Anticipation and Structural Coupling: Two Sides of the Same Coin

Enrique Francisco Soto-Astorga[1] and Juan-Carlos Letelier[2],
[1] Universidad Nacional Autónoma de México [2] Universidad Austral de Chile
astorga@ciencias.unam.mx

Abstract

Two foundational frameworks in theoretical biology, namely Maturana's Structural Coupling and Rosen's Anticipatory Systems, developed in mutual isolation despite addressing the same phenomenon. We argue that Structural Coupling is precisely the historical process by which an organism becomes an anticipatory system: the process by which the regulatory subnetwork's dynamics come into conjugacy with relevant environmental dynamics. We show that the components of Rosen's anticipatory architecture map onto the autopoietic system in a way that satisfies all of Rosen's formal conditions. Furthermore, a conditional probability matrix describing the regulatory subnetwork is actually the observer's statistical redescription of a *bona fide* dynamical system. Because organisms co-constitute rather than merely predict their environments, Rosen's concordance condition must be reformulated as a fixed-point condition. This synthesis carries direct consequences for Artificial Life: current AI systems, including LLMs, are reactive in Rosen's precise sense and incapable of Umwelt construction; open-ended evolution requires true anticipatory agents. A criterion is provided.

Submission type: **Full Paper**

Introduction

Living systems, with their obvious autonomy, their capacity for self-fabrication, and their congruence with the environment containing them, are unique systems that modern science struggles to understand. Collectively, our efforts to explain living systems fall into two unequal traditions: Experimental and Theoretical Biology. Experimental Biology is the undeniable motor of modern biology, and through experiments and observational data, great strides have been made in the last two hundred years. Yet, true understanding of biological organization has not arisen so far from the astronomical amounts of data collected.

As organisms are systems of *organized* processes, one would expect that a *theory of organization* should be developed, or rather, that organization should be studied not as an *explanandum*, but rather as a an *explanans*: a *principle* for Biology. Theoretical Biology, the second tradition involved and the expected guiding light in this discussion, has thus far been a minor actor in the development of Biological thought, even the most recent literature begs for an adequate definition of *organization* for Biology (Mossio, 2024).

After World War II, many researchers were captivated by the new science of Cybernetics (which saw its *point de défaillance* around 1975), as well as by the digital computer metaphor to describe the metabolism and the nervous system. The implicit adoption of this metaphor by mainstream scientists has, in important respects, constrained Biological thinking by privileging input-output descriptions over organizational ones (Rosen, 1991). A sub-tradition then emerged, which disputed the Computer/Information metaphor for failing to reflect a basic condition of living systems: self-fabrication. Among those who advanced this interpretation (working largely in isolation) were Robert Rosen (1934–1998) with $(M, R)-$Systems, Humberto Maturana (1929–2021) & Francisco Varela (1946–2001) with Autopoietic Systems, and Tibor Ganti (1933–2009) with Chemoton Theory.

Maturana and Varela further interpreted living systems not as mere classifiers of reality but as *constructors of reality itself* (Maturana and Varela, 1980). This constructivist notion was termed *Structural Coupling* by Maturana, and *Enaction* by Varela, and it is the conceptual continuation on the notion of *Umwelt* by Jacob von-Uexküll (von Uexkull, 1909). It is important to note that Maturana believed his main contribution to Theoretical Biology was the notion of Structural Coupling and not autopoiesis per se (Maturana, 2012; Maturana et al., 2011). Indeed, his first work in theoretical biology was about how the nervous system *constructs* reality ("Biology of Cognition" Report #9, Biological Computing Laboratory, 1969). Autopoiesis was a translational use of the idea of Structural Coupling to metabolism. Thus, *self-fabrication* and *construction of reality* are interlocking ideas with a common origin.

Coming from the same sub-tradition on the study of self-fabrication, Robert Rosen had an amazing, independent intuition in the late 50s (Rosen, 1958a,b, 1959). His idea can be summarized as follows: the property of self-fabrication can be *logically deduced* from the connectivity of a network of processes that exhibit catalytic closure. He called this entity an $(M, R)-$System and devoted his life to analyzing it. In the initial phase of Rosen's writing about $(M, R)-$Systems (1958-1980), which he regarded as being the correct formal system encoding the causality of living systems, he uses a bare-bone model of metabolic closure that can be subsumed in a simple diagram where the "complete" metabolism is entailed/represented by a single symbol f Letelier et al. (2003, 2006); Soto-Andrade et al. (2011). However, Rosen's discovery has faced severe obstacles in its translation to concrete biological examples(Letelier et al., 2006; Cárdenas et al., 2010; Gatherer and Galpin, 2013; Hofmeyr, 2021).

Rosen also touched upon cognitive systems in his book "*Anticipatory Systems*" (second edition; Rosen (2012)). There, he explored potential connections between self-fabrication and anticipatory behavior. Without actually using the term "cognition", his description of anticipation is necessarily cognitive–the living system builds a model of the relation between it and its environment–as we will explain below.

As the field of Artificial Life faces a turning point in which LLMs and potential open-ended evolution of artificial cognitive systems appear on the horizon as co-potentiators (Karelin et al., 2025), and as interest in Rosen's Anticipatory Systems Theory resurges (Astorga and Rosen, 2025) the moment is right to investigate the parallelism between Structural Coupling (à la Maturana) and Anticipation (à la Rosen). By showing the interconnected contributions of Umwelt, Structural Coupling, and Anticipation, we advance a path forward amalgamating these models as the basis for a non-mechanistic cognitive science in coherence with the Biology of Cognition.

