Supply Chain Analytics 15 (2026) 100229

Contents lists available at ScienceDirect

Supply Chain Analytics

journal homepage: www.sciencedirect.com/journal/supply-chain-analytics

A topological analytics framework for preserving decision structure in supply chain inventory classification

Frank Michael Theunissen *, Shafiq Alam, Aymen Sajjad

School of Management and Marketing, Massey University, QB2.04 Quadrangle B Building, Albany Expressway SH 17, Albany, Auckland 0632, New Zealand

A R T I C L E I N F O

Keywords:
Inventory analytics
Supply chain topology
Criteria selection
Decision structure
Projection loss
Multicriteria modelling

A B S T R A C T

Ineffective inventory classification compromises supply chain resilience when criteria selection methods conflate statistical prominence with decision utility. Standard techniques, specifically Principal Component Analysis (PCA), Recursive Feature Elimination (RFE), and Mutual Information (MI), risk inducing projection loss by excising structurally critical variables. We apply the Topological-Structural Axiomatic Validation (T-SAV) framework, which models criteria sets as simplicial complexes (structures built from points and their higher-order connections) and maps the decision axioms of completeness and non-redundancy to computable topological invariants. Using empirical data from a telecommunications infrastructure provider, we benchmark T-SAV against PCA, RFE, and MI. Results indicate distinct failure modes: PCA fragments the decision manifold into 14 disconnected components (consensus bias), while RFE discards control variables lacking predictive correlation (target dependency). T-SAV instead retains orthogonal keystone criteria that bridge financial and physical operational dimensions. T-SAV achieves a 99.63% variation capture ratio, against PCA (99.24%), MI (99.28%), and RFE (97.39%). Bootstrap resampling ($B = 10{,}000$) indicates these differences are statistically significant ($p < 0.001$), though the margin over PCA and MI is modest and the principal distinction is structural, since T-SAV alone satisfies both topological axioms. T-SAV also shows the lowest resampling variance ($SD = 0.26$), against markedly greater instability in RFE ($SD = 1.84$). In this empirical setting, topological validation can reveal structural weaknesses that statistical covariance and predictive error minimisation may miss. Practically, the framework lets managers check, before deployment, whether a criteria set holds together, lowering the risk of discarding criteria that are statistically quiet but operationally critical.

1. Introduction

Inventory classification underpins working capital optimisation and supply chain resilience. In volatile supply chain environments, ineffective inventory classification directly affects operational and financial performance [1,2]. Concurrently, stockouts attributed to classification errors generate substantial revenue loss. Recent studies [3] estimate that inventory unavailability affects 4% of global retail revenue, equating to annual losses approaching $984 billion. This divergence renders single-criterion methods insufficient and underscores the need for multicriteria approaches that capture the structural complexity of modern supply chain decisions.

Consequently, the validity of Multicriteria Inventory Classification (MCIC) models relies on the structural integrity of criteria sets applied to classify inventory [4]. The selected criteria must capture critical structural dimensions of the decision context to ensure effective classification [5]. However, criteria selection is undermined by the absence of formal validation methods [4]. Furthermore, MCIC has not kept pace with advances in criteria selection within adjacent fields. Unlike domains that employ objective selection methods, MCIC remains reliant on subjective expert judgement or the unvalidated reuse of literature-based criteria [6]. This methodological gap creates an opportunity for principled, data-driven intervention in MCIC criteria selection.

Addressing this methodological deficit requires a shift from subjective heuristics to objective, data-driven selection. Although practitioners favour methods such as Principal Component Analysis (PCA), Recursive Feature Elimination (RFE), and Mutual Information (MI) for their interpretability and efficiency, we argue that these coordinate-dependent approximations introduce distinct theoretical risks when applied to complex inventory manifolds. Rather than eliminating bias, such techniques risk replacing subjective error with systematic algorithmic distortion, specifically the risk of projection loss, a structural

* Corresponding author.
E-mail address: frank.theunissen.1@uni.massey.ac.nz (F.M. Theunissen).

https://doi.org/10.1016/j.sca.2026.100229
Received 26 January 2026; Received in revised form 23 July 2026; Accepted 31 July 2026
Available online 2 August 2026

deficit where statistical projections fail to preserve the topological fidelity of high-dimensional decision manifolds [7,8]. Specifically, PCA decomposes the Covariance Matrix (X^TX), introducing a consensus bias that conflates shared variance with importance. This systematically excises keystone criteria: orthogonal control variables that are structurally critical yet statistically independent [9,10]. Similarly, RFE operates as a predictive wrapper; even when using non-linear estimators, it exhibits target dependency by prioritising features solely on their ability to minimise prediction error, thereby discarding orthogonal control variables that define decision boundaries without explicitly predicting the target variable [11]. Finally, MI calculates probabilities strictly from the Joint Distribution Matrix of variable pairs, imposing a bivariate limitation that fails to detect synergistic dependencies (e.g., XOR interactions) where variables appear individually stochastic but are collectively predictive [21].

Fundamentally, these deficits arise because standard selection methods are coordinate-dependent: they assume that data structure is defined strictly by statistical properties of the feature space, such as covariance, impurity reduction, or pairwise probability. This metric reductionism compromises the structural validity of criteria sets by forcing complex geometric relationships into rigid coordinate systems. Throughout this study we distinguish the machine learning concept of a 'feature' (a statistical variable) from the decision-theoretic concept of a 'criterion' (an evaluative standard). A selected feature qualifies as a criterion only when it carries evaluative meaning within the decision context; algorithmic feature selection therefore does not by itself constitute criteria validation. Our framework makes this condition explicit by admitting a feature as a structurally valid criterion only when it satisfies the topological operationalisation of the completeness and non-redundancy axioms, so that statistical retention and structural validity are assessed separately rather than assumed equivalent. To overcome the limitations of coordinate-dependent projection, we propose a shift towards coordinate invariance, drawing on Topological Data Analysis (TDA) to select and validate criteria sets based on intrinsic connectivity rather than extrinsic axes.

This research applies the Topological-Structural Axiomatic Validation (T-SAV) framework introduced by Theunissen et al. [12] in a comparative case-based criteria selection experiment. T-SAV shifts the basis of selection from magnitude analysis to topological shape analysis. By modelling the criteria set as a high-dimensional simplicial complex, the framework operationalises foundational MCDM axioms of criteria coherence [13,42] as computable topological invariants:

- *Completeness* ($\beta_0 = 1$): Defined as topological connectivity. This invariant ensures the criteria set forms a single, coherent component, preventing the fragmentation of the decision space.
- *Non-Redundancy* ($\beta_1 = 0$): Defined as homological acyclicity. This invariant ensures the absence of cycles or redundant loops ($A \to B \to C \to A$).

We compare the performance of this topological approach against the selected dimensionality and feature selection methods, evaluating whether T-SAV mitigates the distortions inherent in coordinate-dependent models, specifically the pruning of structurally important keystone criteria. Using inventory data from a telecommunications infrastructure logistics provider, this study addresses three research questions:

- *RQ1* (Axiomatic validation): Do standard dimensionality and feature selection methods (PCA, RFE, MI) produce criteria sets that structurally violate the accepted MCDM axioms of completeness and non-redundancy?
- *RQ2* (Operational efficacy): Through a multicriteria inventory classification, does a topologically validated criteria set achieve a higher variation capture ratio (CR_{Var}) compared to sets derived from PCA, RFE, and MI?
- *RQ3* (Case-based evaluation): Does empirical performance provide evidence consistent with the latent risk of projection loss, evaluating whether topological invariants are more reliable indicators of structural necessity than coordinate-dependent methods?

This study makes three primary contributions to Multi-Criteria Decision-Making (MCDM) and supply chain analytics. First, it applies the state-of-the-art topologically based criteria set audit methodology developed by Theunissen et al. [12] in a real-world comparative study using a high-dimensional numerical inventory dataset. Second, it identifies and empirically demonstrates the risk of projection loss in established dimensionality and feature selection methods, providing evidence that statistical and predictive consensus alone is an insufficient proxy for structural necessity in inventory contexts. Third, it evaluates the T-SAV framework as an alternative approach for maintaining the structural integrity of criteria sets in high-dimensional decision spaces.

The paper is organised as follows. Section 2 reviews the evolution of MCIC and the limitations of current feature selection techniques. Section 3 details the mathematical foundations of T-SAV and the topological invariants used for validation. Section 4 describes the experimental methodology and the inventory dataset sourced from a telecommunications infrastructure provider. Section 5 presents the empirical results and a comparative analysis of T-SAV against PCA, RFE, and MI. Section 6 discusses the theoretical and practical implications of the findings and offers directions for future research, and Section 7 concludes the paper.

2. Literature review

This review evaluates the methodological foundations of criteria selection for MCIC. We identify a gap in the literature regarding the rigorous validation of criteria sets. While established dimensionality reduction and feature selection techniques such as PCA, RFE, and MI offer well-documented advantages in managing high-dimensional data, we argue that these techniques introduce theoretical risks when applied to structurally complex inventory decisions. Specifically, we examine whether applying coordinate-dependent approximations to high-dimensional inventory manifolds risks projection loss and the excision of structurally critical constraints.

The following sections analyse these potential failures through the lens of metric conflation, where statistical prominence may erroneously replace decision utility. To address this deficit, we examine axiomatic decision theory as a framework for ensuring criteria set completeness and non-redundancy. Finally, we consider TDA as a methodological alternative capable of preserving the intrinsic geometric fidelity of high-dimensional decision spaces.

2.1. Criteria selection for multicriteria inventory classification

Inventory classification is a fundamental requirement for modern supply chains due to the high volume of stock-keeping units (SKUs) managed by organizations [2,14]. Companies handle thousands to millions of individual SKUs, making uniform management inefficient. Traditional inventory classification, based on annual dollar usage (ADU), assumes that a small percentage of items accounts for most of the value [15,16]. However, this mono-criterion approach is insufficient because it ignores critical criteria such as lead time, criticality, obsolescence, and substitutability, leading to significant financial loss and operational inefficiency [17].

Consequently, classification model success depends heavily on the specific criteria selected [4,14]. If these criteria fail to capture the multidimensional nature of inventory (e.g., a low-cost item that halts production if missing), the system fails regardless of the classification method's sophistication [4,15]. The literature emphasizes that classification must align with specific management objectives rather than serving as a purely statistical exercise [18], yet gaps in formal criteria selection remain an impediment to such alignment. Moreover, criteria selection for MCIC is an underexplored topic, characterised by the lack of formal selection and validation methods [4].

2.2. PCA, RFE and MI in supply chain applications

In the effort to mitigate subjective expert criteria selection with data-driven objectivity, PCA, RFE, and MI are considered foundational methods in supply chain analytics [19,20]. It is essential to acknowledge the established strengths of these methods before detailing the theoretical risks they introduce. These techniques are highly effective when applied to their primary mathematical objectives. PCA excels at compressing high-dimensional spaces and filtering systemic noise by isolating maximum shared variance, making it invaluable for aggregate trend analysis. Wrapper methods such as RFE are exceptionally powerful for building parsimonious models that strictly maximize predictive accuracy [33], a critical requirement for tasks such as demand forecasting. Similarly, MI provides a robust, non-parametric mechanism for capturing complex, non-linear dependencies between variables without relying on rigid distributional assumptions [21].

However, a theoretical misalignment occurs when these algorithms are repurposed for criteria selection. Because these methods are designed explicitly for signal compression, predictive error minimization, or pairwise probability, their application to structurally complex decision problems introduces distinct theoretical risks rooted in metric conflation. By assuming that decision utility is isomorphic to statistical properties, these methods may fail to preserve criteria representing structural constraints within the supply chain. The following subsections analyse three primary modes of this potential failure: consensus bias in variance-based models (PCA), target dependency in predictive wrappers (RFE), and bivariate limitation in probabilistic filters (MI). In each case, we examine how metric reductionism may lead to projection loss, defined as the excision of structurally essential keystone criteria required for manifold connectivity and resilience.

2.2.1. Principal component analysis

PCA is widely valued for its ability to reduce dimensionality while preserving the maximum possible variance, and it remains a cornerstone of exploratory data analysis. However, studies applying PCA to criteria selection exhibit a recurring pattern of linear conflation, characterised by the assumption that decision utility is isomorphic to statistical prominence. This assumption manifests most acutely in the application of PCA to prioritise high-variance features. This consensus bias is frequently exemplified in applications such as pharmaceutical supplier selection [20,22,25]. By explicitly selecting criteria based on explained variance, these models equate the preservation of statistical variance with the preservation of originality. Consequently, this approach systematically excises low-variance structural necessities.

