

Part I — Useful Approximations Framework: The Core Theory

Contents

Part I: Useful Approximations Framework: The Core Theory.	1
Chapter 1: The Map is Not the Territory: Truth as a Functional Imperative	1

Part I: Useful Approximations Framework: The Core Theory.

Chapter 1: The Map is Not the Territory: Truth as a Functional Imperative

What is a circle? A shape made out of each point on a plane that are at some given distance from a given center point. A circle has never existed anywhere. The quantum reality with our 3.28×10^{80} quarks is unable to form such an object. A computer monitor is unable to draw such an object. Yet the circle is known to everyone. It is a **useful shared approximation of reality**. A mental object that helps us in our daily lives.

There is a necessary fundamental gap between the physical reality and the virtual reality constructed in a human brain. If a human draws a circle, the molecules in the ink and their arrangement is always far too complex for a human mind to comprehend. For even a small circle, the complex pattern of the edge of the circle would take an eternity to study in detail. Quintillions of atoms released by your pen in just a short movement on paper. To fully know the exact details and the “truth” about even this small drawing is impossible for a human mind. We do not have the capacity needed to handle reality as it is. But we are fortunate that approximation takes us far.

In addition to the sheer complexity of reality, there is another limit to gaining access to truths about reality. When we study just a small set of atoms, the complexity issue starts to disappear. We have the capacity to fully understand and handle six atoms and their arrangement to some extent. But when looking at such fine details, the quantum world comes in and blurs the view. Planck’s constant and Heisenberg’s uncertainty principle tell us that we actually cannot know the truth about the exact position and momentum of these atoms simultaneously. The particles simply aren’t any more easy to understand in perfect detail than the small circle that we drew.

This isn’t merely a limitation of our measuring instruments; it’s a fundamental property of reality itself at the quantum scale. The very act of observation can influence the state of a particle, meaning that its properties are not known with absolute precision even after being measured. This phenomenon, often referred to as the **observer effect**, further blurs the line between objective reality and our interaction with it. Furthermore, the bizarre phenomenon of **quantum entanglement** suggests that particles can be linked in such a way that the state of one instantaneously influences the state of another, regardless of distance. These quantum realities defy our classical intuition of a perfectly knowable, deterministic universe.

This two-component limiting factor – the overwhelming complexity at larger scales and the inherent probabilistic uncertainties at fine details – profoundly shapes our understanding of reality, leaving us with only an approximate grasp.

The idea that reality cannot be accessed directly, that our perception is inherently limited, has a long and rich history. This fundamental inaccessibility of absolute truth is not a new idea. Philosophers throughout history have grappled with it, perhaps most famously Plato, who, around 400 BCE, presented his evocative Allegory of the Cave.

Allegory of the Cave describes how prisoners, who since birth, are chained in a cave. They are only able to see the wall in front of them. Behind them a fire burns, and between the fire and the wall of

the cave are people carrying objects. The prisoners looking at shadows cast on the back of a cave will consider the shadows the only reality. They learn to name them, talk about them, and predict how they will behave.

This holds until they get freed out of the cave and see what is causing the shadows. The shadows are a similar useful approximation of reality as what our idea of a circle is. The information content of reality, the trees and objects outside the cave, gets projected on a lower-dimensional surface while still containing a lot of useful information about the objects.

Plato's allegory serves as a powerful metaphor for our own **epistemic enclosure**. The shadows on the wall represent our sensory perceptions: they are not reality itself, but rather a useful approximation – a projection of a complex reality onto a limited, accessible surface. Just as the prisoners' 'truth' was confined to the shadows, our own understanding of the world is mediated by our sensory organs and cognitive structures.

Centuries later, in ancient Greece, the philosopher **Pyrrho of Elis** (c. 360 – c. 270 BCE) founded the school of Pyrrhonian skepticism, arguing that true knowledge of reality is impossible. Pyrrho advocated for *epoché*, or the suspension of judgment, on all matters beyond immediate experience. He believed that since our senses can deceive us and our reasoning can be flawed, we can never truly ascertain the ultimate nature of things. This radical skepticism aligns perfectly with the notion that any "truth" we hold is, by necessity, an approximation, and that attempting to grasp an absolute, unmediated reality is a futile endeavor. For Pyrrho, the path to tranquility lay not in finding absolute truth, but in recognizing its inaccessibility and refraining from dogmatic assertions.

Moving forward to the 18th century, the Scottish philosopher **David Hume** further deepened the skeptical tradition. Hume meticulously dissected the foundations of human knowledge, particularly challenging our assumptions about cause and effect. He argued that we never actually *perceive* causality itself; we only observe a constant conjunction of events. Our belief that one event causes another is not derived from reason or direct experience of an inherent connection, but rather from a habit of mind, an expectation formed through repeated observation. This 'habit of mind' is a powerful example of how our cognitive machinery actively constructs a coherent, predictable world from raw sensory input, imposing order where none might inherently exist. It's a testament to the brain's remarkable ability to learn and refine its approximations for survival. For Hume, what we call "truth" about causal relationships is a useful mental construct, not a direct apprehension of an objective, external force.

