

Article

Diagnosing Human–AI Fit in Mid-Market Operating Systems: A Viable-System Alignment Model

Andrej Naraločnik ¹, Andrej Bertoncej ², Maja Meško ^{3,*} and Tine Bertoncej ²¹ Independent Researcher, 1000 Ljubljana, Slovenia; andrej.naralocnik@siol.net² Faculty of Management, University of Primorska, 6000 Koper, Slovenia; andrej.bertoncej@fm.upr.si (A.B.); tine.bertoncej@fm.upr.si (T.B.)³ Faculty of Organizational Sciences, University of Maribor, 4000 Kranj, Slovenia

* Correspondence: maja.mesko@um.si

Highlights

Please indicate how your work links to systems science via your contributions to systems practice, theory, and or methodology.

- The paper treats human–AI fit as a viability problem: Ashby’s variety law and Beer’s recursive functions define what must remain intact when operational judgment is shared with a model.
- It contributes a design-science artifact that operationalizes that theory as four diagnostic dimensions usable in mid-market firms that lack an enterprise control staff.

What are the main findings and/or the implications of the main findings?

- Fit is jointly social and technical. Observability, decision rights, codification, and metabolic load are considered together in diagnosing human–AI fit; improving only the model leaves the viable system uncorrected.
- Each dimension maps onto a Viable System Model (VSM) home. That mapping gives practitioners a systems language for Artificial Intelligence (AI) adoption that generic readiness checklists do not provide.

Abstract

Mid-market firms now share operational judgment with machine-learning systems while remaining too small to staff enterprise-grade governance. Viability theory and recent taxonomies of AI-era organizational pathologies explain why that combination is structurally demanding; they do not yet supply a diagnostic of fit between the human activity system and the model. This article describes the diagnostic. Following a design-science process on the practice-to-theory path, the Viable-System Alignment Model (VSAM) was specified by reanalysis of a theoretically sampled mid-market archive documented between 2003 and 2025. The original archive showed recurring and subsequently stable patterns concerning observation, authority, written judgment, and coordination load, which were reanalyzed as candidate conditions of human–AI fit. The artifact comprises Observability, Decision Rights, Codification, and Metabolic Load, each with a home in the Viable System Model and a socio-technical failure mode. The central proposition is that human–AI fit is a property of the viable system rather than of the software itself. Evaluation at this stage is analytical and descriptive rather than observational; independent field evaluation is reserved for subsequent research.



Academic Editors: Dennis Pepple and Aliyu Sambo

Received: 6 September 2026

Revised: 3 October 2026

Accepted: 16 October 2026

Published: 21 October 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](#)

[Attribution \(CC BY\)](#) license.

Keywords: viable system model; socio-technical systems; human–AI alignment; design science research; mid-market organizations; requisite variety; dynamic capabilities; organizational diagnosis

1. Introduction

An operating system is controllable only to the extent that the variety of the regulator matches the variety of the disturbances it must absorb [1]. That law is not a metaphor. When the disturbance set grows faster than the repertoire of responses, the organization does not become “more agile”; it becomes less able to select a response in time. Mid-market firms now face that condition in a specific form. They have added a second regulator, an artificial intelligence or machine learning layer that classifies, recommends, or, in some cases, acts without adding the specialized staff through which large enterprises absorb residual variety. The firm must continue to grow, maintain the core, and remain intelligible to customers and capital providers. The model increases the speed of some decisions and the opacity of others. The variety gap does not disappear; it changes shape.

Beer’s viable system specifies the recursive functions that have to remain intact if an organization is to survive that load [2]. Operations (System 1) need coordination (System 2), control and audit (System 3 and System 3*), intelligence about the environment (System 4), and a policy function that can still say what the organization is for (System 5). Recursion means that the same five functions appear at more than one level. A mid-market firm that installs a model inside one operational unit without redesigning coordination, audit, and identity does not thereby become more viable. It relocates variety into channels that were never designed to carry it.

That reading is no longer only theoretical. Applied to organizations that embed artificial intelligence, the viable-system architecture already supports a taxonomy of structural, functional, and informational pathologies [3]. Identity can collapse into a vendor tool. Algedonic channels can go silent because nobody is assigned to hear the alarm. Intelligence can be captured by a dashboard that reports what the model can see and not what the environment is doing. Those pathologies name the failure. What they leave open is the diagnosis of fit: a structured answer to whether the human activity system and the model occupy the same viable system, or whether one is running on top of the other as an unassimilated implant.

That gap is the subject of this paper. The organizations considered here are mid-market firms: large enough that managerial intuition alone no longer covers the disturbance set, yet small enough that a dedicated enterprise-level control function is typically not a realistic response. Recent research on AI adoption in small and medium-sized enterprises similarly shows that implementation depends on a combination of technological and organizational conditions, including compatibility, infrastructure, knowledge, resources, culture, competition, regulation, and the surrounding ecosystem [4]. In that band, artificial intelligence is rarely a separate digital transformation program. It arrives as a vendor module inside sales, operations, finance, or service, and it immediately raises four practical questions that existing checklists treat as separate: Can anyone see what the model is doing in time to intervene? Who is allowed to overrule it? Which judgment has been written down, and which still lives in one person? Does the new layer add coordination work that the firm cannot pay for? Those questions are socio-technical, not technical. A more accurate model that nobody may overrule, that nobody can observe, and that consumes the last hours of managerial attention is not an improvement in viability.

In this study, human–AI fit is defined as the degree to which the human activity system and an AI-enabled technical layer are integrated into a single viable system, such that the organization retains timely observability, usable decision rights, codified operational judgment, and sufficient capacity to absorb the resulting coordination load. This differs from task–technology fit, which concerns the match between technology capabilities and the requirements of the tasks it supports [5], and from person–technology fit, which concerns congruence between individual preferences and the characteristics of the technology with which the individual works [6]. It also differs from AI readiness, which assesses the organizational conditions associated with AI adoption [7], and from algorithmic governance, which concerns the organizational rules, accountability structures, and controls governing algorithmic systems [8,9]. VSAM instead diagnoses whether an AI layer entering organizational operations is structurally assimilated into the regulatory functions required for organizational viability. The term alignment is used here in the organizational and regulatory sense of fit within a viable operating system; it does not refer to AI value alignment between model behavior and human values.

The research question is therefore: How can mid-market firms diagnose whether an AI-enabled technical layer fits their human activity system as a viability condition rather than as a local efficiency tool?

The paper designs a named artifact for that problem: VSAM. VSAM is a four-dimensional diagnostic: Observability, Decision Rights, Codification, and Metabolic Load. Each dimension has a home in the viable-system functions above and a failure mode when the social arrangement is left unchanged. The paper makes three contributions. First, it specifies VSAM as a four-dimensional diagnostic artifact mapped onto viable-system homes. Second, it grounds that specification in a longitudinal mid-market archive (2003–2025) re-read for human–AI fit; the archive was not collected as an AI study. Third, it states inspection questions that make the artifact usable in an operating review before any scored protocol is designed.

The argument proceeds as follows. Section 2 places the problem in four theoretical traditions that must be held together: systems method choice, socio-technical joint optimization, dynamic capabilities, and the management of artificial intelligence as distributed judgment. Section 3 states the design-science reanalysis of the archive. Section 4 specifies VSAM, maps each dimension onto the viable system, and supplies inspection questions derived from the process-optimization strand of the same archive. Section 5 locates the artifact against prior diagnostics, states boundary conditions, and identifies what the model deliberately does not claim. Section 6 concludes.

2. Theoretical Background

This section reviews the theoretical foundations of the study. It first discusses viability and organizational pathology, then examines human–AI decision structures, dynamic capabilities, and the design-science perspective that informs the development of VSAM.