The risks of such metric reductionism are evident when PCA is utilised to eliminate features deemed to have a constant influence on alternatives [23]. While mathematically efficient, this excision ignores coordinate invariance. By filtering constraints based on linear correlation, the model masks synergistic dependencies where ostensibly redundant or constant factors may serve distinct topological roles in maintaining decision space connectivity ($\beta_0 = 1$). Theoretical critiques support this view, demonstrating that low eigenvalues often indicate parameter vectors constrained to lower-dimensional affine sets rather than irrelevance [10]. Furthermore, methodical rigor remains inconsistent; literature reviews indicate that up to 56% of studies integrating PCA and MCDM fail to verify foundational assumptions [24]. As a result, PCA-based weighting mechanically favours outliers over managerial reality [25].

The systematic removal of low-variance dimensions for manageability leads to projection loss, meaning the destruction of the decision manifold's shape and boundary conditions. This failure is evident in applications where key criteria are discarded solely due to a lack of variance [26]. Under the T-SAV framework, this represents the excision of a keystone criterion, a structural element essential for system integrity. Removing such criteria fragments the decision space and violates the axiom of completeness ($\beta_0 > 1$). The scale of this projection loss is demonstrated by models that use PCA to discard large proportions of performance indicators to resolve information overload [27]. While statistically efficient, this reduction likely eliminates orthogonal control variables required to detect non-routine events, thereby propagating historical biases.

Furthermore, the literature frequently justifies reductionism through redundancy elimination without accounting for bivariate limitations. Concise hierarchies are frequently prioritised by explicitly removing overlapping objectives [28]. However, this assumes high correlation it implies topological equivalence. In supply chains, highly correlated criteria may form critical fail-safe loops; removing one risks violating the axiom of non-redundancy ($\beta_1 > 0$) and severs essential information loops. Although statistical reduction inherently risks information loss [28], proposed qualitative solutions often lack computational rigour. Supply chain management suffers from an over-representation of traditional unsupervised learning techniques, such as PCA, which are frequently relegated to exploratory roles due to a lack of ground truth [29]. This highlights the gap T-SAV addresses: extending unsupervised reduction from a descriptive tool towards a validated, axiomatic decision framework that seeks to ensure the shape of the decision manifold remains connected.

2.2.2. Recursive feature elimination

Wrapper methods such as RFE offer a principled approach to feature selection by directly optimising predictive performance, and they have proven effective in reducing dimensionality while maintaining model accuracy. However, recent scholarship utilising these methods for criteria selection demonstrates a distinct shift from consensus bias (variance-based) to target dependency (prediction-based). For instance, subjectivity in sustainable supplier selection is often addressed by employing models such as Random Forest combined with Recursive Feature Elimination Cross-Validation (RFECV), which optimise criteria sets based solely on their contribution to classification results [30].

However, maximizing predictive fit often captures chance correlation rather than true domain structure [31]. With high-dimensional descriptors, methods focusing on predictive power frequently fixate on the idiosyncrasies of individual data points, losing sight of the broad picture necessary for generalization [31]. In predictive classification contexts [30], this violates the T-SAV axiom of completeness ($\beta_0 > 1$), as the algorithm fragments the decision space by excising orthogonal variables defining the decision boundary structure without strictly minimising classification error.

Similarly, in supply chain finance, feature selection frameworks utilise RFECV to explicitly remove redundant features to reduce learning costs and predict credit risk [32]. This approach introduces a high risk of structural failure; empirical evidence demonstrates that RFE frequently fails to identify structurally causal variables when correlation is present, prioritising correlated proxies that maximise immediate impurity reduction instead [11]. By applying this logic to credit risk evaluation [32], models likely discard necessary structural buffers that appear statistically redundant but serve as causal fail-safes. Consequently, this creates a non-redundancy violation ($\beta_1 > 0$), ensuring the model lacks the synergistic constraints required for robust risk assessment.

In investigations of Industry 4.0 adoption barriers, RFE is utilised to ensure models ignore predictors with low predictive weights [33]. This exemplifies an operational bias where MCDM considerations are subordinated to predictive accuracy, defining importance strictly as the ability to minimise predictive error [34,35]. Reducing features solely to maximise classification metrics, such as precision and recall, ignores the structural significance of the discarded data [35]. By adopting this accuracy-as-importance paradigm [33], such applications exhibit metric conflation, risking projection loss by compromising the topological shape of the decision manifold.

Collectively, these operationalizations exemplify the flaw of target dependency within modern supply chain analytics. While wrapper methods maximise accuracy, they introduce high computational costs and a significant risk of overfitting, rendering the selected subsets unstable [36]. The limitations of this predictive narrowness are further evident in the fact that global feature importance plots from RFE fail to clarify how features influence specific predictions, necessitating post-hoc explanations [11]. This reliance on data normalisation and accuracy optimisation may alter the topological properties of the dataset, generating models that lack coordinate invariance. Consequently, standard wrappers risk prioritising short-term statistical fit over the structural validation required to detect rare but consequential disruptions.

2.2.3. Mutual information

Recent integrated MCDM frameworks address supplier selection in post-disruption environments by employing techniques such as the Fuzzy Delphi Method (FDM) and Fuzzy Analytic Hierarchy Process (FAHP) to rank qualitative features [19]. While integrated fuzzy models are designed to offer robustness against uncertainty, reliance on expert consensus exemplifies subjective heuristics. Under the T-SAV framework, validating criteria solely through consensus significance value assumes decision utility is isomorphic to subjective experience rather than the objective topological structure of the data. Furthermore, employing FAHP forces the decision manifold into a rigid, linear hierarchy. This linearisation causes projection loss by flattening high-dimensional interactions into simple weighted sums. Consequently, the method offers no guarantee of completeness, potentially fragmenting the decision space and disrupting the connectivity ($\beta_0 = 1$) required for resilient modelling. This risks excising keystone criteria, such as rigid bottlenecks that experts might overlook in favour of popular features.

In contrast to subjective methods, quantitative MI is frequently utilised to analyse complex features, such as those influencing port throughput [37]. MI captures non-linear dependencies without distributional assumptions. However, while MI effectively reduces epistemic uncertainty, its application exemplifies the bivariate limitation. By calculating probabilities strictly from joint distributions of features pairs, the selection is predicated on isolated information gain. This ignores the possibility that supply chain features interact in high-order groups. Consequently, the method may fail to detect synergistic dependencies, such as XOR interactions, where variables appear individually stochastic but become predictive when combined. Theoretical reviews support this critique; Mutual Information Maximization (MIM) criteria ignore dependence structures, as pairwise MI is zero for variables forming an XOR gate even when the 3-way interaction is high [38]. Furthermore, standard redundancy reduction techniques often discard strongly relevant features if they appear statistically redundant, effectively blinding the model to complementarity [21].

These methodological limitations manifest operationally as redundancy conflation, where algorithms excise functional fail-safe mechanisms in the pursuit of parsimony. This is exemplified by applications employing the mRMR algorithm, which explicitly subtracts mutual information between features to avoid information overlap [36,39], and by algorithms proposed to diminish redundancy by treating shared information as a computational defect [39]. In physical supply chains, this mathematical enforcement risks removing necessary feedback loops and backup constraints, potentially violating the axiom of non-redundancy ($\beta_1 > 0$). Finally, the pursuit of variance reduction leads to the excision of keystone criteria. For example, features with minimal explanatory power are frequently explicitly eliminated based on variance thresholds [40]. Under T-SAV, this constitutes a critical failure, as rigid constraints often exhibit zero variance until failure. This metric conflation is evident when higher numerical MI values are equated directly with feature importance, systematically ignoring the topological role of variables exhibiting lower noise and leaving the system vulnerable to decoupling [41].

Table 1 synthesises the methodological approaches of the studies reviewed above, identifying each study's domain, selection method, and whether structural validity of the resulting criteria set was formally assessed. The table reveals a consistent pattern: none of the reviewed studies employing PCA, RFE, or MI assessed the structural coherence of their selected criteria sets against axiomatic requirements. This gap motivates the present study's application of the T-SAV framework, which operationalises completeness and non-redundancy as computable topological invariants.

2.3. Axiomatic decision theory: principles of coherent criteria sets

To address these statistical limitations, criteria selection should be grounded in axiomatic decision theory. A valid decision hierarchy must satisfy established structural requirements [13,42]:

2.3.1. The axiom of completeness (exhaustiveness)

The criteria set must be exhaustive to encompass the full operational context. If two alternatives score identically on all criteria (i.e., $a_i = a_j$), they must be strictly indifferent to all stakeholders [13]. Furthermore, a set is complete only if it adequately indicates the degree to which the overall objective is realized [42]. Omitting essential latent factors renders the analysis invalid regardless of the applied weights, as the criteria set fails to represent the true decision implications.

2.3.2. The axiom of non-redundancy (minimalism)

The set must be minimal to prevent the double counting of preferences. Redundancy distorts weight elicitation; if multiple criteria measure the same underlying concept, that concept receives inflated importance [28]. This dynamic is classified as splitting bias, where the hierarchy structure dictates parameter weights, making validity path-dependent [28]. Consequently, criteria function not as discovered realities, but as constructed keys that assist actors in organizing and communicating their preferences [13].

2.4. Topological data analysis: from description to selection

Topological Data Analysis (TDA) is a mathematical framework rooted in algebraic topology that infers robust qualitative and quantitative information regarding the intrinsic structure of data [43]. Unlike traditional geometric approaches relying on rigid metric properties, TDA characterizes data through its fundamental shape and connectivity, identifying topological features such as connected components, loops, and voids that persist across multiple scales of resolution [43].

A primary motivation for employing TDA is its coordinate invariance; topological properties remain stable under continuous deformations, ensuring that analytical results do not derive from arbitrary coordinate systems or measurement units [44]. Furthermore, TDA methods, most notably persistent homology, exhibit significant robustness to noise and small perturbations [44].

Crucially, TDA offers a distinct analytical advantage over traditional statistical and coordinate-dependent reduction techniques such as PCA or RFE. While coordinate-dependent methods often assume data lies on a low-dimensional affine subspace and conflate linear consensus with feature importance, TDA detects non-linear structures and high-dimensional features that standard methods may overlook [44,45]. For example, datasets with identical summary statistics can possess substantially different topological shapes, a phenomenon where traditional descriptive statistics fail to differentiate between distinct underlying distributions.

Despite these advancements, a critical methodological gap persists. While TDA is established as a diagnostic tool for monitoring systems [44], its application for prescriptive criteria or feature selection remains under-explored relative to established statistical benchmarks. Specifically, rigorous empirical studies benchmarking topologically inspired selection against standard statistical methods like PCA, RFE, or MI for

Table 1
Comparative synthesis of criteria selection approaches in the reviewed literature. The 'Structural Validity' column indicates whether the study formally assessed the completeness or non-redundancy of the resulting criteria set.

Study	Domain	Method	Selection Approach	Structural Validity	Identified Limitation
Modibbo et al. [22]	Pharma supplier selection	PCA	Variance-based reduction	No	Consensus bias: equates variance with importance
Mishra et al. [23]	Manufacturing networks	PCA	Factor loading elimination	No	Removes constant-influence factors without connectivity check
Brint et al. [27]	Automotive SCM	PCA	KPI reduction (28−8)	No	Aggressive reduction eliminates orthogonal control variables
Wu et al. [30]	Sustainable supplier selection	RF + RFECV	Predictive wrapper (19−9)	No	Target dependency: retains only cost-predictive features
Hou et al. [32]	Supply chain finance	RFECV	Predictive wrapper (56−13)	No	Discards causal fail-safes correlated with proxies
Alkhodair and Alkhudhayr [33]	Industry 4.0 adoption	RFE	Low-weight predictor removal	No	Defines importance as predictive accuracy only
Eskafi et al. [37]	Port throughput	MI	Pairwise information gain	No	Bivariate limitation: misses high-order interactions
Nashed et al. [40]	Teacher-support systems	MI + variance	Variance threshold filtering	No	Excises zero-variance constraints critical until failure
Theunissen et al. [12]	Supply chain analytics	T-SAV (TDA)	Topological invariants	Yes	Framework paper; single-case empirical validation pending

criteria selection in MCDM remain scarce. Consequently, the comparative efficacy of TDA-based selection and validation against these statistical baselines warrants investigation.