Reactive Systems

As Rosen explored in diverse literature (Rosen, 1988, 1991, 2000), Western science made great progress on the shoulders of both the Cartesian assumption of fractionability and the Newtonian focus on recursion, which allows systems to be studied in terms of their states and parts. System states $E(t)$ evolve according to

$$E(t + k) = H(E(t), p_1, ..., p_n) \tag{1}$$

or in differential form:

$$dE(t) = F(E(t), t, p_1, \dots, p_n). \tag{2}$$

The future state $E(t + k)$ is thus computed from the present state $E(t)$ and some fixed parameters. This is the world of feedback and reactive systems told from the perspective of *recursive chronicles*[1]: the future of a system is merely some form of computation based on the system's past. This is an assumption in contemporary science[2].

A classical example is the cybernetic loop, where a system's action on its environment is sensed as a feedback stimulus to which the system reacts.

Anticipatory Systems

According to Rosen (first edition (Rosen, 1985), second edition (2012)), an anticipatory system A is one which possesses a model[3] M about itself and/or its environment which is used to proactively determine future states of A. In this definition, which *prima facie* does not appear radically different from the reactive paradigm, systems which are endowed with anticipation are assumed to be *feedforward* and *proactive* systems.

How can general natural systems achieve anticipation? Models in anticipatory systems are not passive and fixed representations of the world, nor are they perfect replicas. They are bounded and actively built through interaction and trial-and-error, and have a purpose. In Rosen's terms (2012):

> As we recall, the decisive feature of an anticipatory system is that, in effect, it employs its model to make predictions about [a given natural system], and more importantly, change of state occurs in [the anticipatory system] at an instant as a function of its predictions. (p. 315)

The model captures only specific, relevant regularities which are important for predictive purposes. Rosen assumes that the model is itself a *formal* counterpart; compact enough to not require the nominal time of the modeled system in order to "*simulate*" its behavior. Thus, the model is updated faster than external time: an essential property that clearly indicates the necessity to construct a theory of *Biological Time*. By "running" its internal model forward in time (a surrogate time) faster than the real world unfolds (nominal time), the system places itself in potential future states and thus determines which action must be taken, **without inducing retrocausality**.

We now establish the conditions for anticipation. Let A and B be systems, then A will be an anticipatory system if

1. A contains a model M of another system B which does not exhaust B: A *interiorizes* **some** relevant meanings of B. In A there will exist a subsystem X containing the aspects of A which mediate the interaction with B and which represent the behaviors of A *if no model* M *existed*. M will not be influenced by X, but X will be influenced

---

[1] A recursive chronicle is a time-indexed sequence of values *f(n)* generated by iterating a mapping *T*, such that each value entails the next: *f(n+1) = T(f(n))* (Rosen, 1991).

[2] As Rodríguez-Vargara & Husbands indicate, we are usually educated to "formalize and categorize" to prove understanding, "which almost inevitably leads to algorithmic theories" (Rodriguez-Vergara and Husbands, 2026).

[3] In the Rosennean sense as explained thoroughly in (Rosen, 2012) via his development of the *modeling relation*.

by M. X will change (pre-adapt) as M anticipates. Ergo, $A : M \times X$.

2. The present state of A at time t must encode in M a future state at time t' of B: the system is not merely reacting to actual states.
3. A must interact with B, using M's predictions of B's behavior, through an effector system E.
4. M preserves some relevant property of A, e. g.: preventing organizational dissipation.

Rosen's original architecture places M, E, and B within the anticipatory A. Input comes from the environment, and output into the same environment is produced by B. In reality, we must assume that X mediates between B and $M \times X$. For Rosen, a true anticipatory system is one for which the model and its predictions are functionally integrated into its own self-preserving organization. The "internal" model M is not really internal in the representational sense, but is embodied in the organization and dynamics of the system. The model may be totally distributed or localized, but the system's constitution is fully involved in its construction. This is an early hint at the potential confluence between the ideas of autopoiesis and Rosennean anticipation.

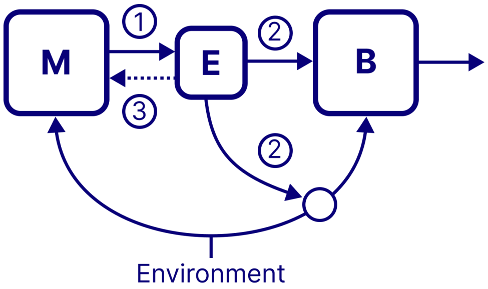

Figure 1: A classical Rosennean Anticipatory System where (3) allows the effector E to *reset* the model M and (2) allows it to change the dynamics of B by acting directly on it or on its environmental inputs based on M's predictions passed on to E by (1).