2.1. Systems Traditions and the Type of Problem

Systems research is not a single method. The same organization can be read as a machine to be optimized, a culture to be interpreted, or a viable system to be kept alive, and the choice of reading determines what will count as a solution [10]. Creative holism is the refusal to apply one of those readings by habit. In this paper, the relevant distinction is between hard and soft systems. Hard systems methods assume a defined purpose and a technical fix. Human–AI operating systems do not present that kind of problem. What the model is for, who may overrule it, and what counts as a valid decision are themselves in dispute [11]. VSAM is therefore a diagnostic for alignment inside a human activity system,

not a procedure for optimizing a machine [11]. Viability theory supplies the functions that must remain intact; it does not, by itself, tell the analyst how to settle contested purpose. That is why a diagnostic artifact is the appropriate product rather than an optimizer. In alignment with this view, recent sociotechnical scholarship increasingly argues that AI outcomes cannot be understood solely as properties of algorithms but must be evaluated within the broader organizational systems in which they operate [12,13].

2.2. Socio-Technical Joint Optimization

Socio-technical research established that technical change without a matching social redesign degrades the work system. Performance depends on joint optimization of the two [14,15]. The founding studies showed that a new technical method does not dictate a single way of organizing work; organizations choose how labor, payment, and control are arranged around the tool, and that choice decides whether variance is handled at its origin or pushed upward as conflict [14]. The same principle has been restated for digital work. When technology changes continuously, joint optimization cannot be a one-time fit of a social system around a fixed machine. Social and technical arrangements have to co-evolve, and variance still has to be handled close to its origin rather than escalated by default [15]. Recent research on human-centered AI extends this socio-technical principle from the individual to the organizational level, arguing that effective AI integration requires keeping not only human users but also organizational practices and structures ‘in the loop’ [16].

In an AI-enabled mid-market firm, the technical element is the model’s training data, its update cycle, and its interface. The social element comprises the rights, observational duties, and escalation paths that enable the model to be used within the viable system. Installing the first without redesigning the second is the contemporary form of the original socio-technical error. VSAM treats that error as diagnosable rather than as a cultural problem to be trained away. This view is consistent with recent sociotechnical AI research that treats machine-learning systems as embedded in broader sociotechnical contexts and shows that vulnerabilities can arise from interactions between technical components, human actors, organizational processes, and the environments in which these systems operate [17].

2.3. Dynamic Capabilities and the Need for Standing Redesign

Fit is not a completed installation. Firms that continue to earn a return as conditions change do so because they can renew ordinary routines rather than freeze them. Teece et al. located that capacity in three elements: processes (how the firm coordinates, learns, and reconfigures), positions (the specific assets it already holds), and paths (the history that opens or closes the next move) [18]. Teece later restated the same capacity as a sequence of microfoundations: sensing, seizing, and reconfiguring as distinct organizational capacities, not slogans for agility [19]. A model that is accurate on last year’s distribution and cannot be revised in this year’s process is a failure of reconfiguration, not a failure of algorithms.

For mid-market operating systems, the implication is sharp. The firm cannot freeze a target operating model around today’s vendor. VSAM, therefore, treats misalignment as a failure to reconfigure the social system at the speed of the technical layer. Observability, decision rights, codification, and metabolic load are the surfaces on which that reconfiguration either occurs or is postponed until variety overload appears as exhaustion.

2.4. Human–AI Management as Distributed Judgment

Automation and augmentation are not alternative strategies that a firm picks once. They are temporally and spatially interdependent. Treating them as substitutes recreates the paradox in which the organization either removes people from the loop and loses the knowledge the model cannot see or keeps people in the loop without redesigning the

work and gains no relief [13]. Recent human–AI teaming research reinforces this view by conceptualizing collaboration as a dynamic socio-technical process in which human and artificial agents contribute complementary capabilities. At the same time, roles adapt through coordination, mutual communication, and shared situational awareness [20]. Managing artificial intelligence is therefore an organizational problem of judgment, data, and accountability, not only a problem of model selection [21]. Recent AI-governance research similarly frames governance as a system of organizational rules, practices, and processes that align AI use with organizational strategies, objectives, values, and stakeholder requirements [8]. This governance problem also has an accountability dimension: AI accountability requires identifiable relations of answerability, including recognized authority, the capacity for interrogation and oversight, and limits on delegated power [9]. In AI-enabled firms, judgment, feedback, and accountability have to be distributed across roles and channels rather than concentrated in a single approver whose calendar becomes the system’s only algedonic path [22]. From an organizational-design perspective, effective human–AI collaboration likewise depends on configuring complementary human and artificial capabilities rather than treating AI primarily as a substitute for human work [23].

That requirement is joint optimization under another name. Observability (who can see in time), decision rights (who may overrule), codification (which judgment is written down), and metabolic load (what the new layer costs in coordination) are the mid-market forms of distributed judgment. VSAM names those four surfaces so they can be inspected before the AI-era failures identified by Perez Rios [3]—identity collapse, silent alarms, and dashboards that replace intelligence—become the firm’s operating reality.

3. Materials and Methods

This section describes the design-science process, the archival dataset, sampling logic, and the analytical procedure used to derive the VSAM dimensions. Consistent with design-science research, the study proceeded through problem identification, definition of design objectives, artifact development, and analytical evaluation. The resulting artifact was the Viable-System Alignment Model (VSAM), developed through theoretical reanalysis of the archive and evaluated through assessment of theoretical coherence, internal consistency, and face utility.

3.1. Design-Science Orientation

Design-science research treats designed artifacts as legitimate scholarly products when they address relevant field problems through rigorous design and evaluation [24,25]. The contribution developed here is such an artifact: a named model with constructs and use principles. The present study draws on the design-science methodology of Pef-fers et al. [25] while focusing on problem identification, the definition of design objectives, artifact development, and descriptive-analytical evaluation. It does not claim to have completed all six DSRM activities in the present study; in particular, demonstration and observational evaluation in live organizational settings are reserved for subsequent field research.

The present artifact-specification study focused on four activities within the broader design-science process. Problem identification established the absence of a viability-oriented diagnostic for human–AI fit in mid-market operating systems. Design objectives translated this problem into the requirement for a parsimonious diagnostic capable of identifying conditions for socio-technical alignment without requiring an enterprise control function. Artifact development used constant comparison of theoretically sampled archive episodes to specify the four VSAM dimensions and subsequently mapped them onto viable-system functions. Evaluation at this stage was analytical rather than observa-

tional: the artifact was assessed for internal consistency and face utility, including whether the four dimensions represented distinct diagnostic conditions and whether the resulting questions could be used in an operating review. Independent field evaluation is reserved for subsequent research.

Two assessment criteria govern the claim [26]. Contribution: the problem is worth solving, and the artifact has pragmatic value for the class of organizations in view. Reasoning: the problem is addressed with internally consistent logic rather than with an accumulation of labels. The work sits on the practice-to-theory path in operations management design science. Field regularities come first; theoretical justification follows [27].

The artifact is derived from recurring mid-market failures in observation, authority, writing down judgment, and coordination load, and then mapped onto the viable system so that those failures can be named as alignment conditions rather than as isolated software defects.

3.2. *Corpus and Reanalysis*

The specification is a reanalysis of an existing mid-market archive: more than 1000 high-growth enterprises documented between 2003 and 2025 through teaching engagements, academic inquiry, and advisory work. Pattern structure in the original archive had stabilized around recurring conditions under which mid-market firms observe, decide, codify judgment, and absorb coordination load. These conditions were established before the current generation of operational AI models. VSAM asks how these pre-existing organizational conditions may matter when a rules-based or statistical layer enters the same operating system. Cases were theoretically sampled [28]. They were chosen because they extended, challenged, or filled a category in the emerging diagnostic, not because they represented a sampling frame of listed companies.

