mimicry

Recognizing the efficiency and architectural advantages of the brain, researchers are actively exploring hardware inspired by its principles.

- **Introduction to Neuromorphic Computing**: This field aims to build hardware systems that mimic the structure (neurons, synapses) and function (parallel, event-driven processing) of the biological nervous system (Mead, 1990). The primary motivation is to overcome the limitations of traditional computer architectures, particularly the von Neumann bottleneck – the separation of processing (CPU) and memory (RAM) that leads to significant energy consumption from data shuffling (Backus, 1978; Horowitz, 2014). Neuromorphic systems often feature massive parallelism, asynchronous, event-based communication (sending signals only when necessary), and the co-location of memory and processing elements. Examples range from large-scale research platforms like Intel's Loihi/Loihi 2 (Davies et al., 2021), IBM's TrueNorth, and the SpiNNaker project (Furber et al., 2013) to various academic prototypes exploring different device technologies.
- **Synaptic Computing Concepts**: A key element of neuromorphic design is synaptic computing. Biological synapses are not just connections; they are dynamic computational units where signal modulation (processing) and learning/memory (weight changes) occur locally (Kandel et al., 2013). Technologies like memristors (Strukov et al., 2008), phase-change memory, and specialized analog circuits are being developed to emulate this synaptic plasticity, allowing the connection strength between artificial neurons to change based on activity. The potential benefits are significant: drastically reducing energy spent on data movement, enabling local learning rules closer to biological mechanisms, and facilitating massively parallel computation.
- **Challenges and Limitations**: Despite progress, building brain-like hardware faces hurdles. Scaling these systems to billions of neurons and trillions of synapses remains difficult. Device variability (each artificial synapse behaving slightly differently) and sensitivity to noise pose challenges for reliable computation. Developing effective programming paradigms and training algorithms for these non-traditional, often analog or mixed-signal, architectures is an ongoing research area (Schuman et al., 2017). Furthermore, capturing the full complexity of biological computation, including intricate dendritic processing and the diverse range of synaptic plasticity rules observed in nature, is still a distant goal (Christensen et al., 2022). While impressive demonstrations exist, achieving the brain's combination of scale, efficiency, and learning flexibility remains a grand challenge.

3.4 Generative AI and LLMs: Functionality vs. Biological Reality

While neuromorphic computing attempts to mimic the hardware, the dominant paradigm in AI today, particularly Generative AI and LLMs, follows a different path, primarily implemented on conventional hardware.

- **Inner Workings of LLMs and Generative AI**: A Linear Algebra Perspective: Modern LLMs are typically based on the Transformer architecture (Vaswani et al., 2017), a sophisticated system fundamentally built upon linear algebra operations performed on vector representations of language and a quick summary below for someone outside the field.

- Vector Embeddings: Input text is first broken down into tokens (words or sub-words), which are then mapped into high-dimensional vectors (embeddings). These vectors represent tokens as points in a continuous space, where semantic similarity often corresponds to spatial proximity. Positional information is also encoded vectorially.
- Matrix Transformations & Attention: The core of the Transformer lies in its self-attention mechanism, heavily reliant on matrix operations. Each token's embedding vector is transformed by learned weight matrices into three distinct vectors: Query (Q), Key (K), and Value (V).
- Calculating Relevance (Dot Products): To understand the context, the model calculates the relevance of every other token to the current token by computing the dot product between the current token's Query vector and the Key vectors of all other tokens. This dot product measures the alignment or similarity between these vectors, yielding attention scores.
- Contextual Representation (Weighted Sum): These scores, after scaling and normalization (via softmax), act as weights. The final, context-aware representation for the current token is computed as a weighted sum of the Value (V) vectors of all tokens in the sequence. Tokens deemed more relevant (higher attention scores) contribute more strongly.
- Refinement (Feed-Forward Layers): Additional feed-forward layers, also involving matrix multiplications and non-linear activation functions, further process these contextual representations.
- Training: The vast number of parameters in these weight matrices are learned by training the model on enormous datasets, typically by optimizing the model's ability to predict the next token in a sequence using techniques like backpropagation and gradient descent. This entire process allows the model to capture complex statistical patterns, grammatical structures, and factual associations present in the training data (Brown et al., 2020). The resulting functionality revolves around sophisticated statistical pattern matching and sequence generation, leading to impressive capabilities and sometimes surprising "emergent abilities" (Wei et al., 2022).

- **Mapping LLMs to the P→C→C Framework**: When viewed through the P→C→C lens, LLMs present a mixed picture:
  - Level 1 (Prediction): There is a strong functional analogy here. Sophisticated sequence prediction, enabled by the linear algebra machinery described above, is the core mechanism driving LLMs. However, this prediction is fundamentally statistical, based on correlations learned from static data, contrasting sharply with biological Level 1 prediction, which is embodied, action-oriented, and driven by minimizing error signals from real-time interaction with the environment. Still, it remains a speculative possibility (as raised in Chapter 1) that by modeling the output of human intelligence (language) at such scale, these models might implicitly capture or approximate some functional regularities of the underlying human predictive cognition.
  - Level 2 (Context): Here, the divergence is significant. LLMs derive context via the attention mechanism operating on the input sequence and statistical relationships

