The OODA Loop as a FRAM Model of Adaptive Orientation

Abstract

This paper argues that the Rivera–Friston interpretation of Boyd's OODA loop is best understood not as a linear decision cycle, but as a bounded adaptive system for maintaining orientation under uncertainty. Drawing on Friston's active inference, predictive processing, Markov blankets, and generative models, it reframes OODA as a recursive process of sensing, inferring, predicting, comparing, acting, and learning. A compact 16-function FRAM model is proposed to represent this structure through four groups — Observe, Orientate, Decide, and Act — with boundary functions preserving the distinction between internal functional coupling and next-cycle environmental feedback. The model shows that Orientation is not a single cognitive step, but a distributed functional achievement involving hidden-cause inference, prediction-error comparison, surprise detection, fast learned response, deliberative simulation, policy selection, and action. The discussion further argues that FRAM can move beyond qualitative mapping by using metadata and equations to carry observed state, predicted state, prediction error, confidence, uncertainty, risk, ambiguity, and expected free energy. In this form, the model becomes a computable architecture for simulating adaptive behaviour across successive cycles, making Boyd's concept of Orientation inspectable, testable, and relevant to digital-twin and sociotechnical system modelling.

Introduction

The starting point for this work was a simple but important question: can the Rivera–Friston interpretation of Boyd's OODA loop be represented as a FRAM model, rather than as a linear decision cycle? The answer is yes, but only if OODA is understood as a bounded, adaptive, functionally coupled system. The usual version of OODA — Observe, Orient, Decide, Act — is too neat. It is useful as a mnemonic, but misleading as a model. It implies sequence where Boyd was concerned with circularity, interaction, hypothesis, action, and adaptation.

The Rivera–Friston conversation (Rivera 2026),is significant because it recovers that richer interpretation. Rivera and McGrath present Boyd's actual sketch as something far more complex than the familiar loop. They emphasise its multiple arrows, its internal interactions, its central Orientation process, and its language of hypothesis and test. Friston's response is especially important because he recognises this as closely aligned with active inference: action is not merely an output after cognition, but a way of soliciting observations that disambiguate the world and test the agent's current hypotheses. This is a very different reading of OODA. It moves the loop away from decision-cycle folklore and towards a theory of adaptive engagement.

That shift also explains why FRAM is a good modelling language for the problem. FRAM is not a component model, and it is not a process flowchart. It represents a system as a set of functions whose outputs are shaped by inputs, controls, preconditions, resources, and time. It is therefore already concerned with the conditions under which work is possible, variable, constrained, amplified, delayed, or adapted. Hollnagel's FRAM was developed specifically to model complex socio-technical systems, where outcomes emerge from functional resonance rather than from simple linear causation (Hollnagel, 2012/2017).

The System Metamodel using the FRAM methodology

The model we developed was deliberately compact. Instead of trying to reproduce every cognitive, neural, cybernetic, or organisational sub-process in the Rivera–Friston discussion, we reduced the structure to a 16-function teaching model. These functions were arranged into four OODA groups: Observe, Orientate, Decide, and Act. Boundary and background functions were retained to represent the external environment, the availability of sensory signals, prior preferences and learned history, and the reception of action effects in the environment. This gives the model enough structure to preserve the active-inference interpretation without making the visual representation unusable.