Each retained case combined at least two lines of evidence. One line was documentary: published accounts, statutory filings, or other records that could not be rewritten after the inquiry. The other was inquiry evidence: structured conversations with executives responsible for operations, technology, or general management. The advisory strand of the same archive includes structured process-optimization engagements with mid-market industrial firms. Those files record activity inventories, cross-functional handoffs, and unused operational knowledge. They supply the incident types against which the four VSAM dimensions were specified. They were not collected as an AI study. Where teaching cohorts produced first drafts, those drafts were reviewed against the documentary line before a case entered the working set. Discrepancies between the narrative and the record were treated as data, not as noise to be smoothed out.

The original program documented recurring scaling frictions across mid-market operating systems. This paper asks a later question of that material: which recurring organizational conditions become relevant to alignment when an AI-enabled technical layer enters the operating system? AI may reproduce, amplify, or redistribute pre-existing organizational failures, while also introducing AI-specific risks that the historical archive could not capture directly.

3.3. *Empirical Base of the Archive*

The phrase ‘more than 1000 high-growth enterprises, 2003–2025’ refers to the archive assembled to address an earlier question: why mid-market firms stall and which operating-system failures recur. That program combined teaching, academic inquiry, and advisory work. Pattern structure in the original program stabilized by 2012–2016; later cases deepened existing categories rather than adding new ones. This article does not recode those 1000 firms as an artificial-intelligence study. It reuses the same material for a later question:

how recurring organizational failures concerning observation, authority structures, codified judgment, and coordination load become relevant when an AI-enabled technical layer enters the operating system.

The 1000-firm record was considered sufficient for the pattern structure of the original research program to stabilize. VSAM was specified on a theoretically sampled part of that record, not on every file in sequence, and not on a second headcount extracted as a list. A case remained in the working set when it extended, challenged, or filled a property of Observability, Decision Rights, Codification, or Metabolic Load. Episodes were excluded from the specification set when they lacked the required combination of documentary–inquiry evidence, concerned organizations outside the defined mid-market operating-system class, or did not extend, challenge, or fill a property relevant to the emerging diagnostic. Specification stopped when further episodes only refined properties already named. The unit of analysis is the regulatory property, not the firm. The archive was not originally designed as a coded dataset with a predefined case ledger. Consequently, the specification process relied on theoretical sampling and specification saturation rather than numerical inclusion targets.

Three institutional strands supplied the volume of teaching and inquiry. The Founder Phase (GEA College, Faculty of Entrepreneurship, and LAB University of Applied Sciences, Finland; 2003–2018) documented early high-growth firms through executive-education cohorts. The Analytical Phase (University of Primorska, Faculty of Management; 2011–2014) applied structured strategic-management criteria to established firms. The Architect Phase (IEDC-Bled School of Management; 2011–2020) added mature mid-market and pre-combination cases. Advisory work in the same window, including industrial process-optimization engagements, supplied activity inventories, cross-functional handoffs, and unused operational knowledge. Those industrial files are the source of the anonymized incident patterns summarized in the respective model dimensions later in the text; they were not collected as an AI study.

The four VSAM dimensions were not derived from a single archive strand. Their specification emerged through comparative reanalysis across organizational advisory engagements, process-improvement and scaling projects, teaching-related organizational inquiries, and academically documented case material. The purpose of this reanalysis was not to compare sectors or source categories, but to identify recurring operating-system conditions concerning observation, authority allocation, codified judgment, and coordination burden. During theoretical sampling, episodes were retained when they contributed to the refinement of an emerging property associated with one of these conditions. Observability, Decision Rights, Codification, and Metabolic Load therefore represent recurring properties identified across multiple archive contexts rather than dimensions attributable to a particular source category.

Organizations in the archive are mid-market operating systems: past startup improvisation and without an enterprise control staff. Ownership is mixed (privately held and publicly listed). Sectors are mixed: high-growth Small and Medium-sized Enterprises (SMEs), industrial and infrastructure firms, manufacturing, retail and distribution, family enterprises, and later digital and venture cases. Geographies include Western and Eastern Europe, Asia, and North America. Teaching and inquiry were conducted primarily in Slovenia and Finland. In this study, “mid-market” is used as a functional and governance-based category rather than as a jurisdiction-specific revenue or employee-size threshold. Inclusion required organizations to have moved beyond founder-level or startup-level coordination yet still lacked a dedicated enterprise-level control architecture capable of absorbing AI governance as a specialized function. The presence of formal audit, risk, or governance arrangements in some listed firms did not, by itself, exclude a case; the

relevant criterion was whether operational coordination, exception handling, and redesign remained dependent on a relatively thin managerial system rather than on a dedicated enterprise control function. This operational definition was applied consistently when selecting episodes for the VSAM specification set.

Each retained teaching case combined two lines of evidence: a documentary line that could not be rewritten after the inquiry (published accounts, statutory filings, or equivalent public record) and a structured conversation with an owner, chief executive, or senior operating manager. Where students produced first drafts, faculty review checked the narrative against the documentary line. Discrepancy was treated as data. No new human-subjects protocol was opened for this article. Client identities are not reported.

The four VSAM dimensions were not hypothesized in advance as an AI checklist. Recurring operating failures in the archive first appeared as observation gaps, concentrated authority, unwritten method, and coordination tax. The reanalysis examined how these recurring organizational conditions could become relevant when an AI-enabled technical layer enters the operating system. Constant comparison held each new episode against the provisional rows. When episodes ceased to add a row, the artifact was specified as four dimensions. Table 1 summarizes the original archive, the specification set used in this study, and the analytical path through which the four VSAM dimensions were derived.

Table 1. Archive versus specification set for this article.

Element	What This Paper Claims
Archive size and window	More than 1000 high-growth/mid-market firms, 2003–2025.
Role of the 1000	Pattern structure stabilized for the original scaling/operating-system question.
Specification set for VSAM	Theoretically sampled cases from the archive until specification saturation of the four dimensions was achieved.
Institutional strands	GEA College and LAB University of Applied Sciences, Finland (2003–2018); University of Primorska (2011–2014); IEDC-Bled School of Management (2011–2020); advisory/industrial process-optimization files.
Organization type	Mid-market operating systems: private and listed; founder-phase through established mid-market; some decline/turnaround cases for contrast.
Sectors and regions	Multi-sector. Western and Eastern Europe, Asia, and North America.
Evidence rule	Document + inquiry. Process maps and activity lists are used only as anonymized incident types.
Analytic path and stop	Reanalysis of recurring operating-system failure patterns → four alignment surfaces. Stop when further episodes refine properties rather than add dimensions.

3.4. Constant Comparison and Saturation

Dimension discovery used the constant comparative method [28]. Each new episode was compared with the provisional set of dimensions. Properties were rewritten when a case could not be placed without stretching a category. Names for those properties were written and rewritten until they could be used without importing a prior codebook [29].

Illustrative derivation example. In one anonymized archive episode, operational reports existed, but exceptions remained visible primarily to a single responsible role. The recurring incident pattern was therefore not an absence of data, but restricted visibility of operationally relevant signals. Through constant comparison, this pattern contributed to the regulatory property of insufficient shared visibility and, subsequently, to the specification of the Observability dimension. Only after this empirical classification was the property theoretically mapped to Systems 2 and 3* of the Viable System Model: System 2 because coordination requires shared visibility across operational units, and

System 3* because independent audit requires access to signals beyond routine reporting. The corresponding VSAM inspection question asks: “Which exceptions are visible without asking a specific person?”

