foundations

MDPI

Article

PCT vs. FEP: A Comparison Between Reorganization Theory and Bayesian Inference

Ty Roachford [1,2,3], **Warren Mansell** [4] and **Rodrigo Pena** [1,2,3,*]

[1] Stiles-Nicholson Brain Institute, Florida Atlantic University, Jupiter, FL 33458, USA; troachford2013@fau.edu
[2] Center for Complex Systems, Florida Atlantic University, Boca Raton, FL 33431, USA
[3] Department of Biological Sciences, Florida Atlantic University, Jupiter, FL 33458, USA
[4] Curtin EnAble Institute, School of Population Health, Curtin University, Perth 6102, Australia; warren.mansell@curtin.edu.au
* Correspondence: penar@fau.edu

Abstract

Perceptual Control Theory (PCT) and the Free Energy Principle (FEP) are two foundational, principle-based frameworks originally developed to explain brain function. However, since their initial proposals, both frameworks have been generalized to account for the behavior of living systems more broadly. Despite their conceptual overlap and practical successes, a mathematical comparison of the two frameworks has yet to be undertaken. In this article, we briefly introduce and compare the philosophical foundations underlying PCT and FEP. We then introduce and compare their experimental and mathematical foundations concretely in the context of bacterial chemotaxis. With these foundations in place, we can use tools from category theory to argue that PCT can be formally understood as a subset of the FEP framework; however, it is worth noting that the mathematical machinery unique to FEP is not required to successfully model bacterial chemotaxis. Finally, we conclude with a proposal for a mathematical synthesis where each framework plays an orthogonal yet complementary role.

Keywords: Perceptual Control Theory; Free Energy Principle; Active Inference; bacterial chemotaxis; category theory; control theory; Bayesian Inference

check for updates

Received: 28 July 2025
Revised: 9 October 2025
Accepted: 23 October 2025
Published: 27 October 2025

Citation: Roachford, T.; Mansell, W.; Pena, R. PCT vs. FEP: A Comparison Between Reorganization Theory and Bayesian Inference. *Foundations* **2025**, *5*, 35. https://doi.org/10.3390/foundations5040035

1. Introduction

1.1. Contextualizing Principle-Based Frameworks

> "There is no such thing as philosophy-free science; there is only science whose philosophical baggage is taken on board without examination."—Daniel Dennett

There is an increasingly large number of scientific frameworks that flip the traditional status quo of creating theories from data on its head; these frameworks focus on generating data that is similar to what one would measure empirically by creating principled models of what could have generated that data. These latter frameworks are what this article calls "principle-based frameworks". Principle-based frameworks empower researchers by using principles to constrain the space of possible models that researchers then construct. In physics, these principles take the form of the *conservation of energy*, the *principle of least action*, and *the principle of maximum entropy*, while in biology, these principles appear as concepts like *autopoiesis*, *homeostasis*, and *natural selection*. Given the growing number of this form of explanation across multiple disciplines and that the systems to which these frameworks purport to explain can overlap, it is pertinent to understand under what circumstances one explanation constitutes a better explanation than another. For example, neural behavior can be modeled at the level of individual molecules following the principles of molecular diffusion or at a higher level of scale using the principles that govern the behavior of electrical circuits. Each of these approaches provides different insights into the phenomena being modeled. However, while these frameworks differ only in scale, frameworks can also vary in kind. Two such principle-based frameworks in the realm of biology that are the subject of this article are Perceptual Control Theory (PCT) and the Free Energy Principle (FEP). Although they are not as fundamental as principle-based frameworks in physics, these frameworks have been rapidly garnering attention since their inception, which calls for a systematic and mathematical comparison of their underlying methodologies.

PCT centers around the principle that for living systems, "Behavior is the control of perception" [1]. When unpacked, this statement asserts that, rather than modeling living organisms as input–output devices—devices that receive an input, or perception, and produce an output, or behavior—living organisms are better modeled as control systems—devices that, like thermostats, produce behavior in service of bringing some physically measurable variable, like temperature, to some particular value, like 90 degrees Fahrenheit. Similarly, the FEP asserts that many, if not all, living systems can be well characterized as minimizing statistical measures of uncertainty called "Variational and Expected Free Energy" [2]. Under the FEP, researchers model organisms as if they were scientists working to infer the causes of their senses in order to bring the environment into their preferred state. This effectively constrains the space of possible models to control system models that are capable of handling uncertainty about the state of the variables they are controlling. In short, both PCT and FEP use principles to cut down the space of models used to model living organisms.

The principles asserted by PCT and FEP have their origins in different places and thus are justified for rather different reasons. PCT uses an experimental methodology called the "Test for the Controlled Variable (TCV)" to determine whether modeling a system as a control system is appropriate by identifying the measurable variables it controls [3]. In short, TCV makes the observation that a measurable variable that is under control will be more difficult to change than if it was not under control. For example, if the temperature of a room is under the control of an air conditioner set at 95 degrees Fahrenheit, then opening a window to let out the hot air would result in less cooling than expected than if this air conditioner was not present. Therefore, the resistance to the change in temperature is evidence of a control system. In this way, the rate of change in measurable variables can be compared against their baseline rates of change (when they are not under control) to identify what measurable variables a system actually controls. In contrast, the FEP relies on whether the researcher is able to characterize the particular living system they are studying as a generative model. In other words, while TCV is about variable selection and system identification, the FEP is only about model specification. Given that these principles do not contradict each other, it is possible to perform TCV on a system and then use the FEP to build a generative model that is nonetheless compatible with PCT.

In this article, we first discuss each of the philosophies underlying PCT and the FEP before comparing their similarities and differences. We then give a summary of the mathematical tools used by PCT and the FEP, namely Control Theory and Bayesian Mechanics, respectively, in addition to examples of their application in modeling bacterial chemotaxis in a simplified setting. We end with a comparison of the performance and form between the two models before concluding our discussion.

1.2. Perceptual Control Theory (PCT)—A Philosophical Summary

Perceptual Control Theory (PCT) was devised by applying the principles of classical control systems engineering to the behavior of living organisms [4–6]. Whilst earlier work had applied control theory to psychology [7,8], PCT made a unique claim at redefining behavior itself as the 'control of perception' [1] and, in turn, the specific organization of components that would achieve this. Thus, rather than assuming that behavior is the output of an organism, PCT proposes that behavior is the controlled input to the system. Take the example of a person balancing an inverted pendulum (see Figure 1). Note that for most limb movements, there is a range of delays and temporal summations of signals within the loop, such that feedback may be intermittent rather than continuous, interspersed among periods of open-loop outputs [9]. Under PCT, balancing an inverted pendulum is not produced by a simple open-loop "input–output" function that maps particular pendulum angles (perception) to particular limb adjustments (action), but rather by a closed-loop function that produces limb adjustments based on the difference between the current pendulum angle and the desired pendulum angle. Therefore, while open loops attempt to balance the pendulum by robotically responding to inputs by producing outputs, closed loops attempt to balance the pendulum by producing outputs that control their inputs (see Figure 1). The dynamic variability of action in the service of control is a characteristic feature of PCT [10], and it emerges because the agent is varying its actions to counteract disturbances, including the unintended effects of the agent's own actions, that would otherwise cause deviations from the agent's goal (in this case, to keep the pendulum vertical). Thus, rather than simply stating a possible way that agents might achieve their goals, PCT provides the mechanism to do so—a closed loop engaged in negative feedback control. In other words, PCT explains how behavior is purposive, in a way analogous to how Einstein's theory of gravity explains how gravity pulls objects to the ground [11].

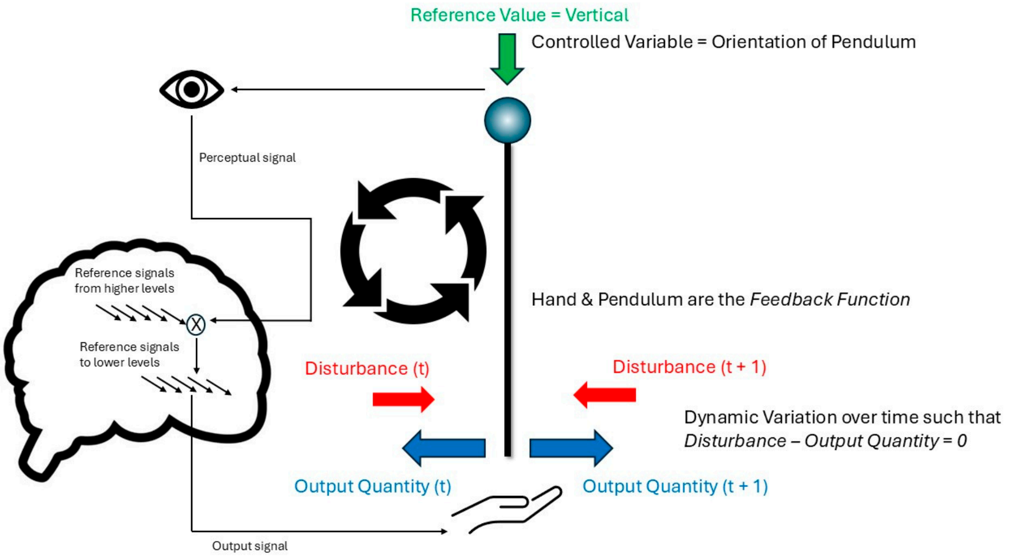

Figure 1. The closed loop of PCT in the case of balancing an inverted pendulum. The diagram illustrates the dynamic variation in hand movement at the bottom of the pole to counteract disturbances to the vertical orientation of the pole, so that the pole is kept close to constantly vertical—the reference value. The source of control comes from higher level control systems in the agent, whose reference values are subtracted from incoming perceptual signals during this task (via the comparator illustrated by a cross in a circle; only one is illustrated for simplicity). The subtraction leads to an error signal that sets lower-level signals for outputs to the muscles of the arm and hand. The diagram is designed to illustrate that this loop is a continuous process that occurs during the performance of this behavior—balancing a pendulum—even in the absence of any learning.

Under PCT, most negative feedback loops that occur in nature and in machines, such as the thermostat, are fixed in the input variable that they control; yet the nervous system allows a variety of variables to be controlled by closed loops at a higher level. Within these hierarchies, closed loops are organized in a branching, hierarchical relationship with one another, so that goals at the top level of the hierarchy are achieved by specifying the subgoals of levels below. At the base of this hierarchy, the organism has a distinct boundary between itself and the environment, and the closed loops here mediate this boundary continuously. Sensors and input functions convert inputs from the environment to perceptual signals that are each compared with their desired perception, or reference signal, to generate an error that is transformed by an output function and actuators, which then drive the activity of feedback functions in the environment (e.g., tools) to counteract the disturbances and keep the error at a minimum. The effect of these actions on the agent's own inputs thereby closes the loop (see Figure 2 for a more detailed diagram).

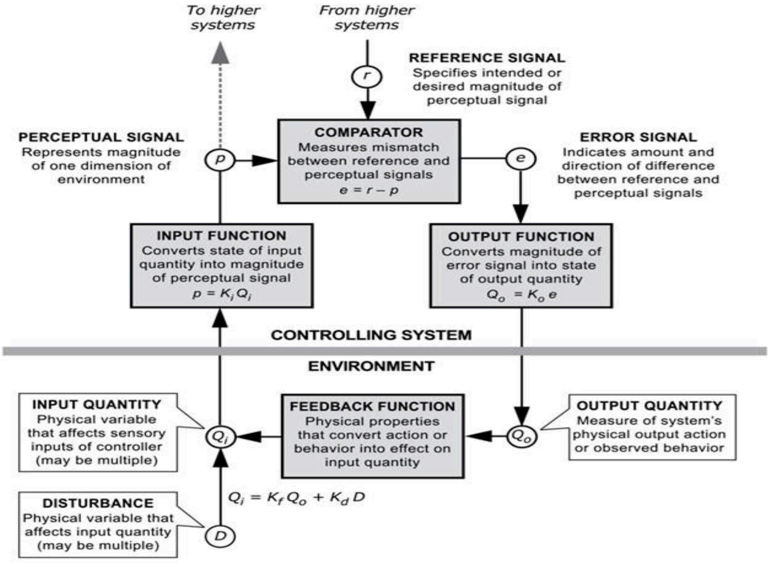

Figure 2. A control unit based on PCT. Reproduced with permission from Dag Forssell. Note: The figure shows the specific components of the closed loop in PCT. Inside the agent, a reference signal specifies the desired magnitude of the perceptual signal, which represents the current magnitude of one aspect of the environment (or self) to be controlled; this is known as the controlled variable (CV); for example, it could be the distance from a target destination. Each aspect of the environment is constructed from an input function, which transforms and integrates signals from lower levels to form the perceptual signal. The difference between the reference signal and perceptual signal is the error signal, which is transformed by the output function within the agent to act through the environment (via the feedback function) to counteract disturbances in the environment that affect the input quantity of the controlled variable.

It is important to note that nothing in the above account describes learning or development—nor how the control units are formed, connected, and optimized. Within PCT, this is achieved via a process known as reorganization. Reorganization describes the in-built tendency for random changes in the properties of control systems, and it is driven by errors generated intrinsically by various subsystems within living control systems. These subsystems monitor deviations in essential variables, such as body temperature or hydration. When intrinsic error increases, the properties of control systems are varied, which is observed as changes in the behavior of the organism, until intrinsic error reduces and the variation pauses (Figure 3).

