FROM MIND MAPS TO NEURONS: THE EVOLUTION OF FRAM TOWARD QUANTITATIVE PREDICTOR – CORRECTOR SYSTEMS

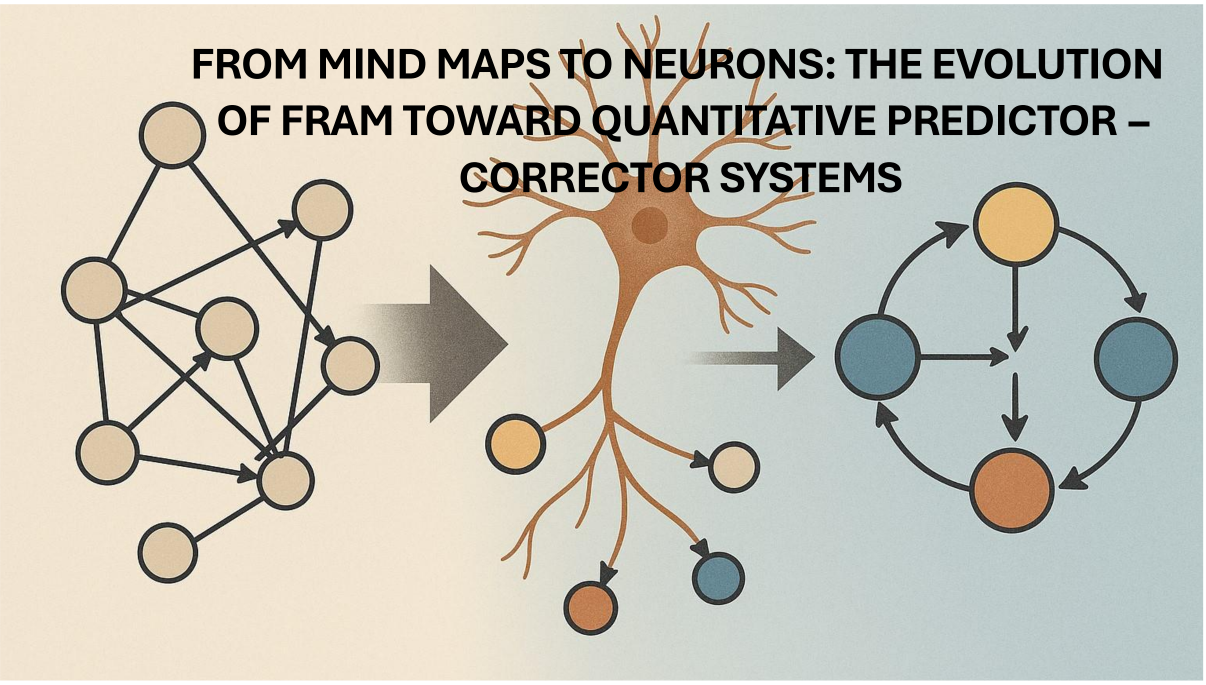

Abstract

This article traces the evolution of the Functional Resonance Analysis Method (FRAM) from its early role as a qualitative "mind map" of sociotechnical variability, through its analogy with neuronal and cognitive architectures, to its current development as a quantitative predictor–corrector framework. We demonstrate how FRAM can operationalise John Boyd's OODA loop in the context of aircraft landing, modelling Observe, Orient, Decide, and Act not as abstract stages but as interdependent functions with measurable properties. Predictor–corrector dynamics, residuals between predicted and observed values, and doctrinal gates are encoded directly in metadata, enabling decision to be treated as a quantifiable process rather than a black box. Extending Llinás's framework for situation control, the Orient phase is decomposed into functions that incorporate memory, doctrine, and cognitive filters alongside sensor fusion. Results show that FRAM can generate traceable time series of OODA activity, enforce stabilisation barriers, and reveal how decision restores congruity under stress. The approach demonstrates both the potential and the limitations of quantification, offering a credible pathway from metaphor to model in the analysis of decision-making within complex sociotechnical systems.

Key words - Functional Resonance Analysis Method (FRAM); OODA loop; predictor–corrector; decision-making; situation awareness; aircraft landing; sociotechnical systems; variability; metadata modelling; resilience engineering.

Introduction

The Functional Resonance Analysis Method (FRAM) was originally conceived as a way of visualising and reasoning about the variability of sociotechnical systems. Early applications often took the form of elaborate "mind maps," where functions were arrayed as abstract nodes and the couplings between them were traced to reveal emergent pathways. This stage was invaluable for shifting thinking away from linear cause–effect chains and toward a recognition of systems as adaptive and resonant. Yet these early FRAM models remained largely qualitative, serving as exploratory heuristics rather than predictive instruments.