In this paper, saturation refers to the artifact’s specification saturation. The comparative coding and specification of properties were conducted by the lead analyst responsible for the archival reanalysis. The study, therefore, does not claim inter-rater reliability or independent coder agreement. Consistent with the design-science purpose of artifact specification, credibility was supported through transparent analytical documentation, theoretical consistency checks, and continuous comparison across heterogeneous archival materials. Independent coding and observational validation remain important directions for future research. A concise property codebook summarizing the properties used in the specification process is provided in Supplementary Table S1, which defines the regulatory property represented by each VSAM dimension, the indicators used to retain an episode as relevant to that property, and the principal boundary that distinguishes it from the other dimensions. These dimension properties are theoretically mapped onto their viable-system functions in Table 2, while Table 3 provides the corresponding archive-derived incident patterns and empirical inspection questions. The operating-system problem space of the original program had already stabilized; later episodes deepened existing categories rather than adding new ones. Reanalysis stopped adding rows to VSAM when further episodes only specified properties already named—for example, how observability fails, how an override is informal rather than absent, how tacit judgment concentrates in one role, or how a new layer taxes coordination. A four-dimensional diagnostic of information-system fit before combination has already shown that such an artifact can be specified from a qualitative mid-market corpus and evaluated by utility and internal consistency [30]. VSAM uses that form on a different problem: whether a model and the human activity system occupy the same viable system. Because the source archive largely predates contemporary AI deployment, specification saturation refers only to the four organizational properties derived through the present reanalysis; it does not imply saturation of AI-specific failure modes. Model drift and update management, automation bias, vendor dependency, technical opacity, and emerging forms of accountability may introduce additional conditions that the historical archive could not surface directly. Some intersect with the present dimensions, for example, drift management with Observability and Codification, and accountability with Decision Rights, but their empirical sufficiency within VSAM remains to be tested in contemporary AI settings.

Table 2. The Viable-System Alignment Model: dimensions, viable-system homes, and socio-technical failure modes.

Dimension	Operational Question	VSM Home [2]	Socio-Technical Failure Mode [14,15]
Observability	Can model behavior be seen in time for intervention?	System 2 (coordination) and System 3* (audit)	The model becomes a black box. Variance is discovered after damage, not at its origin.
Decision rights	Who may overrule the model, and under what rule?	System 3 (control) and System 5 (identity and policy)	Override is informal, concentrated, or absent. Identity collapses into the vendor tool [3].
Codification	Which operational judgment is an organizational asset rather than a person?	System 2 (coordination standards) and System 4 (intelligence)	Tacit knowledge stays in one role. The model trains on what can be logged and ignores what cannot [13].
Metabolic load	What coordination work does the new layer add?	System 3–System 4 homeostat [2]	The firm pays for the model with the last hours of managerial attention. Coordination capacity is consumed by the technical layer.

Table 3. Archive-derived incident patterns and inspection questions for the four VSAM dimensions.

Dimension	Inspection Question	Recurring Anonymized Archive Pattern
Observability	Which exceptions are visible without asking a specific person?	Reports exist; the exception still lives in one head. Data are stored; few are used.
Decision rights	Who may close an exception without an upward loop?	Approval sits at the top. Ownership of a complaint or a non-standard order is undefined.
Codification	Can a second person execute the same case from a written rule?	Method is personal. Forms multiply. Knowledge sits unused.
Metabolic load	What coordination tax does this new layer add to the existing path?	One order crosses many functions. Waiting sits between desks. Analysis is added before activity is removed.

3.5. Case Warrant

The engagements are theoretically sampled cases for artifact specification [31,32]. They are not a probability sample of mid-market firms, and no inferential statistic is claimed. A limited number of information-rich cases is sufficient when the aim is saturation of a diagnostic architecture [33]. Analytical generalization, not statistical generalization, is the warrant [32]: if the four dimensions name mechanisms that already appear in mid-market operating systems and that a model amplifies, those mechanisms can be inspected in further firms of the same class. Their prevalence cannot be estimated from the present design.

Evaluation in this article is therefore analytical. Internal consistency: each dimension has a distinct VSM home and a failure mode that is not a synonym of the others. Face utility was assessed by the authors by determining whether a practitioner could ask the four questions without an enterprise control office. Observational evaluation in new field sites, including inter-rater checks on the four dimensions, is future work and is not reported here.

3.6. Ethics of Archive Reuse

This article reanalyzes an existing archive of teaching, research, advisory, and organizational development. The structured conversations and student-produced analytical drafts were generated as part of routine educational, advisory, and organizational-development activities and were not collected specifically for the present study. The present reanalysis was conducted at the level of organizational patterns and operating-system conditions rather than at the level of individual participants. No personal data, individual participant characteristics, or identifiable responses were analyzed or reported. Identifying details about individuals and organizations, along with commercially sensitive information, were removed or anonymized before the material was retained in the analytical corpus. The article, therefore, reports only anonymized organizational incident patterns and makes no claims about identifiable individuals or organizations.

4. Results

This section presents the resulting artifact and explains the four dimensions of VSAM.

4.1. The Artifact

VSAM is a four-dimensional diagnostic of human–AI fit in mid-market operating systems. It does not classify models by accuracy, vendor, or architecture. The present specification is primarily grounded in AI-enabled decision support and related predictive, classification, recommendation, and statistical systems embedded within organizational operations. Applicability to generative and agentic AI is theoretically plausible but not

established by the present study and requires comparative field testing, as these systems may differ in autonomy, observability, update cycles, and coordination demands. It examines whether the human activity system and the technical layer occupy the same viable system [2]. If they do not, VSAM characterizes the technical layer as an unassimilated regulator rather than as an integrated component of the viable system [1].

Each dimension answers one operational question and occupies one home in the viable-system functions [2]:

- Observability: can anyone see what the model is doing in time to intervene?
- Decision rights: who may overrule it, and is that right usable under load?
- Codification: which judgment has been written down so that coordination and intelligence can use it?
- Metabolic load: what coordination tax does the new layer impose on an already thin management system?

Within the VSAM framework, the four questions are treated as jointly relevant to the diagnosis of human–AI fit. Improving the model while leaving any one of them unaddressed may relocate variety rather than absorb it [1]. That is the socio-technical claim restated in viability language [14,15]: the technical element and the social arrangement have to be diagnosed together.

Table 2 specifies the artifact. Table 3 links each dimension to an inspection question and to a recurring anonymized incident pattern observed in the process-optimization strand of the archive. The inspection questions are practitioner diagnostic prompts, not validated measurement items or a psychometric scale; their potential development into a scored instrument is reserved for subsequent validation research. These patterns provide illustrative empirical grounding for the artifact rather than frequency estimates or statistically representative case evidence. The analytical pathway was sequential: archive evidence was first retained as an anonymized incident pattern, the regulatory property implicated by that pattern was then identified, and only subsequently was it assigned to a VSAM dimension and interpreted through the corresponding VSM function. Table 3, therefore, represents the empirical starting point of the artifact, while Table 2 and Figure 1 represent its subsequent theoretical specification. Accordingly, the archive patterns are observations from the source material, whereas their interpretation through requisite variety and VSM functions is an author-derived theoretical interpretation rather than a directly observed property of the archive.

