Sample Efficiency: A Blueprint for the Next Frontier of Exponential Computing Gains After Moore's Law

Dr Pouya Baniasadi Created: Mar 2026

Last updated: Mar 13th, 2026 - Version 0.92 (draft)

Opportunity Space

The Sample Efficiency Gap (The Gap)

What are the physical limits of how fast an intelligent system can learn with limited data and limited energy? How far are our systems from these limits? The biological brain is the physical proof that intelligence can be orders of magnitude more energy efficient, data efficient, robust, and generalised, than what we currently see in artificial intelligence. We do not need to guess if a better, faster, and cheaper way of processing information exists; evolution has already engineered the proof of concept that is running right now on about 20 watts of power. When you look closely at how a human mind actually acquires, processes, and generalizes information from its environment, it becomes immediately apparent that current AI systems are far from optimal in how they learn. Almost any human-level performance where human level performance has been exceeded by AI, this has been only achieved by many orders of magnitude more data and energy than what the brain requires. The mere existence of biological intelligence proves beyond a shadow of a doubt that the ceiling for computational efficiency is orders of magnitude higher than what our most advanced supercomputers and massive server farms are currently achieving.

To truly understand the sheer magnitude of this algorithmic inefficiency, you only need to look at three stark, quantifiable data points across distinct domains of intelligence. First, consider language acquisition: a 10-year-old child achieves language fluency after being exposed to 40 million words of grounded, sensory context. In stark contrast, a Large Language Model (LLM) like GPT-4 required around 10 trillion words to reach its level of linguistic competence, that is a staggering **250,000x** more samples. Second, look at closed-system logic like chess: chess grandmasters have reached mastery of games with roughly 10,000 hours of chess play, amounting to considering around 10 million positions from the first game to their grandmaster status. AlphaZero required 3 trillion simulated positions in its training to surpass human ability, representing a **300,000x gap** in data efficiency. Finally, in physical world navigation, a human teenager learns to drive to human safety standards with about 100 hours of practice, or roughly 3,000 miles, leveraging intuitive common sense physics. The Waymo Driver, meanwhile, required over 200 million real-world miles after being trained in the simulation, and has not reached level 4 autonomy (no human intervention), and frequently fails in many edge cases. This is almost a **70,000x gap with a pre-trained model**. In almost every single domain, superhuman performance of AI is achieved through many orders of magnitude more experience samples than any human. Indeed, if we control for the amount of data, AI does not reach human level.

This staggering gap, *the sample efficiency gap*, reveals just how much room for improvement remains in the future of intelligent systems.

Urgency: The Exponential Wall of Scaling

Here, I explain the basic dynamics of why investing in sample efficiency is extremely timely, and high-return. The breathtaking commercial and technical success of current foundation models, and the **Scaling Laws** that brought them into existence, is undeniable. For our current deep learning, the core insight of the Scaling Laws was simple and powerful:

> *Model performance improves predictably as you increase compute, data, and model size.*

Scaling laws have been amazing, because essentially even without making fundamental algorithmic breakthroughs, we can buy performance by scaling the hardware infrastructure, and data. This meant that throwing ever more resources at transformer architectures could unlock emergent capabilities that have reshaped the global economy and captivated the public.

But this mainstream may obscure the other side of the story, a painful unavoidable mathematical tap. Specifically, it Scaling laws also imply an **Exponential Scaling Wall** that must be climed:

For each increment of model improvement, we need ***exponentially more*** *data and compute.*

Because achieving merely linear, incremental gains in actual model capability dictates that you must inject an exponentially larger amount of energy, data, and capital into the system. For example, to cut a model's error rate in half, you cannot simply double your data; the math dictates you must multiply your data by eight, ten, or even twenty times to squeeze out that same linear drop in errors.

The industry needs to invest ever-increasing amounts (trillions!) to fight this exponential. While demand for AI products is growing rapidly, demand alone doesn't explain why the industry is investing trillions to build exponentially larger data center capacities. The deeper justification is Scaling Laws!