3. Theoretical framework

This section outlines the mathematical foundation for the T-SAV framework, building upon current derivations of topological decision axioms for criteria selection [12]. We review the shift from coordinate-dependent statistical metrics to coordinate-invariant topological structures. To contextualise the current analysis, we restate the definitions of completeness and non-redundancy and their mapping to homological invariants. Finally, we apply this established framework to define projection loss and predict the failure modes of standard dimensionality reduction techniques.

3.1. The axiomatic basis of criteria selection

To ensure structural validity, criteria selection must be grounded in axiomatic decision theory. A decision hierarchy is valid only if the selected criteria satisfy two necessary conditions [13,42].

3.1.1. The axiom of completeness (exhaustiveness)

The criteria set C must be exhaustive to encompass the full operational context of the objective function. In the context of inventory classification, the objective function requires the simultaneous control of four operational dimensions: *Variation* (D_{vn}), *Value* (D_{vl}), *Volume* (D_{vm}), and *Frequency* (D_f). A criteria set is complete if and only if it spans the fundamental basis of the objective space Ω :

$$C \supseteq \{D_{vn}, D_{vl}, D_{vm}, D_f\}$$

Roy [13] mandates that if two alternatives a and b score identically on all criteria represented as vector equality $g(a) = g(b)$ they must be strictly indifferent to all stakeholders. Keeney and Raiffa [42] reinforce this by noting that omitting essential latent factors renders the analysis invalid regardless of applied weights as the criteria set fails to represent true decision implications.

3.1.2. The axiom of non-redundancy (independence)

The family of criteria must be minimal to prevent the double counting of preferences. If the removal of a criterion c_k from set C does not alter the resulting ranking order R, then c_k is redundant:

$$R(C) \equiv R(C \setminus \{c_k\})$$

A set is non-redundant if the mutual information I between criteria is minimised to prevent statistical correlation and preference splitting:

$$\min \sum_{i \neq j} \mathrm{I}(c_i; c_j)$$

Redundancy manifests as splitting bias, where the criteria hierarchy structure dictates parameter weights [28]. Consequently, criteria must be treated as constructed keys for organizing preferences rather than discovered realities.

These two expressions characterise non-redundancy at different levels. Ranking invariance is an outcome-level property that depends on the weighting and ranking model. The minimisation of mutual information is a pairwise statistical property of the criteria. These conditions are not equivalent, and a set can satisfy one without satisfying the other. In the empirical analysis, we operationalise non-redundancy through topological acyclicity, the condition that the first Betti number is zero, which is the condition the T-SAV algorithm enforces. We adopt acyclicity as the operative definition. We treat ranking invariance and mutual information minimisation as complementary motivations for that definition rather than as equivalent criteria. Acyclicity is a higher-order structural property of the criteria complex, so it can detect cyclic dependencies among three or more criteria that pairwise measures do not capture.

3.2. Projection loss

This study attributes the potential failure of traditional feature selection to preserve topological structure to a fundamental reliance on coordinate-dependent or lower-order metrics. Whether minimising error (RFE), maximising variance (PCA), or maximising information gain (MI), these methods operate strictly on the covariance (Σ), feature space (X), or joint distribution matrix (P_{XY}). This methodological limitation, defined here as projection loss, implies that projecting a complex decision manifold onto a lower-dimensional metric space introduces three critical deficits.

3.2.1. Consensus bias

Standard statistical reduction identifies a signal within the majority consensus (high eigenvalues). However, supply chain operations depend on exceptions and constraints (e.g., safety stock breaches, stockouts, missed deliveries) frequently uncorrelated with routine flow. By discarding low-eigenvalue components, PCA preserves common system behaviour but eliminates independent, orthogonal stabilizers essential for stability.

3.2.2. Target dependency

Predictive models create a circular dependency by retaining only variables correlated with a specific target (e.g., *Cost*). This eliminates orthogonal control variables, namely, factors that define the decision context (e.g., *Operational Noise*) but do not explicitly predict the target value.

3.2.3. Bivariate limitation

Standard information theoretic approaches often define relationships as pairwise scalars, flattening high-order volumetric interactions into dyadic links. This reduction destroys the systemic context necessary to detect synergistic dependencies (e.g., XOR gates) or homological cycles.

3.3. The topological hypothesis: simplicial representation of decision space

Bridging traditional MCDM and the T-SAV framework [12] requires examining the simplicial representation of the decision space. A global criteria set constructs a manifold where each SKU constitutes a point. In this context, the axiomatic requirements of completeness and non-redundancy define the topological shape of this manifold rather than serving as mere logical constraints. To operationalize these parameters, Theunissen *et al.* [12] map foundational MCDM axioms to Betti numbers (β_k), the rank of the k-th homology group in algebraic topology. Intuitively, these invariants describe the shape of the criteria set: the zeroth Betti number counts how many disconnected pieces the set forms, and the first counts how many independent loops it contains. A structurally valid set forms a single connected piece with no redundant loops. This mapping provides the mathematical basis for the T-SAV selection condition.

We treat these mappings as operational topological proxies rather than as proofs of the decision axioms themselves. Topological connectedness is not identical to decision-theoretic completeness, and homological acyclicity is not identical to preference non-redundancy. The proxies are meaningful under two bridging assumptions. First, the criteria distance matrix faithfully encodes the relevant relationships between criteria, so that proximity in the embedding corresponds to shared decision content. Second, the four operational dimensions of *Variation*, *Value*, *Volume*, and *Frequency* span the objective space, so that a single connected component over these criteria corresponds to coverage of the decision context. Under these assumptions, a connected and acyclic complex provides evidence of completeness and non-redundancy. It is not a formal proof of either axiom.

3.3.1. Completeness as $\beta_0 = 1$ (path-connectedness)

In MCDM, the axiom of completeness mandates that the criteria family encompass all relevant problem dimensions.

- *Topological translation*: A complete criteria set must span the conceptual space of the SKU's activity.
- *Mathematical derivation*: In algebraic topology, β_0 (the zeroth Betti number) quantifies connected components. A value of $\beta_0 > 1$ denotes criteria space fragmentation, representing isolated information islands unlinked to a unified metric.
- *The T-SAV rule*: Enforcing $\beta_0 = 1$ establishes that the selected criteria constitute a single connected component. We read this connectedness as an operational proxy for completeness rather than as proof of the axiom. Under the bridging assumptions stated above, a single component indicates continuous information flow across the global set, so that all problem dimensions remain reachable from the core decision objective.

3.3.2. Non-redundancy as $\beta_1 = 0$ (acyclicity)

The axiom of non-redundancy mandates that no criterion duplicates or acts as a linear combination of others.

- *Topological translation*: Data redundancy creates superfluous paths between data points. When two criteria provide identical information, they form a loop, converging on the same conclusion from distinct directions.
- *Mathematical derivation*: β_1 (the first Betti number) quantifies one-dimensional holes or cycles. In persistent homology filtration, a cycle ($\beta_1 > 0$) signifies structural dependency, where multiple features capture the same underlying latent variable.
- *The T-SAV rule*: Enforcing $\beta_1 = 0$ guarantees an acyclic selection. In graph theory, this ensures the selection forms a tree structure (an efficient, non-redundant architecture) rather than a cyclic graph.

3.3.3. The axiomatic selection rule

This framework defines the axiomatic selection rule employed by the T-SAV algorithm. We seek the maximally informative, non-redundant representation of the inventory system. Mathematically, this corresponds to identifying a connected, acyclic subspace. A space where $\beta_0 = 1$ and all higher Betti numbers equal zero ($\beta_k = 0$) is topologically equivalent to a single point. In this context, the algorithm collapses the complex candidate criteria domain into a core singularity of stable, necessary features.

3.4. The topological definition of a keystone criterion

This framework defines the keystone criterion as the primary unit of analysis. While standard inventory classification prioritises criteria based on their contribution to shared variance (statistical consensus), this topological framework categorises features by their structural role within the manifold.

We distinguish between two topological archetypes:

1. *Hubs*: High-consensus features (e.g., *Consumption Value*) exhibiting strong covariance with primary trends. These represent the system's dominant signal and are readily identified by variance-based methods (PCA) via high eigenvalues or by predictive wrappers (RFE) via strong linear correlation.
2. *Keystones*: Orthogonal structural constraints (e.g., *Total Negative Warehouse Variation*) functioning as independent control parameters. These features are frequently statistically obscure, appearing as noise to PCA (orthogonality), irrelevant to RFE (target dependency), or redundant to MI (bivariate limitation).

Formally, a criterion k constitutes a keystone if and only if it exhibits topological insensitivity (statistical obscurity) yet functions as a cut-vertex in the Vietoris-Rips complex. Specifically, the removal of k increases the 0-th Betti number (β_0), fracturing the decision space into disconnected components.

For example, in this study, *Total Negative Warehouse Variation* acts as a boundary constraint rather than a direct profit predictor. Consequently, RFE discards it due to low target correlation and PCA discards it due to a lack of covariance with aggregate volume. Topologically, however, this variable acts as the sole bridge connecting the company's supply chain physical state and financial value vector spaces. Its excision precipitates a structural decoupling of the decision manifold, rendering the model insensitive to latent phantom inventory risks.

3.5. Theoretical mechanisms of model failure

Based on this framework, we theoretically predict the failure modes of comparison models, establishing the basis for the empirical audit (RQ3):

- *PCA (Consensus bias)*: Predicted to violate completeness. Because standardisation equalises variance, PCA penalises statistical orthogonality, discarding independent structural criteria (keystones) that do not correlate with dominant principal components.

- *RFE (Target dependence)*: Predicted to violate completeness. It creates a predictive circular dependency that discards dimensions (e.g., *Variation*) lacking explicit *Consumption Value* prediction capability.
- *MI (Bivariate limitation)*: Predicted to violate non-redundancy. It selects clusters of high entropy synonym variables without detecting cyclic dependencies.

Only a model grounded in simplicial complexes is hypothesised to satisfy both axioms simultaneously by relying on coordinate-free invariants rather than metric-dependent matrix operations. Having established that valid criteria sets must satisfy the topological conditions of path-connectedness and acyclicity ($\beta_1 = 0$), we now apply this theoretical lens to a real-world empirical setting.

4. Methodology

This study validates the T-SAV framework through a comparative experimental design. The topological approach is contrasted against three industry-standard benchmarks: PCA, RFE, and MI. The complete computational protocol is summarised in Algorithm 1 below, which consolidates the four methodology phases: (1) Data Preprocessing, (2) The T-SAV Framework Application, (3) Baseline Model Generation, and (4) Comparative Validation.

Algorithm 1. T-SAV Topological Feature Selection and Validation Framework

Input: $X \in \mathrm{R}^{m \times n}$ (Data), ϵ (Threshold). Output: S_{final} ($T-SAV$ Criteria set), $Audit$ (Axiomatic diagnostics).