The mappings proposed below are author-derived theoretical mappings rather than constructs specified by Beer. Beer's VSM provides the functional architecture; the assignment of the four VSAM dimensions to particular VSM functions is the design contribution of this study. The mappings were made according to the regulatory function implicated by each dimension: Observability concerns coordination and independent audit (Systems 2 and 3*); Decision Rights concerns operational control and policy authority (Systems 3 and 5); Codification supports coordination standards and organizational intelligence (Systems 2 and 4); and Metabolic Load concerns the balance between present operational control and future-oriented intelligence (the System 3–System 4 homeostat). Algedonic channels are not assigned exclusively to any one VSAM dimension. In Beer's VSM, they constitute a distinct escalation pathway from System 1 to System 5; within VSAM, their failure may become diagnostically relevant through different dimensions, including unusable escalation authority under Decision Rights and coordination overload under Metabolic Load. Table 3 provides illustrative incident patterns from the industrial process-optimization strand of the archive; these examples should not be interpreted as the sole evidentiary basis for the four VSAM dimensions, whose specification drew comparatively on the broader archive described in Section 3.3.




Figure 1. Positioning of VSAM dimensions within the Viable System Model. **Note:** System 3* denotes the independent audit channel in the Viable System Model. The asterisk is part of the VSM notation and does not indicate a variable or statistical significance.

Figure 1 places the four dimensions on a single viable-system diagram: Observability on the S2/S3* channels, Decision Rights on the S3/S5 axis, Codification on the S2/S4 standards-and-intelligence loop, and Metabolic Load on the S3–S4 homeostat. Recursion is marked: the same four questions apply at the firm and at the unit that first hosts the model. In practice, the diagnostic process therefore begins at the operational unit where the technical layer operates and is then repeated at the next viable level. A dimension may be locally satisfied but misaligned at a higher level; for example, a unit may possess effective override authority, while firm-level policy constrains its exercise. VSAM is therefore applied iteratively across recursive levels, with contradictions between levels treated as evidence of cross-level misalignment.

4.2. Observability

Observability is the capacity of Systems 2 and 3* to render model behavior inspectable in time for intervention [2]. Coordination needs a shared picture of what the model is doing across operational units. Audit needs an independent view of that picture. Neither function

is satisfied by a vendor dashboard that reports only what the model can already see. The mapping to System 2 does not treat coordination as a general monitoring function. Rather, it reflects the shared visibility required for anti-oscillatory coordination across operational units; independent verification and audit remain the function of System 3*.

In the reanalyzed archive, observability failed in three recurring ways before AI-specific interpretation: latency, where relevant signals arrived after an operational commitment had been made; scope, where available reporting did not capture the conditions required to interpret an exception; and assignment, where no role held an explicit duty to review the relevant signal. When transferred theoretically to an AI-enabled setting, these conditions imply that a model may generate scores or flags without making the conditions under which those outputs are unreliable visible to the organization. The last failure remains social rather than technical: a high-resolution log that no role is required to read does not constitute an audit channel [2,15].

Pattern. Industrial process-optimization files in the same archive recorded the same split in pre-model form: reports and analyses multiplied while only a fraction of the stored data was used. Exceptions remained visible to one role. A later statistical layer inherits that split. It does not invent it.

Observability interacts directly with the other three dimensions because timely signals affect whether decision rights, codified rules, and coordination capacity can be used effectively. An override right that cannot be triggered by a timely signal is a right on paper. Codified rules that cannot be checked against live behavior decay into folklore. Metabolic load rises because managers reconstruct the missing picture through additional meetings.

4.3. Decision Rights

Decision rights are the allocation of override and exception authority across Systems 3 and 5 so that a model cannot become an ungoverned System 1 [2]. System 3 must be able to stop or redirect an operational act that the model has recommended or initiated. System 5 must be able to say what the organization will not automate, because that refusal is part of identity, not an exception log.

The archive pattern was not an absence of managerial judgment, but an absence of usable decision rights. Exception authority was frequently concentrated in a founder or a single specialist whose availability became an operating constraint. Transferred to an AI-enabled setting, a model entering this condition would operate without a readily usable override when that person was unavailable, while exceptions would continue to return to the same individual when that person was present. Distributed judgment requires the opposite arrangement [23]: specified roles, specified thresholds, and a record of the exception that System 3* can later audit.

Pattern. The same archive's industrial process-optimization strand treated authority as a design variable rather than as a personality trait. Measures were identified in the workplace and approved at the top level. Ownership of a complaint or a non-standard order was often undefined. A recurring diagnostic question in that strand concerned which intermediate departments could be bypassed and which responsibilities could be moved to the point of work. That question is the Decision Rights row in operating language. A model that closes cases the unit cannot reopen automates the upward loop that the question was written to remove.

Two pathologies from the AI-era taxonomy attach here [3]. Identity collapse: the firm can no longer state a purpose that is independent of the tool. Silent algedonic channels: the alarm exists, but it has no authorized recipient. Decision rights are the social redesign that prevents both.

Automation and augmentation meet at this dimension [13]. Taking people out of the loop without writing the override is automation by default. Leaving people in the loop without giving them the right that can be exercised under load is augmentation in name only. VSAM does not choose between the two. It asks whether the viable system still contains a controller when the model is running.

4.4. Codification

Codification is the conversion of tacit operational judgment into assets that System 2 can coordinate and System 4 can use for environmental scanning [2]. A rule that lives in one expert cannot be coordinated across units and cannot inform intelligence about the environment. A rule that has been written down can.

The mid-market episodes showed a consistent split between what was formally recorded and what remained tacit. Transactions, tickets, timestamps, and other formal records represented only part of the work, while the reasoning behind accepted exceptions, customer-specific judgments, and relationship-based decisions often remained with one person. Transferred to an AI-enabled setting, a model trained on this organizational state would have access primarily to the recorded representation of the process rather than to the tacit operational judgment behind it. That is the automation–augmentation paradox in operating form [13].

Pattern. Process-optimization files recorded unused operational knowledge and a proliferation of forms that did not transfer judgment. Method remained personal. A second person could not execute the same case from a written rule. A model trained on that state would capture the recorded process but not the tacit operational judgment behind it.

Codification is therefore not documentation for its own sake. It is the joint-optimization move that keeps the social system able to teach the technical system, and able to correct it [14,15]. System 4 depends on it. Environmental intelligence that is fed only by model output will scan a world the model already knows.

4.5. Metabolic Load

Metabolic load is the coordination tax that the new technical layer places on the System 3–System 4 homeostat [2]. Mid-market firms do not keep spare controllers. Every new interface, exception queue, vendor meeting, and reconciliation task consumes the same scarce attention that already coordinates operations and watches the environment.

Metabolic Load is related to, but narrower than, general organizational overload or implementation burden. In VSAM, it refers specifically to the additional coordination, escalation, reconciliation, and exception-handling work required to sustain the technical layer within the operating system. It therefore concerns the regulatory capacity consumed by human–AI coordination rather than workload in general. AI implementations may reduce effort in one activity while increasing coordination requirements elsewhere; the viability question is therefore whether the cumulative coordination burden created by the additional technical layer can be absorbed by the organization.

In the reanalyzed material, the tax appeared as meeting load, dual running of old and new processes, and a rise in escalations that looked like “change management” but were, in fact, unabsorbed variety [1]. When the homeostat is overloaded, System 4 starves first. The firm continues to operate and stops looking outward. That is a viability failure, not a training gap.

Pattern. The industrial files mapped single customer-order paths across many functions, with waiting time sitting between desks. A new analysis was added before the old activity was removed. Indirect departments absorbed the coordination tax first. A model

inserted on that path can increase handoff speed and volume. Metabolic Load asks whether the firm can absorb that coordination tax before the technical layer is expanded further.

Metabolic Load foregrounds a coordination cost that is not consistently made explicit in generic AI-readiness frameworks. They ask whether the firm has data and a sponsor. They do not ask whether the firm can pay the coordination cost of the sponsor's project. VSAM treats that cost as a first-order alignment condition [15,23]. If the load cannot be absorbed, the correct design move is not a better model. It is a smaller scope, a slower insertion, or an explicit decision not to automate that process.