We may also be hitting other obstacles related to the “exponential wall”. Foundation models like the language models have already ingested nearly all of the high-quality, human-generated text that exists in recorded history. Because current models require such astronomical volumes of data, the Data Wall, this exhaustion represents an encounter with the Exponential Wall. Solutions for the data wall are being explored, most prominently synthetic data, feeding AI-generated text back into training; However, here the deep technical problem is that generated data tends to stay within the model's learned distributions rather than representing unexplored information. Regardless, meaningful data must come from interactions with the world, and computation requires physical hardware. Both face the same constraint: exponential scaling in the digital domain eventually collides with finite resources in the physical one.

The eventual need for an industry-wide shift away from *scaling data and compute* to a paradigm of biology-like efficiency is not a philosophical preference or an academic debate, it is a mathematical inevitability. There are reasonable disagreements over whether we are already hitting the exponential wall today or how long that will take, but there can be no disagreement over the fact that we will. Mathematics of physical law, and physical limitations are not negotiable. When a paradigm hits a physical wall, the only way forward is to change the paradigm entirely. The industry will be forced to pivot, not because it wants to, but because it must.

The good news is that the Sample Efficiency Gap is proof that a solution is possible, and it is hidden in the structure of biological computation.

We need Better Learning Architecture not Just Better Hardware

First, what is the solution to the Exponential Scaling Wall? The only mathematical sound way to fight an exponential wall, is an exponential improvement that opposes it.

If hardware improvements continue on *an exponential path*, their impact will certainly be felt on issues of Exponential Scaling Wall. Indeed, if these improvements were keeping up with the demand, we could renew our computers instead of exponentially expanding the computational hardware infrastructure. Moreover, their contribution will not close the sample efficiency gap. This is because hardware and sample efficiency operate on fundamentally different terms: hardware increases the number of available computations you can throw at a problem; sample efficiency decreases the number of required computations to solve it. The point is not that hardware doesn't matter, it is that *an entirely separate axis of exponential improvement exists, and pursuing it is transformative*.

In other words, smaller and faster chips certainly help delay the issues of Exponential Scaling Wall, but they let you run a data-hungry algorithm more quickly without reducing how much data that algorithm needs. The reason a human child learns language so efficiently or to drive so efficiently is not that their neurons use less electricity. It is that their brain's architecture requires fewer examples to extract a more robust representation of the world. The robustness and efficiency are algorithmic advantages, not naturally overcome through faster and more efficient matrix computations.

Furthermore, sample efficiency addresses other fundamental problems that are not related to hardware. Specifically, sample efficiency is the ability to learn more and better from limited data. More elaborately, it is the ability of an intelligent system to extract substantial, generalized learning from a small number of observations. Biology achieves this not by not only accumulating vast statistical correlations across billions of examples, but by building highly *structured internal models* of the world that support prediction.

Sample efficiency acts as a master key that solves a cascade of stubborn downstream problems in AI today. A more sample-efficient system must, by its very nature:

- Achieve higher robustness from fewer observations,
- Represent knowledge at lower degrees of freedom (Occam's razor),
- Grasp the underlying general structures of a problem rather than memorizing arbitrary patterns.

These properties are also needed to enable true zero-shot transfer learning: the ability of an AI to perform complex tasks it has never explicitly been trained on. Improving sample efficiency therefore doesn't address one problem, it simultaneously unlocks progress on generalization, safety, and robustness, the very challenges that currently limit autonomous AI systems.

Why is Sample Efficiency the Second Coming of Moore's Law?

The 1,000,000x Runway

As established in the previous section, relying primarily on scaling to drive AI improvements inevitably requires exponentially more resources—a trajectory that must ultimately contend with the rigid physical laws of data and compute availability. While this presents a structural challenge for the current paradigm, it simultaneously illuminates one of the most significant technological opportunities of our time. Rather than focusing on the constraints of scaling, this dynamic points us directly toward the massive, untapped runway of sample efficiency.

Biological learning clearly demonstrates that highly generalized intelligence can be achieved with a fraction of the data and energy currently utilized by artificial systems. Based on the stark contrast between human cognitive acquisition and the training requirements of current foundation models, we can observe a known, proven efficiency gap ranging from 10,000x to over 1,000,000x. Yet, this biological benchmark merely represents what evolution has managed to engineer; the ultimate theoretical gap between current technology and what is fundamentally physically and mathematically possible may be vastly larger than what biology demonstrates today.

