



Governing AI-Supported Design and Verification Boundaries in Systems Engineering: A Three-Axis Interpretation of the V-Model

**Toshiyasu Watanabe, Dr. Eng.
Waseda University**

Presenter Background



More than 25 years of cross-industry experience across aerospace, automotive, and ICT.



Aerospace

Mitsubishi Heavy Industries

Flight-control R&D and system development for civil aircraft and helicopters



Automotive

Nissan Motor Corporation

Product planning, vehicle testing, and product-attractiveness evaluation



ICT / Consulting

Technology & Management Consulting Firms


ICT, R&D strategy, and social-system planning

Background and Problem

Why engineering boundaries are becoming dynamic



Digital engineering makes design, verification and validation more iterative.




As iteration increases, engineering boundaries shift across levels.



Engineers have to think about vertical and horizontal viewpoints between system boundaries and mission boundaries.

FMEA and FTA are typical examples of this kind of reasoning.




This complexity creates the need for engineering coordinates.



Recently, AI has increasingly supported the interpretation and analysis of engineering information across the V-model.


AI outputs from these activities can influence design assumptions or verification criteria.





However, AI-supported outputs remain engineering candidates until they pass an engineering acceptance gate.

The gate makes the scope, assumptions, traceability, reviewability, and acceptance conditions explicit. Only after this gate can an AI-supported output become engineering evidence.



The issue is not probability itself.


Approach

How Three-Axis structures AI-supported boundary change



Previous research introduced a three-axis framework: System Architecture, Mission Architecture, and Decision Architecture.

These correspond to WHAT, WHEN, and WHO.




These three axes structure engineering status.



This research extends the three-axis framework into 21 reference components. These components provide explicit engineering coordinates.


10 components for SA, 6 for MA, and 5 for DA.



21 components turn the three axes into engineering coordinates.

Three Axes and 21 Reference Components


A common reference space for engineering information





The key point of this approach is to keep unresolved information explicit.

Information that cannot be resolved within the reference space is retained as an engineering residual.



Engineering residual prevents AI from force-fitting unresolved information.




Together with the Three-Axis framework, three viewpoints support the interpretation of decision information: process, hierarchy, and role.

Process traces do not appear in isolation. They are embedded in system levels and organizational role functions.

Three viewpoints

Process Viewpoint: CIO-GoF


What part of the decision process is observed?



- ✓ Context
- ✓ Intelligence
- ✓ Orchestration
- ✓ Governance
- ✓ Feedback

Hierarchical Viewpoint: L0-L3


At what level does the decision information appear?



- ✓ L0 : Strategic
- ✓ L1 : Program
- ✓ L2 : Engineering
- ✓ L3 : Operations

Role Viewpoint: PEAlF

What role function appears?



- ✓ Proposer
- ✓ Evaluator
- ✓ Approver
- ✓ Implementer
- ✓ Feedback Owner

The same decision trace can be interpreted differently by process position, system level, and role function.


Illustrative Case

How the approach applies to flight-control development



This case uses a flight-control system because it is mission-critical and combines digital control with physical response.

The flight-control computer (FCC) processes pilot input and sensed aircraft state to generate an elevator command. The actuator moves the elevator, which produces the aircraft response. Aircraft-state feedback closes the control loop.



Physical response conditions

AoA • dynamic pressure • aerodynamic effectiveness • hinge moment • actuator force/rate • bandwidth / delay

Here, AI supports the engineering analysis, not the flight-control loop.



This case asks how AI can support FMEA and FTA, and how the three-axis framework can structure that support.

Configuration, physical capability, and flight conditions are documented in separate artifacts. Can AI relate this information and help define a traceable verification question?

Baseline: Verification information is distributed

Configuration, physical capability, and flight condition are documented in separate artifacts. Their combined verification question is not explicit.

Case study: Can Three-Axis structure the verification question?

WHAT × WHEN × WHO

Define what to verify, under which conditions, and who is responsible.

Existing engineering artifacts (examples)

Mode definition	Degraded / reconfigured state
FMEA	Degraded actuator performance (e.g., reduced force / rate)
FTA	Insufficient physical authority path under certain conditions
Verification	Demanding condition not explicit in current test coverage

AI-assisted cross-artifact review


Relate information across artifacts

Reconfigured control path (logical)



Control command (e.g., elevator deflection)

Physical response path (condition-dependent)



Aircraft state feedback (pitch rate, attitude, AoA, etc.)

The aim is to support engineering review, not to replace engineering judgment.



This walkthrough applies the three-axis framework to a case of degraded actuator performance.

After reconfiguration, the FCC still generates an elevator command, but actuator capability is reduced. The available verification information does not explicitly cover this degraded configuration under demanding flight conditions.


Baseline: Verification information is distributed

Configuration, physical capability, and flight condition are documented in separate artifacts. Their combined verification question is not explicit.

Case study: Can Three-Axis structure the verification question?

WHAT × WHEN × WHO

Define what to verify, under which conditions, and who is responsible.




Can the degraded system still achieve the required aircraft response under these conditions?



The three-axis framework with 21 reference components structures the analysis from FMEA and FTA to physical verification.