For M to be predictive (condition 2), it must provide A with information regarding possible future states of B. If B evolves via a transformation T_t, then M must contain the corresponding transformation $\tilde{T}_X$ so that $\varepsilon_{B\to A}(T_t(B)) = \tilde{T}_X(\varepsilon_{B\to A}(B))$, which means the present state $s_t(M)$ encodes (ε) a future state $s_{t+i}(B)$. Hence, the evolution of $\tilde{T}_X$ must run faster than that of T_t. (Rosen, 2012). This is the surrogate time referred to above, and as expressed diagrammatically:

$$\begin{tikzcd}
B \arrow[r, "T_t"] \arrow[d, "\varepsilon_{B\to A}"'] & B \arrow[d, "\varepsilon_{B\to A}"] \\
M \arrow[r, "\tilde{T}_X"'] & M
\end{tikzcd}$$

Rosen's ideas on Anticipation, even though developed by some authors (Dubois (2003), Poli (2017), Poli (2019), Poli and Valerio (2019)), have mostly languished in obscurity, largely because of the dense mathematics and new intuitions demanded by Rosen's science[4].

Structural Coupling

The history of Structural Coupling differs sharply from that of Rosen's Anticipation. Maturana's framework began with a simple internal report at Heinz von Foerster's Biological Computing Laboratory (BCL) in 1969. In "*Report 9: Biology of Cognition*," Maturana proposed the following: (Letelier, 2022)

1. The nervous system does not function like a digital computer.
2. The stimulus–response metaphor is fundamentally flawed.
3. An organism's behavior is central to constructing its ecological niche.
4. "Objects" are not independent of the observer but are actively built as sensory–motor correlations.

Maturana and his student Francisco Varela applied idea (4) to cellular metabolism, developing the widely studied theory of Autopoietic Systems (Maturana, 1975). Varela's related formulation, known as Enaction, was introduced in *The Embodied Mind* (Varela et al., 1991), which paved the way for non-mechanistic accounts of cognition.

To present Structural Coupling let us recover Maturana's and Varela's definition of autopoietic machines (Maturana & Varela, 1972):

> An autopoietic machine is a machine organized as a system of processes of production of components concatenated in such a way that they produce components that i) generate the processes (relations) of production that produce them through their continuous interactions and transformations and ii) constitute the machine as a unity in physical space. Consequently, an autopoietic machine continuously specifies and produces its own organization through the production of its own components, under conditions of continuous perturbation

---

[4] Crucially, Rosen never truly bridged his ideas on anticipation and organization represented by $(M, R)-$Systems. Rosen's *opus magnum*, "*Life Itself*" (Rosen, 1991), made almost no mention of Anticipation.

> and compensation of those perturbations (production of components). We can say then that an autopoietic machine is a homeostatic system that has its own organization as the variable that it maintains constant (p. 68, our translation).

Living Systems are thus *devices* that **act**, living in a loop (or spiral) connecting *Perception* to *Action*.

An Autopoietic network's structure changes constantly while its organization remains invariant. We distinguish two key internal subsets:

1. **Organization $\hat{O}$:** The subset of internal processes that define *class identity*, i.e., those that remain mostly invariant under diverse conditions.
2. **Structure $\hat{S}$:** Processes and components that may change without affecting class identity.

Within both sets, we can identify transversal subsets central to the theory:

1. **Sensorium σ:** subsets of Organization and Structure modulated by the biological function of *Perception*.
2. **Effectorium ϵ:** subsets of Organization and Structure modulated by the biological function of *Action*.

Many living systems have specialized sensoria/effectoria: birds can detect magnetic fields; bacteria act primarily by excreting chemicals. Thus, sensoria and effectoria must be understood as complex sensory-motor *geometries* where processes that connect the organism with its medium—its Umwelt—take place.

Internal Signals in Living Systems

When a circumstance triggers the sensorium, internal signals (molecular concentration vectors s_i) propagate through the system, interacting with components of the $\hat{O}$ and/or $\hat{S}$ to produce a contextually coherent action. Critically, there is no one-to-one mapping between external circumstances and the signals they evoke. Although living systems are historical, memory-bearing entities, they are not time-invariant: the same perturbations need not produce the same internal signals. Crucially, this historicity, as we will discuss later, is central to the Biology of Cognition.

Structural Coupling

Every living system (from *Yersinia pestis* to *Ornithorhynchus anatinus* to every member of genus *Homo*) is continuously immersed in a stream of perturbations that drive perception and action. The organism-environment relationship evolves through each successive $(perception - action)$ loop, which also constantly changes $\hat{S}$ (see (Letelier et al., 2002), Fig. 2). The loop works as follows:

1. (t_0): A living system begins recurrent interactions with its environment. Structure and medium follow uncoupled, apparently random dynamics (grey blob in t_0 in figure 2).
2. (t_1): Sustained interaction drives coherent, parallel structural changes in both the living system and its medium (shrinking blob; visible structural change in t_1 in figure 2).
3. (t_2): Structural changes converge further, leaving the living system's structure complementary to its environment (form depicted in t_2 in figure 2). An external observer encountering only this final state might mistakenly interpret it as adaptation to a pre-existing environmental object.

An example of Structural Coupling

Structural Coupling can seem mysterious (as if organisms construct their own environmental objects). A clarifying example is the bacterial production of atmospheric oxygen around 3,200 Mya.

When the first bacteria appeared ($\tilde{3}$,850 Mya), the atmosphere contained no free oxygen: only water, nitrogen, methane, and carbon dioxide. For roughly 1,100 My, all life was anaerobic. Then $\tilde{3}$,200 Mya, one bacterial lineage evolved oxygenic photosynthesis, splitting water molecules to produce oxygen: initially a toxin lethal to competitors that had colonized every available anaerobic niche. Yet this same "poison" eventually drove the evolution of aerobic lineages that came to depend on it. Oxygen thus became a permanent feature of the world from $\tilde{2}$,480 Mya, at the cost of triggering the Huronian glaciation (2,400–2,100 Mya) and nearly extinguishing all life on Earth.