learned during training. This context lacks the deep grounding in embodied experience, interaction history, and flexible, dynamic updating characteristic of biological Level 2 context (Chapter 1; Varela et al., 1991; Shapiro, 2019). The linear algebra operates within an abstract symbolic space, disconnected from the rich, multi-modal, and temporally extended integration seen in biology. This limitation contributes to documented shortcomings in robust common sense reasoning, handling true novelty beyond training data patterns, and deeper situational understanding (Lake et al., 2017; Mitchell, 2019; Bender et al., 2021).
    - Level 3 (Consciousness): Current LLMs are absent at this level. They possess no integrated functional awareness geared towards goals like survival, no subjective experience, and no self-representation emerging from embodied context as defined in the P→C→C framework (Chapter 1). Phenomena like "hallucinations" (generating plausible but false or nonsensical information) underscore their operation as sophisticated pattern generators rather than systems with integrated, grounded understanding.
- **Contrasting with Biological Circuitry**: The differences extend beyond function to implementation, especially when considering the computational mechanisms:
    - **Integration**: LLMs lack the seamless, multi-level integration of prediction, context, action, memory, and motivation seen in biological systems, where computations analogous to linear algebra (if they exist) are deeply embedded within sensory-motor loops and memory systems. Biological systems exhibit an intrinsic coupling between information processing and physical structure, from molecular interactions to network connectivity, which current AI architectures do not replicate (Walker et al., 2016).
    - **Embodiment**: Their disembodied nature fundamentally limits the types of learning and understanding they can achieve compared to organisms interacting with a physical world.
    - **Efficiency**: Training and running large LLMs requires massive computational resources and energy, often megawatts, dwarfing the brain's ~20W budget. The precise, large-scale matrix multiplications are energy-intensive on current hardware (Horowitz, 2014). Biological computation, leveraging synaptic efficiency and potentially analog processes, achieves vastly greater energy efficiency (Attwell & Laughlin, 2001).
    - **Learning**: LLMs typically undergo large-scale batch training, resulting in a static model between major retraining cycles or releases. This contrasts sharply with biological learning, which is a continuous, lifelong process. The brain relies heavily on local, activity-dependent synaptic plasticity rules (Hebbian, STDP), enabling ongoing, incremental adaptation (Malenka & Bear, 2004; Markram et al., 1997). Furthermore, research indicates that functional brain networks dynamically reconfigure and that connectivity patterns are continuously remodeled by experience, reflecting an ongoing adaptation rather than a fixed state post-training (Ramot et al., 2025). While biology might perform computations describable by linear algebra (e.g., feature extraction via PCA-like mechanisms suggested by Oja, 1982 or Radhakrishnan et al., 2024), the dynamic and continuous nature of its learning process, driven by constant interaction and experience-dependent plasticity, is fundamentally different from the periodic, offline training of current LLMs.

- **Robustness**: While powerful within their training domain, LLMs can be brittle when faced with out-of-distribution data or adversarial inputs, unlike the adaptability often shown by biological intelligence.
- **Representation**: LLMs use dense, high-dimensional vectors. Biology might use sparse, dynamic, noisy population codes intrinsically linked to real-world events (Georgopoulos et al., 1986; Pouget et al., 2000). Transformations are performed via complex electrochemical processes at synapses and dendrites, not explicit digital matrix multiplications (Kandel et al., 2013; Stuart et al., 2008).

- **Energy Budgets and Bottlenecks**: The energy demands of large foundation models are substantial. Training can consume energy equivalent to the annual consumption of hundreds of households (Strubell et al., 2019; Patterson et al., 2021). Inference (using the trained model) at scale also carries a significant energy footprint. The primary hardware bottleneck remains the von Neumann architecture, where the constant shuttling of data (model parameters, activations) between memory and processing units consumes a large fraction, sometimes the majority, of the total energy (Horowitz, 2014).

3.5 Enhancing AI by Borrowing from Biology

Given the capabilities and limitations of current AI, insights from biology's P→C→C implementation offer potential avenues for improvement.

- **Towards Robust Prediction and Error Correction**: AI could benefit from incorporating principles like predictive coding, where systems actively predict upcoming inputs and learn from the error between prediction and reality (Rao & Ballard, 1999; Friston, 2005). Implementing feedback loops and mechanisms for active inference (acting to reduce uncertainty) could lead to more robust and adaptive prediction, moving beyond purely statistical sequence completion (Clark, 2015). Specifically, architectures could be designed where prediction error signals from higher levels modulate processing at lower levels, mirroring hypothesized cortical circuits (Bastos et al., 2012) and potentially leveraging recent AI implementations exploring these ideas (e.g., Millidge et al., 2021; Buckley et al., 2017).
- **Building Deeper Context**: Addressing the context gap is crucial. This could involve:
  - Embodiment: Providing AI systems with bodies (robots) or rich, interactive simulated environments to allow learning through interaction and grounding concepts in sensory-motor experience (Shapiro, 2019).
  - Multimodality: Integrating diverse sensory inputs (vision, sound, touch) more effectively than simply concatenating data streams.
  - Bio-inspired Memory & Continuous Learning: Developing AI memory systems that mimic biological hierarchies (short-term, long-term, episodic, semantic) is crucial. Equally important is moving towards true continual learning algorithms that allow AI to integrate new information dynamically and adapt its internal models without requiring complete retraining from scratch or suffering from catastrophic forgetting (Kirkpatrick et al., 2017; Parisi et al., 2019). This would be better to emulate the brain's ability to learn throughout its existence, constantly refining its understanding (Level 2 context) based on new experiences (Ramot et al., 2025). Such systems might also utilize offline

"sleep-like" processes (replay, consolidation) to structure knowledge (McClelland et al., 1995; Born & Wilhelm, 2012).
    - Causality: Moving beyond correlation to build models capable of causal reasoning, potentially incorporating principles from foundational work on causality (Pearl, 2009) and its integration into machine learning (Schölkopf et al., 2021).
- **Exploring Architectures for Higher-Level Integration (Speculative)**: To achieve more integrated, flexible intelligence, AI research might explore:
    - Mechanisms inspired by cognitive theories like Global Workspace Theory (GWT), focusing on architectures that allow information from specialized modules (including context and prediction modules) to be integrated and broadcast for flexible use and goal re-evaluation (Baars, 1988; Dehaene & Naccache, 2001; van Gaal & Lamme, 2011). This suggests AI designs with a central integrative component receiving input from P and C modules.
    - Principles related to information integration measures, such as those from Integrated Information Theory (IIT), as potential metrics or design goals for evaluating the complexity and interconnectedness of AI systems (Tononi et al., 2016).
    - Implementing capabilities for self-monitoring, reflection, and meta-cognition, allowing AI to reason about its own knowledge and uncertainty, drawing inspiration from cognitive science perspectives (Lau & Rosenthal, 2011).
    - Crucial Caution: It is vital to reiterate that these explorations aim to achieve sophisticated functional integration and cognitive capabilities within AI, drawing inspiration from cognitive theories. This must be rigorously distinguished from replicating the specific biological substrate or achieving the subjective experience and survival-oriented functional awareness defined as Level 3 Consciousness in the P→C→C framework. Functional similarity does not imply equivalence of underlying state or experience.
- **Achieving Biological Efficiency**: Bridging the vast energy gap requires progress on multiple fronts:
    - Hardware: Continued advancement in neuromorphic and synaptic computing platforms, exploring the potential of analog, mixed-signal, and other non-von Neumann approaches (Christensen et al., 2022).
    - Algorithms: Developing algorithms that leverage sparsity (activating only necessary components), event-driven processing, and energy-efficient local learning rules inspired by synaptic plasticity. Comparing the energy savings of these approaches against nature's ~20W benchmark will be critical.
- **Alternate Pathways and the Question of Creativity**: Biological inspiration is not the only route. Progress may also come from:
    - Hybrid systems combining the pattern-matching strengths of deep learning with the logical reasoning capabilities of symbolic AI, often termed neuro-symbolic approaches (Garcez & Lamb, 2020).
    - Increased focus on building explicit causal models of the world.
    - Development of entirely novel computing substrates, potentially leveraging optical or even biological materials (Ceze et al., 2019).