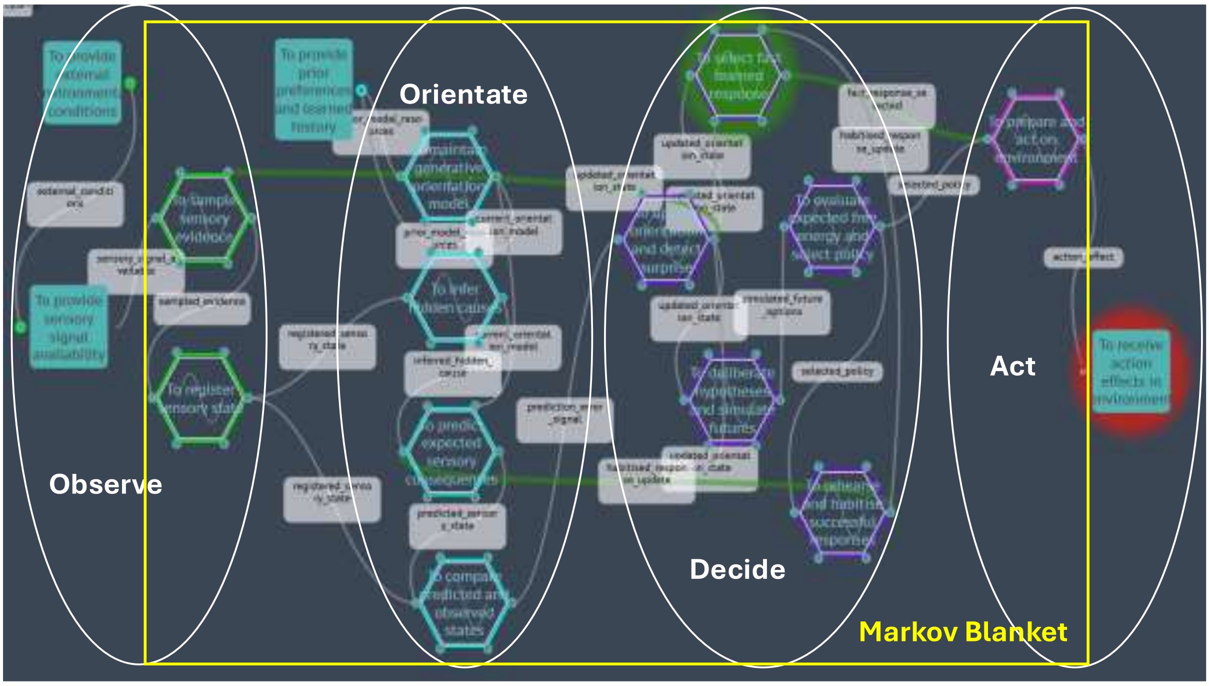

Figure 1 – The System metamodel

The central design decision was to preserve the Markov blanket logic. Friston describes the Markov blanket as the boundary that separates a system from its environment while also coupling it to that environment. It does not seal the system off. It defines the interface through which the system is affected by the world and acts back upon it. Sensory states are influenced by external states and affect internal states. Active states are influenced by internal states and affect the external world. Friston notes that, in systems language, this becomes a formal way of defining inputs and outputs, but he also cautions against treating cognition as a simple flow-through or “sandwich” model. The better interpretation is circular causality around and through the boundary, with internal circular flows shaping what happens at the boundary itself.

That point shaped the final FRAM model. The **Observe** group contains the functions *To sample sensory evidence* and *To register sensory state*. These are intentionally not passive receiving functions. Observation is treated as active sampling. The system does not simply absorb the world; it samples, attends, selects, and registers signals that cross the boundary. This matches Friston’s account of active sensing and active vision, and it also fits the more sophisticated reading of Boyd: observation is shaped by Orientation, not merely fed into it.

The **Orientate** group is the heart of the model. It contains the functions *To maintain generative orientation model*, *To infer hidden causes*, *To predict expected sensory consequences*, *To compare predicted and observed states*, and *To update orientation and detect surprise*. This is where the FRAM model most directly expresses the Rivera–Friston insight. Friston defines a generative model as a probabilistic model of the hidden cause-effect structure producing sensory inputs. Because the causes of sensory evidence are hidden behind the Markov blanket, they cannot be observed directly. They must be inferred from consequences.