The story illustrates the core thesis: organism-environment interactions generate entirely new environmental features, including the selective pressures that will shape subsequent evolution.

Having presented both frameworks independently, we now turn to their integration.

Anticipation and Structural Coupling

We contend that Structural Coupling is the *historical process* by which an organism becomes an anticipatory system. According to Maturana (and von Uexküll before him), every dimension of our sensoriality derives from Structural Coupling (Maturana, 2002). Time is not a simple substance that flows homogeneously and regularly for all observers; it is itself a relational construct, *enacted* through the organism's history of interactions. This makes anticipation—a concept that demands time for its very formulation—an intrinsically complex subject. Rosennean Anticipation, moreover, requires not one but two flows of time: the nominal time of the environment and the surrogate time of the predictive model. As we shall argue, Structural Coupling provides the framework within which both temporalities can be grounded *in the concrete dynamics* of the organism.

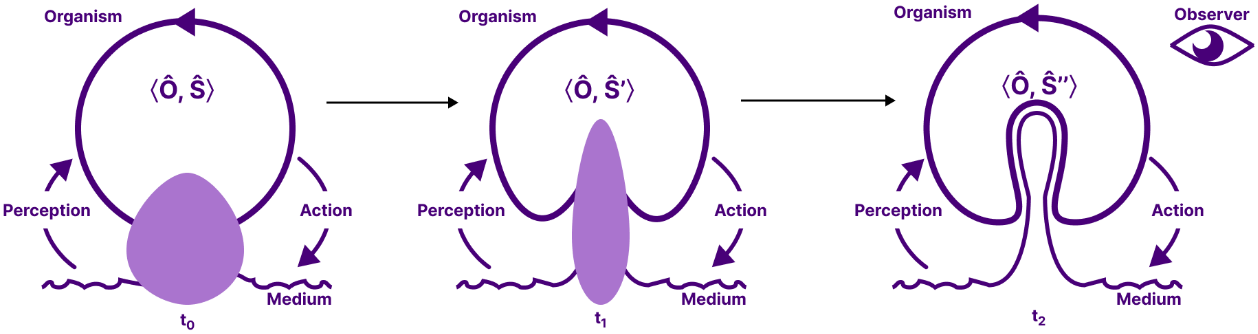

Figure 2: Mechanism of Structural Coupling: At t_0 the autopoietic system, with a given organization $\hat{O}$ and structure $\hat{S}$, confronts a new medium: relevant senso-motor correlations are not established, yet create a diffuse relationship (purple blob). By t_1, new senso-motor correlations are created and the relationship is less diffuse (smaller purple blob). $\hat{O}$ remains invariant, but the $\hat{S}$ has changed to $\hat{S}'$. At t_2, the relationship is clear: no blob, $\hat{O}$ remains invariant but $\hat{S}'$ has changed further to $\hat{S}''$, and the medium presents complementary changes. Naïve external observers erroneously describe the overall process as an *adaptation*.

The *ballistic* character of action

Every organismic action, however simple, has a temporal microstructure. When an action a is initiated, what it entails is a particular sequence of micro-events ($a = b_1 \to b_2 \to \cdots \to b_n$). This sequence, which can be understood as a recursive chronicle, is essentially ballistic: once b_1 is initiated, the subsequent steps must unfold to completion. Consequently, the execution of b_j already *causally entails* both (I) and (II): (I) the future execution of b_{j+1} and (II) the sensory state concomitant with b_{j+1}.

The internal processes that sustain this ballistic trajectory (i.e., signaling cascades, allosteric modulations, regulatory dynamics) constitute a system whose dynamics unfold faster than the corresponding behavioral changes in the medium. Intracellular signaling operates in milliseconds to seconds and gene regulatory responses in minutes; whereas the environmental regularities to which the organism is coupled change in timescales of minutes to hours or longer. This timescale separation is a physical fact. The regulatory subnetwork that coordinates the unfolding of b_1 through b_n is already, in a precise sense, *running ahead* of the medium.

Three levels of internal signaling

We can now develop the concept of *internal signals* introduced above through three levels of description. These levels characterize the three stages in the refinement of the observer's account. The organism always operates as described at Level 2; it is the observer who, in attempting to understand the organism, may begin at Level 0 and progressively refine their description:

Level 0. At the simplest level, we describe the organism as possessing its invariance network $\hat{O}$ and a sensorium that receives signals $s_\alpha = \{a, c, k, m\}$ from the environment. The organism maintains associations between signals and actions: given signal s_α, some action X,Y, or Z is performed with certain regularity. The milieu presents distinct patterns $[X], [Y], [Z]$. But note: in this description, the objects $[X], [Y], [Z]$ *preexist*. They are assumed to be features of a given environment that the organism merely learns to classify. This is the standard cybernetic picture.