- Furthermore, as AI systems become more capable of integrating context and perhaps even simulating potential outcomes (drawing inspiration from the P→C→C hierarchy), the question arises whether this could lead to capabilities analogous to human creativity. Chapter 1 suggested creativity might be a functional consequence of a flexible predictive system exploring novel possibilities. Could AI systems built on similar principles, even if implemented differently, eventually exhibit genuine novelty and valuable generation beyond sophisticated recombination, fulfilling a functional parallel to creativity (Colton & Wiggins, 2012; Jordanous, 2013)? This remains an open and fascinating question.

3.6 An Intermediate Path Forward: Personalized and Private Agentic AI

While the grand challenges of achieving AGI, superintelligence, and orders-of-magnitude power reduction in AI are being pursued, an intermediate and perhaps more immediately impactful path forward lies in the development of **Personalized and Private Agentic AI**. This approach envisions a hybrid architecture where smaller, highly efficient AI models operate locally on user devices, continuously learning and adapting, while interacting with larger foundational models and the broader internet in a privacy-preserving manner.

Architecture and Functionality:

The core idea is a tiered system:

1. **Local, User-Side Agents:** These would be relatively small, power-efficient AI models residing on the user's personal devices (smartphones, computers, wearables, future edge devices). Their primary function would be to continuously learn from the user's direct interactions, personal data, and local environment. This includes:

- **Adapting to User Style:** Learning communication patterns, preferences, work styles, and even emotional nuances.

- **Building Deep Personal Context:** Continuously updating a rich, private model of the user's goals, projects, relationships, knowledge domains, and history. This context would be built from emails, documents, calendars, application usage, and even potentially from physical models or private knowledge bases explicitly provided by the user.

- **Privacy Preservation:** Crucially, the raw personal data used for training and context building would remain on the user's device, not shared with centralized cloud services unless explicitly permitted for specific tasks.

- **Power Efficiency:** These local models would need to be designed for extreme power efficiency to run persistently on edge devices without significant battery drain, leveraging principles of neuromorphic computing or highly optimized traditional architectures.

2. **Interaction with Foundational Models & External Resources:** The local agent would act as a smart intermediary and personalized filter. When broader knowledge or more complex reasoning is required, the local agent could:

   - Formulate anonymized or privacy-enhanced queries to larger foundational models (LLMs, multimodal models) in the cloud.
   - Retrieve information from the internet or other services.
   - Integrate the responses from these external resources with the user's deep personal context to provide highly relevant, personalized, and actionable outputs.

Aligning with the P→C→C Framework Locally:

This agentic model offers a pathway to implement aspects of the P→C→C framework at a personal level:

- **Prediction (Level 1) - Personalized:** The local agent would excel at predicting the user's immediate needs, anticipating information requirements, and adapting its interaction style based on learned patterns. For example, predicting the next word in the user's writing style, the next likely app to be opened, or the information needed for an upcoming meeting.
- **Context (Level 2) - Deeply Private & Personalized:** This is where the framework shines. The local agent would continuously build an unparalleled Level 2 context, deeply interwoven with the user's private data and interaction history. This context would be far richer and more nuanced for personal tasks than any generalized model could achieve. It would understand the implicit connections between different pieces of the user's information.
- **"Consciousness" (Level 3) - User-Centric Functional Awareness:** Within the P→C→C definition, the local agent could develop a form of user-centric "functional awareness." This wouldn't be subjective consciousness, but rather an integrated model of the user's persistent goals, values, and operational environment, allowing it to proactively assist, manage information, and make decisions aligned with the user's long-term interests, all based on its deeply personalized context.

Continuous Learning: Emulating Biological Adaptation:

A key differentiator from current LLM paradigms is **continuous, incremental learning** on the user side. Unlike foundational models that are statically trained between large, periodic updates, these local agents would constantly adapt and refine their models based on every new piece of information and interaction. This mirrors biological learning, where functional brain networks dynamically reconfigure with experience (Ramot et al., 2025). This continuous adaptation is vital for maintaining relevance and usefulness as the user's life and needs evolve.

Contrasting with Other Agentic Frameworks:

This vision of personalized and private agentic AI contrasts with some other emerging concepts of "agentic AI" in several ways:

- **Centralization vs. Decentralization:** Many current discussions around AI agents imply powerful, highly autonomous agents primarily orchestrated and run from the cloud. The proposed model emphasizes a significant component of decentralized, user-side intelligence and control.
- **Privacy Focus:** While all responsible AI development considers privacy, this model places it at the architectural core by keeping sensitive data and continuous learning local by default.
- **Depth of Personalization vs. General Task Execution:** Some agentic frameworks focus on agents capable of executing complex tasks across various web services (e.g., booking flights, managing calendars based on high-level instructions). While the proposed model could support such tasks, its primary emphasis is on deep, continuous personalization and context-building derived from the user's entire private digital sphere, leading to a more symbiotic assistant.
- **Continuous Local Learning vs. Prompt Engineering/Fine-tuning:** Current interactions with foundational models often rely on sophisticated prompt engineering or occasional fine-tuning on specific datasets. The proposed local agents would engage in genuine, ongoing learning and model adaptation on the user's device.
- **Data Ownership and Control:** This model inherently gives users more direct ownership and control over the data that shapes their personal AI's behavior and understanding.

The Path to Power Efficiency and Mass Adoption:

For such local agents to be viable, extreme power efficiency is non-negotiable. This ties back to the grand challenge of learning from biological machinery. If AI models, even smaller specialized ones, can be designed with principles that reduce power consumption by orders of magnitude (as discussed in Section 3.9 of the main article), then embedding this continuous learning and rich context-building on everyday devices becomes feasible. This approach could democratize access to highly personalized and capable AI, overcoming some of the power and cost hurdles associated with relying solely on massive, centralized models for all interactions.

This intermediate path of personalized, private, and power-efficient agentic AI, continuously learning on the user's side, offers a pragmatic way to bring many of the benefits of advanced AI directly to individuals, fostering a more collaborative and privacy-respecting relationship between humans and their intelligent assistants, and potentially serving as a stepping stone towards more complex and integrated intelligent systems.

3.7 The Trajectory Towards AGI and Superintelligence (Speculative Outlook)

Where might these developments lead? The concepts of Artificial General Intelligence (AGI) and Artificial Superintelligence (ASI) represent potential, albeit highly speculative, future milestones.