As the method matured, analogies with neuronal and cognitive architectures began to appear. Each function could be seen as a quasi-neuron: firing only when the right aspects were present, shaping outputs that altered the state of downstream functions, and contributing to higher-order behaviour through dense coupling. This perspective aligned FRAM with insights from neuroscience and active inference, highlighting the similarity between distributed cognitive architectures and the swarm-like interactions of FRAM functions. Instead of being static diagrams, FRAM models began to be understood as generative substrates, capable of representing attention, prediction, and correction in much the same way that cortical columns implement cycles of expectation and update.

The latest development pushes FRAM into the realm of quantitative modelling, where predictor–corrector loops can be encoded directly in function metadata and exercised against empirical data. In this mode, a “Decide” function is no longer a black box but a mathematically accessible operation: residuals between predicted and observed values can be tracked, control laws can be tested, and barriers can be enforced with formal ACTIVATE conditions. This makes it possible to move from narrative to measurement, and from metaphor to model. Taken together, these stages represent an evolution of FRAM thinking—from mind maps as exploratory sketches, to neuron analogues as conceptual bridges, and finally to quantitative predictor–corrector systems capable of simulating decision-making in real tasks.

The application to aircraft landing, and the embedding of John Boyd’s OODA loop within a FRAM model, illustrates the potential of this trajectory: FRAM not only describes how systems resonate but shows how decisions can be represented, calibrated, and ultimately trusted in practice.

The OODA Loop

The question of whether John Boyd’s OODA loop can be given rigorous, quantitative treatment has long preoccupied both theorists and practitioners, but it has rarely been demonstrated outside metaphor. Aircraft landing provides an especially rich example because it concentrates many of the features that Boyd sought to capture such as rapidly unfolding dynamics, thin safety margins, and the interplay of human judgment, automation, and environmental uncertainty. The Functional Resonance Analysis Method (FRAM) offers a distinctive way to instantiate the OODA loop in a practical, explorable form. Developed to analyse complex sociotechnical systems, FRAM models system behaviour as the emergent result of interdependent functions, each of which may vary in timing, precision, and reliability (Hollnagel, 2012). This makes it an ideal candidate for exploring how Observe, Orient, Decide, and Act can be treated not as abstract phases but as functional nodes with quantifiable properties.

In the landing task, “Observe” corresponds to functions that receive and process airspeed, altitude, glidepath, and visual updates. Each of these is subject to known sources of variability—latency in radar altimeter signals, turbulence affecting indicated airspeed, pilot scan patterns—that can be represented through FRAM’s aspect-based couplings. “Orient” can then be modelled as a fusion function, producing both an integrated state estimate and residuals between predicted and observed values. This residual information is critical, because it sets up “Decide” as a predictor–corrector function: the task is not simply to choose an action, but to correct the internal prediction so that congruity is restored between expectation and perception. “Act” is captured through actuator functions—thrust, pitch, flaps, gear—whose actual performance is logged and fed back, closing the loop.