Level 1. At a more refined level, the associations between signals and actions can be described as a **Conditional Probability Matrix** (CPM): a table with signals as columns and actions as rows, whose entries $P(Y, s_\alpha)$ describe the regularities the observer sees in the organism's behavior. This CPM admits an **updating mechanism**: when s_α is produced and action Y is performed and organization is conserved, the observed association is strengthened; i. e. a Hebbian-like description:

$$P(Y, s_\alpha)_{t+1} = P(Y, s_\alpha)_t + \Delta_t \qquad (3)$$

Two things must be emphasized: First, the CPM is the *observer's* statistical description of the effective behavior of the organism's regulatory dynamics. This is consistent with Maturana's insistence that notions of coding, information, and regulation belong to the observer's domain, not to the system's (Rosen had the same intuitions, which he expressed succintly in Rosen (2012)). Second, whatever the CPM describes is realized by concrete metabolic and signaling processes within $\hat{S}(t)$ (for a detailed description see Valdés-Zorrilla (2024)).

Yet, level 1 still assumes pre-existing objects $[X], [Y], [Z]$ as sources of distinct signals. The CPM account maps signals to actions, but the entities that produce the signals are taken for granted.

Level 2. In this level, objects $[X], [Y], [Z]$ do not pre-exist. They are built and continuously maintained by calibrating internal parameters through recurrent interaction. Each "object" is the limit of a trajectory in the $\langle perception, action\rangle$ phase space: a recurrent pattern of sensorimotor interaction that stabilizes through the organism's history of coupling. Different trajectories converge to different invariants; these invariants *are* what the organism treats as objects. At this level, the observer's description finally matches what the organism has been doing all along: co-constituting the regularities it will encounter. This is precisely what an anticipatory system does.

The identification

In Rosen's original architecture, the adaptive anticipatory system can be identified with the organism. This architecture can be redrawn as a $(perception - action)$, allowing for the overlaying of the elements of an autopoietic system on the Rosennean diagram of anticipation. To this end, we produce the following identification which is shown diagrammatically in figure 3:

1. **Model** M is realized by the regulatory and signaling subnetwork of the autopoietic system. This is the dynamical system whose faster-than-medium ballistic timescale makes anticipation possible[5]. At Level 1, the external observer describes its effective behavior as the CPM; but M itself is not a probability table, it is a *bona fide* dynamical system with its own internal entailment $k(M)$, as Louie (2012) requires.
2. **Effectors E** correspond to the effectorium ϵ.
3. X corresponds to the current structural state $\hat{S}(t)$.
4. B is the medium (or something in the medium, like another system or the system modeling itself[6]).
5. **Encoding** $\varepsilon_{B \to A}$ is realized by the sensorium σ.
6. The **property P** that the anticipatory system maintains is the conservation of $\hat{O}$.
7. The **updating of** M **by** E is realized by the structural modifications that the effectorium's activity induces in the regulatory subnetwork (changes in enzyme concentrations, receptor densities, allosteric states, and gene expression profiles).

Conditions previously stated for anticipation are also satisfied:

- Condition 1 (A contains a model M of B) is satisfied by the regulatory subnetwork's conjugacy with the medium dynamics via the sensorium encoding. The same regulatory subnetwork characterizes only the $(perception - action)$ mapping, which means that A does not exhaust B. The autopoietic system possesses an entire metabolic network of processes that are not directly involved in the sensorium-effectorium coupling, which constitute X. The action selected by the regulatory subnetwork modulates the effectorium, which is embedded in the autopoietic network and thereby alters structural degrees of freedom beyond the model itself, which means that X is changed by M.
- Condition 2 (M is predictive) is satisfied by the physical timescale separation between intracellular signaling and environmental change, as argued above.
- Condition 3 (A must interact with B using M's predictions) is satisfied because the effectorium's output is selected by M's faster-timescale, not by the current sensory state. The interaction with B is therefore mediated by M's predictions.
- Condition 4 (M preserves some property) is satisfied because M selects the action before some perturbation propagates to the organizational core, and $\hat{O}$ is conserved proactively rather than by reactive repair. The same timescale separation that grounds Condition 2 also grounds Condition 4.

Our contention is that this identification preserves the Rosennean structure: As Louie (2012) emphasizes, the predictive model in an anticipatory system must not be associated with any kind of probabilistic certainty about the future. Rosen's theory is a general qualitative theory about entailment structures. The regulatory subnetwork satisfies this: its dynamics constitute a genuine inferential entailment i within the organism, conjugate (via the sensorium encoding) to the causal entailment c in the milieu.

There is, however, a critical point of divergence. In Rosen's anticipation, the "environment" is not "*Umwelted*": it exists by itself, with its own dynamics independent of the anticipatory system. The concordance condition (Eq. 6.21 in (Rosen, 2012)) compares the model's prediction against the *actual* future state of an independently evolving system (p. 328):

$$\vec{u}_p(t+h) = \vec{u}(t+h) \tag{4}$$

In Structural Coupling, the organism's action at time t partially determines the medium's configuration at $t+h$. The organism does not predict an independently evolving world; *it co-constitutes the regularity it anticipates.*

---

[5] The ballistic character of the sequence alone does not constitute anticipation since any deterministic causal chain, including a falling rock, entails its own future states. What makes the organism's action anticipatory is the fact that which sequence is initiated depends on the regulatory subnetwork, whose dynamics have already run ahead of the medium.

[6] Rosen permits M to model the system itself, its environment, or both. In the autopoietic case this disjunction is not resolved but dissolved: the medium that M models is partly constituted by the actions M selects, so self-modeling and world-modeling are not separable.