This multi-order-of-magnitude gap must be reframed. It is not a frustrating technical deficit, but rather a profound expanse of untapped economic and technical potential. Bridging this chasm—even by a fraction—has the potential to significantly improve the capabilities, economics, and accessibility of current systems. A runway of this sheer magnitude indicates that the AI industry is not approaching the end of its foundational breakthroughs, but is standing at the threshold of a new vector of growth. For the organizations that successfully unlock the mathematics of sample efficiency, this runway offers decades of compounding returns that could rival or exceed the value created by the initial deep learning boom.

How Exponential Runways Create Revolutions

To truly grasp the scale of this 1,000,000x opportunity, we only need to look at the defining historical precedent of the last half-century: Moore's Law. For fifty years, the continuous, exponential halving of transistor size and cost did not just happen in an academic vacuum; it systematically built the entire modern global economy. The exponential scaling of silicon from massive, room-sized mainframes down to the microchips in our pockets serves as the ultimate proof that traversing a multi-decade technological runway is the most reliable engine for planetary-scale wealth creation.

The core historical lesson of Moore’s Law is that exponential leaps in efficiency do not merely make existing products cheaper or slightly faster; they completely unlock entirely new, previously impossible paradigms. Making a calculator run a thousand times faster didn't just give us a cheaper calculator—it gave us the foundational physics required to invent the internet, mobile computing, and global software ecosystems. The quantitative change in hardware efficiency triggered a qualitative revolution in human capability.

We are now standing at the absolute beginning of the "Algorithmic Epoch," where sample efficiency will play the exact same historical role that silicon scaling played in the 1970s. We must imagine the entirely new, currently impossible creations that will suddenly become reality when AI can learn deeply and rapidly from incredibly sparse data. We are looking at a future of true, fully secure on-device learning, where a highly personalized AI adapts instantly to a single user's nuances on a mobile phone without ever needing to communicate with a massive, energy-hungry data center. We are looking at autonomous robots that can be dropped into entirely novel, unstructured physical environments and learn to navigate them on the fly, just like a biological creature. We are looking at real-time scientific discovery capable of extracting profound ground truths from highly limited, expensive physical experiments. These breakthroughs are physically and economically impossible under the current hardware scaling regime, but they become inevitable once the sample efficiency gap is bridged.

The Asymmetric Potential to Redefine Industrial Leadership

Currently, the artificial intelligence industry operates under a dynamic where market leadership is defined almost exclusively by massive, trillion-dollar infrastructure budgets. The prevailing assumption is that if you do not have the capital to build a gigawatt data center and hoard thousands of GPUs, you cannot compete at the frontier of intelligence. However, the pursuit of sample efficiency introduces a highly disruptive, asymmetric alternative to this arms race. By enabling the development of highly capable intelligent systems at drastically lower computational costs, requiring significantly less training data, and achieving much higher rates of learning, sample efficiency provides a direct path to superior intelligence that simply bypasses the infrastructure monopoly of the incumbents.

This specific technological path creates an environment of extreme capital efficiency for research and development. The foundational mathematical and structural challenges of building sample-efficient architectures do not require you to ingest the entire internet or spin up billion-dollar training runs just to test a hypothesis. Because the core issues of sample efficiency are fundamentally algorithmic, they frequently manifest and can be rigorously proven on fairly modest, highly structured problems. In this domain, organizations are trading massive financial risk for pure technical and mathematical risk. You don't need a billion dollars to prove the math works; you just need the right math.

Because these fundamental algorithmic problems can be solved cheaply in isolated environments and then scaled up to larger domains, we are facing the distinct possibility of rapid, paradigm-shifting changes to the market landscape. Small, highly agile players who focus their resources on innovating in sample efficiency have the legitimate potential to introduce market-rupturing changes virtually overnight. There is an untapped opportunity by attacking the problem of intelligence with fundamentally more efficient technology (fewer data, faster learning) rather than the current exponentially increasing levels of costs, hardware, energy, and diminishing returns of increasing data. Specialized challengers possess the asymmetric potential to completely redefine industrial leadership and topple the established dominances of the current AI giants.