This is also where predictive processing and active inference become relevant. Predictive processing treats perception not as passive registration, but as prediction, comparison, and updating. Clark's account of the predictive brain describes cognition as a situated, embodied process in which the brain uses hierarchical prediction to make sense of the world (Clark, 2013). (

Friston's free-energy principle similarly frames perception, action, and learning in terms of optimisation under uncertainty (Friston, 2010). In our FRAM model, Orientation becomes the functional site where this occurs: the system infers hidden causes, predicts what should be sensed, compares prediction with observation, and updates itself when discrepancy persists.

The **Decide** group was not modelled as a simple decision box. This is important. In the simplified OODA loop, Decide often appears as a discrete stage between Orientation and Action. In the Rivera–Friston reading, decision is better understood as hypothesis and policy selection. We therefore represented Decide as two interacting pathways. The first is a fast, learned, System 1-like route: *To select fast learned response*. This corresponds to Boyd's "implicit guidance and control," where Orientation shapes action without a slow deliberative process. The second is a slower, deliberative route: *To deliberate hypotheses and simulate futures*, *To evaluate expected free energy and select policy*, and *To rehearse and habituate successful responses*.

This distinction is compatible with Kahneman's broad System 1/System 2 distinction, provided it is not over-literalised. Kahneman describes System 1 as fast, automatic, and intuitive, and System 2 as slower and more deliberative (Kahneman, 2011). Friston's account is more mechanistic. He distinguishes fast predictive control, which works when the system can act efficiently in a familiar situation, from more prospective planning, which becomes necessary when the system must evaluate alternative paths into the future. In the transcript, he explains that flow-like behaviour can be fast and efficient, but that more complex situations require deliberation over counterfactual futures.

The **Act** group is represented by *To prepare and act on environment*. This function converts either a fast learned response or a selected deliberative policy into an active output across the boundary. The output is then received by the boundary/background function *To receive action effects in environment*. A crucial correction was made here. We removed a same-cycle coupling from the environmental action effect back to the next environmental input. That correction matters because it preserves the Markov blanket interpretation. Action perturbs the environment, but its consequences become available to the system only through the next cycle of observation. The model therefore distinguishes internal same-cycle couplings from cross-boundary temporal recurrence.

This is not a minor technical detail. It is part of the theory. If the action effect were directly coupled back into the sensory input in the same instantiation, the model would collapse the boundary it is trying to represent. The system would no longer be modelling a self-separated from, but coupled to, an environment. It would instead be modelling a closed internal loop. The Markov blanket requires a distinction between what happens inside the bounded system and what returns from the environment as new evidence in the next cycle.

A Complex Adaptive System

What the FRAM model shows is that OODA is not a loop in the ordinary sense. It is a complex adaptive orientation system. Observation is active sampling. Orientation is generative modelling. Decision is hypothesis and policy selection. Action is an environmental test. The consequences of action change the external world, and those changes become available as new sensory evidence in a later cycle. This is why FRAM fits the Boyd–Rivera–Friston argument so well. FRAM is not constrained to linear causality. It can represent functional dependency, circular causality, boundary coupling, variability, feedback, and emergence.

The model also shows why Orientation should not be treated as a single mental act. It is a cluster of coupled functions that collectively sustain the system's adaptive relation to its environment. The system must infer hidden causes, predict expected consequences, compare expectation with observation, detect surprise, and update its internal model. In FRAM terms, this makes Orientation visible as work. It is not a vague cognitive state; it is a functional achievement.

This has direct relevance to Boyd. Rivera and McGrath argue that Boyd's real sketch has been flattened into a linear cartoon. In their reading, Boyd's Orientation shaped observation, decision, and action through multiple arrows, not through a sequential pipeline. Friston's response strengthens that interpretation because he recognises hypothesis testing as a useful way of understanding active inference. The FRAM model turns that claim into a visible structure. It shows how an adaptive system might move from sensory evidence to inferred cause, from inferred cause to prediction, from prediction to error, from error to update, and from update to either rapid action or deliberative policy selection.