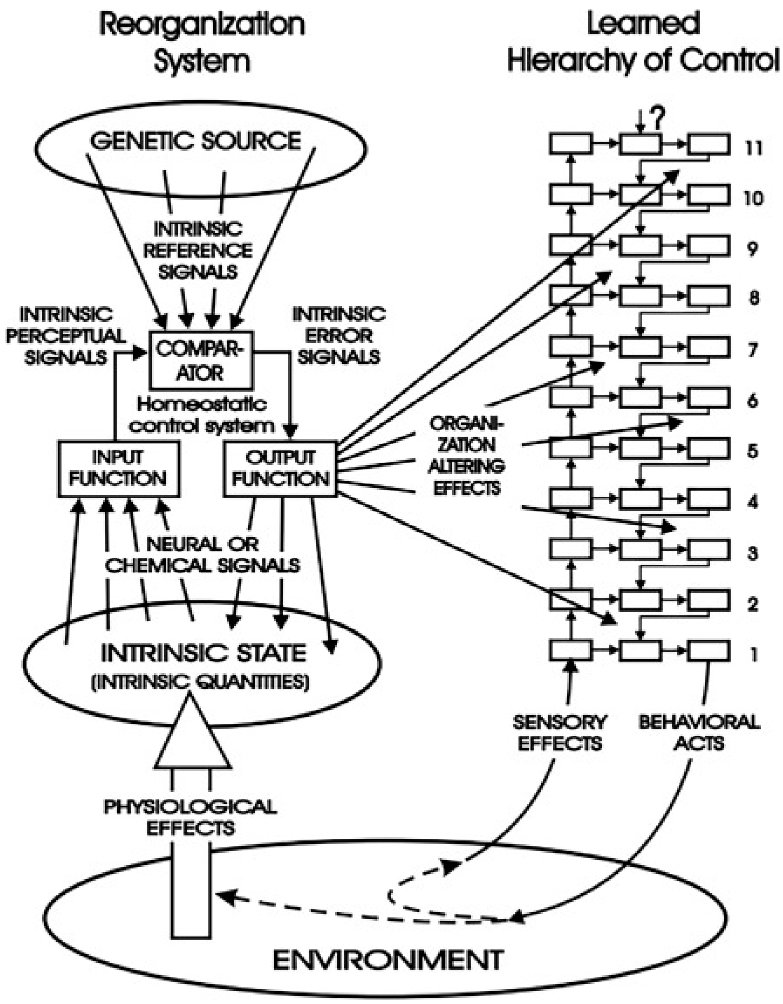

Figure 3. The reorganization system of intrinsic variables and its relation to the learned hierarchy in PCT. The perceptual hierarchy defines and implements goal-directed behavior, whereas the reorganization system prompts trial-and-error changes in the parameters (e.g., gains) of the perceptual hierarchy when intrinsic variables shift from their reference values. Reproduced with permission from Living Control Systems Publishing.

According to PCT, it is not the case that the inside of the organism mirrors or models the outside environment; rather, the organism develops, by contributing its own selected variation, whichever internal structures (e.g., input functions and connections to other units) it requires to control the variables that support its survival. Indeed, from the PCT perspective, objective reality exists, and organisms adapt to reality, not by testing their beliefs, but by constructing and reorganizing internal functions to interact with reality in a manner that meets their needs, i.e., reduces intrinsic error [11].

1.3. Free Energy Principle (FEP)—A Philosophical Summary

The Free Energy Principle (FEP) is often conceptualized as arising from two complementary explanatory perspectives: The "Low Road" and the "High Road". The low road draws from the tradition of Helmholtz's theory of unconscious inference, which proposes that perception is a process of hypothesis testing, where the brain infers the most probable causes of its sensory inputs [2]. For example, one might predict that a cat is nearby upon hearing a "meow". This view has been further elaborated through the Bayesian brain hypothesis, which posits that the brain continuously generates predictions about incoming sensory data and updates these predictions based on the mismatch between expected and actual input—i.e., prediction errors [2,12]. In other words, if one did not see a cat upon hearing a "meow" then one would update their internal model, or "beliefs", such that one's future predictions are more accurate given one's model. Under this perspective, the primary function of perception and cognition is to reduce uncertainty about the hidden causes of sensory signals and is sufficiently described by Bayes' rule. Bayes' rule is a mathematical way to update "beliefs" about a set of hypotheses in a way that systematically reduces uncertainty about a particular observation given a particular belief structure, or internal model of the world. In this setting, beliefs are represented by probability distributions following all the normal rules of probability theory. However, Bayes' rule is often intractable to solve by computers, and the same is expected to be true for organisms. Therefore, the FEP leans into an approximation to Bayes' rule called "Variational Bayesian Inference" which minimizes a statistical quantity called variational free energy (VFE). VFE effectively measures the uncertainty that an agent with a particular belief structure has about the state of their world, given what they see [2,12,13].

In contrast, the "High Road" begins with a more ontological question: "What does it mean for a system to persist over time?"—from this perspective, the FEP is not initially concerned with brains, but rather with the formal conditions under which any "thing" can be identified as existing at all [2,14,15]. The FEP suggests that identifying a system as a coherent entity (i.e., something that maintains its identity over time) implies that it is sparsely coupled to its environment. In practical terms, this means that its internal dynamics are not entirely swamped by environmental fluctuations and that an observer can consistently identify its characteristic features. In other words, an observer is only able to pick out features of a "thing" when that "thing" is sufficiently causally separated from the stuff that is around it. As a toy example, imagine a universe consisting of a single bucket with a divider such that red liquid fills the left side and blue liquid fills the right. The divider allows an observer to distinguish the two liquids as separate "things." If we poke a hole in the divider, red and blue begin to mix, and what was once two distinguishable systems dissolves into one purple liquid. The more holes in the divider, the less causally separated the liquids are, and the faster this indistinguishability happens. This illustrates how identity, persistence, and recognizability of any "thing" rely on a certain degree of causal separation, or statistical independence, from the surroundings (Figure 4).

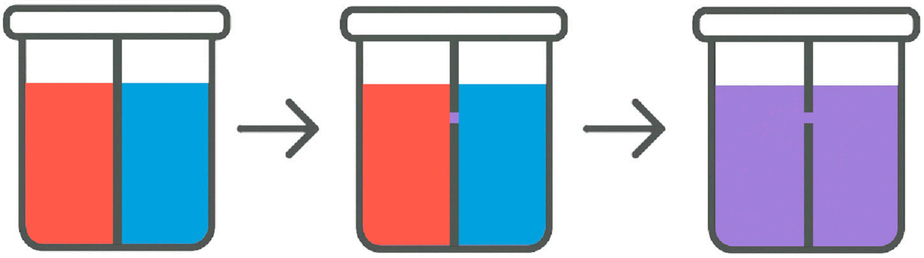

Figure 4. An analogy illustrates how the identity of objects depends on their distinguishability, which itself depends on some degree of causal separation. Here, each liquid is a different "object", while the divider represents the degree of causal separation. The more holes in the divider, the smaller the degree of separation, and the faster the two "objects" become one, an indistinguishable object.

Given that a system maintains a certain degree of causal separation over time, the FEP says that this system can be described as embodying a statistical generative model, over which the minimization of variational free energy allows them to do so. This is performed by selecting a particular system and partitioning it into "internal states", "active states" and "sensory states", while its environment is designated as "external states". This process is known as the construction of a Markov blanket, which mathematically formalizes the sparse coupling between a "thing" and its environment [2,14,15]. In particular, external states are statistically independent of internal states given active and sensory states. In other words, the environment (external state) can only affect the internal state of the organism by acting on that organism's "sensory state". Dually, the internal state of the organism can only affect the environment through acting on its "active state". These active and sensory states, which mediate the interaction between organism and environment, are called "blanket states". Therefore, the FEP is a way to characterize any system that maintains causal separation from its environment by treating action, planning, and learning as Bayesian Inference (Figure 5).

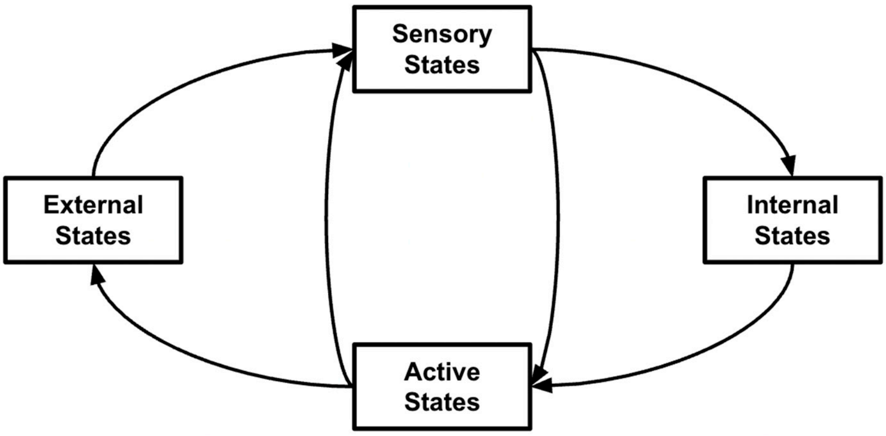

Figure 5. Figure illustrating the Markov blanket construction broadly, where external states constitute the environment while internal states constitute organismal variables that are causally separated from the environment, given sensory and active states. The sensory and active states constitute the Markov blanket mediating interaction between the environment and the agent.

However, as Mel Andrews [16] emphasizes, the FEP is best understood not as a theory with ontological commitments, but as a model structure—a formal mathematical framework that is only helpful when combined with the right interpretation. From this stance, called the instrumentalist stance, the FEP does not claim what "things" really are—it only provides a flexible modeling framework for systems that can be described as minimizing variational free energy. Misunderstanding this distinction has led to conceptual reification, where mathematical terms are mistaken for literal features of the world. Critics such as Thomas [17] and Bruineberg et al. [18] have pointed out the tension between realist and instrumentalist readings of the FEP—especially when it is interpreted as implying that all systems must construct internal models of their environments. Our analysis adopts the instrumentalist view, where the FEP provides tools for describing and simulating systems, without requiring metaphysical commitments about the intrinsic nature of the systems being modeled. To attribute internal models to systems without careful consideration and experimentation would be a category error.

1.4. PCT vs. FEP Philosophical Comparison

Echoing Kennaway [19], we agree that PCT and FEP are widely different in philosophical terms because of their starkly contrasting epistemology. Yet surprisingly, they converge towards analogous structural processes—namely, closed sensorimotor loops of bounded agents that maintain their existence despite external disturbances. The connection between the approaches is illustrated by Friston's references to PCT as "…a landmark theory in our journey of self-understanding". Furthermore, both approaches address fundamental assumptions regarding the nature of living organisms—but each one constructs a closed loop between system and environment to do so. PCT closed loops are constructed with equations describing the mechanics of the parts inside each loop. In contrast, FEP loops are constructed with equations describing an inferential process, which may have no mapping to the physical parts that the closed loop is made of. In other words, PCT models are justified only when they detail precisely what physically measurable variables the closed loop is controlling, in addition to the physical mechanism of how those variables are controlled, whereas FEP models are justified only when they allow researchers to predict, or recapitulate, the data generated by such control. In that respect, PCT achieves a mechanistic definition of behavior by focusing on a ubiquitous physical feature that living and some non-living things must have—namely, control through mechanistic negative feedback. This is in contrast to the FEP, which achieves a statistical definition of behavior by ubiquitously interpreting all things as minimizing free energy over a generative model. Whilst PCT redefines behavior as the control of perception, FEP redefines behavior as the reduction in uncertainty about states expected to bring about preferred perceptions. To put it even more simply, PCT regards control as fundamental and as a precondition for probability estimation, whereas FEP regards information as fundamental and as a precondition for control. At their core, PCT systems have no internal model to maintain and update, and require no information for probability theory to utilize, whereas an internal world model updated by reducing uncertainty through new information is fundamental to FEP.

Despite the above contrast, PCT does utilize information to specify the variables that are perceived and controlled and their relationships with one another—this has been termed control information to distinguish it from the symbolically coded information within information theory. To the degree that this specification is imperfect, for example, when an organism is learning a new variable for control, then an understanding of probability is relevant. Indeed, the process of reorganization in PCT was originally based upon the cybernetic concept of negentropy—the process of forming new structures through minimization of error, a process somewhat parallel to the minimization of surprise in FEP. Thus, PCT and FEP are somewhat complementary rather than contradictory when described in this way; PCT clarifies that the closed loop of perceptual control is a fundamental mechanical process that accounts for behavior itself, whereas it could be argued that FEP clarifies that the closed loop of surprise minimization is a fundamental process that accounts for the development of a structured perceptual system that specifies the appropriate variables for control. Potentially, a synthesis of PCT and FEP would regard both information and control as orthogonal, fundamental, complementary constructs within living agents. Control without information is exemplified by machines like the steam engine governor that are often described as operating purely through physical forces with no internal capacity for using information to alter what it controls [20]. Information without control is exemplified by digital computers that are often described purely in terms of their information processing capacity with no direct capacity for physical control. Neither PCT nor FEP alone is equipped to explain these extreme cases, and although each begins with opposing fundamentals, they converge on the closed-loop structure with insufficient explanatory power—PCT with insufficient consideration of the fundamentals of information, and FEP with insufficient consideration of the fundamentals of control.

Both PCT and FEP are essentially functionalist and hierarchically reductionist accounts. They share the assumption that purposive living systems can be explained and modeled through how simpler functional components are organized, and this lends them both to mathematical modeling that does not incorporate the material substrates (e.g., cells, molecules, bodies) upon which these functions operate and depend. However, readers need to be aware that the paradigm of Choice Causation challenges this view and points to the apparent pre-existence of purposive selection of alternatives within living organisms at the molecular level, which in itself requires an explanation [21].