Medium-term Opportunity: World Models

This section presents a high-level path toward dramatically improving sample efficiency. It does not propose a fully detailed or mathematically complete architecture. Instead, it explains why today's dominant frameworks are unlikely to deliver meaningful improvements and identifies the conditions necessary for progress in the alternative architectures. The scope of this section is therefore to inform strategic decisions about where to place, or avoid, future bets of their time and resources, rather than a clear recipe for how to build sample efficient AI.

The perspective advanced here remains contrarian: fundamental mathematical limits will prevent today's heavily funded, end-to-end architectures from ever achieving the sample efficiency improvements. Current end-to-end generative models must therefore give way to more promising methods. The anticipated gains are so substantial that paradigms unable to match them will simply fail to compete, similar to newer faster hardware architectures giving way to newer more efficient hardware during the age of Moore's law. While this stance is gaining broader acceptance as a path towards more capable forms of intelligence (for example evidenced by the shift toward world models championed by figures such as Yann LeCun and Richard Sutton), the direction advocated here goes further. It offers a stricter, mathematically grounded vision for the future of AI: one built on taking more advantage of more special mathematical structures that enable sample efficiency, or in other words, Compact Representations for Abstract Reasoning (CRAR).

The Contrarian View: Abandoning End-to-End Generative AI

Current generative models have undeniably mastered some domains,most notably sequence prediction when learning a distribution over a finite number of discrete outcomes; for example, the next work (or token) in a sequence of text. When operating within the bounded rules of text tokens or distinct categorical choices or labels, these architectures excel at predicting the next logical step in a sequence.

On the other hand, the limitations of the current paradigm are well documented. It is however arguable what the limits of the capabilities are, and which limits are fundamental rather than temporary state of affairs. This article is on the side of arguing that some of the limits are fundamental, and these approaches face a hard limitation when applied to advanced physical AI and complex reasoning environments. The physical world is fundamentally not sequential; it is highly parallel, continuous, and dynamic. The possibility spaces governing real-world physics and abstract reasoning are vastly too large to be effectively modeled as simple finite sequences, rendering pure sequence prediction inadequate for general intelligence.

The foundational inefficiency of these models is rooted in their end-to-end objective: mapping full-resolution inputs directly to full-resolution outputs. Forcing a system to predict every single pixel or fine-grained detail wastes massive computational energy on elements that are either entirely irrelevant to the core task or fundamentally unpredictable based on the available information.

We can clearly illustrate this wastefulness using the example of a rotating camera in a physical room. In an end-to-end generative paradigm, the model is forced to predict the high-granularity details of the objects hidden behind the camera as it pans, a prediction that is unreasonable based on the starting information. This is highly illogical; that specific visual information is neither available in the preceding frames nor relevant to understanding the core geometric transformation taking place.

Furthermore, generative architectures often rely on successive generations or actions that rely on a previous generation. Given there is always a chance \epsilong for error, even if this \epsilong is small, successive generations inherently compound their errors exponentially with every step. Because each output becomes the input for the next cycle, a microscopic error early on rapidly degrades performance, making the system incredibly brittle unless each move is kept completely independent from the last. This is one of the main reasons why generative models start to become dehoherent as the number of successive generations becomes high.

In contrast, biological systems seem to reason and predict in the abstract space, rather than end-to-end. There are many instances of scientific studies documenting this, but this is even accessible through a mental introspection to a reader of this article. For example, imagine an object like a banana, and then attempt to rotate the imagined object it in your head. Note that you do not render the microscopic texture of its skin or the precise reflection of ambient light. You render only the abstract, geometric granularity required to perform the transformation, gaining a massive computational advantage by instinctively stripping away the irrelevant details that current AI models waste energy trying to predict.