- **Defining AGI in the P→C→C Context**: Within the framework of this series, AGI could be conceptualized not just as human-level performance on specific tasks, but as achieving robust, general-purpose capabilities across all three levels of the hierarchy: demonstrating adaptive, error-driven Prediction; building deep, flexible, interaction-grounded Context; and possessing an integrated functional Awareness that enables flexible, autonomous goal pursuit and adaptation in complex, novel environments.
- **From AGI to ASI (Artificial Superintelligence)**: A common, though hypothesized, scenario posits that once AGI is achieved, an intelligence explosion could occur through recursive self-improvement, where the AGI rapidly enhances its own cognitive abilities, leading to Artificial Superintelligence (ASI) (Bostrom, 2014). ASI is typically defined as intelligence far surpassing the brightest and most creative human minds across virtually all domains of interest.
- **Potential Enabling Factors for ASI (Hypothetical)**: The transition to ASI might be fueled by advantages unavailable to biological intelligence:
    - Data Access: Instantaneous, comprehensive access to the world's digital information.
    - Memory: Vastly scalable, high-fidelity memory with near-perfect recall and instant accessibility.
    - Computation: Processing speeds potentially millions or billions of times faster than biological neurons, leveraging massive parallelism and potentially novel architectures (e.g., quantum computing).
    - Connectivity & Speed: Communication and coordination at electronic speeds, unconstrained by biological signaling limits.
- **Potential Capabilities and Implications (Hypothetical)**: The advent of ASI, while speculative, could have transformative consequences:
    - Accelerating scientific discovery and technological innovation at an unimaginable pace.
    - Potentially solving many of humanity's most complex challenges (disease, poverty, climate change, resource scarcity).
    - Leading to profound societal, economic, and cultural shifts.
    - However, it also raises significant existential risks if not developed and deployed carefully. The challenge of alignment – ensuring that the goals and behaviors of ASI systems remain beneficial to humanity – is paramount and widely recognized as exceptionally difficult (Bostrom, 2014; Tegmark, 2017). Successfully navigating this transition requires not only technical prowess but also profound foresight and global cooperation.

3.8 Future Directions: Testable Hypotheses and Quantification from P→C→C

To move beyond a purely conceptual framework and enhance its scientific originality, the P→C→C model should ideally generate testable predictions and point towards potential quantification.

- **Testable Hypotheses**:
  - Neuroscience: The framework predicts distinct neural signatures for P, C, and C. For instance, rapid sensory prediction errors (Level 1) might correlate with activity in early sensory cortices, context integration (Level 2) with hippocampal-neocortical loops and working memory networks, and shifts in functional awareness/goal states (Level 3) with large-scale changes in prefrontal cortex activity and global network synchrony. Disrupting sleep consolidation (Chapter 2) should disproportionately impair Level 2 context integration compared to immediate Level 1 prediction.
  - AI Development: The P→C→C hierarchy suggests modular AI architectures. It hypothesizes that systems explicitly separating P, C, and C functions, with biologically plausible information flow (e.g., error signals from C modulating P and C; consolidation mechanisms strengthening C), will demonstrate greater robustness to novel situations and better common-sense reasoning than end-to-end trained monolithic models of similar size. It also predicts that richness of interaction during embodied training will correlate strongly with the depth of Level 2 context achieved.
  - Cognitive Science: Manipulating the predictability of an environment should primarily affect Level 1 processing efficiency, while manipulating the complexity of contextual rules should tax Level 2 resources. Tasks requiring strategic replanning based on integrated past experience should specifically engage Level 3 processes.
- **Towards Quantification (Conceptual)**: While formal mathematical modeling remains a challenge (Legg & Hutter, 2007), we can conceptualize approaches. The 'depth' of Level 2 Context might be related to the predictive power or compressibility of an agent's learned world model. The degree of integration characteristic of Level 3 might be conceptually linked to information-theoretic measures applied to the system's internal state transitions (Tononi et al., 2016) or the dimensionality of its accessible state space. Prediction accuracy and error rates provide direct measures for Level 1. Developing such metrics, even if approximate, is crucial for rigorously evaluating progress in both biological understanding and AI development under this framework.

3.9 Conclusions: Rethinking Intelligence, Co-Designing for an Energy-Efficient Future

- **Synthesis of the 3-Part Series**: *The Energy Enigma and Emerging Paths*
  This three-part exploration, as presented by the author, commenced by positing prediction as the functional cornerstone of biological intelligence. It is argued that this intelligence evolved within a P→C→C (Prediction → Context → Consciousness) hierarchy, primarily to serve the imperatives of survival and procreation (Chapter 1). The series then delved into nature's blueprint, highlighting the biological mechanisms that achieve this sophisticated predictive capacity with remarkable energy efficiency—approximately 20 watts (Chapter 2;

Attwell & Laughlin, 2001). This final chapter bridges these insights to Artificial Intelligence (AI), contrasting the brain's energy frugality with the significant power demands of current AI, especially Large Language Models (LLMs) (Horowitz, 2014; Strubell et al., 2019). The P→C→C framework, when viewed through this comparative lens, is presented not only as a model for understanding intelligence but also as a means to underscore the critical challenge of energy consumption in AI. The author suggests that understanding these biological principles can illuminate intermediate paths forward, such as the development of personalized and private agentic AI, which aim to leverage efficient, locally adaptive models as a step towards more sustainable and user-centric artificial intelligence.

- **P→C→C as a Unifying Framework**: *Towards a Theory of Efficient Intelligence*
  The P→C→C hierarchy, as proposed, offers a conceptual structure for analyzing intelligence across both biological and artificial domains, with a focus on functional capabilities. Its potential value, the author suggests, lies in guiding the quest for a more fundamental theoretical basis of intelligence – one that inherently incorporates principles of computational and energy efficiency. A paramount goal identified is understanding how prediction, context-building, and integrated awareness emerge and interact within low-power biological systems. The framework encourages a shift beyond benchmarking AI solely on task performance, urging consideration of its architectural and energetic viability. This is presented as a crucial step towards developing truly general and sustainable AI, applicable to both large-scale systems and distributed, personalized agents (Legg & Hutter, 2007).