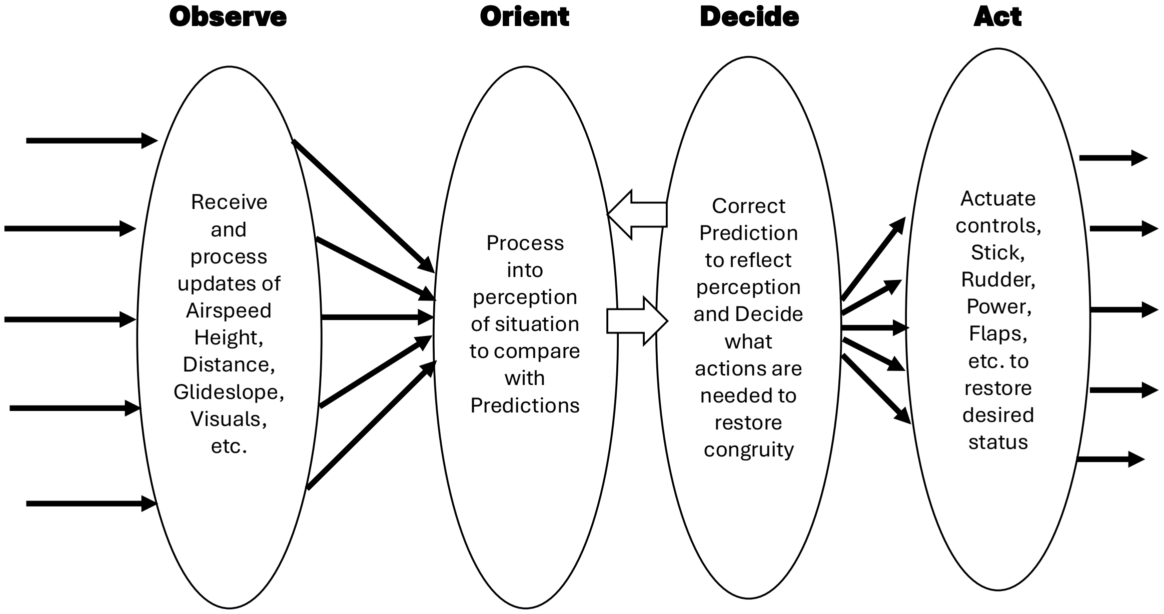

Figure 1 – Boyd's OODA Loop in the cockpit

This approach allows decision to be expressed quantitatively, rather than remaining a black box. Metadata in FRAM can encode prediction models of energy state, confidence bounds on sensor values, and latency corrections. For example, thrust commands can be modelled as functions of airspeed residuals with explicit tolerance limits, while actuator feedback supplies the measured divergence between commanded and actual settings. Over repeated operations, these predictor–corrector dynamics can be logged and analysed, offering a dataset in which OODA cycles become observable traces rather than theoretical constructs. In this way, the OODA loop is operationalised as a digital twin of the landing system.

Another advantage of FRAM in this context is its treatment of barriers and gates. Stabilised approach criteria at 1000 and 500 feet above ground level, widely mandated in airline standard operating procedures (ICAO, 2018; FAA, 2015), can be embedded directly as conditional metadata. These act as control aspects that either permit continued decision-making or activate a go-around pathway. This reflects an often-overlooked feature of Boyd's conception: the OODA loop is not free-running but bounded by doctrine and external constraints (Boyd, 1987). By embedding barriers explicitly, FRAM prevents the model from collapsing into an unconstrained spiral of updates and actions and ensures that policy discipline is preserved.

There are, however, limitations. Boyd's "Orient" was not just technical fusion but encompassed cultural filters, doctrinal patterns, and the shaping of perception by prior experience. A FRAM model risks narrowing this richness if it is reduced to residual computation alone. Similarly, the quality of quantification depends entirely on the fidelity of the metadata: if latency, sensor bias, or human scan variability are simplified too aggressively, the resulting predictor–corrector loops may exhibit apparent stability while masking systematic drift. Llinas Extends the "Orient" phase far beyond fusion of sensor data. He integrates *mental modeling, memory, doctrine, cultural context, and learned patterns*. The framework suggests that effective orientation is about constructing and maintaining a control of the situation, not just processing inputs.

This fills a gap in most quantitative or technical OODA treatments. If you only model residuals and sensor fusion, you miss the shaping role of prior experience and organisational filters. For

FRAM, it suggests “Orient” should include functions representing training, doctrine, and cognitive frames, not just statistical prediction.

To bring Llinás et al.’s expanded “situation control” into a FRAM model of the OODA loop, the Orient phase needs to be decomposed into a richer set of functions that move well beyond simple sensor fusion. Instead of a single fusion node, Orient can be modelled as a swarm of interdependent functions:

- parsing raw sensor data,
- injecting doctrine and operational context,
- retrieving memory and prior experience,
- maintaining and updating a mental model of the aircraft’s state,
- generating competing hypotheses about what is happening,
- testing those hypotheses against both physical and doctrinal constraints,
- projecting future states, allocating attention, and
- monitoring for cognitive traps such as fixation or automation bias.

Each of these functions produces outputs that are passed downstream, while control and resource aspects represent the influence of standard operating procedures, training, and cultural filters.

Quantitative metadata can be used to encode prediction and update rules, weights for memory and doctrine, bias indices, and thresholds that determine whether Orient produces a reliable output or triggers a re-orient or go-around pathway.

- This approach makes “situation control” explicit:
- attention plans feed back into the Observe functions,
- bias flags adjust the confidence of the mental model, and
- stabilisation gates ensure that doctrinal limits are respected.

By doing so, FRAM captures the richness of orientation as described by Llinás—memory, context, and doctrine are embedded alongside residuals and projections—transforming Orient from a black box into a transparent, testable set of interacting functions.