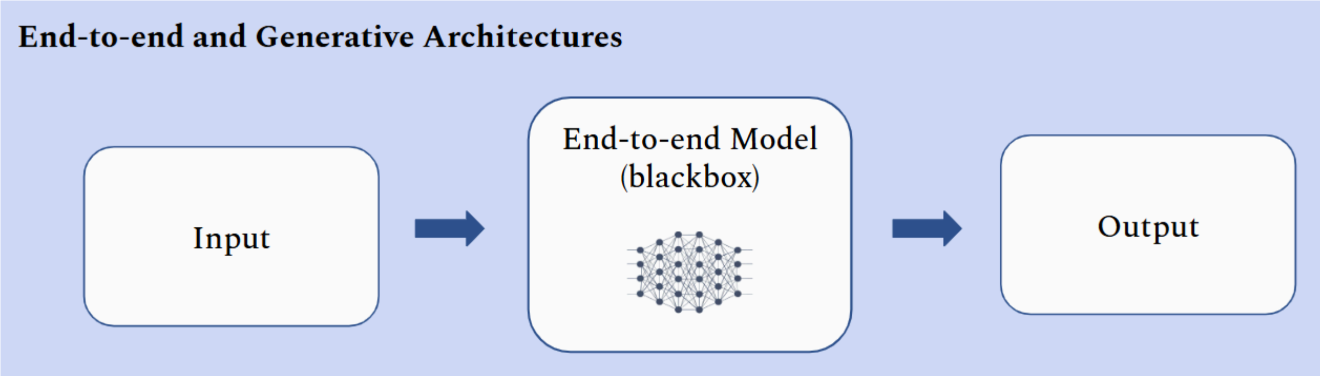

The Next Shoe to Drop: Predictive Abstract World Models

The transitional architectural shift currently taking place across the vanguard of the industry involves decoupling the inputs and outputs, from the minimal representations needed for good predictions. Instead of mapping an input directly to a fully realized output, these emerging systems predict an abstract, latent state of the world. A separate, specialized decoder is then utilized to translate that abstract state back into full resolution only if explicitly required by the user or the downstream task.

The mechanical advantage of this decoupled shift is profound. Because the core cognitive reasoning and future-state prediction happen purely within this lower-dimensional abstract space, the system is structurally forced to ignore irrelevant inputs, noisy artifacts, and unpredictable granular details. The architecture isolates the computational work, focusing the model's energy exclusively on the most critical, causally relevant information driving the environment.

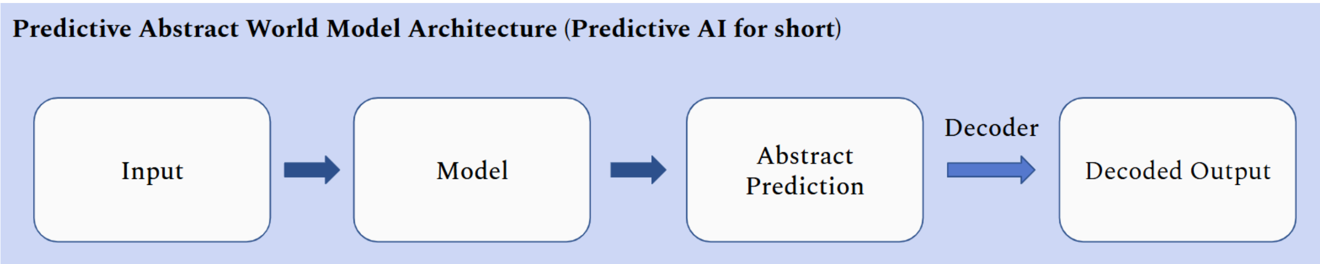

This transitional shift fundamentally redefines the economics of learning and prediction. Operating purely within lower-dimensional, abstract spaces opens the possibility to aggressively slashes the computational burden across the entire lifecycle of the model, simply by avoiding computations that are not necessary outside of the encoder and decoder. During the training phase, predicting compact abstractions requires vastly lower dimensional parameter updates and significantly less data to reach convergence than attempting to map high-dimensional pixel arrays. During inference, processing lightweight, conceptual representations requires a fraction of the energy and memory bandwidth, enabling highly complex reasoning to occur rapidly and cost-effectively.