This requires a reformulation of the concordance condition. Let μ denote the function describing how the organism's action transforms the medium, σ the sensorium mapping, and π the action selected by the regulatory subnetwork. The organism-medium loop is defined by

$$s(t+h) = \sigma(\mu(\pi(s(t)), m(t))). \tag{5}$$

The medium that M models is partly constituted by the actions M selects, so forcing a choice between $B = \hat{S}(t)$ and $B = medium$ would artificially impose a separation that structural coupling dissolves. The autopoietic case is thus not a special instance of Rosen's architecture but a generalization of it. Structural coupling is achieved when this loop possesses stable recurrent patterns: when M selects actions that generate medium configurations consistent with M's own entailments. Rather than a matching condition between prediction and an independently given reality, we have a *fixed-point condition*: the model must select actions that produce a world compatible with its own dynamics, subject to the constraint that $\hat{O}$ is conserved. The existence, uniqueness, and stability of such fixed points are open mathematical questions; the sheer fact that organisms achieve and maintain structural coupling serves as an existence proof. This fixed-point condition constitutes the core of the formal research program implied by the synthesis we have proposed.

This fixed-point structure captures what is absent from Rosen's framework: the organism is not merely a predictor but a co-author of the regularity it predicts. It also captures what is absent from Maturana's framework: the co-authorship operates through a specific mechanism (the faster-than-nominal-time dynamics of the regulatory subnetwork) that gives the organism's structural trajectory its anticipatory character.

Still, the medium is open: it is subject to perturbations not encoded in the regulatory subnetwork. As the medium changes beyond the scope of the model, adequacy of M progressively fails. This is Rosen's temporal spanning: the interval during which property P is maintained is finite because the model is necessarily an incomplete closed system encoding an open system (Rosen, 2012).

In autopoiesis, the equivalent phenomenon is the loss of structural coupling: the organism's structure no longer admits the perturbations presented by the medium, not because any component has broken but because the organism's structural trajectory has diverged from the medium's. Rosen emphasizes that this constitutes a *global failure* (no individual component malfunctions; what fails is the *relational conjugacy* between model and environment). In autopoiesis, the loss of adaptation is not a local event but a breakdown of congruence between the organism's changing structure and the medium's changing configuration.

Unlike Rosen's treatment, where temporal spanning leads inexorably to system failure, the Structural Coupling framework admits a dynamic of *re-coupling*. If $\hat{O}$ is not lost, the regulatory subnetwork can undergo further structural modifications that restore commutativity. This is what an external observer would describe as *learning*. Learning is re-coupling as re-anticipation, a resource that Rosen's framework, focused on the degradation of models, does not naturally provide.

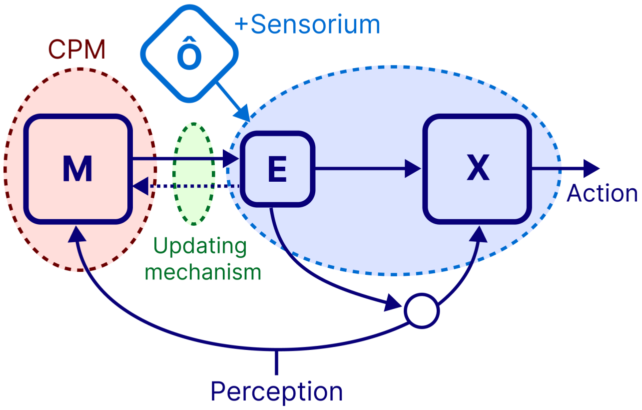

Figure 3: The diagram of an anticipatory system with autopoietic aspects superimposed. The corresponding object-system B is replaced by internal X.

Strong & Weak Anticipation

Structural Coupling also bridges the gap between weak and strong modes of anticipation as distinguished by Dubois (2003). Weak anticipation employs a decoupled internal model, whereas strong anticipation demands a dynamical structure within the metabolism itself, isomorphic to the organism-medium interaction (in a mechanistic sense: a real analogue computer).

The mechanisms of Structural Coupling correspond to strong anticipation precisely because the regulatory subnetwork functioning as M is not a detachable computational component but an integral part of $\hat{S}(t)$ whose dynamics are metabolic processes with a faster timescale arising from the physical properties of $\hat{S}(t)$'s molecular constituents. The conjugacy between model and environment is achieved not by engineering a separate representation but by the history of structural coupling itself reshaping the regulatory network until its internal entailment tracks the relevant environmental entailment.

Nova Synthesis

Anticipation theory introduces conceptual dimensions not explicitly articulated in Structural Coupling. Conversely, Structural Coupling foregrounds mechanisms of object and world-constitution absent from Rosen's treatment, where the environment is taken as independently given. Moreover, Structural Coupling provides what Rosen's framework lacks: a historical account of how the encoding functor is established, maintained, and restored after degradation, along with a strong account of cognition.

Since the 1920s, numerous researchers have expressed serious concerns that something fundamental is missing from our theoretical frameworks for understanding organisms. For Rashevsky, it was a theory of organization. For Rosen, Maturana, Varela, and Gánti, it was a theory of self-fabrication. For others, it is the incorporation of probability theory or information-theoretic principles.