The above synthesis may necessarily be a simplification, as there are additional unique features of both PCT and FEP that are not discussed here. However, given the multiple components of both theories, there are likely to be a variety of situations that favor the explanatory capacity of one principle-based framework (PBF) over another. Nonetheless, no direct mathematical comparisons appear to have been published, and so the first step is to use a simple test case to compare the fundamental features of both theories in a way that exemplifies both their overlaps and distinctions.

1.5. Why Bacterial Chemotaxis?—Motivation

To compare the practical implications of modeling under FEP vs. PCT, we chose to model bacterial chemotaxis because it is a phenomenon that is both relevant across many areas of biology and conceptually simple. More generally, chemotaxis is the movement of an organism in response to a chemical stimulus—this movement can be attractive (towards the chemical), or repellent (away from the chemical). Organisms typically accomplish this using temporal sensing, whereby they essentially compare the concentration of some chemical at their previous position with the concentration of the chemical at their current position [22,23]. This concentration difference induces a chemical gradient that informs the movement of organisms towards and away from particular chemicals. If the concentration difference is positive, the organism is moving towards the chemical, and if the concentration difference is negative, the organism is moving away from the chemical.

While bacterial chemotaxis is one of the most studied and well-understood examples, the relevance of chemotaxis extends far beyond microorganisms. In multicellular organisms, chemotaxis plays essential roles in early development (e.g., the directed movement of cells during embryogenesis) [24], immune function (e.g., the recruitment of leukocytes to sites of infection or injury) [25], and reproductive processes (e.g., sperm navigating toward the egg via chemical cues) [26]. Chemotactic signaling is also fundamental to wound healing [27], angiogenesis [28], and neuronal migration [29]. Thus, although our modeling focuses on the bacterial case, chemotaxis provides a powerful window into fundamental biological control processes that scale from subcellular mechanisms to whole-organism behaviors.

Bacteria, such as *E. coli*, cannot "control" the direction in which they swim in the typical sense. They operate with only two options: (1) continue swimming in their current direction (a "run"), or (2) trigger their flagella to rotate in a way that disrupts coordinated motion, causing the bacterium to randomly reorient in space (a "tumble"). Even when choosing to continue "running", *E. coli* are unable to maintain a straight path for more than a few seconds due to rotational diffusion, a physical process that gradually randomizes their orientation over time [30]. As a result, *E. coli* must rely on stochastic sampling of the environment, combined with probabilistic modulation of tumbling frequency, to bias their movement toward favorable chemical gradients.

This apparent limitation demonstrates that even with such minimal control, only a binary choice between tumbling and not tumbling, bacteria are capable of goal-directed behavior. The simplicity and robustness of this behavioral loop make bacterial chemotaxis an ideal test case for evaluating the explanatory power of different models of biological control. Specifically, modeling chemotaxis under both the Perceptual Control Theory and the Free Energy Principle allows us to examine how different formalizations of goal-seeking, prediction, and feedback manifest in a system with extreme constraints on action and perception; however, it still succeeds in adaptively shaping behavior to meet environmental demands.

2. Materials and Methods

2.1. Perceptual Control Theory (PCT)—A Mathematical Model of Bacterial Chemotaxis

To begin our study, we reimplemented a PCT-based model of bacterial chemotaxis, first created by Dr. Richard Marken and Dr. Bill Powers [31]. As a PCT-based model, the *E. coli* model is thought of as a control system featuring an input function, comparator, internal reference signal, and output function. The input function represents the mode by which *E. coli* senses its environment, while the output function represents the mode by which *E. coli* modifies its environment. The comparator, aptly named, compares the input signal provided by the input function to the internal reference signal generated by the organism. The comparator then sends this comparison, known as the error signal, to the output function. Through reciprocal interaction with the environment via this action–perception loop, *E. coli* shape their environment such that their input signal matches their internal reference signal. This model of behavioral organization is illustrated in the following diagram in Figure 6.

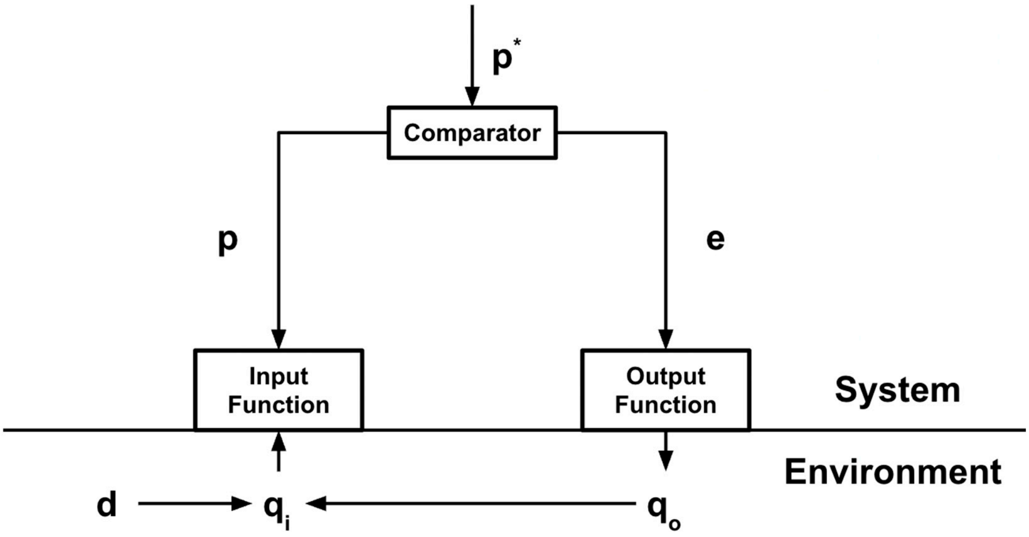

Figure 6. A simplified reproduction of Figure 2, showing a control unit based on PCT. From the diagram we have, **qi** that represents the actual environmental stimulus, like light, or sound; **p** for the "representation" of the environmental stimulus (**qi**) produced by the input function—known as the perceptual signal; d that represents a disturbance to the environmental stimulus (**qi**) produced by environmental factors, like an obstacle that blocks one's view, or earmuffs that block one's sense of hearing; **p*** for the internal reference signal, or signals, generated by the organism; and e for the error generated by the comparator after comparing both the representation of the environmental stimulus (**p**) and the internal reference signal (**p***).

In Marken and Powers, *E. coli* is modeled as a point in a continuous finite 2D space that moves at a constant velocity with respect to some other fixed "target" point representing the attractive chemical stimulus. Internally, *E. coli* is modeled as a control system whose perceptual signal is a function of how quickly the *E. coli* is moving (its velocity), the difference between the swimming angle of the *E. coli* and the target, and the magnitude of the radial gradient. The velocity and magnitude of the radial gradient are kept fixed; therefore, *E. coli's* perceptual signal consists only of its orientation with respect to the target. It is assumed that a substance accumulates inside real *E. coli*, which results in it tumbling. This is modeled as a tumble in response to accumulated error by the comparator. As such, Marken and Powers allow for their *E. coli* model to move in a straight line, thus simulating an environment free of the rotational diffusion normally experienced by real *E. coli*.

As mentioned in the philosophy of PCT section, the process of reorganization under PCT is responsible for the in-built tendency of control systems to change their properties in response to errors generated by various subsystems. As such, the PCT *E. coli* model here, which is only a single control system, follows only an analogous version of reorganization rather than that as exactly coined by [1,4,5]. To quote Powers:

> "In *E. coli* what is reorganized is its direction of swimming, a measure of behavior, which is not the sort of reorganization meant here. But the bacterium's direction of swimming through space can be treated as analogous to the direction of altering

> a set of control system parameters in some abstract multidimensional space, and a tumble can be treated as an alteration in the proportions by which the parameters are being changed, so the direction of metaphorical swimming in this multidimensional space changes."

In other words, tumbling is what happens when *E. coli* is reorganizing to randomly pick a new direction of swimming. Each new direction of swimming is really the result of a collection of control systems, perhaps one controlling each cilia, and each of these systems undergoes a reorganization process on its own and contributes to the final swimming direction of the *E. coli*.

One final simplification to consider explicitly is with regard to how the perceptual signal is defined. In this model, the velocity of *E. coli* and the difference between the target angle and the swimming angle of *E. coli* are used directly. In reality, these variables are calculated in analog using receptors embedded in the *E. coli*'s membrane that detect chemical concentration directly and trigger various molecules that diffuse across the intracellular space, thereby driving its chemotactic behavior. Therefore, our use of angle calculation as a perceptual signal acts as a surrogate function that stands in for the biophysical mechanisms of the organism.

We recreated the Marken and Powers chemotaxis controller line-for-line in TypeScript and instrumented it with additional bookkeeping required for our comparative study. We collected the average distance from the center and the amount of "food" collected for all runs. The "food" is accumulated as a function of the distance from the target once the agent is within a distance threshold, as will be visually illustrated later. The code and a live demonstration can be accessed here: https://github.com/Tyfoods/PCT-vs-FEP (accessed on 24 October 2025). Below is a table summing up the original Marken and Powers implementation (Table 1).

While our recreation is in lockstep with Marken and Powers, for added analysis, we added a circular boundary around the target whose area indicates where "food" can be accumulated. Once inside the boundary, food is accumulated as a real number from 0 to 1 linearly based on the distance from the center. The following illustration (Figure 7) shows the result of a PCT *E. coli* model, briefly after reaching the target, along with its path taken.

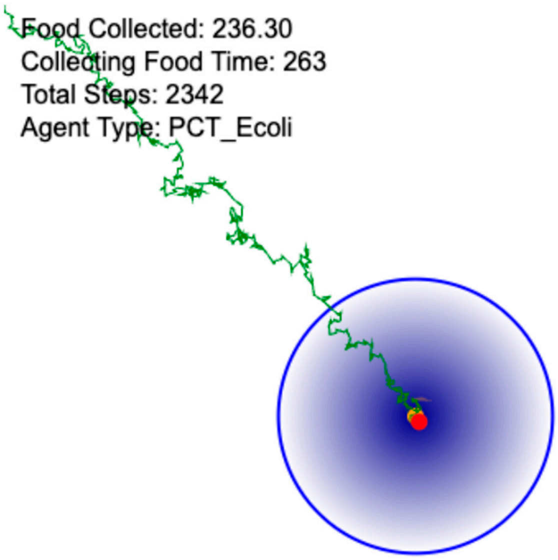

Figure 7. An image from a visualization of a PCT *E. coli* model's trajectory to the target attractant. The red dot indicates the position of the *E. coli*. The yellow dot indicates the position of the target. The green trail indicates the path taken by the *E. coli*. The blue boundary represents the start of the area in which "food" is accumulated.

Table 1. This table summarizes the PCT *E. coli* model algorithm. The variables and parameters are defined.

Step	Algorithm	Description
1	$x := x + v \cdot \cos(a)$ $y := y + v \cdot \sin(a)$	Swim a fixed distance v in the current heading a.
2	$\varphi := \mathrm{atan2}(y_t - y, x_t - x)$	Compute the angle φ to the attractant source.
3	$p := G{\cdot}v{\cdot}\cos(\varphi - a)$	Input function: convert motion through the gradient into a perceptual signal p (temporal derivative of concentration).
4	$e := p^* - p$	Comparator: difference between reference p* and perception p.
5	$dq := K{\cdot}e$	Output function: convert error to the increment dq of an internal "messenger" substance.
6	$q := q + dq$	Accumulate the substance.
7	if $q > Q$ then $q := 0$ $a := \mathrm{random}(0, 2\pi)$	When accumulation exceeds threshold Q, trigger a tumble (pick a new heading uniformly at random).

- p, p*, e—same as previously defined;
- v—velocity that the *E. coli* can move at—Kept Fixed;
- a—angle that the *E. coli* is heading;
- G—scaling constant for the perceptual signal—Kept Fixed;
- K—scaling constant for the error function;
- dq—Increment of accumulated "substance";
- q—total "substance" accumulated;
- Q—tumble threshold;
- Φ—angle to the target, or attractant source;
- (x_t, y_t)—location of the target;
- (x, y)—location of the *E. coli*.

- Parameters for this figure:
  - ○ Velocity = 1 – Fixed;
  - ○ Tumble Threshold = 100 – Fixed;
  - ○ Scaling Constant = 32 − Dynamically Found Via Parameter Search;
  - ○ Reference Signal = 1.49 − Dynamically Found Via Parameter Search.

To find parameters capable of fulfilling the task, a parameter search was conducted, varying the scaling constant and reference signal. For this parameter search, both the target and starting position of the *E. coli* remained fixed.

2.2. Free Energy Principle (FEP)—A Mathematical Model of Bacterial Chemotaxis

We construct our model in the style of a Partially Observable Markov Decision Process (POMDP) to match the previously discussed PCT *E. coli* model. The POMDP framework is used to model agents that do not have access to the full, or true, state of the environment, where the state of the environment is discrete and depends only on its previous state and the agent's previous action [32,33]. First, we will briefly run through the basics of Bayes Theorem and explain how its mathematical structure can often be intractable to calculate. This will lead us to approximate Bayesian Inference, which is more commonly used than exact Bayes when modeling things as engaging in inference. This section can be skipped if one is already familiar with the mathematics behind Bayesian Inference, but reading is encouraged. In the following section, we will detail the structure of our POMDP *E. coli* model in the context of Bayesian Inference.