1: // **Phase 1: Embedding & Calibration**

2: $X' \leftarrow \ln(1+|X|)$; $Z \leftarrow (X' - \mu)$; $\mathrm{D} \leftarrow 1 - \mathrm{CosineSim}(Z^T)$

3: $\epsilon \leftarrow$ LSI_Optimization(D) $\triangleright$ *Find longest stable interval*

4: // **Phase 2: Centrality Ranking ($\mathbf{C_T}$)**

5: $PD_{full} \leftarrow \mathrm{PersistDiag}(D)$; for each $c_i \in Z$ do

6: $C_T(c_i) \leftarrow W_\infty\left(PD_{full}, \mathrm{PersistDiag}(D \setminus \{c_i\})\right)$

7: $C_{ranked} \leftarrow \mathrm{SortDesc}(C_T)$

8: // **Phase 3: Axiomatic Pruning (T-SAV)**

9: $S \leftarrow \{c_{hub}\}$; for each c_i in C_{ranked} do

10: $(\beta_0, \beta_1) \leftarrow \mathrm{Homology}\big(VR(S \cup \{c_i\}, \epsilon)\big)$

11: if $(\beta_0 == 1 \wedge \beta_1 == 0)$ then $S \leftarrow S \cup \{c_i\}$

12: $S_{final} \leftarrow S$

13: // **Phase 4: Comparative Baselines**

14: $S_{RFE} \leftarrow RFECV(\mathrm{Z})$; $S_{PCA} \leftarrow \mathrm{PCA}(Z, 0.95)$; $S_{MI} \leftarrow \mathrm{MutualInfo}(Z)$

15: // **Phase 5 & 6: Audit & Validation**

16: for each model $M \in \{S_{RFE}, S_{PCA}, S_{MI}\}$ do $\mathrm{Audit}[M].\ \beta \leftarrow \mathrm{Homology}\big(VR(M, \epsilon)\big)$

17: for each set $S \in \{S_{final}, S_{RFE}, S_{PCA}, S_{MI}\}$ do

18: $S' \leftarrow \mathrm{MinMax}(S)$; $W \leftarrow \mathrm{Entropy}(S')$; $Rank \leftarrow \mathrm{TOPSIS}(S, W)$

19: $\mathrm{Results}[S].\ C_R^{Var} \leftarrow \sum_{i \in ClassA} |Var_i| / \sum_{All} |Var|$

20: return $S_{final}, Audit, Results$

4.1. Phase 1: data acquisition and geometric embedding

Implementation: Algorithm 1 (Lines 1–3)

Empirical validation was conducted using data from a telecommunications infrastructure fulfilment provider specialising in fibre provisioning, network buildsand fault restoration. The organisation operates a fleet of over 100 technicians whose Service Level Agreement (SLA) performance depends on high inventory availability. Since 2024, the company has undertaken a supply chain transformation program to mitigate material flow variation and avoid significant financial penalties for unaccounted inventory. Despite process improvements, the organisation faces critical challenges in inventory variation management. Discrepancies between procured, issued, and returned materials generate hoarded stock and incorrect material utilisation. Consequently, the primary operational objective employs MCIC to identify specific SKUs driving this systemic variability.

The dataset contains transaction activity metrics for the entire SKU catalogue over a 12-month period. The raw data matrix comprises SKUs and candidate criteria (operationally defined as statistical features), spanning consumption volume, variability, procurement frequency, stock holding, cost, and transactional activity. The final dataset comprised $N = 2810$ SKUs and $n = 54$ criteria. To prepare the data for topological analysis, distributions were normalised using a log transform of the magnitude of each criterion, $\ln(1 + |x|)$, and Z-score standardisation detailed in Algorithm 1 (Phase 1). We transform on magnitude because, for these criteria, the sign records the direction of stock movement rather than a difference in activity: a recorded value of -1000 or $+1000$ reflects the same 1000 units of activity. Directionality is preserved through the dedicated positive and negative variation criteria, which enter the analysis as separate variables. This ensures variance is equalised ($\sigma^2 = 1$) before computing the cosine distance matrix D.

All computational procedures were implemented in Python 3.11. Persistent homology was computed using the Ripser library, with Vietoris-Rips complex construction and Betti number extraction. Scikit-learn was used for the baseline models (PCA via SVD decomposition, Random Forest for RFECV, and k-NN mutual information estimation). Distance matrices, entropy weighting, and TOPSIS ranking were implemented using NumPy and SciPy. The Random Forest regressor for RFECV was initialised with a fixed random seed (random_state = 42) to ensure reproducibility. The dataset used in this study contains commercially sensitive operational data and is not publicly available; however, the complete analytical pipeline and parameter configurations are described in sufficient detail to permit replication with comparable datasets.

4.2. Phase 2: The T-SAV framework (Selection Protocol)

Implementation: Algorithm 1 (Lines 4–12)

The T-SAV framework isolates a criteria subset adhering to the dual topological axioms of completeness and non-redundancy.

Topological calibration and ranking

The filtration value ϵ was established via Longest Stable Interval (LSI) analysis (Line 3). Prior to selection, criteria were ranked by their topological centrality (C_T), calculated as the bottleneck distance between persistence diagrams (Lines 5–7). This identifies keystone criteria whose removal induces the most significant structural collapse.

Axiomatic pruning

We seed the selection with a single hub criterion, defined as the criterion ranked first by topological centrality (Algorithm 1, Line 9), that is, the criterion whose removal induces the greatest change in the persistence diagram as measured by the bottleneck distance. In this dataset the hub criterion is *Warehouse Transfers In Units*, the same criterion identified in Section 5.2 as the topological keystone bridging the operational dimensions. The core mechanism employs the iterative forward-selection loop defined in Algorithm 1 (Lines 8–12). A candidate criterion is accepted if and only if it maintains a single connected component ($\beta_0 = 1$) without introducing essential loops ($\beta_1 = 0$).

4.3. Phase 3: comparative baseline models

Implementation: Algorithm 1 (Line 14)

To benchmark the topological approach against industry standards, alternative criteria sets were generated using three distinct dimensionality reduction techniques as executed in Algorithm 1 (Phase 4):

- *Recursive Feature Elimination*: To test for target dependency, a Random Forest Regressor (100 estimators) was employed with *Consumption Value* as the target proxy. The Random Forest estimator was selected for its capacity to capture non-linear feature interactions, while 100 estimators provide a standard balance between ensemble stability and computational cost. 5-fold cross-validation (RFECV) was integrated to identify the subset size that maximised the negative mean squared error, ensuring the wrapper selects features based on their aggregate predictive contribution rather than univariate statistics.
- *Principal Component Analysis*: To test for consensus bias, Singular Value Decomposition (SVD) was performed on the standardised feature matrix. The number of components k required to explain 95% of the variance was selected, recovering original criteria importance via absolute loadings. The 95% explained variance threshold is a widely adopted convention in applied PCA, balancing dimensionality reduction against information retention. Criteria were retained if they exhibited high absolute loadings on the selected components, ensuring traceability from principal components back to the original feature space.
- *Mutual Information*: To test for bivariate limitation, the information gain between candidate criteria and the target was estimated using a k-nearest neighbour (k-NN) estimator. Selection followed a mean thresholding strategy where the selected set S_{MI} comprises criteria with MI scores exceeding the global average. The mean thresholding strategy provides an objective, data-driven cut-off that avoids arbitrary parameter tuning: criteria contributing above-average information gain are retained, while those below are pruned. The k-NN estimator was chosen for its non-parametric flexibility, as it does not impose distributional assumptions on the feature space.

Two of the three benchmarks are supervised. RFE and MI both use *Consumption Value* as the target, whereas T-SAV and PCA are unsupervised. The RFE and MI selections are therefore conditional on this target choice. We treat target dependency as a property of supervised selection relative to the chosen target rather than an intrinsic defect of either method. A supervised method optimised for a value target will tend to discard criteria, such as variation measures, that do not predict that target. The comparison reported here is accordingly conditional on the *Consumption Value* target, and we return to this conditionality when interpreting the results in Sections 5.3 and 6.2.2.

4.4. Phase 4: comparative validation protocol

Implementation: Algorithm 1 (Lines 16–20)

The final phase subjects the generated criteria sets (S_{TSAV}, S_{RFE}, S_{PCA}, S_{MI}) to a dual-stage audit designed to answer the specific research questions.

Comparative topological audit (RQ1)

To evaluate structural validity, each baseline set was projected back into the geometric embedding space and their Betti numbers were

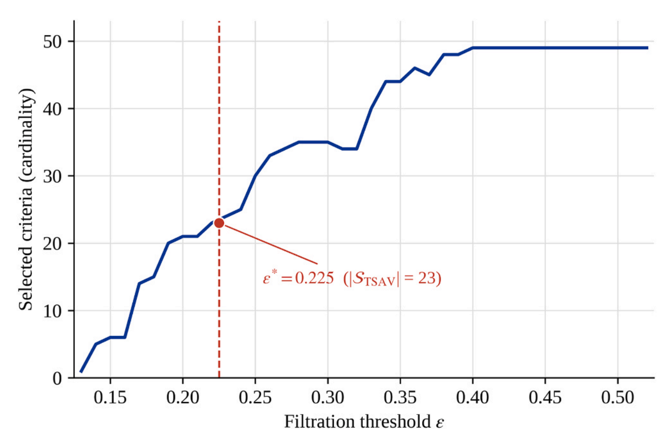

Fig. 1. T-SAV threshold sensitivity. The plot maps the cardinality of the T-SAV criteria set as a function of the filtration threshold. The set grows from a small, fragmented core at strict thresholds toward near-total inclusion at permissive thresholds. The red dashed line marks the selected threshold at the onset of the connected regime.

computed (Line 16). The audit evaluates each benchmark set against the topological axioms that T-SAV is constructed to satisfy. We therefore treat it as an internal structural consistency test rather than a neutral external benchmark, and we rely on the variation capture ratio (Section 5.6), which is computed independently of these axioms, as the operationally independent measure. Under this framing the following were measured:

- *Completeness* ($\beta_0 = 1$): Detecting fragmentation ($\beta_0 > 1$) where a method fails to capture the full scope of the inventory system.
- *Non-redundancy* ($\beta_1 = 0$): Detecting circular loops ($\beta_1 > 0$) where a method selects correlated features that introduce noise.

Operational validation via MCIC (RQ2)

To evaluate practical utility, a comparative MCIC was conducted. To eliminate weighting bias, weights were calculated using Shannon Entropy, and SKUs were ranked using the TOPSIS method (Lines 17–19). Consistent with the Pareto principle, the top 20% of SKUs were designated as Class A. Efficacy was quantified using the variation capture ratio (CR_{Var}):

$$CR_{Var} = \frac{\sum_{i \in A} v_i}{\sum_{all} v_i}$$

Operationally, the variation capture ratio (CR_{Var}) quantifies the proportion of total systemic inventory volatility concentrated within the primary inventory tier (Class A). In MCIC, isolating high-value items is insufficient if the model fails to capture highly volatile SKUs that introduce disproportionate operational risk. Consequently, CR_{Var} serves as a direct measure of a model's operational efficacy. A high ratio demonstrates that the selected criteria successfully prioritise high-risk, unstable SKUs for strict managerial control. Conversely, a lower ratio indicates structural failure, revealing that critical physical variation has leaked into lower-priority inventory tiers where it remains unmanaged.

Methodological Output

The execution of this protocol yields two distinct sets of performance metrics: the topological Betti profiles for structural compliance and the variation capture ratios for operational impact. These quantitative results form the basis of the empirical analysis presented in Section 5.

5. Results

The topological sensitivity analysis, visualised in Fig. 1, maps the cardinality of the selected criteria set as a function of the filtration parameter (ϵ).

The resulting persistence curve exhibits three distinct topological phases:

- *Fragmentation phase* ($\epsilon < \epsilon_{\min}$): At strict thresholds, the semantic graph remains disconnected. The system identifies minimal structure ($n \leq 7$), representing only the most trivial correlations.
- *Transition interval* ($\epsilon_{\min} \leq \epsilon \leq \epsilon_{\max}$): A sharp phase transition marks the coalescence of core features into a connected backbone. In this window the admissible backbone grows rapidly as the primary semantic clusters merge.
- *Saturation interval* ($\epsilon > \epsilon_{\max}$): The admissible backbone expands further between $\epsilon_{\max}$and ϵ_{sat}, approaching the full candidate set. Beyond ϵ_{sat}, the system enters a saturation regime, admitting weak correlations and stochastic noise as the count slowly drifts toward greater than 50.

To prioritise semantic fidelity, the structural threshold was fixed at ϵ_{opt}. Situated within the transition interval, this threshold enforces a rigorous standard of cohesion. We select this value at the onset of the connected regime, the smallest scale at which the criteria form a single connected, acyclic backbone, which retains the most parsimonious set that spans the decision space while minimising dependence on subjective parameter choice.

Table 2
The T-SAV topological criteria set (S_{TSAV}). The selection demonstrates multidimensional structural balance, distinctively retaining critical instability metrics (Dimension: *Variation*) such as *Total Nett Warehouse Variation* that define the boundaries of the decision manifold.