4.6. How the Four Dimensions Work Together

The dimensions are not a maturity staircase. A firm can have clear decision rights and still be blind. It can see the model and still be unable to overrule it. It can codify rules and still exhaust the homeostat if maintaining those rules creates excessive coordination demands. Diagnosis, therefore, starts with all four questions at the recursion where the model actually sits, usually one operational unit, and only then at the firm.

That sequence follows from requisite variety and recursion [1,2]. The local system must be viable before the embedding system can treat the model as an asset. Installing intelligence at the firm level, even when the unit cannot observe or overrule the same model, is the structural pathology already identified in AI-era organizations [3].

The four dimensions are analytically distinct because they diagnose different regulatory failures. Observability concerns access to timely information about model behavior; Decision Rights concerns authority to act on that information; Codification concerns whether operational judgment exists in transferable organizational form; and Metabolic Load concerns whether the coordination required to sustain the human–AI arrangement can be absorbed. A firm may therefore satisfy one or more dimensions while failing another. The four dimensions should consequently be interpreted as theoretically distinct diagnostic propositions rather than as empirically validated latent constructs.

For the purposes of the present specification, the artifact is operationally usable when a practitioner can ask the four questions in Table 2, locate the answers on the viable-system diagram, and identify whether the next intervention is social, technical, or both. That is the utility criterion used in Section 3. Internal consistency requires that no two rows in Table 2 are synonymous. Observability is not Decision Rights. Codification is not Metabolic Load. Each failure mode points to a different repair. A single incident may implicate more than one dimension when it contains multiple regulatory failures. Classification is therefore based on the regulatory property being diagnosed rather than on mutually exclusive assignment of an incident to a single dimension; for example, an invisible exception may indicate an Observability failure, while the absence of authority to resolve the same exception constitutes a separate Decision Rights failure.

What the artifact does not produce is a score, a ranking of vendors, or a price. Those outputs belong to other instruments. VSAM's result is a structured description of fit. Illustrative operating-review scenario. Consider a mid-market operating unit introducing an AI-enabled recommendation system for non-standard customer orders. In an operating review, the practitioner first asks whether exceptions generated by or relevant to the model are visible in time to intervene (Observability). The review then identifies who may override a recommendation and under what conditions (Decision Rights), whether the judgment required to handle non-standard cases is represented in transferable organizational rules rather than held by one specialist (Codification), and whether reviewing exceptions, reconciling outputs, and coordinating with the vendor adds work that the existing management system can absorb (Metabolic Load). The scenario does not generate

a score or rank the four dimensions; it demonstrates how the inspection questions structure an operating review.

5. Discussion

This section interprets the findings and positions VSAM within the broader literature on organizational viability and human–AI systems. It discusses the theoretical and practical implications of the proposed artifact, highlights its limitations, and identifies opportunities for future validation and refinement.

5.1. What the Artifact Adds

VSAM does not replace the established VSM diagnostic tradition, socio-technical design, or the literature on managing artificial intelligence. Rather, it extends the diagnostic use of the VSM to the specific problem of human–AI fit in mid-market operating systems. Beer established the VSM as a diagnostic architecture for organizational viability [2,34], and subsequent VSM scholarship has further developed its use for organizational diagnosis and redesign [35,36]. VSAM builds on this tradition by specifying four inspection surfaces for the organizational assimilation of an AI-enabled technical layer.

That is the theoretical contribution. Human–AI fit is stated as a viability condition with four observable surfaces. That move changes what counts as evidence. Accuracy of the model is not evidence of fit. The presence of a sponsor is not evidence of fit. Within VSAM, human–AI fit is conceptualized through four regulatory conditions: Systems 2 and 3* can render relevant model behavior visible in time; Systems 3 and 5 retain a usable override; judgment required for coordination and intelligence is codified; and the System 3–System 4 homeostat can absorb the coordination load.

The same claim is a socio-technical claim. Joint optimization is the finding that a new technical method does not dictate a single social arrangement, and that variance handled far from its origin becomes a source of conflict [14,15]. Socio-technical design later made participation and quality of working life explicit requirements of that finding [37]. VSAM retains this principle while operationalizing it as a diagnostic rather than a participatory design process. Within VSAM, the four dimensions define the core social arrangements to be inspected when assessing whether a model can be integrated into a mid-market operating system without reproducing the original socio-technical error.

5.2. Relation to Adjacent Artifacts and Studies

A recent systematic review of human–AI collaboration in information systems research identifies limited theoretical integration and the absence of unified evaluation frameworks as persistent gaps in the field [38]. VSAM addresses a narrower part of this problem by providing a theoretically grounded diagnostic of organizational fit rather than a general framework for human–AI collaboration.

Organizational AI-readiness research assesses whether a firm possesses the technological, informational, infrastructural, data, participant, process, customer, and environmental conditions required to adopt AI [7]. VSAM addresses a different question: once a model enters an operating system, it diagnoses whether the resulting human–AI arrangement preserves the regulatory functions required for organizational viability.

Table 4 clarifies that VSAM does not replace these adjacent approaches; it addresses a distinct diagnostic question concerning the structural assimilation of an AI-enabled technical layer into the regulatory functions of a viable operating system.

Table 4. Positioning VSAM relative to adjacent approaches to human–AI and organizational fit.

Approach	Primary Concern	Unit/Focus	Difference from VSAM
Task–Technology Fit [5]	Match between task requirements and technology capabilities	Task/user	VSAM examines the viability of the wider human–AI operating system
Person–Technology Fit [6]	Congruence between person and technology	Individual	VSAM focuses on organizational regulatory integration
AI Readiness [7]	Conditions enabling AI adoption	Organization	VSAM diagnoses fit once AI enters operations
AI/Algorithmic Governance [8,9]	Rules, accountability, oversight, and control	Governance system	VSAM diagnoses specific viable-system alignment conditions
Human–AI Collaboration/Teaming [20,23,38]	Complementary human and AI capabilities	Team/work system	VSAM provides a viability-oriented organizational diagnostic
ISOFAM [30]	Information-system fit before combination	IS/organization	VSAM examines a human activity system and an AI layer within one operating system

A four-dimensional diagnostic of information-system fit before combination used the same methodological form: named dimensions, theoretical sampling, and evaluation based on utility and internal consistency [30]. Information Systems–Organizational Fit Alignment Model (ISOFAM) asks whether two information systems can be combined. VSAM asks whether a human activity system and a model already inside one firm occupy the same viable system. The shared form is methodological. The problem class is not.

Knowledge-creation research clarifies why Codification is a viable surface and not a documentation project. Organizations that convert tacit judgment into explicit assets can coordinate and can teach; organizations that leave judgment in persons cannot scale that judgment across units [39]. In an AI-enabled setting, the conversion has a second audience. The model learns what is explicit. What remains tacit is progressively excluded from the work’s official representation [13]. Codification is therefore both an organizational-learning act and an alignment act.

Execution research clarifies Metabolic Load and Decision Rights. A strategy that is clear on the vertical axis still fails when horizontal coordination is unreliable [40]. Mid-market firms live on that horizontal axis. They do not have departments whose only job is to reconcile exceptions. A model that increases the exception rate without a specified recipient recreates the coordination failure Sull and colleagues measured, this time inside the technical layer.

Penrose’s managerial limit is the economic name of the same constraint [41]. Growth is bounded by how fast existing management can absorb complexity. A model sold as relief from that limit can tighten it if the model’s coordination work exceeds the capacity it frees. VSAM treats that possibility as a diagnostic question rather than as an implementation footnote.