The relevance to Friston is equally clear. The model respects the distinction between sensory and active coupling, and it preserves the idea that a system is separated from its environment but not isolated from it. It also reflects the active-inference view that sentient behaviour can be understood in terms of probabilistic inference across perception, planning, and action (Parr, Pezzulo and Friston, 2022). The FRAM model does not reproduce the mathematics of active inference, but it gives those ideas an operational, inspectable form.

The relevance to FRAM is perhaps the most important practical point. FRAM gives us a way to represent Orientation as a distributed functional achievement. The model does not simply say "the agent decides." It shows how sensing, inference, prediction, comparison, surprise detection, fast response, deliberative simulation, policy selection, action, and learning are coupled. It therefore offers a bridge between Boyd's strategic insight and Friston's formal account of adaptive perception and action.

This also brings Ashby back into the discussion. Friston explicitly links the problem of environmental complexity to requisite variety: a system that must engage with a complex world needs sufficient internal degrees of freedom to model and regulate that world. Ashby's law of requisite variety remains one of the foundational insights of cybernetics: only variety can absorb variety (Ashby, 1956). The FRAM model expresses this by refusing to reduce Orientation to a single box. Orientation requires enough internal functional variety to represent ambiguity, novelty, competing hypotheses, fast response, deliberation, learning, and action.

From representation to computation

The model should not be understood merely as a qualitative mind map of the Rivera–Friston–Boyd OODA loop. It is better understood as a computable functional architecture. The visual model shows the structure of adaptive orientation, but metadata can allow each function to carry numerical state information. That is what moves the model from explanation towards simulation.

The key step is to treat each FRAM function as a state-bearing unit. A function does not simply exist as a labelled activity. It can carry metadata describing what it observed, what it predicted, how confident it was, how uncertain or ambiguous the signal was, and whether the function was activated in the current cycle. The Observe functions can carry values for signal strength, signal quality, sensory precision, and sampling priority. The Orientation functions can carry values for belief state, predicted state, observed state, prediction error, confidence, uncertainty, novelty, and Bayesian surprise. The Decide functions can carry values for fast-response readiness, deliberative need, expected information gain, ambiguity, risk, and expected free energy. The Act function can carry values for action confidence, action intensity, and the expected effect of action in the next cycle.

This changes the status of the FRAM model. It is no longer only a diagram of functional relationships. It becomes a structured hypothesis about how an adaptive system updates itself over time. In each cycle, observed sensory data can be imported into the model and compared with what the generative orientation model predicted. The difference between predicted and observed state becomes a prediction error. If the error is small, the model is reasonably calibrated. If the error is large or persistent, the model must update its belief state, reduce confidence, increase uncertainty, trigger further sampling, or move from fast learned response to slower deliberation.

This gives the model a direct connection with active inference. Friston's account treats prediction error as the "newsworthy" part of sensory evidence: the difference between what the system predicted and what it actually sensed. That error drives Bayesian belief updating. In the FRAM model, the equivalent computation can be represented by metadata fields such as OBSERVED_STATE, PREDICTED_STATE, PREDICTION_ERROR, PRECISION, UNCERTAINTY, and UPDATED_BELIEF. These need not be perfect neuroscientific variables at the first stage. They can be defensible proxies that allow the model to simulate adaptive behaviour.

The Decide group can then use the updated orientation state to choose between response modes. If the situation is familiar, prediction error is low, and confidence is high, the model may select a fast learned response. If ambiguity, novelty, risk, or surprise is high, the model may activate a deliberative pathway, simulate possible futures, evaluate expected free energy, and select a policy. Friston's work on active inference and learning is relevant here because expected free energy can be decomposed into risk and ambiguity, and behaviour can be understood as balancing pragmatic and epistemic value: acting to achieve preferred outcomes while also seeking information where uncertainty matters (Friston et al., 2016).