2.2.1. Bayes Theorem and Variational Bayesian Inference

Bayes Theorem is a mathematically optimal way to update beliefs in light of observations when beliefs are encoded as probability distributions (Figure 8). The theorem is defined as follows:

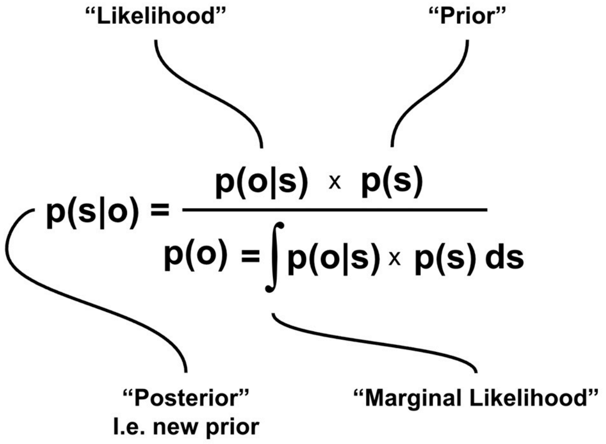

Figure 8. The equations for the exact form of Bayesian Inference are where $p(s \mid o)$, or the "posterior", quantifies the optimal new beliefs about hidden state "s" after one observes an event, or "evidence", which is denoted "o" for "observations". In English, this is usually read "the probability of the hidden state given the observations". In contrast, $p(o \mid s)$—is the "probability of observations given the hidden state". This is called the "likelihood" and is an inversion of the posterior. Therefore, the likelihood quantifies beliefs about what hidden states are most likely to cause a particular observation. $p(s)$ are "prior", or current beliefs about the hidden state, while $p(o)$, also called the "marginal likelihood" of the observation—that is, the probability that a particular observation will occur, averaged over all possible hidden states according to their prior probabilities ($p(s)$). Importantly, while $p(s \mid o)$, $p(s)$, and $p(o)$ are 1-Dimensional Probability Distributions, $p(o \mid s)$ is matrix such that the product of $p(o \mid s)$ and $p(s)$ is the joint probability distribution $p(s,\ o)$ which is referred to as the agent's "generative model".

Using Bayes' Theorem exactly requires one to calculate the marginal likelihood $p(o)$, which is often intractable when the space of hidden states is large or continuous, because it requires integrating over all possible hidden states, weighting each by its prior probability, and multiplying by its likelihood.

$$p(o) \ = \ \int p(o \mid s) \cdot p(s)\, ds \tag{1}$$

This computational difficulty motivates the use of approximate inference methods, such as the variational inference this article employs, which replaces exact Bayesian Inference with optimization over a simpler family of distributions. Variational Bayesian Inference avoids intractable marginal likelihood calculations by introducing an approximate posterior distribution which is found in a way that minimizes a statistical quantity called "Variational Free Energy". Concretely, this is performed by taking the minimizing the Kulback-Liebler Divergence between the approximate posterior distribution $q(s)$, and the true posterior distribution $p(s \mid o)$. Mathematically, this is written as follows:

$$KL[q(s) \ \ || \ \ p(s \mid o)] = \int q(s) \log\left(\frac{q(s)}{p(s \mid o)}\right) ds \tag{2}$$

However, since the true posterior $p(s \mid o)$ requires computing the marginal likelihood $p(o)$ we rewrite this using Bayes' Rule:

$$KL[q(s) \ || \ p(s \mid o)] = \int q(s) \log\left(\frac{q(s)}{\frac{p(o|s) \cdot p(s)}{p(o)}}\right) ds \tag{3}$$

Rearranging, we get

$$KL[q(s) \ || \ p(s \mid o)] = \int q(s) \log\left(\frac{q(s)}{p(o \mid s) \cdot p(s)}\right) ds + \log p(o) \tag{4}$$

where $\int q(s) \log\left(\frac{q(s)}{p(o|s) \cdot p(s)}\right) ds$ is the variational free energy, and $\log p(o)$ is the log marginal likelihood, or model evidence, and when written as $-\log p(o)$, is called the "surprise". This surprise is precisely the same quantity defined by Claude Shannon, which in this case quantifies how improbable observations are, given the generative model [2,34]. Rearranging the equations further we see the following definition:

$$-\log p(o) = \int q(s) \log\left(\frac{q(s)}{p(o \mid s) \cdot p(s)}\right) ds - KL[q(s) \ || \ p(s \mid o)] \tag{5}$$

Since the KL Divergence is greater than or equal to zero, Equation (5) becomes the following:

$$0 \leq \int q(s) \log\left(\frac{q(s)}{p(o \mid s) \cdot p(s)}\right) ds + \log p(o) \tag{6}$$

And, therefore, Equation (6) becomes the following:

$$-\log p(o) \ \leq \int q(s) \log\left(\frac{q(s)}{p(o \mid s) \cdot p(s)}\right) ds \tag{7}$$

Therefore, variational free energy is an upper bound on surprise, and minimizing the variational free energy minimizes the KL divergence, which results in an approximate posterior distribution that closely matches the true posterior.

To obtain an intuition for this, consider the following reorganized form of Bayes' Theorem made to reflect an unknown marginal likelihood (Figure 9):

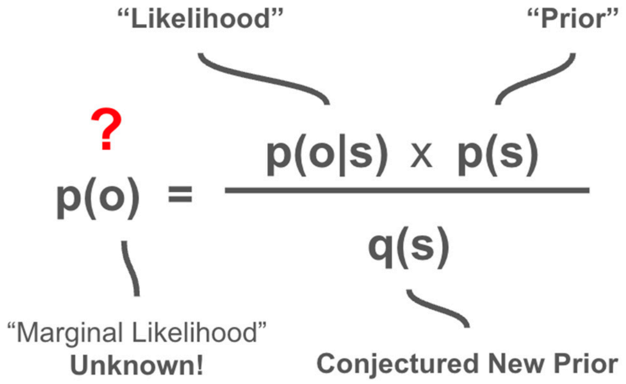

Figure 9. Rearrangement of Bayes' Theorem to reflect an unknown, or intractable marginal likelihood.

This equation illustrates that if we were able to guess the correct posterior, which here takes the form q(s), then this would allow us to calculate the marginal likelihood. Variational free energy minimization formalizes this procedure by allowing for adjustments based on a statistical gradient. In practice, minimizing variational free energy requires gradient-based optimization of the parameters of the approximate posterior distribution

$q(s)$; While a full exploration of these techniques is beyond the scope of this article, numerous algorithms exist for performing this minimization. Notably, some of these algorithms lend themselves more readily to biological plausibility, which has prompted investigation into how physical systems might implement approximate Bayesian Inference in a neural or otherwise materially grounded manner [32].

2.2.2. A POMDP Model of Bacterial Chemotaxis

The POMDP *E. coli* Model consists of four different components named the likelihood (A), transition matrix (B), preference (C), and prior (D). Together, these components characterize the agent's generative model, which describes how the agent learns, plans, and takes actions against what it expects will satisfy its preferences. Each of these components has the following definition:

- **A (likelihood matrix):** The probability of a particular observation, given a particular environmental state—$p(o \mid s)$.
- **B (transition matrix):** The probability of the next environmental state, given the current environmental state, and the action taken—$p(s' \mid s,\ a)$.
- **C (preferences):** The agent's log prior preference vector over observations—C.
- **D (prior):** The probability that the environment is in any given state—$p(s)$.

To make the POMDP *E. coli* model as comparable as possible to the PCT *E. coli*, the POMDP model employed the same "Input Function" over the same "Input Quantity". In other words, the input quantity here was the angle of travel with respect to some arbitrary reference point, while the perceptual signal was generated by the aforementioned equation in Table 1:

$$p := G{\cdot}v{\cdot}\cos\varphi - a \tag{8}$$

The perceptual signal quantifies how quickly the model is moving towards the target. However, since it is a continuous variable, it was discretized into three distinct categories. Since the velocity is constant, and the domain of the cosine function is between -1 and +1, the value of the input function is maximal when the angle between the model and target is zero. Therefore, we use a piecewise function to split the domain of the cosine function into three categories that map intuitively to the angle between the model and target. Values between 1 and $^1/_3$ represent angles between 0 and 120 degrees, values between $^1/_3$ and $-^1/_3$ represent angles between 120 and 240 degrees, and values between $-^1/_3$ and -1 represent angles between 240 and 360 degrees (see Figure 10). In this way, observation labels (0–120, 120–240, and 240–260) were constructed in a way that matches the perceptual signal in PCT models.

The following tables (Tables 2–9) are what the components of the generative model look like as parameterized by the code attached to this paper:

Table 2. Log preferences C for each possible label in the model. Note: While preferences are expressed as a vector, they are not probability distributions; therefore, their values need not sum to 1 or be bounded between 0 and 1. This is because preferences do not encode beliefs about the world but rather reflect relative desirability. For example, a preference vector like [−10, 0, 5] implies that state 0 is strongly avoided, state 1 is neutral, and state 2 is highly preferred. What matters are the relative magnitudes, not absolute probabilistic interpretations.

Observation Label	Log Preferences
0–120	0.00
120–140	−8.48
240–360	−8.48

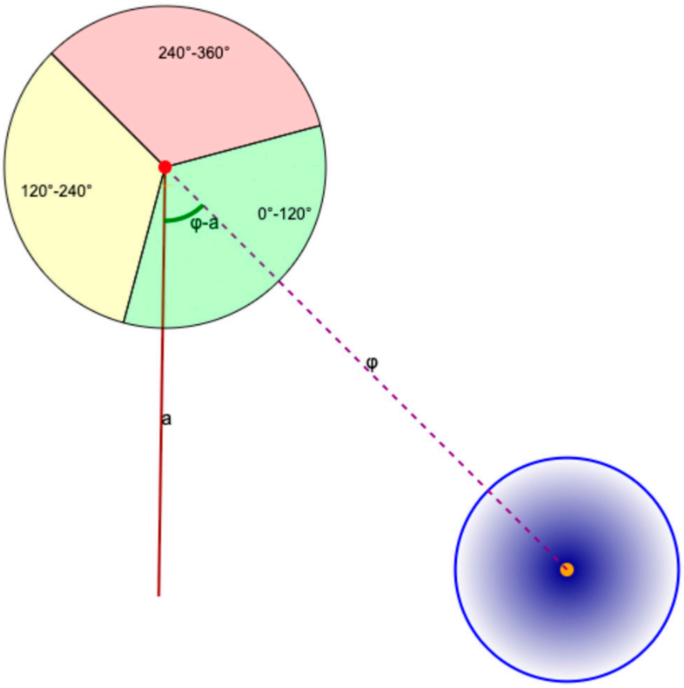

Figure 10. This figure illustrates the discretization of the perceptual signal into buckets that faithfully represent the difference between the angle the *E. coli* is facing and the angle between the *E. coli* and the target. The code reproducing the figure can be found attached to this article.

Table 3. Probability values for each hidden state in the prior $p(s)$—each with an equivalent value, which maximizes the entropy of the distribution and can be interpreted as having no bias towards being some distance away from the target.

State Label	Prior State Probability
Near (to target)	33%
Mid (from target)	33%
Far (from target)	33%

Table 4. The likelihood matrix, which details the probability of being near, mid, or far from the target given the observed angle between the agent and the target. This can be interpreted as having sharp biased beliefs, where the agent believes it is near the target with 100% certainty if the angle it observes is 0–120.

	Near (to Target)	Mid (from Target)	Far (from Target)
0–120	100%	0%	0%
120–140	0%	100%	0%
240–360	0%	0%	1.00%

Table 5. Transition Matrix “Part 1” that defines the probability of the next state given the previous state and the action taken. In this case, the action being taken is “approach”—Therefore, this can be interpreted as the agent believing that if it is near the target, then “approach” will result in being near with a 90% probability.

	Near (to Target)	Mid (from Target)	Far (from Target)
Near (to target)	90%	5%	0%
Mid (from target)	10%	90%	10%
Far (from target)	0%	5%	90%

Table 6. Transition Matrix "Part 2". In this case, the action being taken is "tumble". Therefore, this can be interpreted as the agent having no biased beliefs about where it will be in relation to the target after tumbling.

	Near (to Target)	**Mid (from Target)**	**Far (from Target)**
Near (to target)	33%	33%	33%
Mid (from target)	33%	33%	33%
Far (from target)	33%	33%	33%

Table 7. "hi" corresponds to high probability mass, which is equal to the likelihood sharpness. "lo" corresponds to low probability mass and is calculated with the following function: $\frac{(1-hi)}{2}$ This results in a proper probability matrix where rows sum to 1.

	Near (to Target)	**Mid (from Target)**	**Far (from Target)**
0–120	hi	lo	lo
120–140	lo	hi	lo
240–360	lo	lo	hi

Table 8. This illustrates how the nonlocality bias (nb) is used as a scalar hyperparameter to tune the transition matrix for the "approach" action of the model.

	Near	**Mid**	**Far**
Near	$1 - nb$	nb	0%
Mid	$\frac{nb}{2}$	$1 - nb$	$\frac{nb}{2}$
Far	0%	nb	$1 - nb$

Table 9. This illustrates how the log prior preference vector was encoded with a single scalar hyperparameter.

Observation Label	**Log Preferences**
0–120	0.00
120–140	log(x)
240–360	log(x)

To conduct a parameter search only the likelihood matrix, "Approach" transition matrix, and preferences were allowed to vary. However, given that the POMDP model uses probability vectors and matrices, rather than scalars, these parameters were themselves made tunable by a scalar hyperparameter for a more standardized comparison.