We can point to Yann LeCun's JEPA (Joint Embedding Predictive Architecture) as prominent, high-level validation of this concept. The core idea of a "joining architecture" that aligns abstract representations rather than generating pixels confirms that the industry's top minds are already recognizing the profound inefficiency of end-to-end generation. The field is steadily moving toward architectures that enable and prioritize abstract prediction as the primary mechanism for understanding the world.

Long-term Blueprint: Reasoning with Compact Representations

The Ultimate Evolution: Compact Representations for Abstract Reasoning (CRAR)

Driven by the unforgiving economics of sample efficiency, architectures will inevitably be compelled to evolve far beyond basic abstract prediction. The ultimate evolution of this trajectory is what we define as Compact Representations for Abstract Reasoning (CRAR). The critical leap in this paradigm is that the "Abstract Reasoning" step does not merely occur inside a standard, black-box latent space. Instead, the reasoning occurs inside highly structured, "well-behaved" mathematical spaces that have been intentionally engineered for computational elegance.

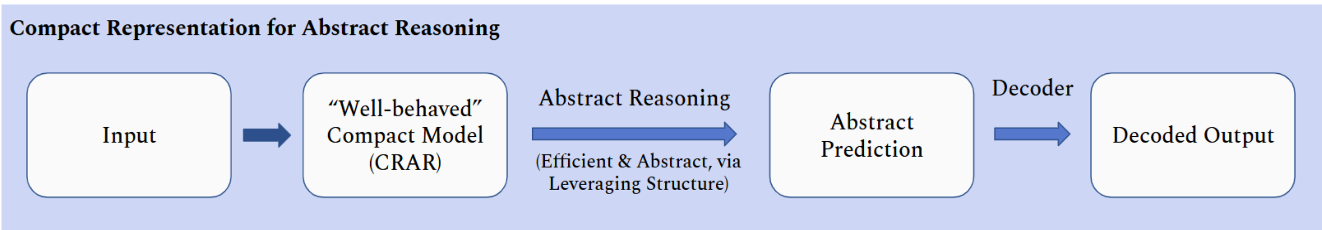

By endowing these compact spaces with specific, targeted geometric and algebraic properties, we allow the system to execute computations and transformations that are vastly superior to fully general-purpose methods. Mathematics and optimization theory clearly demonstrate that highly expressive, structured algebraic systems can make calculations far more efficient than unstructured ones. We can move beyond relying solely on the generalized, unguided search of standard gradient descent, leveraging the inherent geometry of the space to arrive at correct predictions with radically less data.

We must conclude by acknowledging the inherent trade-off of this ultimate evolution. Imposing these structured algebraic and geometric systems makes the model slightly less "general purpose" and less of a true, unconstrained blank slate. However, in exchange for relinquishing this absolute generality, we make the system exponentially more efficient, sample-sparing, and mathematically capable of solving wide classes of complex problems that are currently intractable for today's massive, resource-heavy architectures.

Why Bio-Inspired Learning is a Blueprint

Proceeding from the established reality that pure scaling trajectories will ultimately collide with hard resource walls, sample efficiency stands out as the absolutely necessary path forward for the industry. If we accept that we must drastically reduce the data and compute required to achieve advanced intelligence, we must look for a structural blueprint. The most logical, scientifically sound, and historically proven blueprint for designing the next generation of foundational architectures is to study the only existing, physically proven model of hyper-efficient intelligence in the universe: biological intelligence.

The Cross-Disciplinary Roots of AI

Historically, the most foundational and paradigm-shifting leaps in artificial intelligence did not originate purely from within the silo of computer science; they occurred when researchers combined deep technical backgrounds with profound inspiration from biology and cognitive science. The founding fathers of the current era perfectly illustrate this dynamic. Geoffrey Hinton, widely recognized as the father of deep learning, leveraged a deep background in cognitive psychology to pioneer neural networks. Richard Sutton, the father of reinforcement learning, drew heavily from the study of animal behavior and classical conditioning. Yann LeCun, the father of convolutional neural networks, built his visual processing architectures by directly modeling the mammalian visual cortex. Every one of these pioneers explicitly cites their deep knowledge of cognitive science and biology as a fundamentally important factor in their technical developments.