Code	Criterion	Dimension	Code	Criterion	Dimension
C_1	Std Dev Monthly Usage	Variation	C_{15}	Warehouse Transfers In Value	Value
C_3	Warehouse Transfers In Units	Volume	C_{18}	Total Warehouse Variation Count	Frequency
C_4	Van Warehouses Without Stock Count	Frequency	C_{19}	Warehouses Without Stock Count	Frequency
C_5	Avg Monthly Usage	Volume	C_{24}	Total Negative Warehouse Variation Units	Variation
C_6	Adjustment Transaction Count	Frequency	C_{25}	Total Qty On Hand	Volume
C_7	Direct To Warehouse Purchase Order Units	Volume	C_{27}	Total Stock Value	Value
C_8	Total Combined Positive Variation Units	Variation	C_{31}	Total Positive Warehouse Variation Units	Variation
C_9	Total Combined Negative Variation Units	Variation	C_{34}	Total Nett Warehouse Variation	Variation
C_{11}	Warehouse Transfers In Transaction Count	Frequency	C_{37}	Consumption Value	Value
C_{13}	Total Combined Nett Variation	Variation	C_{39}	Total Nett Van Variation	Variation
C_{14}	Purchase Order Units	Volume	C_{42}	Physical Warehouses Without Stock Count	Frequency
			C_{45}	Van Warehouses With Stock Count	Frequency

5.1. Sensitivity analysis of the filtration threshold (ϵ)

To assess parameter sensitivity, the topological stability of the criteria set was evaluated across adjacent thresholds:

- *Under-filtration* ($\epsilon < \epsilon_{opt}$): At the lower bound of the transition interval, the connectivity radius is too strict and the candidate graph stays fragmented, so the T-SAV set retains only a small core of 14 criteria that under-represents the decision space. At this scale the candidate graph violates completeness ($\beta_0 > 1$).
- *Selected threshold* ($\epsilon = \epsilon_{opt}$): The T-SAV forward-selection returns a connected, acyclic backbone of 23 criteria spanning all four operational dimensions. This threshold sits at the onset of the connected regime, yielding the most parsimonious backbone that still spans the decision space.
- *Over-filtration* ($\epsilon > \epsilon_{opt}$): At the upper saturation boundary, the connectivity radius admits weakly correlated criteria and the backbone

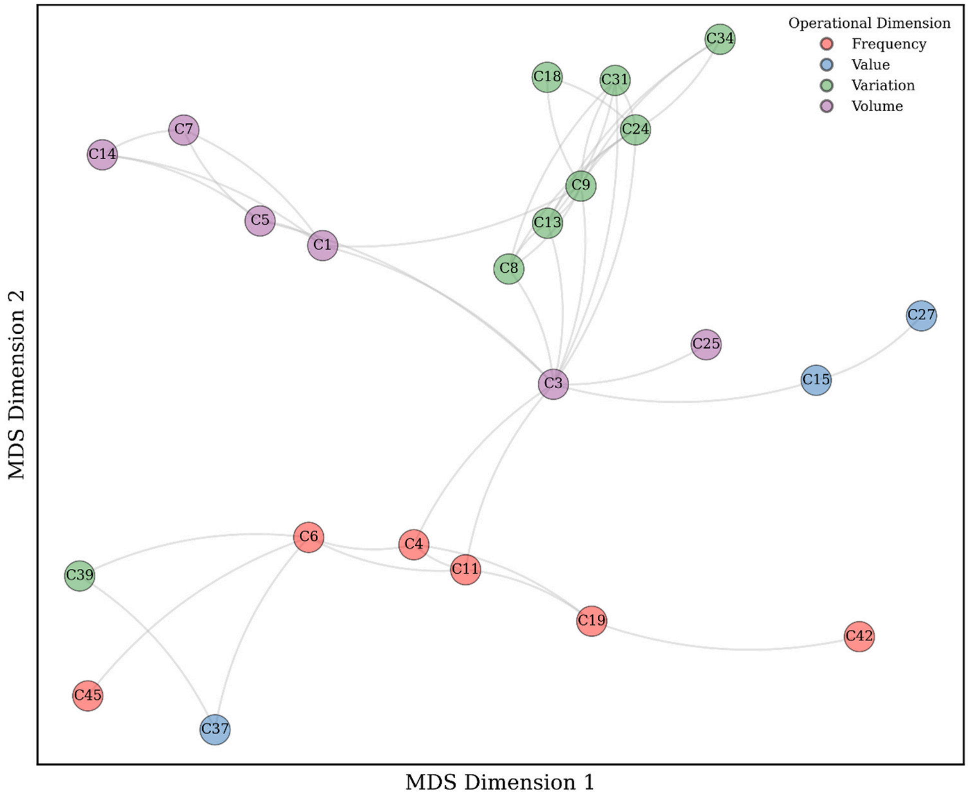

Fig. 2. Empirical topological manifold projection of the selected criteria backbone (S_{TSAV}). Each node is identified by its index code and colour-coded by its semantic operational dimension. Curved edges represent significant topological links where the correlation distance satisfies the stability threshold (ϵ_{opt}) identified via persistence analysis.

inflates toward near-total inclusion, reaching 49 criteria by the upper boundary. At this scale the candidate graph develops redundant cycles ($\beta_1 > 0$).

Therefore, the selected threshold (ϵ_{opt}) sits at the onset of the connected regime, the smallest scale at which the selection forms a single connected, acyclic backbone spanning the decision space. Unlike the saturated regime ($\epsilon > \epsilon_{opt}$), ϵ_{opt} isolates a compact topological backbone (S_{TSAV}) satisfying completeness and non-redundancy axioms without compromising distinctiveness. Table 2 details the composition of this topological backbone. The criteria are indexed (C_T) according to their topological centrality. Notably, the set retains a dual portfolio: it includes high-consensus Hub metrics (e.g., *Consumption Value*) alongside orthogonal keystone criteria (e.g., *Variation Count*) which operate independently of the main transaction flow. See Appendix A for the complete criteria list.

5.2. T-SAV selection

The T-SAV axiomatic pruning process reduced the initial candidate pool to a minimal structural backbone of 23 criteria. This selected subset satisfies the dual requirements of topological completeness (connectivity) and non-redundancy (acyclicity). The final criteria set, categorised by operational dimension, is detailed in Table 2.

The structural validity of this selection is visualised in Fig. 2. The MDS projection confirms that the selected criteria form a cohesive, interconnected manifold rather than a fragmented set of isolated features. The geometric arrangement of nodes, determined purely by empirical correlation distances (D), aligns closely with the semantic operational definitions, providing strong evidence for the validity of the proposed taxonomy. As illustrated in Fig. 2, the criteria naturally organise into distinct topological clusters that correspond to specific operational behaviours:

- *Variation (Green nodes)*: Criteria related to stock variation (e.g., *Total Nett Van Variation*, *Total Combined Positive Var Units*) form a tightly coupled cluster in the upper-right quadrant. The density of connections within this group indicates that variation metrics are highly co-dependent; a fluctuation in one variation metric is almost invariably mirrored by changes in others.
- *Volume (Purple nodes)*: High-volume stock metrics (e.g., *Avg Monthly Usage*, *Total Qty On Hand*) occupy the upper-left region. Their separation from the variation cluster empirically confirms that value and variation are distinct operational phenomena, i.e., a high-volume SKU is not inherently variable.
- *Frequency (Red nodes)*: Procurement and usage frequency metrics (e.g., *Adjustment Transaction Count*, *Warehouses Without Stock*) form a distinct chain along the lower axis. This linear arrangement suggests a continuous progression in frequency behaviour, distinct from the stochastic nature of the variation cluster.
- *Value (Blue nodes)*: Financial criteria (e.g., *Total Stock Value*, *Consumption Value*) are notably positioned at the periphery of the manifold. Their loose connectivity to the central core suggests that unit cost and value density act as exogenous variables, behaving independently of the physical flow dynamics (volume and frequency).

A critical finding is the centrality of the node *Warehouse Transfers In Units*, which acts as the topological keystone of the system. Situated at the intersection of the *Volume*, *Variation*, and *Frequency* clusters, this node exhibits the highest betweenness centrality. Operationally, this identifies it as the primary bridging criterion; it is the single criterion most predictive of the overall system state, physically linking the otherwise distinct dimensions of *Volume* and *Variation*. Conversely, the

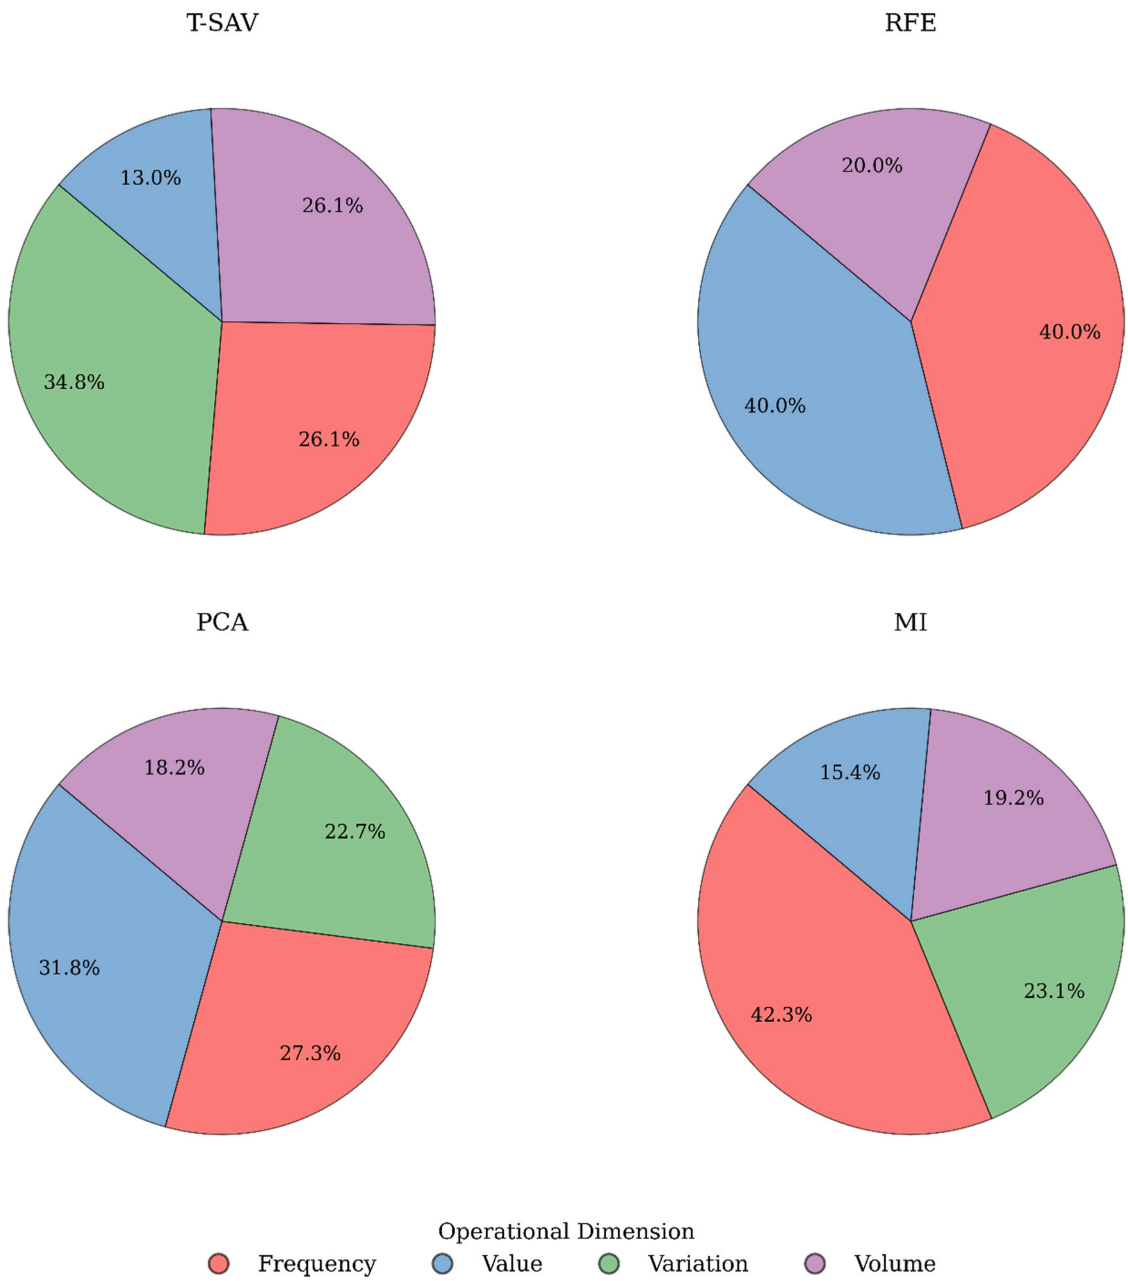

Fig. 3. Comparative dimensionality profiles of selected feature sets. T-SAV (S_{TSAV}) exhibits topological sensitivity to instability, allocating the plurality of selection to *Variation* (34.8%). Conversely, RFE (S_{RFE}) demonstrates target dependency, optimising for *Value* and *Frequency* (40.0% each) while entirely excising the variation dimension (0.0%). PCA (S_{PCA}) reflects consensus bias toward Value (31.8%), whereas MI (S_{MI}) illustrates the bivariate limitation, retaining a disproportionate density of Frequency metrics (42.3%).

peripheral placement of *Value* criteria implies that financial value is a poor predictor of operational risk compared to these core flow metrics.