The 3 × 3 likelihood matrix A was parameterized by a single scalar hyperparameter called the "likelihood sharpness". This concentrates the probability mass along the diagonal elements as illustrated in the following table:

The placements of "hi" vs. "lo" constitute different hypotheses about the relationship between angle to target and the value of the hidden state. The value of "hi" therefore constitutes the degree of confidence in that belief structure. Similarly, the transition matrix for approaching was constructed with a single parameter called the "Nonlocality Bias (nb)". The transition matrix is set up so that the way probability mass is concentrated respects the adjacency of different states. In other words, if the current state is "far" we expect the next state to be "near" with a low probability because "mid" must come first. Therefore, the nonlocality bias represents a sort of uncertainty about locality. This is illustrated in the following table:

Finally, the preference vector was encoded simply with a scalar value x between 0 and 1, as illustrated below:

To find parameters capable of fulfilling the task of navigating to the target, a parameter search was conducted varying the likelihood, "Approach" transition matrix, and log prior preference vector. As for the PCT *E. coli* models, both the target and starting position of the *E. coli* remained fixed during the search. The following illustration shows the result of a POMDP *E. coli* model briefly after reaching the target along with its path taken (Figure 11).

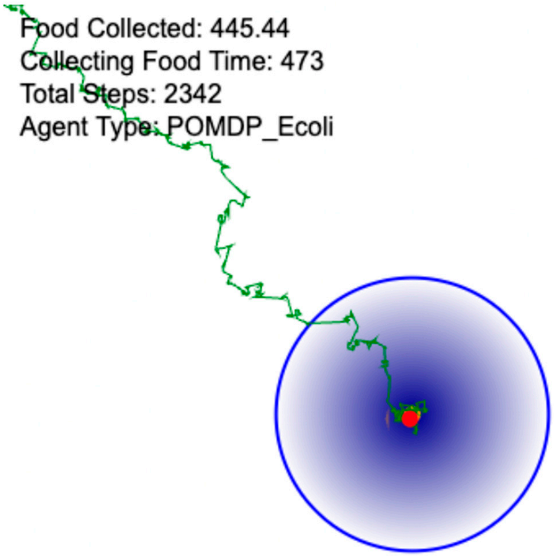

Figure 11. A trajectory of a POMDP *E. coli* model towards the center of a target attractant.

- Parameters for this figure are as follows:
    - Velocity = 1− Fixed;
    - Likelihood Sharpness = 1− Found dynamically via parameter search;
    - Approach Bias = 0.0625− Found dynamically via parameter search;
    - Preference Strength = $\overline{0.36}$ − Found dynamically via parameter search.

2.3. PCT vs. FEP Categorical Formulation of PCT and FEP Models

2.3.1. Category Theory—A Brief Introduction

Category theory is a meta-mathematical framework created by Saunder’s Maclane and Samuel Eilenberg to solve problems in algebraic topology that require a language capable of well-describing both the mathematical structures found in geometry and in algebra [35–37]. However, due to its expressiveness and capacity to describe phenomena in many different areas of science, category theory is increasingly becoming something of a lingua franca for the sciences. Here, we apply category theory to describe both the PCT and FEP mathematical models of agents by constructing categorical diagrams of each framework and conjecturing a mapping from the FEP-agent category to the PCT-agent category. We begin with a brief introduction to the mechanics and terminology of category theory.

A category consists of two main components, objects and morphisms. Objects correspond to mathematical structures like Boolean values, real numbers, or probability distributions. Morphisms correspond to mappings between objects, like operators, functions, or transformations. An object is represented as a capital letter like A, B, or C while morphisms are represented as arrows $\to$ between objects, where the source object is at the tail of the arrow, and the target object is at the head of the arrow. Trivially, each arbitrary object in a category, say X has an identity arrow $X \to X$, which is typically not drawn in categorical diagrams but assumed to exist. In category theory, composition is at the core of how mathematical structures are understood. Given a category with objects A, B, C, D and morphisms $f : A \to B$, $g : B \to C$, and $h : C \to D$, then there exists a composition $g \circ f : A \to B \to C$ that satisfies associativity. In other words, one can bracket the operations however they would like as long as the order of compositions is preserved. More concisely this is written: $h \circ (g \circ f) = (h \circ g) \circ f$.

For a concrete example of a simple category, consider a category with objects $A = People, B = Integers, C = Boolean$ and morphisms:

- $f : A \rightarrow B -$ How many cats does this person have?
- $g : B \rightarrow C -$ Do they have more than one cat?

The composition $g \circ f$ is then:

- $g \circ f : A \rightarrow C = A \rightarrow B \rightarrow C$—Is the answer to the question "Does this person have more than one cat?"

The diagram is illustrated in the following way (Figure 12):

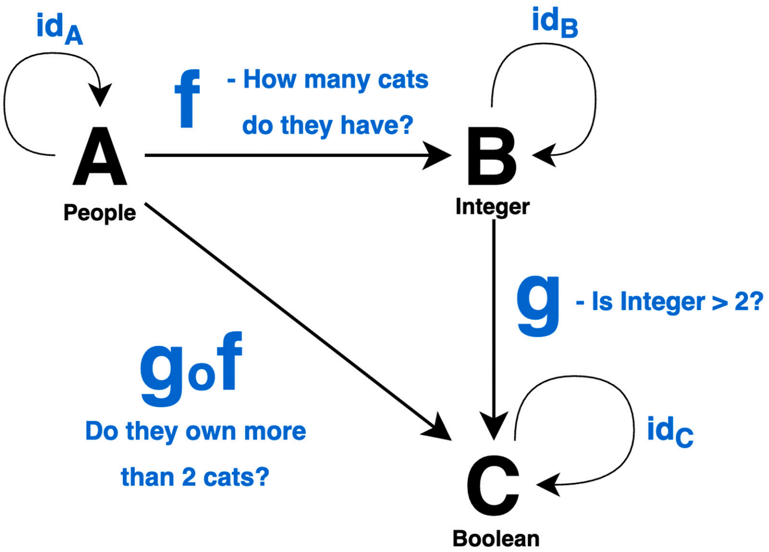

Figure 12. Diagram of a simple category where all morphisms have blue labels that describe how they map their source and target objects. The diagram includes identity morphisms for each object, which are normally not drawn.

While there are morphisms between objects within a category, there are also morphisms between categories called functors. Functors map the objects and morphisms in one category, to the objects and morphisms in another category in a way that preserves composition and associativity. Functors give us a way to formally write the possible ways in which two categories can be related. While the application of functors is where this article ends, beyond functors, there are natural transformations, which are morphisms between functors. This gives one a way to understand the ways in which relationships between categories are themselves related.

2.3.2. Agent Categories and a Functor Between Them

Both the PCT and FEP agent categories were constructed in a way that broadly reflects the core way that these agents are implemented mathematically without becoming overly detailed. This happens in two ways: first, we compare only categorical diagrams of the most basic set up for each framework. In practice, both PCT and FEP diagrams can consist of complex hierarchical structures and additional variables conveying novel parameters that modify the capabilities of those models. Second, we avoid a full categorification of each framework. A full categorification would require working inside an established background category such as **Set**, **Vect**, or **Stoch**, whose objects and morphisms are already well-defined mathematical structures. While there is active research in this direction (see, e.g., St Clere Smithe's doctoral thesis "Mathematical Foundations for a Compositional

Account of the Bayesian Brain" [38]), engaging with those technical foundations is beyond the scope of this article.

Even still, applying the categorical framework at a high level allows us to reason systematically about information flow and structural correspondence between PCT and FEP frameworks. Therefore, the following diagrams are to true categorical diagrams as pseudocode is to code written in a particular programming language. As a result, many of the morphisms in the following diagrams do not typically have common names apart from their associated mathematical symbols; therefore, these morphisms have been given arbitrary names that at least roughly reflect the meaning of the associated operation(s).

First, we introduce the PCT (Figure 13) and FEP agent categories (Figure 14), and show a diagram of how they interact (Figure 15). Then we construct a functor from the FEP agent category to the PCT agent category (Figure 16), which can be shown to be composite of three subfunctors (Figures 17–19).

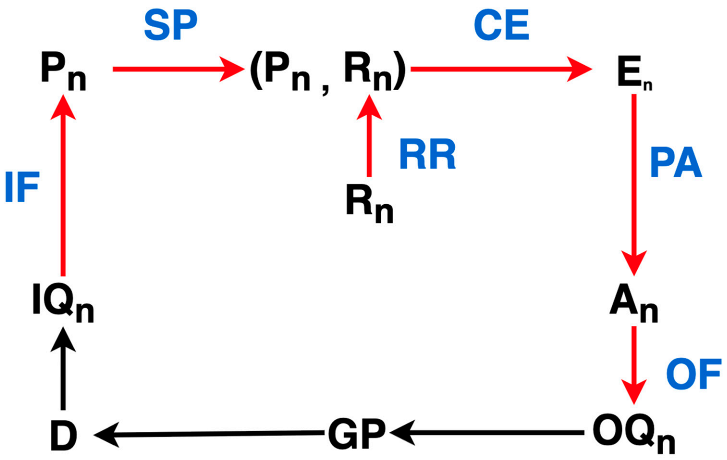

Figure 13. The PCT Agent Category. Morphisms that are internal to the category, that is, actually modeled explicitly by the PCT paradigm, are colored in red. Black morphisms indicate morphisms whose function may not be known and are located in the "environment category". The blue symbols are the abbreviated names of all morphisms internal to the category which are not necessarily named as such in the literature.

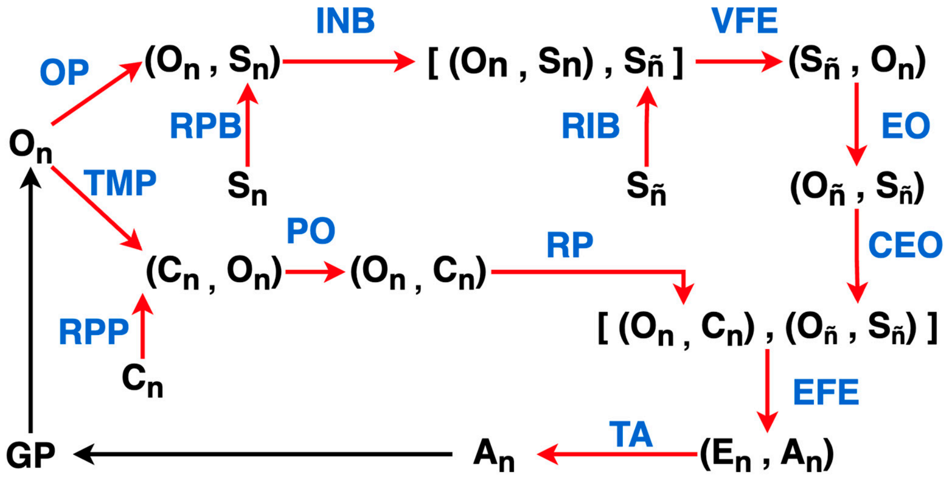

Figure 14. The FEP Agent Category. Morphisms that are internal to the category, that is, actually modeled explicitly by the FEP paradigm, are colored in red. Black morphisms indicate morphisms whose function may not be known and are located in the "environment category". The blue symbols are the abbreviated names of all morphisms internal to the category, which are not all necessarily named as such in the literature.

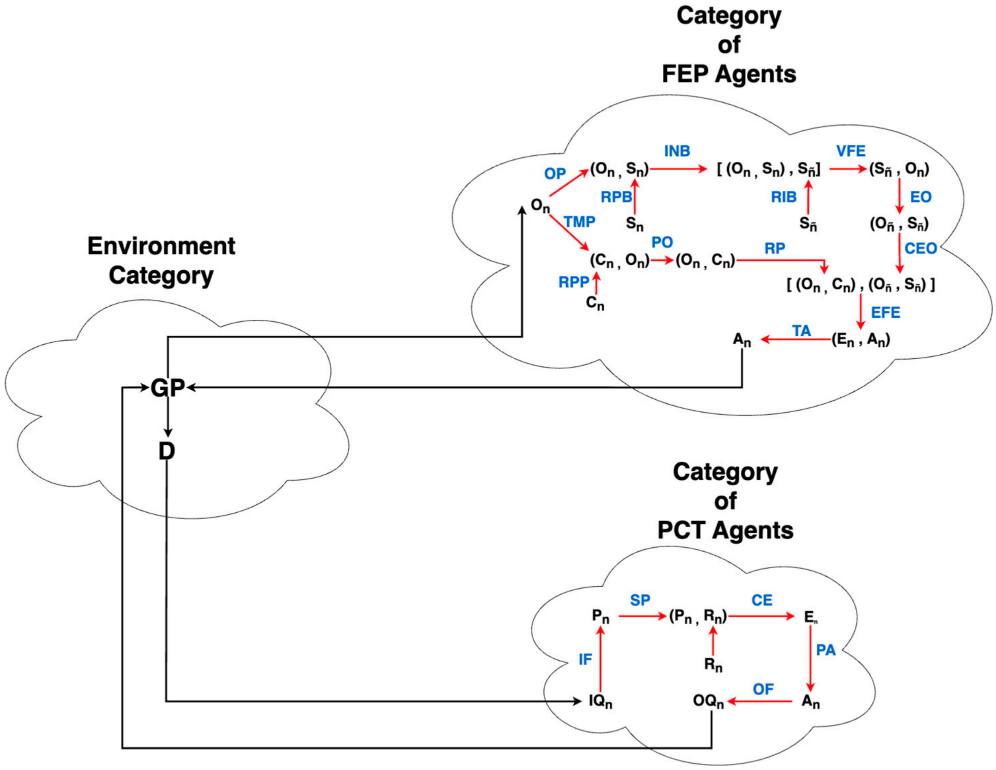

Figure 15. A diagram representing how one would couple interactions between PCT and FEP agents. The environment $GP_n \rightarrow D_n$ that affects the input quantity $D_n \rightarrow IQ_n$.