We hypothesize that a comprehensive theory of biological organization can be achieved through a *novel synthesis* of existing frameworks. Two steps are critical. First, Structural Coupling and Anticipation must be integrated, as together they address niche construction more comprehensively than Umwelt as originally introduced. This integrated framework provides the proper foundations for a Biology of Cognition urgently needed for contemporary inquiries—such as in multi-species systems like holobionts, where recent work on entangled agencies (Sepúlveda-Pedro et al., 2026) has highlighted the need to account for individual-level agency within collective sympoietic dynamics.

Second, any rigorous theory of biological organization must begin with these foundations. Cognition is not an emergent feature of autopoietic networks but a condition of their very existence: anticipation is built into the regulatory dynamics that enable self-fabrication. Thus, cognition and biological organization mutually presuppose each other and are co-extensive.

Conclusions

The synthesis just presented carries concrete consequences for the enterprise of Artificial Life. The dominant paradigm in contemporary AI (LLMs, deep reinforcement learning, generative networks) produces systems that are, in Rosen's precise terminology, reactive. However sophisticated their behavior, these systems operate within the reactive paradigm described by Equations (1) and (2): their future states are computed from past states via fixed (or gradient-updated) parameters. They do not contain internal predictive models whose dynamics run on a faster timescale; they contain statistical compressions of training data[7]. As Louie (2012) emphasizes, Rosen's theory of anticipation has nothing to do with probability. A conditional probability distribution over next tokens, however large and however useful, is not a predictive model in the Rosennean sense: it is a summary of past correlations, not a dynamical system with its own internal entailment running ahead of reality.

Simple systems do not make mistakes. An error presupposes an internal organizational norm against which behavior can be evaluated as inadequate. In anticipation, that norm is P: when M selects an action that fails to preserve $\hat{O}$, genuine organizational errors occur. No such norm exists in reactive systems; its absence is structural. Rosen's famous result on the uncomputability of organisms has been contested, but analysis under the Sieg-Gandy account of computability (a generalized form of Turing's analysis) suggests that the processes constitutive of genuine anticipatory organization lie outside of what any machine can realize (Astorga, 2025).

The distinction between reactive and anticipatory has a central consequence: a reactive system cannot co-constitute its environment. It can only respond to inputs it receives. An anticipatory system, by contrast, acts on the basis of its model's predictions and thereby transforms the medium it will subsequently encounter. The fixed-point condition we derived in (5) captures something that no reactive system can achieve: the generation of stable sensorimotor invariants (objects, niches, *Umwelten*) through the organism's own anticipatory action. This is precisely what is missing from current artificial systems: they classify a given world but do not construct one.

The artificial life community has long sought open-ended evolution, and even now LLMs appear to be great candidates for open-ended evolving systems (*Karelin et al. (2025)*). The framework developed here suggests a reason why this goal has proven so elusive: open-ended evolution requires agents that are not merely self-reproducing but *anticipatory and structurally coupled*. Consider this: in most artificial evolving platforms, the "environment" is specified independently of the agents. The fitness landscape is given; the agents adapt to it. The environment is not Umwelted. But if organisms are Umwelt constructors then the fitness landscape is not fixed; it is continuously reshaped by the very agents that are adapting to it. This is niche construction understood not as a secondary effect but as a constitutive feature of living organization. Open-ended evolution, from this perspective, requires agents whose anticipatory models and whose environments are coupled in a fixed-point relation: each shapes the other. Without this coupling, evolution will always converge because everything behaves like a simple system.

This Level progression defines a criterion: Level 0 artificial evolving systems cannot achieve open-ended dynamics. Level 1 systems may exhibit longer transients but will eventually exhaust the novelty of a fixed environment. Only Level 2 systems, where the environment is a co-product, have the structural prerequisites for genuinely open-ended dynamics.

An artificial life system worthy of its name must both self-fabricate and Umwelt-construct. Up until now, an LLM's umwelt is not self-constructed but derived from others' umwelten, leaving it without any clear biological congruity.

---

[7] These dynamics are computational dynamics over token sequences, not dynamical models of a natural system, and thus are not conjugate to environmental dynamics via an encoding functor on a faster timescale: no computational system has dynamics that are *physically conjugate* to environmental dynamics, because computation is substrate-neutral.

Acknowledgments

The authors express their gratitude to Arturo González-Bretón for his work on the Tikz code used in this paper and to Diane Greenstein for editorial assistance.

References

Astorga, E. F. S. (2025). Of organisms and machines: Robert rosen's mechanisms and machines under the light of sieggandy computability. Master's thesis, National Autonomous University of Mexico, CDMX, Mexico.

Astorga, E. F. S. and Rosen, J. (2025). International Conference on Anticipatory Systems and Rosennean Complexity (ICAS) @ UNAM-MX - Book of Abstracts.

Cárdenas, M. L., Letelier, J.-C., Gutiérrez, C., Cornish-Bowden, A., and Soto-Andrade, J. (2010). Closure to efficient causation, computability and artificial life. *Journal of Theoretical Biology*, 263(1):79–92.

Dubois, D. M. (2003). Mathematical Foundations of Discrete and Functional Systems with Strong and Weak Anticipations. In Goos, G., Hartmanis, J., Van Leeuwen, J., Butz, M. V., Sigaud, O., and Gérard, P., editors, *Anticipatory Behavior in Adaptive Learning Systems*, volume 2684, pages 110–132. Springer Berlin Heidelberg, Berlin, Heidelberg. Series Title: Lecture Notes in Computer Science.