The presence of a continuous backbone of edges connecting all four dimensions visually validates the completeness axiom. There are no isolated islands of data; every selected criterion can be reached from any other through a sequence of significant correlations ($\epsilon \leq \epsilon_{opt}$). This indicates that the T-SAV framework has captured a holistic model of the company's supply chain, integrating *Value*, *Volume*, *Variation* and *Frequency* operational dimensions into a single unified framework.

5.3. Comparative feature selection profiles

To validate the operational robustness of the T-SAV framework, its selection profile was benchmarked against three established feature selection methodologies: MI, PCA, and RFE. The comparative distribution of selected criteria across operational dimensions is summarised in Fig. 3. See Appendix B for a detailed record of criteria selected by model.

5.3.1. Behavioural archetypes and model bias

Analysis indicates four distinct selection behaviours rooted in the underlying algorithmic logic:

- *Consensus bias* (PCA): PCA exhibited a balanced preference for *Frequency*, *Volume*, and *Value*. However, it marginalized the *Variation* dimension, which constituted only a minor portion of its selection. Notably, PCA was the only model to prioritise external *Value* metrics (e.g., *Direct To Site Purchase Order Value*), confirming a bias toward global signal magnitude over internal operational stability.

- *Target dependence* (RFE): RFE exhibited the most aggressive selection. It allocated 80% of its selection to *Value* (*Buy Price*, *Warehouse Transfers In Value*) and *Frequency* metrics combined. Critically, RFE completely excluded *Variation* criteria. This confirms that predictive wrappers optimised for the *Consumption Value* target are blind to multi-dimensional systemic risk if that risk does not strictly correlate with the target variable.
- *Bivariate limitation* (MI): MI yielded the most expansive set. Lacking a redundancy filtration mechanism, MI skewed heavily toward *Frequency* (∼42%), retaining clusters of collinear count-based variables (e.g., *Consumption Transaction Count*, *Job Issue Incident Transaction Count*). While it captured some variation, its focus remained on transaction velocity rather than structural risk.
- *Topological Preservation* (T-SAV): T-SAV achieved the most balanced profile. Unique among the models, T-SAV prioritised *Variation* (∼35%) as its largest dimensional component, followed by *Volume* and *Frequency*. It notably de-prioritised *Value*, retaining only essential financial nodes. This suggests T-SAV is the only model to explicitly identify *Variation* as a primary structural feature of the company's supply chain.

5.3.2. Analysis of unique selections & blind spots

A granular analysis of criteria selected exclusively by a single model illustrates specific operational blind spots. PCA uniquely retained *Direct-To-Site* metrics. While mathematically dominant due to high variance, these features are operationally irrelevant to internal warehouse availability. MI uniquely retained niche transaction counts (*Job Issue Incident Transaction Count*), inflating dimensionality with low-level frequency noise. T-SAV was the sole model to retain negative variation (*Total Negative Warehouse Variation Units*) and Inbound Volume (*Purchase Order Units*).

The retention of *Total Negative Warehouse Variation Units* is operationally critical; it is the direct proxy for stockouts and inventory loss. The fact that RFE, PCA, and MI all discarded this metric suggests that standard statistical methods systematically undervalue negative signals (risks) in favour of positive signals (consumption). By uniquely preserving the inbound flow (*Purchase Order Units*), T-SAV ensures the model can correlate downstream shortages with upstream replenishment failures, a capability lost in the benchmark models.

5.4. Empirical validation of the keystone criterion

The theoretical framework posits a distinction between Hubs (high-consensus signal carriers) and keystones (low-consensus structural bridges). To validate this distinction empirically, the selection overlap was analysed for the study's specific keystone candidate: *Total Negative Warehouse Variation Units* (*TNWV*).

5.4.1. Divergence on statistical orthogonality

As predicted by the topological definition, *TNWV* exhibited statistical orthogonality, leading to its systematic rejection by all benchmark models (as confirmed in Appendix B). *TNWV* was discarded by PCA, confirming that inventory risk (variation) is orthogonal to the dominant eigenvectors of *Volume* and *Value*. RFE excluded *TNWV*, validating that negative inventory signals (shortages) often lack a strong predictive correlation with the positive target of *Consumption Value*. MI failed to select *TNWV*, likely due to its weak bivariate dependence on transaction counts.

5.4.2. Topological necessity (cut-vertex effect)

In contrast, T-SAV was the sole model to retain *TNWV*. This unique retention validates the cut-vertex property defined in Section 3.4. While standard models prioritised Hubs such as *Warehouse Transfers In Value* (selected by all 4 models), T-SAV identified that removing *TNWV Units* would topologically disconnect the physical state of warehouse operations activity from the financial impacts of its activity. By retaining this feature, T-SAV preserves the manifold's structural integrity, ensuring the model remains sensitive to transactional risks that do not carry a strong statistical signal but act as critical failure points in the company's supply chain.

The exclusion of *TNWV Units* by RFE, PCA, and MI empirically demonstrates the topological insensitivity of statistical feature selection in this context. These models successfully optimised for the centre of the data distribution (Hubs) but failed to map the boundaries (keystones). T-SAV's retention of the *Variation* dimension demonstrates its capacity to operationalise the keystone criterion, distinguishing true structural constraints from statistical noise.

Table 3
Topological audit results. Evaluation of model selections against structural axioms at the stable threshold (ϵ_{opt}). β_0 denotes connected components (Ideal = 1); β_1 denotes coverage loops/redundancies (Ideal = 0).

Model	Criteria Count	β_0 (Completeness)	β_1 (Redundancy)	Result
T-SAV	23	1	0	PASS
RFE	5	4	0	FAIL (Fragmented)
PCA	22	14	0	FAIL (Fragmented)
Mutual Info	26	4	1	FAIL (Both)

5.5. The topological audit

To validate structural integrity (RQ1), persistent homology was applied to each criteria set. The results are presented in Table 3.

The topological audit demonstrates that traditional statistical models fail to preserve the system's structural invariants:

- *T-SAV* (Benchmark): T-SAV was the sole model to satisfy the audit criteria ($\beta_0 = 1$, $\beta_1 = 0$), confirming the selection forms a single, non-redundant set.
- *PCA* (Fragmentation): Despite selecting a comparable feature count ($n = 22$), PCA exhibited significant fragmentation ($\beta_0 = 14$). By maximising statistical consensus, PCA selected features from disjoint high-loading clusters while discarding the bridge variables required for connectivity. This yielded a set of statistically significant but topologically isolated information islands.
- *RFE* (Sparse Disconnect): RFE failed the completeness test ($\beta_0 = 4$), generating a sparse model capable of *Consumption Value* prediction but lacking the structural connectivity required to explicate upstream variability.
- *MI* (Homological Redundancy): MI failed both axioms. Its high selection count introduced topological loops ($\beta_1 > 0$), indicating circular redundancies where multiple features describe identical latent phenomena without adding unique structural value.

Taken together, these audit results reveal a structural pattern: higher criteria counts do not guarantee a connected, acyclic selection. PCA selected 22 criteria yet produced 14 disconnected components, while MI selected 26 criteria but introduced cyclic redundancy. This indicates that the failure mode is not one of insufficient quantity but of misaligned selection logic. PCA fragments the manifold because it selects from disjoint high-loading clusters without ensuring inter-cluster connectivity. MI inflates dimensionality because it evaluates information gain dyadically, admitting collinear *Frequency* metrics that form topological loops without contributing new structural information. RFE, by contrast, achieves parsimony ($n = 5$) but at the cost of extreme fragmentation, retaining only the narrowly predictive core. The T-SAV framework avoids all three failure modes by conditioning selection on the global topological signature of the combined set rather than on the individual statistical properties of each criterion.

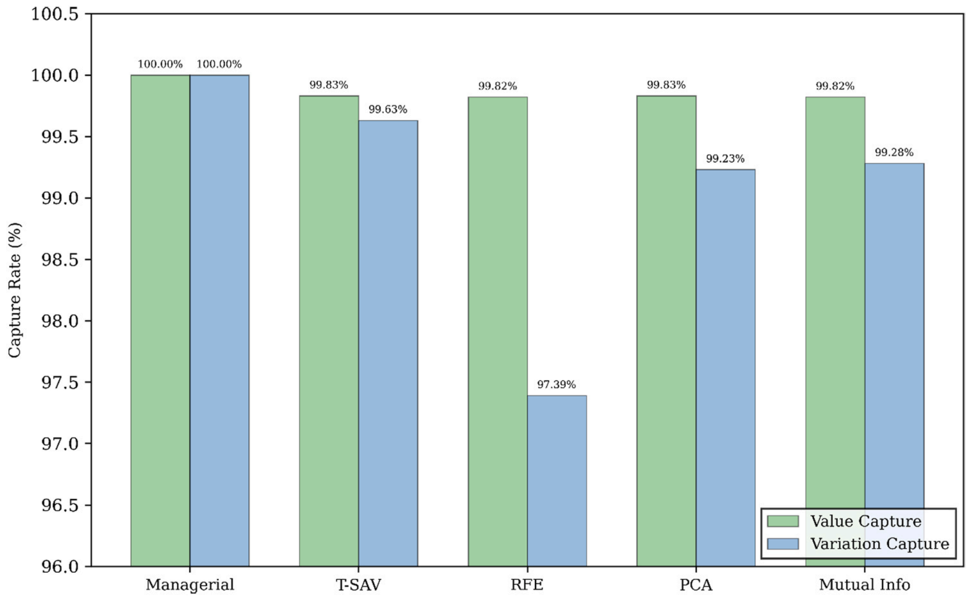

Fig. 4. Comparative classification performance. Bars show value and variation capture for each of the four selection models; the dashed line marks the four-criterion managerial heuristic set (100%).

5.6. Operational validation: multi-criteria inventory classification (RQ2)

To quantify operational efficacy (RQ2), a validation simulation was conducted using an entropy-weighted TOPSIS classification model. Performance was benchmarked against the theoretical upper bound using the value capture ratio (CR_{Val}) and the variation capture ratio (CR_{Var}). Fig. 4 illustrates these results.

While all four methods achieved near-perfect convergence on the value capture ratio (CR_{Val}), significant divergence emerged in their ability to explain system variation (CR_{Var}). Consistent with the target dependency hypothesis, RFE exhibited the most significant structural failure ($CR_{Var} = 97.39\%$). The 2.61% residual variance leakage suggests that predictive wrappers discard keystone criteria lacking explicit *Consumption Value* prediction utility. Statistical filter methods (PCA and MI) outperformed RFE but did not match the topological benchmark (PCA: 99.24%, MI: 99.28%). The residual variance gap indicates that reliance on covariance (PCA) or pairwise entropy (MI) inevitably filters out a subset of independent structural constraints. T-SAV achieved the highest variation capture rate (99.63%), minimising residual variance to 0.37%.

This performance differential, representing a 2.24 %age point improvement over RFE and a 0.40 %age point improvement over PCA and MI, suggests that topological invariants provide a more reliable mechanism for capturing systemic risk than either predictive or statistical correlations. While the absolute differences between the filter-based methods (PCA and MI) and T-SAV are modest, the critical distinction lies in the topological audit results: T-SAV is the only method that simultaneously satisfies both structural axioms, indicating that the variation captured by T-SAV reflects structurally coherent criteria rather than statistically convenient proxies.