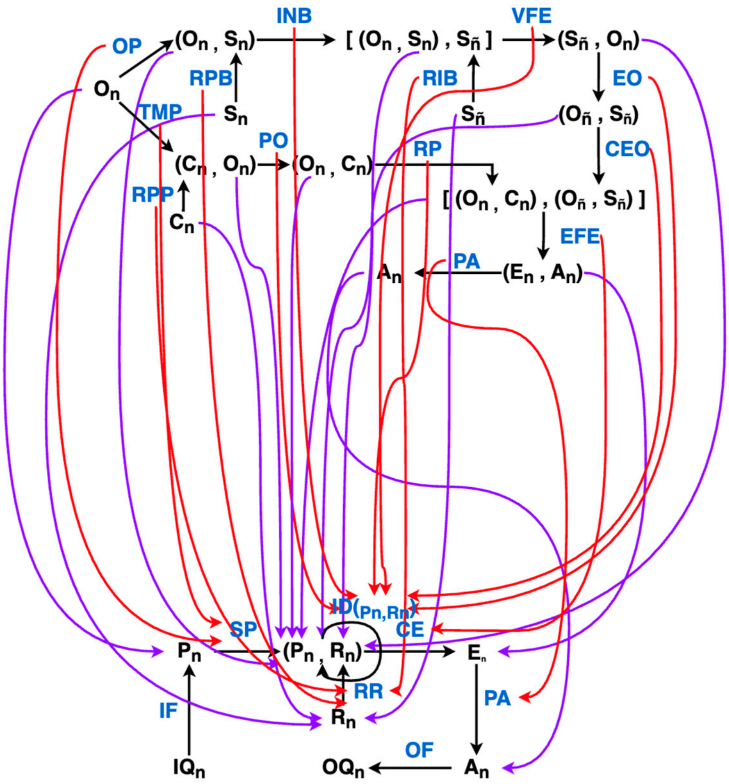

Figure 16. The $FEP\ Agent \rightarrow PCT\ Agent$ functor. Here, morphisms internal to categories are black. The functor consists of mappings from objects and morphisms in the FEP category to objects and morphisms in the PCT category. Environmental morphisms and their associated objects have been removed from the diagram for clarity.

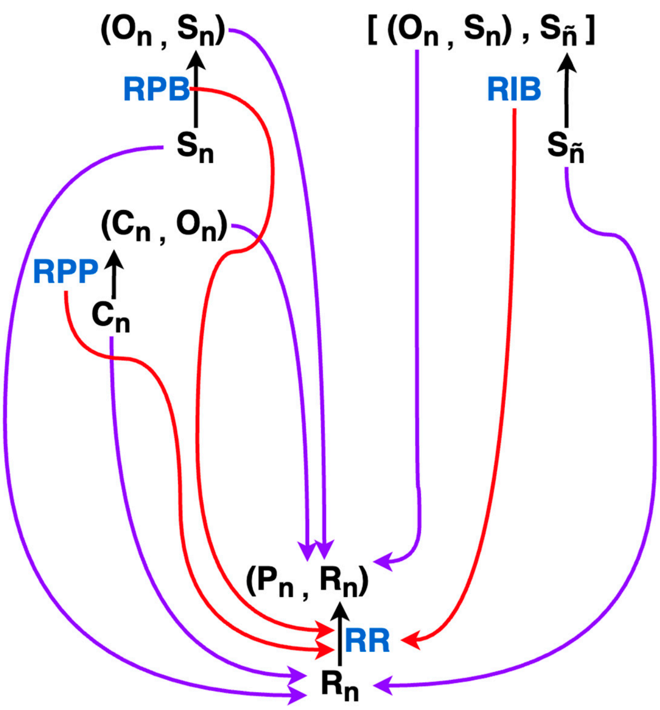

Figure 17. An illustration of the Internal Variable Subfunctor. Here, the source objects of each morphism in each subcategory correspond to the parameters of that agent category. In this case, the prior (S_n), approximate new prior distribution (S) and log prior preference (C_n) are FEP parameters, while PCT parameters consist only of the reference variable R_n. The target objects are outputs of the operations represented by their respective morphisms.

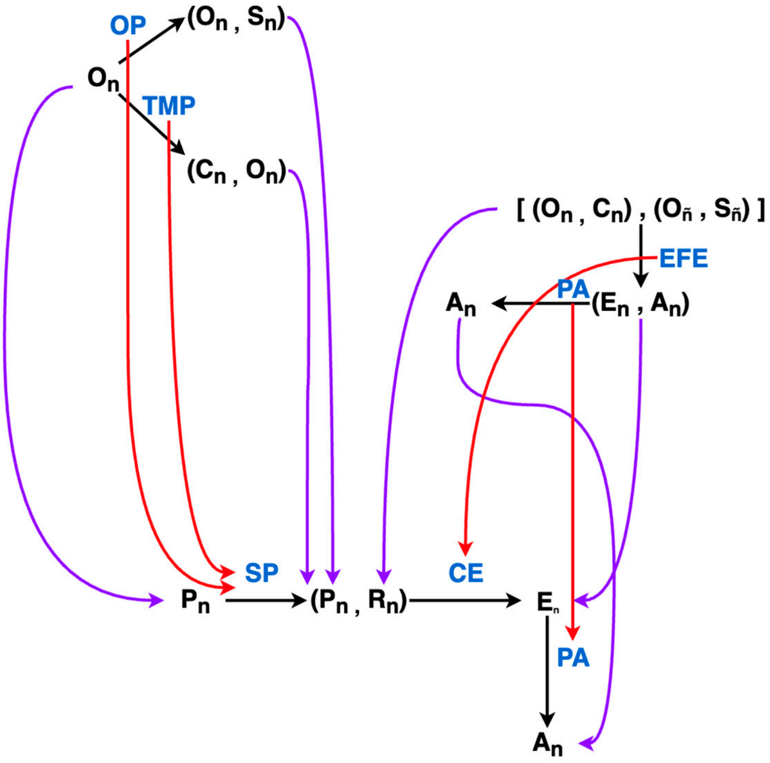

Figure 18. An illustration of the Perception Action Subfunctor. Here, each of the subcategories consists only of objects and morphisms corresponding to how perception is received and how action is performed.

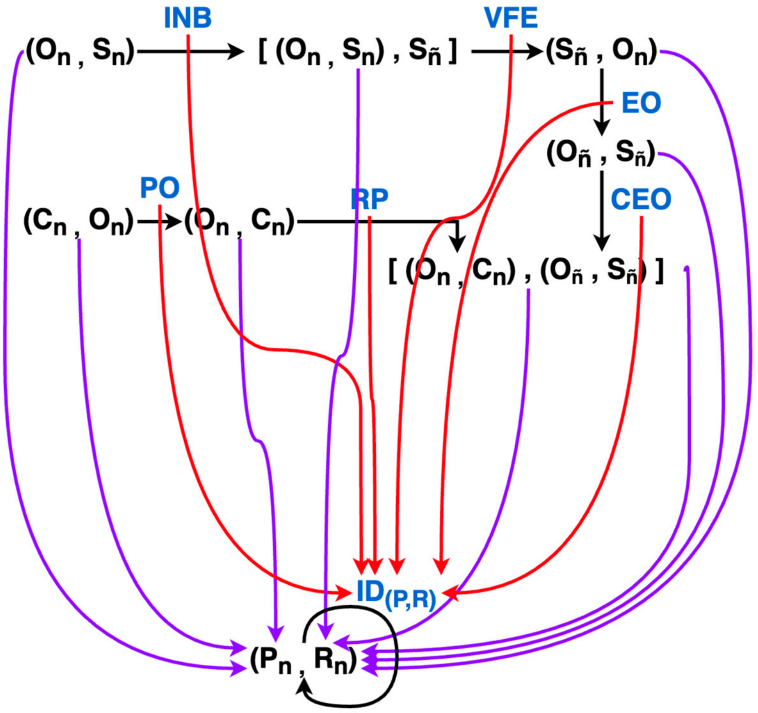

Figure 19. An illustration of the Uncertainty–Certainty Subfunctor. Here, each of the subcategories consists only of objects and morphisms concerned with how integrated perceptions are prepared such that they can drive action.

- Objects
    - ○ IQ_n—Input Quantity;
    - ○ P_n—Perceptual Signal;
    - ○ R_n—Reference Signal;
    - ○ E_n—Error Signal;
    - ○ A_n—Action;
    - ○ OQ_n—Output Quantity.
- Morphisms
    - ○ $IQ_n \rightarrow P_n$—Input Function (IF);
    - ○ $P_n \rightarrow (P_n, R_n)$—"Remember Perception" (RP);
    - ○ $R_n \rightarrow (P_n, R_n)$—"Remember Reference" (RR);
    - ○ $(P_n, R_n) \rightarrow E_n$—"Calculate Error" (CE);
    - ○ $E_n \rightarrow A_n$—Pick Action (PA);
    - ○ $A_n \rightarrow OQ_n$—Output Function (OF);
    - ○ $D_n \rightarrow IQ_n$—Disturb Input Quantity;
    - ○ $GP_n \rightarrow D_n$—Produce Disturbance.
- Objects
    - ○ O_n—Observation;
    - ○ S_n—Prior beliefs about hidden state GP_n;
    - ○ Cn—Log prior preference about hidden state GP_n;
    - ○ S—Approximate new prior S_n;
    - ○ O—Expected Observation;
    - ○ E_n—Error Signal;
    - ○ A_n—Action.

- Morphisms
  - ○ $O_n \ \rightarrow \ (O_n \ , \ S_n)$—"Observe Perception" (OP);
  - ○ $S_n \rightarrow \ (O_n \ , \ S_n)$—"Remember Prior Beliefs" (RPB);
  - ○ $S \ \rightarrow \ [(O_n \ , \ S_n) \ , \ S \ ]$—Remember Imagined Beliefs (RIB);
  - ○ $(O_n \ , \ S_n) \ \rightarrow \ [(O_n \ , \ S_n) \ , \ S \ ]$—"Imagine New Beliefs" (INB);
  - ○ $[(O_n \ , \ S_n) \ , \ S \ ] \rightarrow (S \ , \ O_n)$—Pick A New Prior Belief (Variational Free Energy —VFE);
  - ○ $(S \ , \ O_n) \ \rightarrow \ (O \ , \ S)-$ Produce Expected Observation (EO);
  - ○ $O_n \ \rightarrow \ (C_n \ , \ O_n)$—Trigger Memory of Preference (TMP);
  - ○ $C_n \ \rightarrow \ (C_n \ , \ O_n)$—Remember Prior Preferences (RPP);
  - ○ $(C_n \ , \ P_n) \ \rightarrow \ (O_n \ , \ C_n)$—Produce Preferred Observation (PPO);
  - ○ $(O \ , \ S \ ) \ \rightarrow (O_n \ , \ C_n) \ , \ (O \ , \ S)$—"Consider expected observation" (CEO);
  - ○ $(O_n \ , \ C_n) \ \rightarrow \ [(O_n \ , \ C_n) \ , \ (O \ , \ S)]$—"Remember preference" (RP);
  - ○ $[(O_n \ , \ C_n) \ , \ (O \ , \ S)] \ \rightarrow \ (E_n \ , \ A_n)$—Produce Error (Expected Free Energy—EFE);
  - ○ $(E_n \ , \ A_n) \ \rightarrow \ A_n$—Take Action (TA).

Since the morphisms that are part of the environment are typically not modeled uniquely by either framework, we exclude them for the construction of a functor from the FEP agent category to the PCT agent category. Here, internal morphisms are black, while functor arrows that map objects and morphisms are purple and red, respectively.

- Object Mappings
  - ○ $O_n \rightarrow P_n;$
  - ○ $[(O_n \ , \ C_n) \ , \ (O \ , \ S \ )] \rightarrow (P_n \ , \ R_n);$
  - ○ $[(O_n \ , \ S_n) \ , \ S \ ] \rightarrow (P_n \ , \ R_n);$
  - ○ $(S \ , \ O_n) \rightarrow (P_n \ , \ R_n);$
  - ○ $(C_n \ , \ O_n) \rightarrow (P_n \ , \ R_n);$
  - ○ $(O_n \ , \ C_n) \rightarrow (P_n \ X \ R_n);$
  - ○ $(O_n \ X \ S_n) \rightarrow (P_n \ X \ R_n);$
  - ○ $(O \ , \ S) \rightarrow (P_n \ , \ R_n);$
  - ○ $C_n \rightarrow R_n;$
  - ○ $S_n \rightarrow R_n;$
  - ○ $S \rightarrow R_n;$
  - ○ $A_n \rightarrow A_n;$
  - ○ $(E_n \ , \ A_n) \rightarrow E_n.$
- Morphism Mappings
  - ○ $CEO \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $EO \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $EO \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $PO \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $RP \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $VFE \rightarrow ID_{(P_n \ , \ R_n)};$
  - ○ $EFE \rightarrow CE;$
  - ○ $OP \rightarrow SP;$
  - ○ $PA \rightarrow PA;$
  - ○ $RIB \rightarrow RR;$
  - ○ $RPB \rightarrow RR;$
  - ○ $RPP \rightarrow RR;$
  - ○ $TMP \rightarrow SP.$

The $FEP\ Agent \rightarrow PCT\ Agent$ functor can be understood as a combination of three "subfunctors", which are structure-preserving decompositions of the functor such that

1. The source and target categories of each subfunctor are subcategories of the source and target categories in the original functor;
2. The disjoint union of the subfunctors is the functor itself—In other words, "gluing" all subfunctors together yields the original functor.