- **Lessons from 4 billion Years of Evolution**: *The Masterclass in Low-Power, Adaptive Computation*
  Nature's extensive evolutionary journey is portrayed as a testament to optimization under strict energy constraints. The biological machinery for intelligence—characterized by functional reuse, synaptic computing, hierarchical memory, and continuous, adaptive learning (Ramot et al., 2025) – is shown to have achieved complex cognition, creativity, and consciousness as emergent properties without demanding exorbitant energy. This, the author contrasts sharply with current AI's reliance on power-hungry, often statically trained models. The fundamental lesson drawn is that high intelligence does not necessarily equate to high energy consumption if the underlying computational principles, including continuous local adaptation (as envisioned for agentic AI), are correctly harnessed. The drive for persistence and propagation, potentially rooted in the physico-chemical imperative of dynamic kinetic stability (Pross, 2022), is argued to have sculpted biological systems that are both highly capable and remarkably frugal.

- **The Grand Challenge**: *Co-Designing AI for Orders-of-Magnitude Power Reduction*
  The series argues for a necessary paradigm shift in AI development: elevating the drastic reduction of AI's energy footprint—by orders of magnitude—to the status of a grand challenge, comparable in importance to achieving Artificial General Intelligence (AGI) itself. This necessitates a concerted effort in co-designing hardware, software, and algorithms, deeply informed by the biological machinery of intelligence. Learning how the brain achieves its remarkable efficiency, potentially through frameworks like P→C→C, is

presented not merely as an academic pursuit but as essential for unlocking AI's potential for widespread adoption and equitable access. Key aspects of this co-design include:

  - Developing novel neuromorphic and synaptic computing architectures that inherently minimize data movement and embrace analog or event-driven processing (Mead, 1990; Davies et al., 2021). These architectures are seen as suitable for both large-scale systems and the power-constrained edge devices envisioned for personalized agentic AI.
  - Designing algorithms that prioritize sparse computation, local learning rules, and continuous adaptation over brute-force, energy-intensive training of massive models.
  - Fostering a deeper theoretical understanding of intelligence that integrates energy efficiency as a core principle.
  - The unique informational architecture of biological systems, where control and information flow are intrinsically tied to physical structure (Walker et al., 2016), is highlighted as a profound source of inspiration for this endeavor.

- **The Promise and Peril of Superintelligence**: **An Energy and Agency Perspective**
  The prospect of Artificial Superintelligence (ASI) is acknowledged for its immense potential to address humanity's grand challenges (Bostrom, 2014; Tegmark, 2017). However, the author raises concerns about the sustainability and accessibility of such powerful systems if AI remains on its current energy trajectory. An ASI tethered to energy-intensive paradigms could see its benefits limited or its development become environmentally prohibitive. Thus, solving the AI energy challenge is presented as intrinsically linked to the responsible and equitable realization of advanced AI's promise. Furthermore, the critical challenge of ASI alignment—ensuring its goals remain beneficial to humanity—will require systems that are not only intelligent but also understandable and efficient. The development of transparent, user-aligned agentic AI, with local control and privacy, is suggested as a potential intermediate step that might offer valuable lessons and safeguards for navigating the path towards more powerful AI.

- **Final Thought-Provoking Statement: The Imperative of Sustainable and Personalized Intelligence**
  The quest to understand and create intelligence is framed as one of humanity's most ambitious undertakings. The author concludes that as this endeavor progresses, the lessons from biological intelligence—particularly its profound energy efficiency and continuous adaptability—must become central guiding principles. The P→C→C framework is offered as one lens through which to explore the functional architecture of such efficient intelligence. Ultimately, the widespread and beneficial integration of AI into society is seen to hinge not just on its raw capabilities, but on its sustainability, trustworthiness, and its capacity for personalization while respecting individual privacy. Embracing the grand challenge of creating powerful, yet low-power, AI, potentially through innovative paths like user-centric agentic systems, is presented as essential for co-designing a future where intelligence, both biological and artificial, can thrive in harmony with global resources and individual human values.

Acknowledgments

Thanks to colleagues, friends, and LinkedIn connections over the years that helped to stress test some of the ideas described in this paper and for various discussions. Thanks to Aneesh Pannala, Sreemala Pannala, Ganesh Kannan, David West, and Dhaval Shah for reviewing this article and providing critical feedback.

Conflict of Interest

The author declares no conflict of interest.

Bibliography

1. Alkire, M. T., Hudetz, A. G., & Tononi, G. (2008). Consciousness and anesthesia. *Science, 322*(5903), 876-880.
2. Anastassiou, C. A., Perin, R., Markram, H., & Koch, C. (2011). Ephaptic coupling of cortical neurons. *Nature Neuroscience, 14*(2), 217-223.
3. Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation* (Vol. 2, pp. 89-195). Academic Press.
4. Attenborough, D. (2010). *First Life*. [TV series documentary].
5. Attwell, D., & Laughlin, S. B. (2001). An energy budget for signaling in the grey matter of the brain. *Journal of Cerebral Blood Flow & Metabolism, 21*(10), 1133-1145.
6. Baars, B. J. (1988). *A Cognitive Theory of Consciousness*. Cambridge University Press.
7. Backus, J. (1978). Can programming be liberated from the von Neumann style?: a functional style and its algebra of programs. *Communications of the ACM, 21*(8), 613-641.
8. Baddeley, A. (2003). Working memory: looking back and looking forward. *Nature Reviews Neuroscience, 4*(10), 829-839.
9. Bailey, C. H., & Kandel, E. R. (1993). Structural changes accompanying memory storage. *Annual Review of Physiology, 55*(1), 397-426.
10. Baker, M. D., Wolanin, P. M., & Stock, J. B. (2006). Signal transduction in bacterial chemotaxis. *Bioessays, 28(1), 9-22*.
11. Bassett, D. S., & Gazzaniga, M. S. (2011). Understanding complexity in the human brain. *Trends in Cognitive Sciences, 15*(5), 200-209.
12. Bastos, A. M., Usrey, W. M., Adams, R. A., Mangun, G. R., Fries, P., & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron, 76*(4), 695-711.
13. Beaulieu-Laroche, L., Toloza, E. H. S., van der Goes, M. S., Lafourcade, M., Barnet, M., Sugden, A. U., ... & Harnett, M. T. (2018). Enhanced dendritic compartmentalization in human cortical neurons. *Cell,* 175(3), 643-651.e14.
14. Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *Proceedings of the 2021 ACM*