To assess the robustness of these performance differences, a bootstrap resampling analysis was conducted ($B = 10{,}000$, resampling SKUs with replacement). In each replicate the four criteria sets were held fixed at their full-sample selections, and only the entropy-weighted TOPSIS classification and the capture ratios were recomputed on the resampled SKUs. The analysis therefore assesses the stability of the operational metric under sample variation, conditional on the selected criteria, rather than the stability of the selection procedures themselves. The 95% confidence intervals for the variation capture ratio were: T-SAV [98.87%, 99.87%]; RFE [92.11%, 98.92%]; PCA [97.64%, 99.71%]; MI [97.82%, 99.77%]. Pairwise Wilcoxon signed-rank tests confirmed that T-SAV significantly outperformed all three benchmark methods ($p < 0.001$ in each case). Beyond statistical significance, the bootstrap distributions reveal notable differences in structural stability. T-SAV exhibited the smallest variance across resamples (SD = 0.26), indicating that its criteria set captures variation consistently regardless of sample composition. By contrast, RFE exhibited the largest variance (SD = 1.84), with its 95% confidence interval spanning nearly 7 %age points. This instability is consistent with the topological fragmentation identified in Table 3: a criteria set with four disconnected components is inherently sensitive to which subset of SKUs is evaluated, because no structural bridge connects the retained dimensions. These results indicate that the direction of the T-SAV advantage is stable across resamples and not attributable to sampling variability.

Statistical significance and practical magnitude must be distinguished here. With 10,000 resamples, a small difference can reach significance even when its operational relevance is limited, so we also report effect sizes for the paired comparisons against T-SAV. We use the matched-pairs rank-biserial correlation, computed from the same bootstrap distributions, together with the median paired difference in the variation capture ratio expressed in percentage points (pp). Against RFE, the median difference is 2.32 pp (95% CI [0.86, 7.03]) with a rank-biserial correlation of 1.00, and T-SAV captures more variation in every resample. Against PCA, the median difference is 0.39 pp (95% CI [0.05, 1.47]) with a rank-biserial correlation of 0.98. Against MI, the median difference is 0.28 pp (95% CI [−0.02, 1.36]) with a rank-biserial correlation of 0.93. The rank-biserial values are large in all three cases, which reflects a highly consistent ordering across resamples and explains the small p-values. The magnitude of the advantage over the two filter methods is nonetheless modest, under half a percentage point at the median, and for MI the difference interval includes zero at its lower bound. We therefore do not read the variation capture margin over PCA and MI as a large operational effect. The substantive separation between the methods rests on the topological audit, where T-SAV is the only method to satisfy both structural axioms, rather than on the size of the capture-ratio gap.

5.7. Empirical validation of projection loss (RQ3)

The comparative audit is consistent with the projection loss hypothesis, indicating that in this case study the statistical and predictive models excised structural criteria. PCA's operational failure ($CR_{Var} = 99.24\%$) coupled with significant topological fragmentation ($\beta_0 = 14$) is consistent with consensus bias. Despite standardisation, the algorithm prioritised high-covariance hubs over statistically independent (orthogonal) keystones, indicating that statistical variance is not isomorphic to structural necessity. The residual variance gap (2.61%) is consistent with the risk of target dependency within RFE. By optimising for the Consumption Value target, the wrapper discarded orthogonal control variables, suggesting that predictive utility is not isomorphic to structural necessity. Topological loops detected in the MI set ($\beta_1 > 0$) demonstrate the bivariate limitation. Despite high feature retention, MI failed to detect circular redundancies, confirming that pairwise metrics cannot identify high-order homological cycles. Consequently, T-SAV mitigates projection loss, minimising dimensionality while preserving the topological manifold. This empirical performance addresses RQ3, indicating that topological invariance predicts systemic criticality more reliably than statistical consensus or predictive error minimisation in this dataset. The following section interprets the theoretical and practical implications of these findings.

6. Discussion

The primary objective of this study was to determine if topological invariants could serve as a more reliable basis for criteria selection than traditional statistical feature selection methods. Empirical evidence indicates a fundamental divergence between metric reduction and structural preservation. This discussion interprets those findings, demonstrating that widely used techniques, specifically those relying on variance (PCA), predictive error (RFE), or pairwise entropy (MI), systematically compromise the decision manifold by violating the axioms of completeness and non-redundancy. The following sections detail how T-SAV resolves these violations, re-establishing the link between mathematical form and operational function in supply chain analytics.

6.1. Principal findings

This study demonstrates that the application of topological invariants provides a criteria selection method for MCDM that addresses the limitations of traditional coordinate-dependent methods. Results are consistent with the primary hypothesis: in this case study, standard dimensionality reduction and feature selection techniques, specifically PCA, RFE, and MI, exhibit projection loss by systematically violating structural decision axioms.

Addressing RQ1 and RQ3, the topological audit reveals that while T-SAV satisfied the axioms of completeness ($\beta_0 = 1$) and non-redundancy ($\beta_1 = 0$), every comparative model failed to preserve the decision manifold. PCA produced a highly fragmented criteria set ($\beta_0 = 14$). Because standardisation neutralised magnitude effects, this failure is consistent with consensus bias: the algorithm prioritises high-covariance hubs while excising keystone criteria statistically orthogonal to dominant principal components. This suggests that statistical covariance is not a valid proxy for structural necessity. MI failed to detect homological cycles ($\beta_1 > 0$). This confirms that pairwise entropy metrics remain blind to high-order redundancy, retaining clusters of synonymous variables that topological filtration successfully collapses. Regarding RQ2, RFE's predictive optimisation caused significant variation leakage. While T-SAV achieved a variation capture ratio (CR_{Var}) of 99.63%, RFE captured only 97.39%. This supports the target dependency hypothesis: by optimising exclusively for a cost proxy, RFE discards orthogonal control variables essential for explaining systemic variance.

In this case study, these findings indicate that topological connectivity can identify structural weaknesses that statistical covariance, pairwise entropy, and predictive error minimisation may overlook.

6.2. Interpretation and contextualisation with literature

This study was premised on a fundamental theoretical tension identified in the literature: the divergence between statistical significance and decision-theoretic validity. While contemporary supply chain analytics increasingly relies on metric dimensionality reduction to manage high-dimensional data, empirical results provide confirmation that these methods rely on the assumption that decision utility is isomorphic to metric properties such as variance, predictive correlation, or pairwise entropy (i.e., metric conflation).

By subjecting standard feature selection methods to a topological audit, this study moves beyond qualitative critiques of statistical feature reduction to provide quantitative evidence of projection loss. The following subsections interpret these findings through the lens of the axiomatic framework established in Section 3.

6.2.1. Evidence consistent with consensus bias in variance-based models (PCA)

The topological audit of the PCA-derived criteria set is consistent with consensus bias. While the literature frequently employs PCA to prioritise high-variance features under the assumption that variance equates to structural information [20,22,25], empirical results demonstrate that this approach systematically excises structural necessities. Despite standardisation neutralising magnitude differences, PCA failed the completeness audit ($\beta_0 > 1$), producing a fragmented decision space. This empirical fragmentation aligns with theoretical critiques indicating that low eigenvalues often represent constraints on affine sets rather than noise [10]. Operationally, this manifested in the excision of 37.5% of the Variation criteria identified by T-SAV. By prioritising noisy variables that covary with aggregate volume (Hubs), PCA discarded orthogonal keystone criteria, such as Total Negative Warehouse Variation, which are essential for defining system boundaries. This confirms prior warnings in the literature: removing variables based strictly on variance equates to removing structural axioms required for system integrity [26].

6.2.2. Target dependency (RFE)

The empirical performance of RFE is consistent with the structural limitations associated with target dependency in supply chain analytics. Although wrapper methods provide a mechanism to minimise subjectivity during criteria selection [30,32], results demonstrate that this approach introduces predictive narrowness that compromises model robustness. As noted in the literature, wrapper algorithms are inherently model-oriented, evaluating feature relevance strictly through their impact on prediction error [46]. However, predictive accuracy is an insufficient proxy for structural validity. Models optimized solely for prediction frequently overfit to spurious statistical associations within the training data that fail to generalize to production environments [47]. By prioritizing isolated statistical correlation over structural necessity, wrapper methods obscure the mechanistic transparency required to validate algorithmic logic, ultimately compromising model reliability in mission-critical supply chain operations.

Demonstrating this predictive narrowness within the dataset, RFE achieved high efficacy in predicting financial value (CR_{Val}) yet demonstrated the highest residual variance leakage (CR_{Var}) and a total exclusion of variation criteria (0.0%). This empirical evidence substantiates theoretical caveats identified in the literature: the unconstrained maximisation of predictive fit risks incorporating spurious correlations while rendering the model insensitive to latent causal structures orthogonal to the immediate error minimisation function [11,31]. By retaining strictly those features with direct financial utility, RFE systematically excised the operational control parameters essential for resilient classification. This confirms that the reduction of criterion utility to a simple function of predictive accuracy yields criteria sets characterised by structural instability [35].

6.2.3. The bivariate limitation of Mutual Information

The topological failure of the MI set provides quantitative evidence for the bivariate limitation identified in prior information-theoretic frameworks [21,38]. Because estimating true conditional mutual information in high-dimensional spaces is subject to the curse of dimensionality, practical MI algorithms are forced to rely on truncated, lower-order dependencies. Consequently, although MI retained the largest feature set ($n = 26$), the detection of persistent homological loops confirms that pairwise probability metrics neglect multi-information and fail to recognise cyclic redundancy among strongly dependent variables.

While splitting bias is frequently identified in manual selection [28], MI introduces an algorithmic equivalent. By assessing information gain dyadically rather than topologically, the model retained clusters of highly collinear frequency variables. This behavior reflects the information bottleneck principle, wherein compressing a signal to maximize target predictive relevance inherently sacrifices the global structural topology of the data [38]. Ultimately, this confirms that without a high-order structural check, specifically simplicial complex filtration capable of encoding k-way interactions, probabilistic filters inevitably inflate dimensionality without genuinely increasing the manifold's structural information content.

6.2.4. Topological invariance as the basis for axiomatic validity

In contrast to these statistical failures, the success of T-SAV validates the application of established axioms of criteria set coherence [13,42] to automated feature selection. By defining criteria selection through topological invariants, specifically connectedness ($\beta_0 = 1$) and non-redundancy ($\beta_1 = 0$), rather than metric thresholds, the model identified the keystone criteria that statistical methods discarded. The retention of *Total Negative Warehouse Variation* is the primary proof of concept. This variable, statistically obscure to PCA (orthogonality) and RFE (low target correlation), was identified by T-SAV as a topological bridge connecting financial value to physical variation. Its inclusion allowed the T-SAV model to achieve the highest variation capture ratio (99.63%), minimizing the decoupling risk that pervades modern inventory management. This result bridges established methodological gaps [29], transforming unsupervised reduction from a descriptive exploratory tool into a rigorous, theoretically grounded selection mechanism that ensures the shape of the decision manifold remains connected.

6.3. Theoretical implications

6.3.1. Shifting from consensus to topology

This study extends the domain of TDA beyond the diagnostic applications identified in existing literature [44]. By demonstrating that topological invariants can guide ex-ante criteria selection, this framework addresses specific methodological gaps regarding prescriptive selection [45].

Beyond immediate operational utility, these findings prompt a re-evaluation of the definition of noise in decision science. Traditional coordinate-dependant ontologies, particularly PCA, conflate low eigenvalues (low shared variance) with statistical insignificance [9]. Consequently, criteria that fail to align with the global variance signal are mathematically suppressed. Conversely, this research demonstrates that within complex supply chain manifolds, low-eigenvalue features frequently function as topological bridges, serving as critical edges that connect distinct operational dimensions (e.g., connecting financial value to physical variation).

Because standardisation in the methodology neutralised magnitude effects, the systematic exclusion of these features by PCA is attributable to statistical orthogonality rather than scale. In this case study, this is consistent with consensus bias: the tendency of statistical models to maximise covariance (the majority vote) while discarding unique, orthogonal signals. The validation of the T-SAV framework argues that criterion importance must be redefined by its contribution to connectivity rather than its alignment with consensus. This recontextualises criteria selection from a filtering task (outlier removal) to a manifold learning task (structural preservation).