Gatherer, D. and Galpin, V. (2013). Rosen's (M,R) system in process algebra. *BMC Systems Biology*, 7:128.

Hofmeyr, J.-H. S. (2021). A biochemically-realisable relational model of the self-manufacturing cell. *Biosystems*, 207:104463.

Karelin, G., Nakajima, K., Soto-Astorga, E. F., Carr, E., James, M., and Froese, T. (2025). Indeterminism in large language models: An unintentional step toward open-ended intelligence. Submitted to Artificial Life journal.

Letelier, J. (2022). Cybernetics in chile: a history with unexpected chapters. *AI and SOCIETY*, 37:1105–1113.

Letelier, J. C., Marín, G., and Mpodozis, J. (2002). *Computing with Autopoietic Systems*, pages 67–80. Springer London, London.

Letelier, J.-C., Marín, G., and Mpodozis, J. (2003). Autopoietic and (M,R) systems. *Journal of Theoretical Biology*, 222:261–272.

Letelier, J.-C., Soto-Andrade, J., Guíñez Abarzúa, F., Cornish-Bowden, A., and Cárdenas, M. L. (2006). Organizational invariance and metabolic closure: Analysis in terms of (M,R) systems. *Journal of Theoretical Biology*, 238(4):949–961.

Louie, A. H. (2012). Anticipation in (M,R)-systems. *International Journal of General Systems*, 41(1):5–22.

Maturana, H. (2002). Autopoiesis, structural coupling and cognition: A history of these and other notions in the biology of cognition. *Cybernetics Human Knowing*, 9(3-4):5–34.

Maturana, H. (2012). Reflections on my collaboration with francisco varela. *Constructivist Foundations*, 7(3):155–164.

Maturana, H., Paucar-Caceres, A., and Harnden, R. (2011). Origins and implications of autopoiesis. preface to the second edition of *De Máquinas y Seres Vivos*. *Constructivist Foundations*, 6(3):293–306.

Maturana, H. R. (1975). The organization of the living: A theory of the living organization. *International Journal of Man-Machine Studies*, 7(3):313–332.

Maturana, H. R. and Varela, F. J. (1980). *Autopoiesis and Cognition: The Realization of the Living*, volume 42 of *Boston Studies in the Philosophy of Science*. D. Reidel, Dordrecht.

Mossio, M. (2024). Introduction: Organization as a Scientific Blind Spot. In Mossio, M., editor, *Organization in Biology*, volume 33 of *History, Philosophy and Theory of the Life Sciences*, pages 1–22. Springer International Publishing, Cham.

Poli, R. (2017). *Introduction to Anticipation Studies*, volume 1 of *Anticipation Science*. Springer International Publishing, Cham.

Poli, R., editor (2019). *Handbook of Anticipation: Theoretical and Applied Aspects of the Use of Future in Decision Making*. Springer International Publishing, Cham.

Poli, R. and Valerio, M., editors (2019). *Anticipation, Agency and Complexity*, volume 4 of *Anticipation Science*. Springer International Publishing, Cham.

Rodriguez-Vergara, F. and Husbands, P. (2026). Is Every Cognitive Phenomenon Computable? *Mathematics*, 14(3):535.

Rosen, R. (1958a). A relational theory of biological systems. *Bulletin of Mathematical Biophysics*, 20:245–260.

Rosen, R. (1958b). The representation of biological systems from the standpoint of the theory of categories. *Bulletin of Mathematical Biophysics*, 20:317–341.

Rosen, R. (1959). A relational theory of biological systems II. *Bulletin of Mathematical Biophysics*, 21:109–128.

Rosen, R. (1985). *International federation for systems research*. Elsevier.

Rosen, R. (1988). Effective processes and natural law. In Herken, R., editor, *The universal Turing machine: a half century survey*, Oxford science publications, pages 485 – 498. Oxford Univ. Press, Oxford.

Rosen, R. (1991). *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press, New York.

Rosen, R. (2000). *Essays on Life Itself.* Columbia University Press, New York.

Rosen, R. (2012). *Anticipatory Systems: Philosophical, Mathematical, and Methodological Foundations*, volume 1 of *IFSR International Series in Systems Science and Systems Engineering*. Springer New York, New York, NY. Original publish date: 1985.

Sepúlveda-Pedro, M. A., Ramírez-Vizcaya, S., Hernández-Ochoa, J. L., Vergara-Silva, F., and Gastelum-Vargas, M. (2026). Entangled Agencies: The Enactive Approach and Eco-Evo-Devo in the Scientific Study of Sympoiesis as a Dialectics of Life. *Adaptive Behavior*, page 10597123251411884.

Soto-Andrade, J., Jaramillo-Riveri, S., Gutiérrez, C., and Letelier, J.-C. (2011). Ouroboros avatars: A mathematical exploration of self-reference and metabolic closure. In *Proceedings of the 11th European Conference on Artificial Life (ECAL 2011)*, pages 763–770. MIT Press.

Valdés-Zorrilla, A. (2024). A categorical approach to autopoietic structural coupling. Master's thesis, Universidad de Chile, Santiago, Chile.

Varela, F. J., Thompson, E., and Rosch, E. (1991). *The Embodied Mind: Cognitive Science and Human Experience*. MIT Press, Cambridge, MA.

von Uexkull, J. (1909). *Umwelt und Innenwelt der Tiere*. Julius Springer, Berlin.