AI-assisted analysis traces information across FMEA and FTA artifacts. The selected components help define what to verify, under which conditions, and who evaluates the evidence. Physical evidence then supports engineering reassessment.




21 Reference Components Specify the Verification Scope



The selected reference components make the verification scope explicit across WHAT, WHEN, and WHO.


They identify the affected system elements, operating conditions, timing constraints, evidence requirements, and decision responsibilities. In the illustrative simulation, the degraded configuration meets the objective under nominal conditions, but not under demanding conditions.





WHAT, WHEN, and WHO define the engineering response.

Engineers use additional verification to decide what action to take. They update the affected artifacts and verify any changes. If the evidence is insufficient, they return to verification instead of closing the loop.



**AI-assisted analysis identifies a candidate concern.
Engineers verify, decide, update, and close the loop.**


Discussion and Conclusions

The role of the framework and the key takeaways



The three-axis framework can complement these methods by making engineering context explicit.

The 21 reference components structure context before AI processing and provide reference points for output review. Integration with these methods remains future work.





The three-axis framework helps keep AI-supported engineering traceable across changing boundaries.

1

**Explicit
engineering status**



AI-supported output
requires explicit
engineering status.

2

**Independent
coordinates**



WHAT, WHEN, and WHO
provide independent
coordinates for
engineering use.

3

Explicit residual



Unresolved information
remains explicit as a
residual.

Explicit status, coordinates, and residuals help keep AI-supported contributions reviewable.



Thank you

Toshiyasu Watanabe, Dr. Eng.

Waseda University

watanabet@aoni.waseda.jp




- [1] International Council on Systems Engineering (INCOSE), Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities, 5th ed., INCOSE-TP-2003-002-05, Wiley, 2023.
- [2] S. Friedenthal, A. Moore, and R. Steiner, A Practical Guide to SysML: The Systems Modeling Language, 3rd ed., Morgan Kaufmann, 2015.
- [3] SAE International, ARP4761A: Guidelines for Conducting the Safety Assessment Process on Civil Aircraft, Systems, and Equipment, SAE International, 2023.
- [4] T. Watanabe, “Decision Architecture: An Independent Third Dimension for Socio-Technical Systems in the AI Era,” presented at CSER 2026, Arlington, VA, Apr. 2026.
- [5] T. Watanabe, “Decision Architecture: An Observation Framework for Structuring Decision Information in Systems Engineering,” INCOSE International Symposium, vol. 36, no. 1, pp. 1827–1846, 2026. doi: 10.1002/iis2.70224.
- [6] SAE International, ARP4754B: Guidelines for Development of Civil Aircraft and Systems, SAE International, 2023.
- [7] ISO/IEC/IEEE 42010:2022, Software, Systems and Enterprise — Architecture Description, 2022.

Appendix




Engineering use of AI output requires three questions: what it affects, when it is valid, and who can accept it.

These questions make engineering boundaries, applicable conditions, and acceptance authority explicit.






The purpose of this research is to define a governance process that makes AI-supported contributions reviewable and verifiable.




Three Axes and 21 Reference Components

A case activates only a subset



Three Axes and 21 Reference Components


A case activates only a subset





The V-model works not only through formal artifacts, but also through implicit coordination by the project decision owner.

This role connects both sides of the V-model and also aligns decisions across organizational levels, from strategy to operations.




**The role exists in practice.
The decision context around it often remains implicit.**



Can Decision be described at the same conceptual level as System and Mission?

ISO/IEC/IEEE 42010 already recognizes architecture decisions and rationales, but decision context often remains difficult to observe, compare, and trace across SE activities.




This is the entry point for Decision Architecture.

Source: T. Watanabe, "Decision Architecture: An Observation Framework for Structuring Decision Information in Systems Engineering," INCOSE International Symposium 2026



Decision Architecture (DA) adds a *Who* axis to observe decision intent around *What* (system) and *When* (mission).

The goal is not to replace existing SE models. The goal is to make decision information around them observable.




**“WHO” means decision intent, responsibility, and role position
— not a person name.**




A decision process is guided by priorities, criteria, responsibility, and intent — not only by actions.

A decision process is not directly observable as an internal state; it becomes observable through external expressions and process traces — an outside-in approach.

SE activities as external traces



Process view of decision-making flow




Decision flow reveals what is being preserved, changed, accepted, or rejected.




CIO-GoF makes the projection process explicit by separating the processing flow from its governance.

Context (C), Intelligence (I), and Orchestration (O) form the processing flow. Governance of Feedback (GoF) provides reference conditions and supports refinement.






A landscape view maps representative AI-related technologies to engineering activities across the V-model.



- Illustrative mapping based on the current landscape review; not exhaustive.




The three-axis framework applies the same three governance views to AI-supported contributions across technologies and engineering activities.





The engineering activity changes. **WHAT, WHEN, and WHO** remain the same governance views.

Each activity applies the three views to its own engineering objects, conditions, and authority.



Same governance views. Different engineering objects.



- **Extend empirical evaluation across engineering tasks**
- **Expand application across domains and lifecycle activities**
- **Integrate Three-Axis context with MBSE, RAG, and knowledge graphs**
- **Develop metrics for projection fit, residuals, and output consistency**