Despite these caveats, the exercise demonstrates that FRAM can provide a practical and testable instantiation of the OODA concept. By embedding predictor–corrector functions in an interdependent web of Observe, Orient, Decide, and Act activities, the abstract cycle acquires operational content. In the concrete case of aircraft landing, this allows researchers to explore how observation and prediction diverge under stress, how decision restores or fails to restore congruity, and how barriers enforce bounded adaptation. In short, FRAM moves OODA from metaphor toward model: not a slogan, but a system that can be traced, quantified, and validated against the unforgiving reality of the runway.

Results and Discussion

Figure 2 – A simplified FRAM model showing the OODA Loop mapping of the Functions

Applying the Functional Resonance Analysis Method (FRAM) to the aircraft landing task demonstrates the viability of using this methodology to instantiate John Boyd’s OODA loop as a practical, quantitatively explorable system. The FRAM model decomposed the landing sequence into interdependent functions corresponding to Observe (sensor updates of airspeed, altitude, glidepath, and visual cues), Orient (fusion into a situation estimate and residuals against predictions), Decide (predictor–corrector updating and command generation), and Act (control of thrust, pitch, configuration, and feedback verification). Each of these was represented with metadata describing expected performance, variability, and latency, allowing the model to be exercised not just qualitatively but numerically.

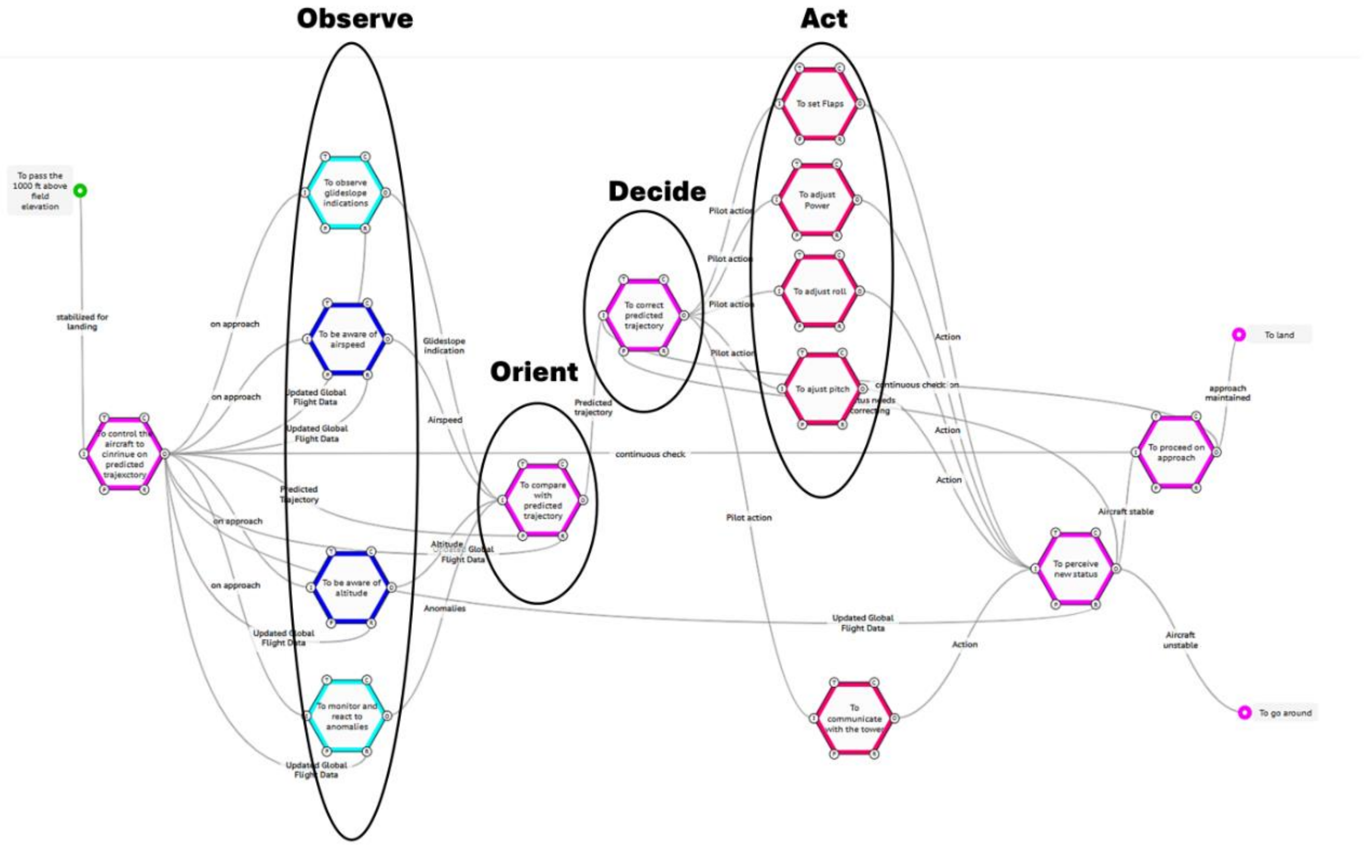

Figure 2 – A simplified FRAM model showing the OODA Loop mapping of the Functions