Therefore, to achieve the next major, multi-order-of-magnitude breakthrough in sample efficiency, the field of artificial intelligence must actively return to this interdisciplinary well of biological inspiration. Simply taking the existing, established computer science models and scaling their parameter counts by another order of magnitude will not yield fundamentally new architectural efficiencies. True structural innovation requires us to look outside the echo chamber of current engineering practices and systematically translate the mechanisms of biological learning into our mathematical and computational frameworks.

Areas to Learn from Biological Intelligence

When we analyze biological intelligence, several core advantages immediately present themselves. These traits highlight the specific functional mechanisms that current models lack.

Transfer Learning

In biological systems, transfer learning is the highly efficient cognitive ability to take a concept, rule, or insight acquired in one specific context and seamlessly apply it across radically different, previously unseen domains. A human does not need to relearn the basic physics of gravity and momentum when switching from catching a baseball to catching a set of keys; the underlying knowledge is transferred instantly, vastly accelerating the learning process for the new task.

Robustness and Recovery

A second defining characteristic of biology is Robustness and Recovery. Biological intelligence possesses the remarkable ability to navigate extreme edge cases, noisy inputs, and completely unexpected environments gracefully. When a human encounters a novel physical obstacle or even suffers localized cognitive disruption, the system does not experience the catastrophic, system-wide failures and bizarre hallucination states that are so common in current artificial intelligence models. Biology degrades gracefully and recovers dynamically.

Generalization and Adaptation

Furthermore, biological systems excel at Generalization and Adaptation. This is the innate capacity of an intelligent agent to operate reliably and make accurate inferences entirely outside the specific, narrow distribution of its previous training experiences. While current AI often fails spectacularly when the test data diverges even slightly from the training data, biological agents continuously adapt to shifting environmental distributions without requiring massive, formalized retraining regimes.

Abstraction

At the core of these capabilities is Abstraction. Abstraction is biology's definitive method for processing the universe; it is the ability to build highly structured, causal, and conceptual internal models of the world. Rather than relying on flat, high-dimensional statistical mappings that simply correlate inputs to outputs, biological brains compress complex environments into elegant abstract representations, allowing them to reason about the world at a fundamental, geometric level.

Learning from Experience vs. Human Knowledge

Crucially, biology achieves all of this by learning more from direct experience rather than relying on the transfer of pre-packaged human knowledge. Current artificial intelligence methods largely succeed by transferring massive amounts of limited, static, human-labeled data into machine weights. In contrast, biological intelligence learns dynamically, interactively, and autonomously through direct physical and cognitive experience. Because high-quality, human-labeled data is a strictly finite resource that fundamentally cannot scale with our technological ambitions, future architectures must inevitably shift away from static knowledge transfer and toward experiential, biological-style learning to continue advancing.

The Unifying Umbrella: Sample Efficiency

It is critical to understand that Transfer Learning, Robustness, Generalization, and Abstraction are not separate, isolated features that can be bolted onto an architecture one by one. Rather, they are the exact, interconnected functional mechanisms that directly and indirectly create sample efficiency. They set the mathematical and structural standards for what it means to learn efficiently. An architecture cannot be truly sample efficient if it lacks these traits; they are the fundamental engines of data reduction.

We can explicitly link every single one of these biological traits to the drastic reduction of data requirements. Abstraction dramatically reduces the dimensionality of the problem space, exponentially decreasing the amount of data needed to map it. Transfer Learning allows a system to reuse past insights, completely eliminating the need to start from scratch and gather basic data for every new task. Generalization minimizes the massive requirement to constantly collect vast, specialized datasets just to cover rare edge cases. Finally, robustness prevents the incredibly data-heavy cycle of constantly retraining a model with massive new data injections every time it encounters an unexpected gap in its knowledge.

In conclusion, intentionally engineering new architectures to structurally possess these specific biological traits is not merely an academic exercise; it is the direct, pragmatic path toward closing the sample efficiency gap. By translating these cognitive mechanisms into formal algorithmic structures, we can fundamentally lower the mathematical requirement for data. This biological blueprint sets the precise stage and the required specifications for the mathematical frameworks that must be developed next.