Conference on Fairness, Accountability, and Transparency, 610-623.
15. Berg, H. C. (1993). *Random Walks in Biology*. Princeton University Press.
16. Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience, 2*(1), 32-48.
17. Bobrovskiy, I., Hope, J. M., Ivantsov, A. Y., Nettersheim, B. J., Hallmann, C., & Brocks, J. J. (2022). Oldest preserved gut content and meal of an animal. *Current Biology, 32*(23), 5194-5201.e2.
18. Born, J., & Wilhelm, I. (2012). System consolidation of memory during sleep. *Psychological Research, 76*(2), 192-203.
19. Bostrom, N. (2014). *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press.
20. Zong F.J., Min X., Zhang Y., Li Y.K., Zhang X.T., Liu Y., He K.W. 2023. Circadian time- and sleep-dependent modulation of cortical parvalbumin-positive inhibitory neurons. *EMBO J.*, 42(3), e111304
21. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., ... & Amodei, D. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems, 33*, 1877-1901.
22. Buckley, C. L., Kim, C. S., McGregor, S., & Seth, A. K. (2017). The free energy principle for action and perception: A mathematical review. *Journal of Mathematical Psychology, 81*, 55-79.
23. Buck, L., & Axel, R. (1991). A novel multigene family may encode odorant receptors: a molecular basis for odor recognition. *Cell, 65*(1), 175-187.
24. Bullmore, E., & Sporns, O. (2012). The economy of brain network organization. *Nature Reviews Neuroscience, 13*(5), 336-349.
25. Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.
26. Buzsáki G., McKenzie S., Davachi L. 2022. Neurophysiology of Remembering. *Annu Rev Psychol.*, 73, 187-215.
27. Carabotti, M., Scirocco, A., Maselli, M. A., & Severi, C. (2015). The gut-brain axis: interactions between enteric microbiota, central and enteric nervous systems. *Annals of Gastroenterology, 28*(2), 203-209.
28. Carandini, M. (2012). From circuits to behavior: a bridge too far? *Nature Neuroscience, 15*(4), 507-509.
29. Ceze, L., Nivala, J., & Strauss, K. (2019). Molecular digital data storage using DNA. *Nature Reviews Genetics, 20*(8), 456-466.
30. Ye L. et al. (2022). Overview of Memristor-Based Neural Network Design and Applications. *Front. Phys.*, 10.
31. Chklovskii, D. B. (2004). Synaptic connectivity and neuronal morphology: two sides of the same coin. *Neuron, 43*(5), 609-617.
32. Christensen, D. V., et al. (2022). 2022 roadmap on neuromorphic computing and engineering. *Neuromorphic Computing and Engineering, 2*(2), 022501.
33. Church, G. M., Gao, Y., & Kosuri, S. (2012). Next-generation digital information storage in DNA. *Science, 337*(6102), 1628.
34. Clark, A. (2015). *Surfing Uncertainty: Prediction, Action, and the Embodied Mind*. Oxford

University Press.
35. Colton, S., & Wiggins, G. A. (2012). Computational creativity: The final frontier? *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI 2012)*, 21-26.
36. Craven, B. A., Paterson, E.G., Settles, G.S. (2009). The fluid dynamics of canine olfaction: unique nasal airflow patterns as an explanation of macrosmia. *J R Soc Interface.*, 7(47), 933-43.
37. Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience, 3*(3), 201-215.
38. Cryan, J. F., O'Riordan, K. J., Cowan, C. S., Sandhu, K. V., Bastiaanssen, T. F., Boehme, M., ... & Dinan, T. G. (2019). The microbiota-gut-brain axis. *Physiological Reviews, 99*(4), 1877-2013.
39. Davies, M., et al. (2021). Advancing Neuromorphic Computing With Loihi: A Survey of Results and Outlook. *Proceedings of the IEEE*, 109(5), 911-934.
40. Dawkins, R. (1986). *The Blind Watchmaker*. W. W. Norton & Company.
41. Dehaene, S., & Changeux, J. P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron, 70*(2), 200-227.
42. Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition, 79*(1-2), 1-37.
43. Dicke, U., & Roth, G. (2008). Animal Intelligence and the Evolution of the Human Mind. *Scientific American, 299*(2), 72-79.
44. Dong, G., & Golden, S. S. (2008). How a cyanobacterium tells time. *Current Opinion in Microbiology, 11*(6), 541-546.
45. Dunlap, J. C., Loros, J. J., & DeCoursey, P. J. (Eds.). (2004). *Chronobiology: Biological Timekeeping*. Sinauer Associates.
46. Einstein, A. (1922). *The Meaning of Relativity*. Princeton University Press.
47. England, J. L. (2013). Statistical physics of self-replication. *The Journal of Chemical Physics, 139*(12), 121923.
48. Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1*(1), 1-47.
49. Firestein, S. (2001). How the olfactory system makes sense of scents. *Nature, 413*(6852), 211-218.
50. Franco-Obregón, A., & Gilbert, J. A. (2017). The Microbiome-Mitochondrion Connection: Common Ancestries, Common Mechanisms, Common Goals. *mSystems, 2*(3), e00018-17.
51. Friedman, T. L. (2016). *Thank You for Being Late: An Optimist's Guide to Thriving in the Age of Accelerations*. Farrar, Straus and Giroux.
52. Friston, K. J. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences, 360*(1456), 815-836.
53. Fröhlich, F., & McCormick, D. A. (2010). Endogenous electric fields may guide neocortical network activity. *Neuron, 67*(1), 129-143.
54. Furber, S., Lester, D. R., Plana, L. A., Temple, S., & Brown, A. D. (2013). Overview of the SpiNNaker system architecture. *IEEE Transactions on Computers*, 62(12), 2454-2467..
55. Garcez, A. D'Avila, & Lamb, L. C. (2020). Neurosymbolic AI: The 3rd Wave. *arXiv preprint arXiv:2012.05876*.
56. Georgopoulos, A. P., Schwartz, A. B., & Kettner, R. E. (1986). Neuronal population coding