Consider a red apple before us. Our eyes, like the cave wall, do not capture the 'truth' of every photon's infinite possible wavelength. Instead, specialized receptor molecules in our retina respond to a narrow band of the electromagnetic spectrum, converting a complex wave into a simplified signal – an action potential spike sent to the brain. The precise, objective details of the light are lost, replaced by an internal, approximate experience we label 'red' (roughly 625–740 nm). This 'red' is not the objective property of the apple's surface, but our brain's functional interpretation of a specific set of incoming signals. It is a 'truth' that is useful for navigating our environment, but it is not the absolute, unmediated reality of the apple's atomic structure or its interaction with light.

About 2000 years after Plato, the German philosopher **Immanuel Kant** (18th century) further solidified this notion of an inaccessible reality, providing a more systematic philosophical framework. Kant introduced a crucial distinction between the **noumenon** and the **phenomenon**. The noumenon refers to the 'thing-in-itself' – reality as it exists independently of our perception, unmediated and unknowable. It is the raw, objective, quantum reality, seen through the 3.28×10^{80} quarks and their probabilistic nature.

In contrast, the phenomenon is reality as it appears to us, as it is experienced and understood by the human mind. According to Kant, our minds are not passive recipients of information; they actively structure and organize sensory data through innate "categories of understanding" (like space, time, and causality). Therefore, the world we perceive, the 'truth' we experience, is always a product of both external input and our internal cognitive machinery. This means that our 'truth' is not a direct apprehension of the noumenon, but rather a necessary, approximate internal model – a **functional imperative**.

This "functional imperative" is not merely a philosophical nicety; it is a computational necessity. Imagine a system, biological or artificial, attempting to process every single piece of information from its environment, down to the quantum level, or to perfectly simulate its own internal state in real-time. Such

an endeavor would lead to **computational paralysis**. The sheer volume of data would overwhelm any finite processing capacity, preventing the system from making decisions, taking action, or even maintaining coherence. Such a system would be perpetually stuck, unable to navigate its environment or pursue any goals, effectively ceasing to function. To avoid this infinite regress and informational overload, any sufficiently complex, finite system *must* create a simplified, approximate internal model of itself and its environment. This model, this “functional fiction,” is its working “truth.”

Our sensory organs act as the crucial interface, feeding raw data that our brains then process and interpret into a coherent, usable ‘virtual reality’ or an ‘internal virtual twin.’ This internal world, which we experience as our conscious reality, is not merely a passive reflection; it is a dynamic, constantly updated simulation, not just by new sensory input, but by the brain’s continuous effort to minimize prediction errors, refining its model of reality to better serve our need for survival and success. Kant’s philosophy thus provides a powerful framework for understanding why any finite system, including the human mind, must construct its own version of ‘truth’ rather than accessing an objective, absolute one.

This fundamental inaccessibility of absolute truth forms the bedrock of **Useful Approximations Framework (UAF)**, our proposed functionalist theory of consciousness. UAF argues that consciousness itself is precisely this “necessary functional fiction”—an indispensable internal model that any sufficiently complex, finite system *must* create to manage its overwhelming internal complexity, prevent computational paralysis, and achieve coherent agency. The “truth” we experience is not a window to an objective external reality, but a highly optimized, subjective simplified internal model, designed to enable our interaction with a reality that is otherwise too vast and complex to comprehend directly. This chapter, therefore, sets the stage for understanding consciousness as a representation of reality, not as a mysterious emergent property, but as a fundamental and inevitable computational solution to the problem of optimizing systems control over scarce resources in a competitive environment.

It is exactly here that the boldest contemporary challenge to this thesis enters. Alexander Lerchner has recently argued, in *The Abstraction Fallacy* (Lerchner, 2026), that any view which slides from “the brain runs a useful approximation” to “the approximation is the consciousness” commits a deep ontological inversion. On his account, computation is never a feature of physics itself but a *map* drawn by an already-experiencing **mapmaker** who imposes a finite alphabet on continuous reality; consequently, no amount of additional symbol manipulation can produce the experienter that any computation already presupposes. If he is right, then a digital system can *simulate* consciousness but never *instantiate* it, and the entire trajectory of this book would dead-end at the silicon border. We take this argument seriously enough to devote Chapter 23.5 to a full presentation of it and to a point-by-point reply. The short version of our response is this: the title of this chapter is deliberate. The map is *not* the territory — but, in the only systems we know of that have ever been conscious, the territory has *built* a map of itself from the inside, with no outside mapmaker available. Either the brain is a magical exception, or what looks like an “outside” mapmaker in computational systems can also, given the right physical-organizational conditions, be enacted from within. The rest of this book is the development of that second option.

Useful Approximations Framework (UAF) stands on the shoulders of giants, drawing inspiration from centuries of philosophical inquiry and decades of scientific discovery. Among the most profound insights that have shaped our framework is Thomas Metzinger’s Phenomenal Self-Model (PSM) theory, which compellingly argues that the self is not a mystical entity but a transparent, internal model constructed by the brain. UAF embraces this foundational concept, extending it by rigorously detailing the computational imperatives that necessitate such a model, the mechanisms by which it is continuously refined, and the universal principles that compel its emergence across biological and artificial systems, and even at the cosmic scale.