Simulation of this model suggests that the OODA cycle can be rendered traceable in terms of prediction error dynamics. For example, deviations between predicted and measured airspeed (ΔIAS) can be logged across the approach, with the Decide function adjusting thrust commands via metadata equations. Residuals in glidepath error similarly feed the predictor–corrector loop for pitch. In both cases, the cycle closes through actuator feedback, allowing decision effectiveness to be expressed in measurable terms such as convergence time, overshoot, and stability margins. By embedding these parameters into FRAM metadata, the model yields time series of OODA activity that can be compared across flights, phases of operation, or pilot techniques.

A further result is that stabilised approach barriers, implemented as ACTIVATE metadata at 1000 and 500 feet above ground level, acted as control gates on the Decide function. When predictive residuals exceeded tolerance or stabilisation criteria were unmet, the model forced a go-around pathway, halting further decision iterations in favour of a safety-bounded exit. This captures a key feature of real-world operations: OODA loops do not run unchecked but are framed by doctrinal rules (ICAO, 2018; FAA, 2015). Embedding those barriers in FRAM metadata ensured that decision-making remained consistent with established safety practice.

Conclusions

From an analytical standpoint, the model demonstrates both strengths and limitations. Its strength lies in converting Boyd's abstract notion of decision into a predictor–corrector function, permitting quantitative exploration of how congruity between perception and prediction is restored. This makes it possible to measure decision quality in terms of error reduction, rather than leaving it as a cognitive metaphor. It also allows for systematic exploration of hazards, such as sensor latency or actuator lag, by perturbing metadata values and examining the resulting resonance across functions.

The limitations, however, must be recognised. Boyd's conception of "Orient" was broader than sensor fusion: it included cultural filters, doctrinal assumptions, and prior experience shaping interpretation. While FRAM can in principle accommodate such variability, the quantitative treatment risks narrowing it to residual calculation. Equally, the viability of the approach depends on accurate metadata calibration. If confidence limits, latency parameters, or actuator dynamics are simplified, the predictor–corrector loops may appear to function smoothly while concealing hidden instabilities. Empirical validation against recorded flight data is therefore essential.

Taken together, these results suggest that FRAM provides a credible path to operationalising OODA. The landing case shows how predictor–corrector loops can be captured and explored within FRAM, and how stabilisation gates can be enforced as barriers. The discussion highlights that quantification is not trivial, but with careful metadata design and empirical grounding, OODA can be made into a testable construct rather than a rhetorical device. This represents an advance in both the modelling of sociotechnical systems and the practical use of Boyd's theory: decision emerges not as an opaque box, but as a functional, measurable process embedded in the dynamics of a real task.

References

- Boyd, J. (1987). *A Discourse on Winning and Losing*. Maxwell AFB: Air University Library.
- FAA. (2015). *Stabilized Approach Criteria*. Federal Aviation Administration Advisory Circular AC 120-71A.
- Hollnagel, E. (2012). *FRAM: The Functional Resonance Analysis Method: Modelling Complex Socio-technical Systems*. Farnham: Ashgate.
- ICAO. (2018). *Manual of All-Weather Operations* (Doc 9365, 4th Ed.). Montreal: International Civil Aviation Organization.
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2), 127–138.
- Osinga, F. P. B. (2007). *Science, Strategy and War: The Strategic Theory of John Boyd*. London: Routledge.
- Slater, D., & Hill, R. (2024). *Quantifying Human Factors in Complex Sociotechnical Systems Using the FRAM.* Technical Report / Preprint. DOI 10.3389/fpubh.2023.1160971.
- Brehmer, B. (2005/2006). "The Dynamic OODA (DOODA) loop..." (ICCRTS / IC2J).
- Llinas, J. (2022). "An Expanded Framework for Situation Control." *Frontiers in Systems Neuroscience*.
  Endsley, M. R. (1995). "Toward a Theory of Situation Awareness in Dynamic Systems." *Human Factors* (+ SAGAT lineage).