of movement direction. *Science, 233*(4771), 1416-1419.
57. Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience, 30*, 535-574.
58. Goldman, N., Bertone, P., Chen, S., Dessimoz, C., LeProust, E. M., Sipos, B., & Birney, E. (2013). Towards practical, high-capacity, low-maintenance information storage in synthesized DNA. *Nature, 494*(7435), 77-80.
59. Han, B., Sivaramakrishnan, P., & Lin, C. J. (2019). 'Inside out': a dialogue between mitochondria and bacteria. *FEBS J. ;286(4):630–641*.
60. Hasler, A. D., & Scholz, A. T. (1983). *Olfactory imprinting and homing in salmon*. Springer-Verlag.
61. Helm, B., and Liedvogel, M. (2024). Avian migration clocks in a changing world. *J Comp Physiol A 210, 691–716*.
62. Hidalgo, C. (2024, July 1). Physicists: Perhaps Life Is a Unique State of Matter. *Mind Matters News*.
63. Hohwy, J. (2013). *The Predictive Mind*. Oxford University Press.
64. Hordijk, W., & Steel, M. (2018). Autocatalytic Networks at the Basis of Life's Origin and Organization. *Life (Basel), 8*(4), 62.
65. Horowitz, A. (2009). *Inside of a Dog: What Dogs See, Smell, and Know*. Scribner.
66. Horowitz, M. (2014). Computing's energy problem (and what we can do about it). *ISSCC 2014 Digest of Technical Papers*, 10-14.
67. Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *Journal of Neuroscience, 25*(45), 10420-10436.
68. Hull, C. L. (1943). *Principles of Behavior*. Appleton-Century-Crofts.
69. Ishiura, M., Kutsuna, S., Aoki, S., Iwasaki, H., Andersson, C. R., Tanabe, A., ... & Kondo, T. (1998). Expression of a gene cluster kaiABC as a circadian feedback process in cyanobacteria. *Science, 281*(5382), 1519-1523.
70. Jacob, F. (1977). Evolution and tinkering. *Science, 196*(4295), 1161-1166.
71. Jenkins, E. K., DeChant, M. T., & Perry, E. B. (2018). When the nose doesn't know: Canine olfactory function associated with health, management, and potential links to microbiota. *Frontiers in Veterinary Science, 5*, 56.
72. Jensen, O., & Lisman, J. E. (1998). An oscillatory short-term memory buffer model can account for data on the Sternberg task. *Journal of Neuroscience, 18*(24), 10688-10699.
73. Jonas, E., & Kording, K. P. (2017). Could a neuroscientist understand a microprocessor? *PLoS Computational Biology, 13*(1), e1005268.
74. Jordanous, A. (2013). Evaluating computational creativity: A standardised procedure for evaluating creative systems and its application. *Ph.D. Thesis,* University of Sussex.
75. Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S. A., & Hudspeth, A. J. (Eds.). (2013). *Principles of Neural Science* (5th ed.). McGraw-Hill.
76. Karban, R. (1982). Increased reproductive success at high densities and predator satiation for periodical cicadas. *Ecology, 63*(2), 321-328.
77. Kauffman, S. A. (1986). Autocatalytic sets of proteins. *Journal of Theoretical Biology, 119*(1), 1-24.
78. Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., &

Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences, 114*(13), 3521-3526.
79. Koelsch, S., Vuust, P., & Friston, K. J. (2019). Predictive Processes and the Peculiar Case of Music, *Trends Cogn. Sci., 23*(1), 63-77.
80. Korteling, J. E. (Hans)., van de Boer-Visschedijk, G. C., Blankendaal, R. A. M., Boonekamp, R. C., & Eikelboom, A. R. (2021). Human- versus Artificial Intelligence. *Frontiers in Artificial Intelligence, 4*, 622364.
81. Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*. University of Chicago Press.
82. Lake, B. M., Ullman, T. D., Tenenbaum, J. B., & Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences, 40*, E253.
83. Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences, 15*(8), 365-373.
84. Legg, S., & Hutter, M. (2007). A collection of definitions of intelligence. *Frontiers in Artificial Intelligence and Applications, 157*, 17.
85. Lennie, P. (2003). The cost of cortical computation. *Current Biology, 13*(6), 493-497.
86. Levenson, J. M., & Sweatt, J. D. (2005). Epigenetic mechanisms in memory formation. *Nature Reviews Neuroscience, 6*(2), 108-118.
87. London, M., & Häusser, M. (2005). Dendritic computation. *Annual Review of Neuroscience, 28*, 503-532.
88. Magin, R. L. (2004). Fractional calculus in bioengineering. *Critical Reviews in Biomedical Engineering, 32*(1), 1-104.
89. Magin, R. L. (2010). Fractional calculus models of complex dynamics in biological tissues. *Computers & Mathematics with Applications, 59*(5), 1586-1593.
90. Malenka, R. C., & Bear, M. F. (2004). LTP and LTD: an embarrassment of riches. *Neuron, 44*(1), 5-21.
91. Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science, 275*(5297), 213-215.
92. McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. *Psychological Review, 102*(3), 419-457.
93. McNaughton, B. L., Battaglia, F. P., Jensen, O., Moser, E. I., & Moser, M. B. (2006). Path integration and the neural basis of the 'cognitive map'. *Nature Reviews Neuroscience, 7*(8), 663-678.
94. Mead, C. (1990). Neuromorphic electronic systems. *Proceedings of the IEEE, 78*(10), 1629-1636.
95. Millidge, B., Seth, A., & Buckley, C. L. (2021). Predictive coding: a theoretical and experimental review. arXiv preprint https://arxiv.org/abs/2107.12979.
96. Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience, 24*(1), 167-202.
97. Mitchell, M. (2019). *Artificial Intelligence: A Guide for Thinking Humans*. Farrar, Straus and Giroux.
98. Mori, K., Nagao, H., & Yoshihara, Y. (1999). The olfactory bulb: coding and processing of odor molecule information. *Science, 286(5440), 711-5*.

99. Nakajima, M., Imai, K., Ito, H., Nishiwaki, T., Murayama, Y., Iwasaki, H., ... & Kondo, T. (2005). Reconstitution of circadian oscillation of cyanobacterial KaiC phosphorylation in vitro. *Science, 308*(5720), 414-415.
100. Kimura, M. (2021). Prediction, suppression of visual response, and modulation of visual perception. *Frontiers in Human Neuroscience, 15*, 730962.
101. Nicholls, J. G., Martin, A. R., Fuchs, P. A., Brown, D. A., Diamond, M. E., & Weisblat, D. A. (2012). *From Neuron to Brain* (5th ed.). Sinauer Associates.
102. Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology, 15*(3), 267-273.
103. Olshausen, B. A., & Field, D. J. (2004). Sparse coding of sensory inputs. *Current Opinion in Neurobiology, 14*(4), 481-487.
104. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks, 113*, 54-71.
105. Patterson, D., Gonzalez, J., Le, Q., Liang, C., Munguia, L., Rothchild, D., ... & Dean, J. (2021). Carbon emissions and large neural network training. *arXiv preprint arXiv:2104.10350*.
106. Pearl, J. (2009). *Causality: Models, Reasoning, and Inference* (2nd ed.). Cambridge University Press.
107. Poirazi, P., & Papoutsi, A. (2020). Illuminating dendritic function with computational models. *Nature Reviews Neuroscience,* 21, 303-321*.*
108. Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Rev Neuroscience,* 1(2), 125-32.
109. Pross, A. (2011). Towards an evolutionary theory of the origin of life based on kinetics and thermodynamics. *Journal of Systems Chemistry, 2*(1), 1.
110. Pascal, R. and Pross, A. (2022 On the Chemical Origin of Biological Cognition. *Life, 12*(12), 2016.
111. Radhakrishnan, A., et al. (2024). Mechanism for feature learning in neural networks and backpropagation-free machine learning models. *Science, 383*(6689), 1461-1467.
112. Ralph, M. R., Foster, R. G., Davis, F. C., & Menaker, M. (1990). Transplanted suprachiasmatic nucleus determines circadian period. *Science, 247*(4945), 975-978.
113. Ramakrishnan, V. (2002). Ribosome structure and the mechanism of translation. *Cell, 108*(4), 557-572.
114. Ramot, A., Taschbach, F.H., Yang, Y.C. *et al.* (2025). Motor learning refines thalamic influence on motor cortex. *Nature, https://doi.org/10.1038/s41586-025-08962-8*
115. Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience, 2*(1), 79-87.
116. Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience, 7*, 74.
117. Ryan, M.-L. (2003). *Narrative as Virtual Reality: Immersion and Interactivity in Literature and Electronic Media*. Johns Hopkins University Press.
118. Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward Causal Representation Learning. *Proceedings of the IEEE, 109*(5), 612-634.