We have decomposed this functor into three subfunctors, each called the following:

- The Internal Variable Subfunctor;
- The Perception Action Subfunctor;
- The Uncertainty–Certainty Subfunctor.

3. Results

Our results include quantifications of model performance and interpretations of the categorical diagrams that were constructed in the previous section.

3.1. PCT vs. FEP Model Performance Comparison

To make a quantitative comparison between PCT and FEP models, we compared the behavior of each mathematical model against a random walker in terms of metrics like the average distance from the target and the amount of "food" collected. However, we emphasize that direct comparisons between these two models remain difficult for a number of reasons that will be addressed in the discussion section.

A parameter search was conducted to find both PCT and FEP models that could fulfill the task of reaching their target in under 5000 time steps. Parameter ranges and step sizes for the PCT and FEP models were selected such that the percentage of successful models was approximately equal with values of 96% and 96.1%, respectively, out of 2880 total model parameter variants for each. For PCT, the resultant parameter range and step sizes were as follows:

- Tumble Threshold:
  - Range: 50–100;
  - Step Size: 5.55.
- Scaling Constant:
  - Range: 2–40;
  - Step Size: 4.22.
- Reference Signals:
  - Range: 0.85–3;
  - Step Size: 0.07.

For the FEP model, the resultant parameter ranges and step sizes were as follows:

- Likelihood Sharpness:
  - Range: 0.8–1;
  - Step Size: 0.022.
- Nonlocality Biases:
  - Range: 0–0.72;
  - Step Size: 0.08.
- Preference Strengths:
  - Range: 0–10;
  - Step Size: 0.323.

Figures 20 and 21 plot the distribution of average distances from the target and the distribution of time spent collecting food over each of the aforementioned model's parameter spaces:

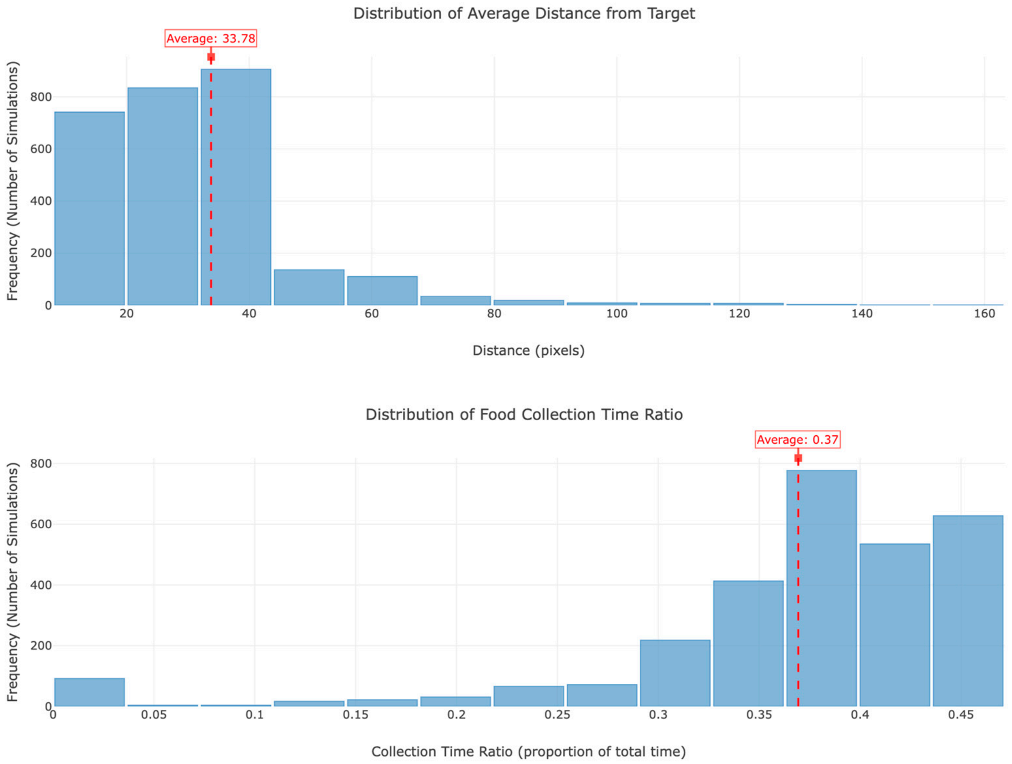

Figure 20. An illustration of the distribution of average distance from the center and the ratio of time spent collecting food for 2880 different PCT *E. coli* Model Parameters.

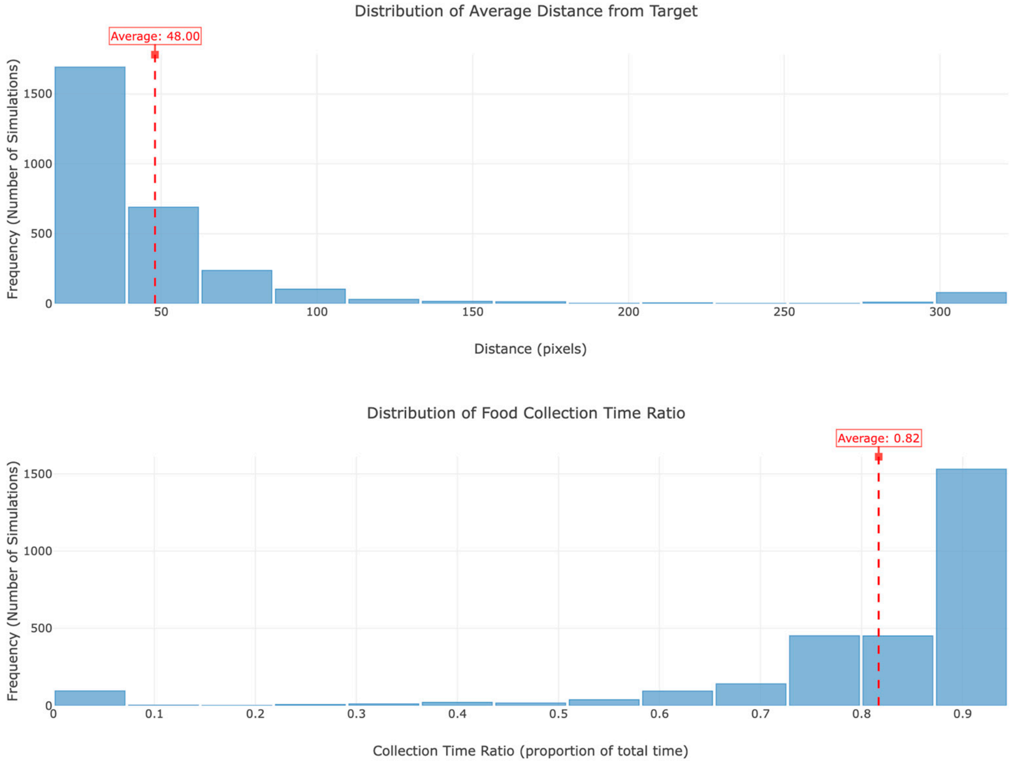

Figure 21. An illustration of the distribution of average distance from the center and ratio of time spent collecting food for 2880 different FEP *E. coli* Model Parameters.

3.2. PCT vs. FEP Mathematical Comparison

The PCT and FEP agent categories that were constructed illustrate the core way that each framework mathematically implements its modeling scheme. However, in practical cases, each framework entails additional sophistication in their diagrams when modeling living systems in more detail. For example, since PCT typically centers on dealing with equations describing the dynamics of the physical materials that embody the living system, additional sophistication comes from mapping the objects and morphisms in the PCT agent diagram to the category containing those equations. In other words, the PCT agent category of a particular agent is incomplete without an accompanying functor to the physical dynamics. For example, neurons are often modeled with a set of differential equations that describe how the voltage of a neuron changes in response to ionic and synaptic currents [39]. From the perspective of PCT, neurons can be considered as agents that control their membrane potential voltage. The reversal potentials driving each current act as the physical reference variables for the control system, while the magnitudes of the current represent the physical error signals whose physical dynamics push the voltage towards the neurons resting voltage potential. The following diagram illustrates such a functor from the “Category of Hodgkin–Huxley Neurons” to the PCT Agent Category (Figure 22):

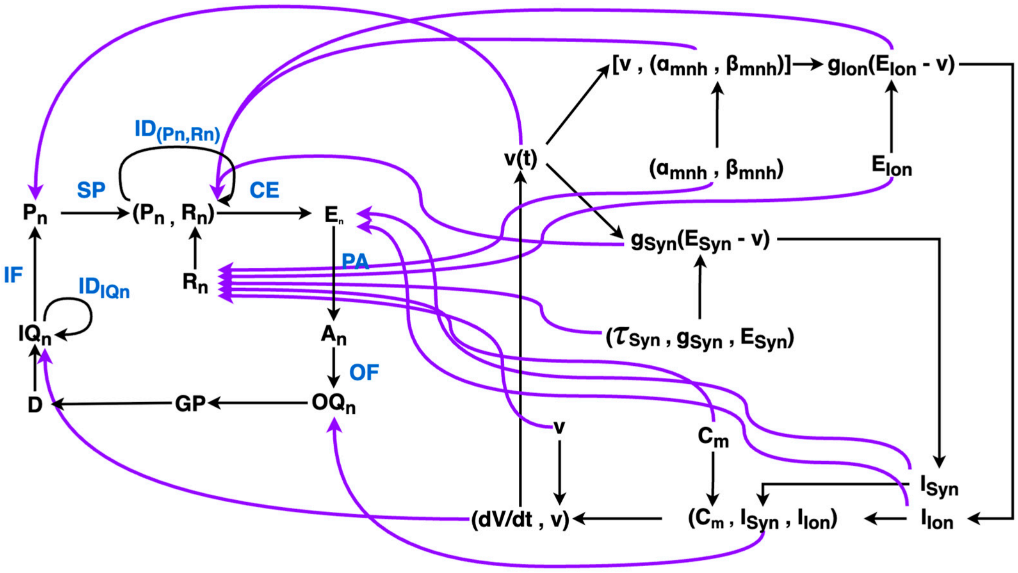

Figure 22. An illustration of a functor from a “Category of Hodgkin–Huxley Neurons” to the PCT Agent Category is illustrated in Figure N. For brevity, only object mappings for the functor are pictured. Many of these morphism mappings can be inferred given their source and target object mappings. When source and target objects are mapped to the same object in the target category, the morphism between them must be mapped to the identity of that target object. Note that $C_m\ X\ I_{Syn}\ X\ I_{Ion}$ is the total electrical state of the cell, which maps to a pair of the change in voltage and current voltage $\left(\frac{dV}{dt}, v\right)$ with the following morphism: $\left(C_m\ ,\ I_{Syn}\ ,\ I_{Ion}\right) \rightarrow \left(\frac{dV}{dt}, v\right)$. This is isomorphic with the composite morphism in the PCT category: $OQ_n \rightarrow GP \rightarrow IQ_n$. In other words, the total electrical state of the neuron is its output quantity, and the voltage of the cell is its input quantity.

This approach contrasts with the FEP, which deals with equations describing the mechanics of inference over generative models. Since FEP-based models are typically not concerned with the physical dynamics of living systems, additional sophistication comes from extending the FEP agent category by adding objects and morphisms to it. Therefore, implementations of the FEP agent category are more likely to be independent of the physical system that they model than are PCT agent implementations. Due to this, FEP agent categories are more easily comparable to each other than PCT agent categories.

FEP agent categories, like the example detailed here, are only subsets of more complicated categories. However, fully comparing PCT agent categories involves using functors to formalize not only how PCT agent categories map to physical equations, but how those functors may be related by natural transformations.

While the structure of the FEP agent category does not preclude a mapping to the physical equations governing a living system, the PCT agent category presents a simpler way to achieve this mapping. Nonetheless, it is likely that the FEP agent category may be more parsimonious when it comes to predicting the behavior of living systems when the physical equations describing that system are intractable or unknown. This is due to the enormous number of parameters that can be tweaked in a generative model. However, while generative models are not as entangled in their representations of the physical system they are modeling as machine learning techniques like artificial neural networks, it remains clear that the FEP agent category presents a lot of extra overhead for researchers hoping to create realist, rather than instrumentalist, models.

However, a potential synthesis emerges when considering living systems where the physical equations comprising a living system are thought to be known, but whose organization and/or specific parameter values are unknown. Rather than building generative models of agents and the beliefs they have, this article suggests explicitly building generative models that parameterize the uncertainty that researchers have about model architecture and unknown parameters. Diagrammatically, this would be illustrated as an FEP agent category where probability distributions represent beliefs about sets of functors from PCT agent categories to physical equation categories, given data. Indeed, in the literature, this is known as the “dual” formulation of the FEP, where the modeling focus is flipped; Instead of attributing beliefs to the system, we model our beliefs about the system’s measurable properties and their dynamics [40]. It is in this sense that PCT and the FEP meet at a deep level: both acknowledge that the system itself may be ‘information-free,’ but that modeling inevitably involves information and uncertainty.