Because the Viable System Model is recursive, VSAM may be applied at multiple organizational levels. A human–AI arrangement that appears aligned at one level may generate pathology at another. For example, a local operational unit may exhibit adequate observability and decision rights while simultaneously creating coordination overload at a higher level of recursion. Conversely, organizational policies may appear coherent while operational units lack the information or authority required for effective action. The diagnostic should therefore begin at the level where the model directly influences activity and subsequently extend across higher recursive levels. Contradictions between levels should be interpreted as potential indicators of misalignment rather than anomalies. In

this sense, VSAM is intended not as a single-point assessment but as a recursive diagnostic heuristic consistent with the logic of the Viable System Model.

5.3. Practical Implications

For a mid-market operator, the artifact is a sequence rather than a scorecard.

First, locate the recursion. Ask the four questions in the unit where the model already runs. A firm-level policy that the unit cannot see or override is misaligned.

Second, repair the social surface before enlarging the technical one. If Observability is missing, a more accurate model makes the firm faster at being wrong. If Decision Rights are missing, augmentation is a slogan. If Codification is missing, the model trains on a partial world. If the metabolic load is already high, additional coordination demands may themselves become part of the pathology.

Third, treat identity as a System 5 question [2,3]. What the firm will not automate is part of what the firm is. That sentence belongs in a policy, not in an exception comment.

Fourth, do not staff the diagnosis as an enterprise program. The four questions are designed for organizations that cannot hire a control office. They can be asked in an operating review. They cannot be outsourced to the vendor that supplied the model.

5.4. Limitations and Further Work

Five limitations follow from the method.

First, the evaluation in this article is analytical. The present evaluation addresses internal consistency and face utility. Observational evaluation in new sites, including inter-rater checks on the four dimensions, is the next study.

Second, the archive is theoretically sampled from one mid-market research-and-advisory program assembled to study operating-system failure, not model adoption. Analytical generalization to other mid-market firms is the warrant [32]. Statistical generalization is not. Firms outside the mid-market band and firms whose AI layer is a separate digital business rather than an embedded layer in the operating core fall outside the scope of the present claim.

Third, exploration–exploitation remains a background tension [42]. A firm that can see and overrule a model may still allocate attention only to exploitation. That allocation is the dynamic-capability problem already stated in Section 2 [18,19]. VSAM inspects the surfaces that enable reconfiguration; it does not measure how attention is split between exploration and exploitation.

Fourth, VSAM diagnoses whether the human activity system and the model occupy the same viable system. It does not diagnose why a named person refuses to share judgment. If the override fails because a leader will not release it, the Decision Rights row records the failure; the repair is not a finer diagnostic. The causes of individual leadership behavior fall outside the scope of the present diagnostic.

Fifth, the limitation is that the study retrospectively reinterprets an archive originally assembled for a broader operating-system research program, which introduces potential author interpretive influence and heterogeneity across sectors and geographical contexts. Much of the archive also predates widespread organizational AI deployment; its use here is therefore a theoretical transfer proposition: earlier cases are not treated as AI implementations, but as evidence of operating-system failure mechanisms that an AI-enabled layer may reproduce or amplify. The present findings should consequently not be assumed to transfer unchanged to generative or agentic AI systems without direct field evaluation.

The archival material predates many contemporary AI applications. The study, therefore, does not claim that historical cases involved AI systems in their present form. Rather,

the archive was used to identify recurring operating conditions relating to observation, authority, codification, and coordination. The argument is that contemporary AI systems amplify, expose, or redistribute these conditions rather than replace them. Accordingly, the transfer from historical organizational pathologies to present human–AI contexts should be understood as a theoretical reinterpretation rather than direct empirical observation of contemporary AI implementations.

Further work should take three paths. Field evaluation: apply the diagnostic dimensions and inspection questions in Tables 2 and 3 to a small number of mid-market sites that are introducing decision-support models, and report whether independent observers consistently classify the same unit across the four dimensions. Instrument work: if those studies hold, a scored protocol can be designed. It is not designed here. Comparative work: test whether the four dimensions remain sufficient when the technical layer is generative rather than classificatory. Requisite variety does not change with the model's architecture [1]. The properties of Observability and Metabolic Load might.

6. Conclusions

The research question asked how a mid-market firm can diagnose whether an AI-enabled technical layer fits its human activity system as a viability condition rather than as a local efficiency tool. The answer offered here is a four-dimensional diagnostic. Mid-market firms that share operational judgment with a model face a viability problem, not a software problem. Requisite variety still has to be matched [1]. The recursive functions of the viable system still have to remain intact [2]. When they do not, the failures already have names in the AI-era pathology taxonomy [3]. What those accounts did not supply is a diagnostic of a fit small enough to use without an enterprise control staff.

VSAM is that diagnostic. Within the analytical scope of this study, VSAM conceptualizes human–AI fit through four jointly considered dimensions: Observability, Decision Rights, Codification, and Metabolic Load. These dimensions represent social arrangements that must co-evolve with the technical layer [14,15], with each dimension mapped to a distinct function within the viable system. The present study establishes their theoretical coherence and face utility rather than their statistical sufficiency or predictive validity.

The artifact was specified on the practice-to-theory path of design science [24,27] by reanalysis of a theoretically sampled mid-market archive documented between 2003 and 2025. The specification of the four rows stopped when further episodes refined properties rather than adding dimensions. Evaluation in this paper is analytical, focusing on internal consistency and face utility. Observational evaluation of VSAM in live organizational settings is the next study.

Two implications follow. For systems practice, AI adoption in the mid-market should be inspected as an alignment of regulators, not as a procurement event. In systems methodology, a small, named diagnostic is a legitimate design-science product when the problem class is real, and the reasoning is explicit [26]. VSAM helps the operator identify alignment surfaces that warrant attention before the technical layer is extended or further models are introduced.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/systems14101333/s1>, Table S1: Concise Property Codebook for VSAM Dimensions.

Author Contributions: Conceptualization, A.N.; methodology, A.N.; formal analysis, A.N.; writing—original draft preparation, A.N.; writing—review and editing, A.B., T.B. and M.M.; supervision, T.B. and M.M. All authors have read and agreed to the published version of the manuscript.

Funding: The University of Maribor, Faculty of Organizational Sciences, and the authors acknowledge the financial support from the Slovenian Research and Innovation Agency (research core funding P5-0018).

Institutional Review Board Statement: This article reanalyzes an existing teaching, research, and advisory archive. No new human-subject recruitment was undertaken.

Informed Consent Statement: Informed consent for the present secondary analysis was not obtained because the archival materials were originally generated through routine educational, advisory, and organizational development activities rather than specifically collected for this study. The present reanalysis was restricted to anonymized organizational patterns and operating-system conditions; no personal data, identifiable responses, or individual participant characteristics were analyzed or reported.

Data Availability Statement: The underlying engagement files are confidential client and teaching records and are not publicly archived. The diagnostic specification and the inspection questions used in the article are contained in Tables 2 and 3.

Acknowledgments: During the preparation of this manuscript, the authors used a large language model assistant for outline control, first-appearance checking of numbered citations, and language editing. The authors have reviewed and edited every sentence and take full responsibility for the scientific content of this publication. No model output was used as empirical data, and no figure or table was generated as a primary result.