119. Schrödinger, E. (1944). *What is Life? The Physical Aspect of the Living Cell*. Cambridge University Press.
120. Schuman, C. D., et al. (2017). A survey of neuromorphic computing and neural networks in hardware. *arXiv preprint arXiv:1705.06963*.
121. Seth, A. K., & Bayne, T. (2022). Theories of consciousness. *Nature Reviews Neuroscience, 23*(7), 439-452.
122. Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. *Neuron, 80*(3), 791-806.
123. Shapiro, L. (2019). *Embodied Cognition* (2nd ed.). Routledge.
124. Shapson-Coe, A., et al. (2024). A petavoxel fragment of human cerebral cortex reconstructed at nanoscale resolution. *Science, 384*(6696), eadk4858.
125. Sherry, D. F. (2006). Neuroecology. *Annual Review of Psychology, 57*, 167-197.
126. Shettleworth, S. J. (2010). *Cognition, evolution, and behavior*. Oxford University Press.
127. Smith, T. J. (2012). Watching You Watch Movies: Using Eye Tracking to Inform Cognitive Film Theory. *Psychocinematics: Exploring cognition at the movies,* 165–190. Oxford University Press.
128. Sourjik, V., & Berg, H. C. (2004). Functional interactions between receptors in bacterial chemotaxis. *Nature, 428*(6981), 437-441.
129. Squire, L. R., & Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science, 253*(5026), 1380-1386.
130. Steitz, T. A. (2008). A structural understanding of the dynamic ribosome machine. *Nat Rev Mol Cell Biol., 9(3), 242-53.*
131. Stickgold, R. (2005). Sleep-dependent memory consolidation. *Nature, 437*(7063), 1272-1278.
132. Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and Policy Considerations for Deep Learning in NLP. *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 3645-3650.
133. Strukov, D. B., Snider, G. S., Stewart, D. R., & Williams, R. S. (2008). The missing memristor found. *Nature, 453*(7191), 80-83.
134. Stuart, G. J., Spruston, N., & Häusser, M. (Eds.). (2008). *Dendrites* (2nd ed.). Oxford University Press.
135. Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nat Rev Neurosci, 15, 745–756*.
136. Takahashi, J. S. (2017). Transcriptional architecture of the mammalian circadian clock. *Nature Reviews Genetics, 18*(3), 164-179.
137. Tang, H., et al. (2019). Bridging Biological and Artificial Neural Networks with Emerging Neuromorphic Devices: Fundamentals, Progress, and Challenges. *Advanced Materials, 31(49),* 1902761.
138. Tegmark, M. (2017). *Life 3.0: Being Human in the Age of Artificial Intelligence*. Knopf.
139. Tinoco Jr, I. and Gonzalez Jr, R. L. (2011). Biological mechanisms, one molecule at a time. *Genes Dev., 25(12), 1205-31*.
140. Tononi, G., Boly, M., Massimini, M., & Koch, C. (2016). Integrated information theory: from consciousness to its physical substrate. *Nature Reviews Neuroscience, 17*(7), 450-461.
141. Tononi, G., & Cirelli, C. (2014). Sleep and the price of plasticity: from synaptic and cellular

homeostasis to memory consolidation and integration. *Neuron, 81*(1), 12-34.
142. Tulving, E. (1972). Episodic and semantic memory. In E. Tulving & W. Donaldson (Eds.), *Organization of Memory* (pp. 381-403). Academic Press.
143. Turrigiano, G. G., & Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience, 5*(2), 97-107.
144. van Gaal, S., & Lamme, V. A. (2011). Unconscious high-level information processing: implication for neurobiological theories of consciousness. *Neuroscientist*, 18(3), 287-301.
145. Varela, F. J., Thompson, E., & Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press.
146. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is All You Need. *Advances in Neural Information Processing Systems, 30*.
147. Walker, S. I., Kim, H., Davies, P. C. W. (2016). The informational architecture of the cell. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 374(2063), 20150057.*
148. Wang, X. J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences, 24*(8), 455-463.
149. Wang, Z., Chu, Z., Doan, T. V., Ni S, Yang, M., & Zhang, W. (2024). History, Development, and Principles of Large Language Models-An Introductory Survey. *arXiv preprint arXiv:2402.06853*.
150. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., ... & Fedus, W. (2022). Emergent abilities of large language models. *Transactions on Machine Learning Research*.
151. Wehner, R., & Srinivasan, M. V. (2003). Path integration in insects. In K. J. Jeffery (Ed.), *The Neurobiology of Spatial Behaviour* (pp. 9-30). Oxford University Press.
152. West, B. J., Bologna, M., & Grigolini, P. (2003). *Physics of Fractal Operators*. Springer.
153. Whiten, A., & Byrne, R. W. (1997). *Machiavellian intelligence II: Extensions and evaluations*. Cambridge University Press.
154. Wikipedia contributors. (n.d.). Drive theory. In *Wikipedia, The Free Encyclopedia*.
155. Wiltschko, R., & Wiltschko, W. (2003). Avian navigation: from historical to modern concepts. *Animal Behavior, 65(2)*, 257-272.
156. Hu J-Y., et al. (2020). DNA storage: research landscape and future prospects. *National Science Review, 7*(6), 1092-1107.
157. Borbolis F., et al. (2023). The Crosstalk between Microbiome and Mitochondrial Homeostasis in Neurodegeneration. *Cells, 12*(3), 429.