The model can therefore be run as a cycle-by-cycle simulation. In cycle N, the system samples sensory evidence, registers the observed state, infers hidden causes, predicts what should be happening, compares prediction with observation, and updates its orientation. It then decides whether to act, deliberate, or gather more evidence. The action output crosses the boundary and perturbs the external environment. In keeping with the Markov blanket interpretation, that action does not feed back into the same cycle. Its consequences become part of the external conditions available to the next cycle. In cycle N+1, new sensory evidence allows the model to test whether the previous action produced the expected effect.

This makes the model suitable for testing behaviour, not just describing it. Signal quality can be reduced. Ambiguity can be increased. Prior confidence can be strengthened or weakened. Learned response habits can be made more or less dominant. Time available for deliberation can be shortened. The model can then show whether the system tends towards rapid skilled action, cautious deliberation, repeated evidence-seeking, overconfident misclassification, delayed response, or maladaptive persistence in a wrong hypothesis. In FRAM terms, variability can be introduced into inputs, controls, resources, timing, and outputs. Metadata then allows us to examine how that variability propagates through the OODA system.

The important point is that computation does not require the first version to implement the full mathematics of active inference. A practical first version can use proxy variables: observed state, predicted state, prediction error, confidence, uncertainty, risk, ambiguity, information gain, and expected free energy. These values can be calculated in Excel, FRAIXL, FMI, or another execution environment, then written back into the FRAM metadata. FMV can then act as the visual explanation layer, showing which functions are active, uncertain, surprised, confident, overloaded, or driving action.

This is what turns the model into a practical epistemic digital twin. The FRAM structure defines the functions and couplings. The metadata defines the state of each function. The equations define how those states change over time. The cycle logic defines how the system learns from discrepancy. The visual model then shows not only what the OODA system is, but how it behaves under different conditions. It can be used to compare fast and slow response, calibrated and miscalibrated orientation, exploration and exploitation, adaptive and maladaptive action.

Conclusion

The result is not merely a new diagram of OODA. It is a way of making Orientation computationally inspectable. Boyd provides the insight that adaptive action depends on Orientation. Friston provides the formal logic of prediction, uncertainty, surprise, and action. FRAM provides the functional architecture. Metadata and equations then turn that architecture into a runnable model that can simulate, test, and compare behaviours across successive cycles.

The claim should therefore be made carefully but confidently. FRAM is not active inference. It is not a model of neurons. It is not a substitute for Friston's mathematics. But it is a powerful practical system model for representing the functional organisation that active inference describes: boundary, sensing, hidden-cause inference, prediction, comparison, surprise, action, and learning. That makes it highly relevant to the Rivera–Friston interpretation of Boyd. It shows that OODA is not best understood as a four-step loop. It is better understood as a functional architecture for adaptive orientation under uncertainty.

David Slater

References

Ashby, W. R. (1956). *An Introduction to Cybernetics*. Chapman & Hall. (ashby.info)

Clark, A. (2013). "Whatever next? Predictive brains, situated agents, and the future of cognitive science." *Behavioural and Brain Sciences*, 36(3), 181–204. (Edinburgh Research)

Friston, K. (2010). "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, 11, 127–138. (Nature)

Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. and Pezzulo, G. (2016). "Active inference and learning." *Neuroscience & Biobehavioural Reviews*, 68, 862–879. (ScienceDirect)

Hollnagel, E. (2012). *The Functional Resonance Analysis Method: Modelling Complex Socio-technical Systems*. Ashgate. (erikhollnagel.com)

Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus and Giroux / Penguin. (Penguin)

Parr, T., Pezzulo, G. and Friston, K. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. MIT Press. (direct.mit.edu)

Rivera, B., McGrath, M. and Friston, K. (2026). *Karl Friston Decodes the Real OODA Loop: Active Inference and What Boyd Got Right*. (https://www.youtube.com/watch?v=e8R7KicFCSs Transcript supplied)