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

References

1. Ashby, W.R. Requisite variety and its implications for the control of complex systems. *Cybernetica* **1958**, *1*, 83–99.
2. Beer, S. *The Heart of Enterprise*; Wiley: Chichester, UK, 1979.
3. Perez Rios, J. The viable system model and the taxonomy of organizational pathologies in the age of artificial intelligence (AI). *Systems* **2025**, *13*, 749. [\[CrossRef\]](#)
4. Schwaewe, J.; Peters, A.; Kanbach, D.K.; Kraus, S.; Jones, P. The new normal: The status quo of AI adoption in SMEs. *J. Small Bus. Manag.* **2025**, *63*, 1297–1331. [\[CrossRef\]](#)
5. Goodhue, D.L.; Thompson, R.L. Task-technology fit and individual performance. *MIS Q.* **1995**, *19*, 213–236. [\[CrossRef\]](#)
6. Tomer, G.; Mishra, S.K. Understanding person-technology fit and analyzing its influence on work-related outcomes among IT professionals. In *Proceedings of the 18th Pacific Asia Conference on Information Systems (PACIS 2014), Chengdu, China, 24–28 June 2014*; Association for Information Systems (AIS): Atlanta, GA, USA, 2014; Paper 113.
7. Tehrani, A.N.; Ray, S.; Roy, S.K.; Gruner, R.L.; Appio, F.P. Decoding AI readiness: An in-depth analysis of key dimensions in multinational corporations. *Technovation* **2024**, *131*, 102948. [\[CrossRef\]](#)
8. Birkstedt, T.; Minkinen, M.; Tandon, A.; Mäntymäki, M. AI governance: Themes, knowledge gaps and future agendas. *Internet Res.* **2023**, *33*, 133–167. [\[CrossRef\]](#)
9. Novelli, C.; Taddeo, M.; Floridi, L. Accountability in artificial intelligence: What it is and how it works. *AI Soc.* **2024**, *39*, 1871–1882. [\[CrossRef\]](#)
10. Jackson, M.C. *Systems Thinking: Creative Holism for Managers*; Wiley: Chichester, UK, 2003.
11. Checkland, P. *Systems Thinking, Systems Practice*; Wiley: Chichester, UK, 1981.
12. Berente, N.; Gu, B.; Recker, J.; Santhanam, R. Managing artificial intelligence. *MIS Q.* **2021**, *45*, 1433–1450. [\[CrossRef\]](#)
13. Raisch, S.; Krakowski, S. Artificial intelligence and management: The automation–augmentation paradox. *Acad. Manag. Rev.* **2021**, *46*, 192–210. [\[CrossRef\]](#)
14. Trist, E.L.; Bamforth, K.W. Some social and psychological consequences of the longwall method of getting coal. *Hum. Relat.* **1951**, *4*, 3–38. [\[CrossRef\]](#)
15. Pasmore, W.; Winby, S.; Mohrman, S.A.; Vanasse, R. Reflections: Sociotechnical systems design and organization change. *J. Change Manag.* **2019**, *19*, 67–85. [\[CrossRef\]](#)
16. Herrmann, T.; Pfeiffer, S. Keeping the organization in the loop: A socio-technical extension of human-centered artificial intelligence. *AI Soc.* **2023**, *38*, 1523–1542. [\[CrossRef\]](#)
17. Dobbe, R.; Wolters, A. Toward Sociotechnical AI: Mapping Vulnerabilities for Machine Learning in Context. *Minds Mach.* **2024**, *34*, 12. Correction in *Minds Mach.* **2024**, *34*, 35. [\[CrossRef\]](#)
18. Teece, D.J.; Pisano, G.; Shuen, A. Dynamic capabilities and strategic management. *Strateg. Manag. J.* **1997**, *18*, 509–533. [\[CrossRef\]](#)

19. Teece, D.J. Explicating dynamic capabilities: The nature and microfoundations of (sustainable) enterprise performance. *Strateg. Manag. J.* **2007**, *28*, 1319–1350. [[CrossRef](#)]
20. Berretta, S.; Tausch, A.; Ontrup, G.; Gilles, B.; Peifer, C.; Kluge, A. Defining human-AI teaming the human-centered way: A scoping review and network analysis. *Front. Artif. Intell.* **2023**, *6*, 1250725. [[CrossRef](#)]
21. Ivaschenko, A.; Diyazitdinova, A.R.; Nikiforova, T. Optimisation of the Rational Proportion of Intelligent Technologies Application in Service Organisations. *Organizacija* **2021**, *54*, 147–160. [[CrossRef](#)]
22. Alibašić, H. Systemically mediated leadership in AI-enabled organizations: A socio-technical systems theory of distributed judgment, feedback, and accountability. *Systems* **2026**, *14*, 984. [[CrossRef](#)]
23. Kolbjørnsrud, V. Designing the Intelligent Organization: Six Principles for Human-AI Collaboration. *Calif. Manag. Rev.* **2024**, *66*, 44–64. [[CrossRef](#)]
24. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design science in information systems research. *MIS Q.* **2004**, *28*, 75–105. [[CrossRef](#)]
25. Peffers, K.; Tuunanen, T.; Rothenberger, M.A.; Chatterjee, S. A design science research methodology for information systems research. *J. Manag. Inf. Syst.* **2007**, *24*, 45–77. [[CrossRef](#)]
26. Dimov, D.; Maula, M.; Romme, A.G.L. Crafting and assessing design science research for entrepreneurship. *Entrep. Theory Pract.* **2023**, *47*, 1543–1567. [[CrossRef](#)]
27. Gauss, L.; Lacerda, D.P.; Siluk, J.C.M.; Romme, A.G.L. Design science in operations management: A review and synthesis of the literature. *Int. J. Manag. Rev.* **2025**, *27*, 221–237. [[CrossRef](#)]
28. Glaser, B.G.; Strauss, A.L. *The Discovery of Grounded Theory*; Aldine: Chicago, IL, USA, 1967.
29. Charmaz, K. *Constructing Grounded Theory*, 2nd ed.; Sage: Thousand Oaks, CA, USA, 2014.
30. Naraločnik, A.; Bertoneclj, A.; Bertonecl, T. Information systems in pre-combination M&A: Developing an ISOFAM. *Systems* **2025**, *13*, 565. [[CrossRef](#)]
31. Eisenhardt, K.M. Building theories from case study research. *Acad. Manag. Rev.* **1989**, *14*, 532–550. [[CrossRef](#)]
32. Yin, R.K. *Case Study Research and Applications: Design and Methods*, 6th ed.; Sage: Thousand Oaks, CA, USA, 2018.
33. Flyvbjerg, B. Five misunderstandings about case-study research. *Qual. Inq.* **2006**, *12*, 219–245. [[CrossRef](#)]
34. Beer, S. *Diagnosing the System for Organizations*; John Wiley & Sons: Chichester, UK, 1985.
35. Espejo, R.; Reyes, A. *Organizational Systems: Managing Complexity with the Viable System Model*; Springer: Berlin/Heidelberg, Germany, 2011. [[CrossRef](#)]
36. Schwaninger, M. Design for viable organizations: The diagnostic power of the viable system model. *Kybernetes* **2006**, *35*, 955–966. [[CrossRef](#)]
37. Mumford, E. The story of socio-technical design: Reflections on its successes, failures and potential. *Inf. Syst. J.* **2006**, *16*, 317–342. [[CrossRef](#)]
38. Seini, A.B.; Adam, I.O.; Preko, M. Human-AI collaboration in information systems research: A systematic literature review and future research directions. *Hum.-Intell. Syst. Integr.* **2026**. [[CrossRef](#)]
39. Nonaka, I.; Takeuchi, H. *The Knowledge-Creating Company*; Oxford University Press: New York, NY, USA, 1995.
40. Sull, D.; Homkes, R.; Sull, C. Why strategy execution unravels-and what to do about it. *Harv. Bus. Rev.* **2015**, *93*, 57–66.
41. Penrose, E.T. *The Theory of the Growth of the Firm*; Blackwell: Oxford, UK, 1959.
42. March, J.G. Exploration and exploitation in organizational learning. *Organ. Sci.* **1991**, *2*, 71–87. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.