4. Discussion

A comparison of PCT and the FEP has resulted in a recognition of a conceptually sound approach to synthesizing the two frameworks. Both frameworks contribute to formally handling different kinds of scientific knowledge, controlling information, and uncertainty about control information, respectively. Quantitatively comparing PCT and FEP models of bacterial chemotaxis reached an impasse because these models were both able to achieve the task of navigating to the target attractant with similar levels of success. The percentage of successful models for each type was highly dependent on the parameter space. Although the parameter space for FEP models was reduced, it proved difficult to justify that the parameter spaces selected for each model were truly comparable. We settled on choosing a parameter space where the percentage of models that succeeded on the task was about equal; however, we emphasize that even still, it is the parameter spaces chosen that account for the variation in measured distributions. Qualitatively, both models had a variety of trajectories taken from their starting points to the target attractant. However, the variety in these trajectories was also determined by the parameter space chosen. A possible approach to finding truly comparable parameter spaces may be to train an unsupervised machine learning model to categorize trajectories, so that patterns between parameter selection and model trajectory behavior can be identified. In this way, parameters that result in model trajectories that are in the same category can be directly compared.

Although there was no clear difference in performance between models, there remains a question as to the relative parsimony and redundancy of the models themselves. Classically, it is proposed that when comparing two models with equal evidence, the more parsimonious model is preferred [41]. In our simplified case, the PCT model clearly provides the more parsimonious option; however, further investigation across diverse and more complex scenarios is necessary to clarify when, where, and if ever, FEP models are more parsimonious than PCT models. In particular, our simulation omitted the dynamic and relatively unpredictable disturbances in chemical concentration over time and location that would occur within a real-life fluid substrate that is subject to the Brownian motion of particles, currents, and non-local disturbances from other agents. It is possible that introducing these disturbances, together with identifying parameterized PCT and FEP models in the same model trajectory category, would lead to stronger conclusions concerning when one of these models, and or frameworks, provides an advantage over the other in terms of performance and parsimony.

Furthermore, elaborations of both PCT and FEP models, including hierarchical structures and additional elements, may help shed light on how a concrete mathematical synthesis can be achieved. In both frameworks, hierarchies are not only practical modeling tools but deeply tied to how agency, goal-directedness, and self-regulation are conceptualized. For example, bacterial chemotaxis decides not just what to move towards, but what to move away from, namely toxins, predators, and other environmental repellents. To capture these behaviors, more detailed models of bacterial chemotaxis should be constructed. These models should specify cell properties such as membrane receptor distributions and intracellular mechanisms, in addition to environmental properties like fluid viscosity and dynamic chemical gradients. At the molecular level, computations are stochastic, and the signal-to-noise ratio of functions implemented by these molecules must be sufficient to reliably drive goal-oriented behaviors in the right contexts. Hierarchical PCT models incorporating these details or more biophysically inclined surrogate functions may shed light into how conflict among multiple sub-control systems are resolved mechanically; hierarchical FEP models may shed light into either (1) how the cellular dynamics might be efficiently simulated when cast as inference or (2) of how multiple sub-control systems can be organized such that they are able to meet a variety of goals in different contexts. The investigation of the structured mathematical relationships between these models may form the basis for future comparative work or theoretical unification.

This article argues that category theory offers the right level of abstraction for both rigorously formalizing such comparisons and for constructing a unified framework. In particular, functors from FEP agent categories to PCT agent categories provide a precise way of illustrating how generative models can be reinterpreted as expressing beliefs about the parameters and structure of physical control systems, rather than about the beliefs held by the agents themselves.

Importantly, this categorical approach allows us to formally distinguish between epistemic uncertainty about a system and ontological claims about what a system is. That distinction is often blurred in discussions of the FEP, where generative models are sometimes ambiguously attributed to organisms, experimenters, or the models themselves. By tracing functors between categories, we gain a way to clearly separate models of inference about systems from models of inference within systems. Furthermore, the development of natural transformations between functors representing different modeling perspectives may offer a formal mechanism to translate between generative models that are instrumentally useful and physical models that are mechanistically explanatory.

Ultimately, both PCT and the FEP share a commitment to explaining purposeful behavior, though they do so from different epistemological standpoints. By recasting both within the language of category theory, we not only clarify their similarities and differences but also reveal a space of possible hybrid models that integrate mechanistic structure with inferential flexibility. Such hybrid models may prove essential for advancing our scientific understanding of life, agency, and cognition.

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

Funding: This research received no external funding.

Data Availability Statement: The code used to produce many of the figures and results shared in this article can be found at this link: https://github.com/Tyfoods/PCT-vs-FEP (acessed on 24 October 2025).

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

Abbreviations

The following abbreviations are used in this manuscript:

PCT	Perceptual Control Theory
PBF	Principle Based Framework
VFE	Variational Free Energy
EFE	Expected Free Energy
FEP	Free Energy Principle
HH	Hodgkin Huxley

References

1. Powers, W.T. *Behavior, the Control of Perception*; Aldine: New York, NY, USA, 1981; ISBN 978-0-202-25113-4.
2. Parr, T.; Pezzulo, G.; Friston, K.J. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*; The MIT Press: Cambridge, MA, USA, 2022; ISBN 978-0-262-36997-8.
3. Marken, R.S. The Dancer and the Dance: Methods in the Study of Living Control Systems. *Psychol. Methods* **1997**, *2*, 436–446. [CrossRef]
4. Powers, W.T.; Clark, R.K.; McFarland, R.L. A General Feedback Theory of Human Behavior. Part I. *Percept. Mot. Ski.* **1960**, *11*, 71–88. [CrossRef]
5. Powers, W.T.; Clark, R.K.; McFarland, R.L. A General Feedback Theory of Human Behavior. Part II. *Percept. Mot. Ski.* **1960**, *11*, 309–323. [CrossRef]
6. Mansell, W.; Marken, R.S. The Origins and Future of Control Theory in Psychology. *Rev. Gen. Psychol.* **2015**, *19*, 425–430. [CrossRef]
7. Ashby, W.R. *Design for a Brain*; Wiley: New York, NY, USA, 1954.
8. Wiener, N. *Cybernetics or Control and Communication in the Animal and the Machine*; The M.I.T. Press: Cambridge, MA, USA, 1948.
9. Lubashevsky, I.; Lubashevskiy, V. Towards Naturalized Phenomenology: Dynamics of Space–Time Clouds and Power Law of Working Memory. *Cogn. Syst. Res.* **2025**, *92*, 101374. [CrossRef]
10. Arocha, J.F. Scientific Realism and the Issue of Variability in Behavior. *Theory Psychol.* **2021**, *31*, 375–398. [CrossRef]
11. Cochrane, T.; Nestor, M.J. Control at the Heart of Life: A Philosophical Review of Perceptual Control Theory. *Curr. Opin. Behav. Sci.* **2025**, *63*, 101526. [CrossRef]
12. Friston, K. The Free-Energy Principle: A Unified Brain Theory? *Nat. Rev. Neurosci.* **2010**, *11*, 127–138. [CrossRef]
13. Jaynes, E.T. *Probability Theory: The Logic of Science*; Cambridge University Press: Cambridge, UK; New York, NY, USA, 2003; ISBN 978-0-521-59271-0.
14. Friston, K.; Da Costa, L.; Sakthivadivel, D.A.R.; Heins, C.; Pavliotis, G.A.; Ramstead, M.; Parr, T. Path Integrals, Particular Kinds, and Strange Things. *Phys. Life Rev.* **2023**, *47*, 35–62. [CrossRef]
15. Ramstead, M.J.D.; Sakthivadivel, D.A.R.; Heins, C.; Koudahl, M.; Millidge, B.; Da Costa, L.; Klein, B.; Friston, K.J. On Bayesian Mechanics: A Physics of and by Beliefs. *Interface Focus* **2023**, *13*, 20220029. [CrossRef]
16. Andrews, M. The Math Is Not the Territory: Navigating the Free Energy Principle. *Biol. Philos.* **2021**, *36*, 30. [CrossRef]
17. van Es, T. Living Models or Life Modelled? On the Use of Models in the Free Energy Principle. *Adapt. Behav.* **2021**, *29*, 315–329. [CrossRef]
18. Bruineberg, J.; Dołęga, K.; Dewhurst, J.; Baltieri, M. The Emperor's New Markov Blankets. *Behav. Brain Sci.* **2022**, *45*, e183. [CrossRef] [PubMed]
19. Kennaway, R. Perceptual Control Theory and the Free Energy Principle: A Comparison. *Curr. Opin. Behav. Sci.* **2025**, *65*, 101575. [CrossRef]

20. Corning, P.A. Control Information Theory: The "missing Link" in the Science of Cybernetics. *Syst. Res. Behav. Sci.* **2007**, *24*, 297–311. [CrossRef]
21. Abel, D.L. Selection in Molecular Evolution. *Stud. Hist. Philos. Sci.* **2024**, *107*, 54–63. [CrossRef]
22. Wadhams, G.H.; Armitage, J.P. Making Sense of It All: Bacterial Chemotaxis. *Nat. Rev. Mol. Cell Biol.* **2004**, *5*, 1024–1037. [CrossRef]
23. Ridley, A.J.; Schwartz, M.A.; Burridge, K.; Firtel, R.A.; Ginsberg, M.H.; Borisy, G.; Parsons, J.T.; Horwitz, A.R. Cell Migration: Integrating Signals from Front to Back. *Science* **2003**, *302*, 1704–1709. [CrossRef]
24. Kurosaka, S.; Kashina, A. Cell Biology of Embryonic Migration. *Birth Defects Res. Part C Embryo Today Rev.* **2008**, *84*, 102–122. [CrossRef]
25. Wong, C.H.Y.; Heit, B.; Kubes, P. Molecular Regulators of Leucocyte Chemotaxis during Inflammation. *Cardiovasc. Res.* **2010**, *86*, 183–191. [CrossRef]
26. Eisenbach, M. Sperm Navigation in Humans: A Concerted Action of Multiple Means. *Commun. Biol.* **2025**, *8*, 923. [CrossRef]
27. Schneider, L.; Cammer, M.; Lehman, J.; Nielsen, S.K.; Guerra, C.F.; Veland, I.R.; Stock, C.; Hoffmann, E.K.; Yoder, B.K.; Schwab, A.; et al. Directional Cell Migration and Chemotaxis in Wound Healing Response to PDGF-AA Are Coordinated by the Primary Cilium in Fibroblasts. *Cell. Physiol. Biochem.* **2010**, *25*, 279–292. [CrossRef] [PubMed]
28. Michaelis, U.R. Mechanisms of Endothelial Cell Migration. *Cell. Mol. Life Sci.* **2014**, *71*, 4131–4148. [CrossRef] [PubMed]
29. Aberle, H. Axon Guidance and Collective Cell Migration by Substrate-Derived Attractants. *Front. Mol. Neurosci.* **2019**, *12*, 148. [CrossRef] [PubMed]
30. Drescher, K.; Dunkel, J.; Cisneros, L.H.; Ganguly, S.; Goldstein, R.E. Fluid Dynamics and Noise in Bacterial Cell–Cell and Cell–Surface Scattering. *Proc. Natl. Acad. Sci. USA* **2011**, *108*, 10940–10945. [CrossRef]
31. Marken, R.S.; Powers, W.T. Random-Walk Chemotaxis: Trial and Error as a Control Process. *Behav. Neurosci.* **1989**, *103*, 1348–1355. [CrossRef]
32. Smith, R.; Friston, K.J.; Whyte, C.J. A Step-by-Step Tutorial on Active Inference and Its Application to Empirical Data. *J. Math. Psychol.* **2022**, *107*, 102632. [CrossRef]
33. Heins, C.; Millidge, B.; Demekas, D.; Klein, B.; Friston, K.; Couzin, I.D.; Tschantz, A. Pymdp: A Python Library for Active Inference in Discrete State Spaces. *J. Open Source Softw.* **2022**, *7*, 4098. [CrossRef]
34. Shannon, C.E. A Mathematical Theory of Communication. *Bell Syst. Tech. J.* **1948**, *27*, 379–423. [CrossRef]
35. Spivak, D.I. *Category Theory for the Sciences*; The MIT Press: Cambridge, MA, USA, 2014; ISBN 978-0-262-02813-4.
36. Yanofsky, N.S. *Monoidal Category Theory: Unifying Concepts in Mathematics, Physics, and Computing*; MIT Press: Cambridge, MA, USA, 2024; pp. i–viii. ISBN 978-0-262-38078-2.
37. Cheng, E. *The Joy of Abstraction: An Exploration of Math, Category Theory, and Life*, 1st ed.; Cambridge University Press: Cambridge, UK; New York, NY, USA; Port Melbourne, Australia; New Delhi, India; Singapore, 2024; ISBN 978-1-108-70844-9.
38. Smithe, T.S.C. Mathematical Foundations for a Compositional Account of the Bayesian Brain. *arXiv* **2022**, arXiv:2212.12538. [CrossRef]
39. Gerstner, W.; Kistler, W.M.; Naud, R.; Paninski, L. *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*; Cambridge University Press: Cambridge, UK, 2014; ISBN 978-1-107-06083-8.
40. Sakthivadivel, D.A.R. A Worked Example of the Bayesian Mechanics of Classical Objects. In Proceedings of the Active Inference; Buckley, C.L., Cialfi, D., Lanillos, P., Ramstead, M., Sajid, N., Shimazaki, H., Verbelen, T., Eds.; Springer Nature: Cham, Switzerland, 2023; pp. 298–318.
41. Epstein, R. The Principle of Parsimony and Some Applications in Psychology. *J. Mind Behav.* **1984**, *5*, 119–130.