6.3.2. Reframing objectivity in decision science

The literature frequently frames the transition from expert heuristics to quantitative feature selection as an advancement toward objectivity [6], a view frequently championed in contemporary MCDM models [48]. However, the identification of consensus bias challenges this premise.

The systematic exclusions observed in PCA and RFE establish that mathematical formulation does not guarantee structural neutrality; rather, it often codifies algorithmic biases (e.g., variance prioritisation or predictive circularity). Consequently, the construct of objectivity in MCDM necessitates reconceptualisation: it requires transcending the mere mitigation of decision-maker subjectivity to ensure the elimination of algorithmic distortion. We postulate a theoretical hierarchy wherein structural objectivity, predicated on coordinate-invariant topological properties, constitutes a necessary antecedent to statistical objectivity. A decision model cannot be deemed objective if the criteria selection process compromises the topological integrity of the problem space ($\beta_0 > 1$) ex-ante.

6.3.3. The computability of qualitative axioms

Historically, foundational decision axioms governing criteria selection, such as completeness [13,42] and exhaustiveness [13], functioned as qualitative, normative guidelines. By translating these requirements into computable topological invariants ($\beta_0 = 1$ for completeness; $\beta_1 = 0$ for non-redundancy), this study quantifies the topological integrity and axiomatic consistency of criteria sets.

This translation bridges a historic divide between Soft Operations Research (problem structuring) and Hard Operations Research (optimisation). Problem structure validity is transposed from a qualitative normative ideal to a computable structural invariant. This allows for the verification of logical consistency prior to the computation of weights, ensuring that MCDM models satisfy necessary axiomatic conditions before resources are committed to optimisation.

6.3.4. Decoupling covariance from utility

Finally, the empirical failure of PCA challenges the uncritical application of eigenvalue maximisation in engineered systems. Standard statistical theory posits that utility is a function of shared variance (consensus). However, supply chains are engineered for stability; therefore, control often manifests as a lack of variance until a failure occurs.

This study establishes the Topological Keystone: a criterion that is statistically orthogonal to principal components (low consensus) yet structurally central to the manifold (high connectivity). The retention of *Total Negative Warehouse Variation* by T-SAV illustrates a case in which covariance and structural utility diverge, so that a criterion of low statistical prominence carries high structural importance. We present this as a case-specific observation rather than a general law, since it rests principally on a single keystone criterion within one dataset. Because standardisation equalised variance, the exclusion of this criterion by the benchmark methods is attributable to signal uniqueness rather than weak magnitude.

6.4. Managerial implications

This study cautions supply chain practitioners and data scientists against the uncritical application of automated feature selection. While recent literature notes an increasing reliance on algorithms like RFE and PCA, results demonstrate that these methods may excise critical structural information. To mitigate unmodeled variation, practitioners should adopt three protocols:

- *Revaluation of statistical independence*: Practitioners must reject the assumption that statistical independence equates to noise. Standardised criteria that are uncorrelated with dominant clusters (e.g., *Negative Warehouse Variation)* likely function as topological keystones. Eigenvalue-based filtration of these variables effectively removes system boundary constraints.
- *Differentiation of redundancy and connectivity*: Contrary to heuristics that strictly purge correlated variables [36,39], this study confirms that certain high-correlation links are structural necessities. Practitioners must distinguish between circular redundancy (which distorts weights) and structural connectivity (which bridges financial and physical dimensions). Indiscriminate removal based on pairwise correlation risks model fragmentation.
- *Adoption of topological audits*: The systemic lack of ex-ante validation identified in the literature [4] is addressable by operationalising completeness ($\beta_0 = 1$) and non-redundancy ($\beta_1 = 0$). These invariants should function as a mandatory validation protocol, ensuring the criteria set forms a single, coherent component prior to MCDM deployment.

To facilitate practical adoption, a five-step implementation workflow is outlined for practitioners seeking to apply the T-SAV framework: (1) Assemble the candidate criteria set from domain experts, operational databases, or literature, and structure the raw data matrix with SKUs as rows and criteria as columns. (2) Apply log-transformation and *Z*-score standardisation to normalise distributions, then compute the pairwise cosine distance matrix between criteria. (3) Perform persistent homology on the full criteria distance matrix and identify the stable filtration threshold via Longest Stable Interval analysis, selecting the value at the onset of the topological plateau. (4) Rank criteria by topological centrality (bottleneck distance impact) and execute the iterative forward-selection loop, accepting each candidate only if it maintains a single connected component without introducing cycles. (5) Audit the resulting set by verifying the Betti profile satisfies completeness and non-redundancy, then proceed to weight elicitation and MCDM ranking. This workflow requires standard scientific computing libraries for persistent homology computation and can be executed on a conventional workstation for datasets with fewer than 100 candidate criteria.

Regarding computational complexity, the dominant cost in the T-SAV pipeline is the construction of the Vietoris-Rips complex and the computation of persistent homology. For a criteria set of dimension n, the theoretical worst-case complexity scales exponentially at $\mathrm{O}(2^{n})$ if computing the full simplicial complex. However, because T-SAV operates on the criteria distance matrix (D) rather than the full data matrix (X) and requires only lower-dimensional homology (β_0 and β_1), modern libraries bypass the construction of the full complex. Consequently, for the present study ($n = 54$), exploiting matrix sparsity at the targeted threshold (ϵ_{opt}) allowed the full pipeline to execute in under 60 s on standard hardware.

6.5. Limitations

While this study establishes the efficacy of topological invariants in criteria selection, distinct limitations frame the generalisability of the findings.

First, as discussed in Section 6.4, the computational cost of persistent homology can grow steeply in the worst case with the number of candidate criteria (approx. $O(2^n)$. While T-SAV proved efficient for the analysed criteria dimension ($n = 54$), its application to ultra-high-dimensional datasets may require approximation algorithms or landmark selection to remain computationally feasible. The dominant cost is not a single homology computation but the leave-one-out structure of the audit, which recomputes the persistence of the criteria complex once for every candidate criterion to measure that criterion's structural contribution. The number of these recomputations grows linearly with the criteria count, while each recomputation grows with the number of simplices admitted at the homological dimensions retained. For the present problem ($n = 54$ criteria) this remained tractable, but criteria sets numbering in the hundreds would make the repeated full-complex computation the binding constraint. Practical routes to scale include sparse or approximate Vietoris-Rips constructions, witness or landmark complexes that summarise the point cloud with a reduced vertex set, and restricting the audit to the lowest homological dimensions that carry the axiomatic signal, namely connectedness and the first cycle structure, so that higher-order simplices need not be enumerated [43,44].

Second, the study employed a static topological audit. Supply chain ecosystems are dynamic systems where relationships between variables may drift over time. This cross-sectional analysis presumes that the topological shape of the decision manifold remains stable; however, structural breaks (e.g., supply shocks) could alter the topology, potentially necessitating dynamic re-calibration of the filtration parameter (ϵ_{opt}).

Finally, the validation relies on the cosine distance metric to define the embedding space. While appropriate for normalised inventory vectors, topological features are sensitive to the choice of metric. Using Euclidean or Mahalanobis distances could yield divergent Betti profiles, suggesting that the shape of the decision problem is partially dependent on the geometric lens applied.

Several simplifying assumptions also warrant acknowledgement. The framework assumes that criteria relationships are adequately captured by pairwise cosine similarity at a single point in time and that the resulting topological features generalise beyond the observed sample period. The four-dimensional operational taxonomy (*Variation*, *Value*, *Volume*, *Frequency*) is defined a priori based on domain knowledge; in other contexts, the appropriate dimensional structure may differ. Furthermore, this study validates the framework using a single organisational dataset from the telecommunications infrastructure sector. While the dataset is operationally complex (54 criteria, 2810 SKUs), the generalisability of the specific findings, particularly the identification of negative variation as a keystone criterion, is bounded by this single-case design. Future studies should replicate the T-SAV framework across multiple industries, product types, and supply chain configurations to assess the stability and transferability of the topological approach.

A further limitation concerns preprocessing. Sixteen of the 54 criteria carry a directional sign, which we treat as magnitude on the grounds that, for these criteria, the sign records the direction of stock movement rather than a difference in activity. The three nett-variation criteria are therefore represented by their absolute size, and their net direction enters the analysis only through the separate positive and negative variation criteria rather than being modelled directly.

6.6. Future research

The T-SAV framework opens a new avenue for research at the intersection of Algebraic Topology and Decision Science: Topological Model Validation. Future scholarship should move beyond feature selection to address the structural validity of the decision models themselves. A complementary methodological extension is a nested bootstrap that reruns each selection procedure within every resample, which would assess the stability of the selection procedures themselves rather than the operational metric alone (Section 5.6).

6.6.1. Topological consistency in AHP

A critical area for future research is the application of homological algebra to consistency measurement. Currently, frameworks like the Analytic Hierarchy Process (AHP) rely on linear algebra (eigenvalue consistency ratios) to detect logical errors. However, intransitivity in pairwise comparisons (e.g., A > B > C > A) naturally forms geometric cycles. Future research could model the pairwise comparison matrix as a directed simplicial complex, using the first Betti number ($\beta_1 > 0$) to detect and locate specific consistency holes or logical loops within expert judgment. This would allow for a precise, topological audit of preference structures before weights are calculated.

6.6.2. Topological stability of rankings

Furthermore, future studies should expand beyond criteria selection and explore Topological Sensitivity Analysis for ranking methods like TOPSIS or VIKOR. By applying time-varying topology (Zigzag Persistence) to MCDM rankings under weight perturbation, researchers could visualize the stability of a ranking order. This would allow analysts to define the topological stability region of a decision model, i.e., the specific range of inputs within which the decision manifold (and consequently the ranking) preserves its homological shape. This offers a robust, coordinate-free alternative to standard linear sensitivity checks, identifying rank reversal not as a statistical artefact, but as a topological fracture in the decision space.

7. Conclusion

This study addressed a methodological deficit in MCIC: the reliance on coordinate-dependent feature selection methods that may compromise structural model validity. By bridging TDA with foundational MCDM axioms, this research evaluated whether PCA, RFE, and MI introduce systematic projection loss when applied to criteria selection in a telecommunications inventory context. The empirical results provide answers to the three research questions posed. Regarding RQ1, the topological audit confirmed that all three benchmark methods violated the structural axioms: PCA produced 14 disconnected components, RFE produced 4, and MI introduced cyclic redundancy, while T-SAV was the sole method to satisfy both completeness and non-redundancy simultaneously. Regarding RQ2, T-SAV achieved the highest variation capture ratio (99.63%), outperforming RFE (97.39%), MI (99.28%), and PCA (99.23%) on the MCIC validation task. Regarding RQ3, the convergence of topological failure and operational deficit across all three benchmarks supports the projection loss hypothesis: methods that violate structural axioms produce measurably lower variation capture, confirming that coordinate-dependent metrics are unreliable proxies for structural necessity in this context.

The introduction of the T-SAV framework provides a mathematically rigorous alternative. By operationalising completeness as connectedness ($\beta_0 = 1$) and non-redundancy as acyclicity ($\beta_1 = 0$), T-SAV isolated a criteria backbone that improved variation capture by 2.24 %age points compared to the predictive wrapper. These findings indicate that low-variance, orthogonal criteria can serve as structural bridges within a decision manifold, a role that standard statistical methods fail to detect. For practitioners, this supports a transition from criteria selection based on statistical magnitude to selection based on topological connectivity, with the T-SAV framework providing a computable protocol for validating criteria set integrity prior to MCDM deployment. In practical terms, this gives managers a structural check on a criteria set before it drives classification, flagging cases where a statistically minor criterion is in fact holding the decision structure together. These conclusions rest on a single telecommunications inventory dataset, so the specific criteria identified should be read as illustrative rather than universal, and confirmation across other sectors and datasets remains necessary.

CRediT authorship contribution statement

Frank Michael Theunissen: Writing – original draft, Visualization, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Aymen Sajjad:** Writing – review & editing, Supervision, Conceptualization. **Shafiq Alam:** Writing – review & editing, Supervision, Conceptualization.

Declaration of Competing Interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Frank Michael Theunissen reports a relationship with the case company that includes: employment. The co-authors declare that they have no known competing financial or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

Global Criteria Index The complete list of candidate criteria ranked by topological centrality. These codes (C1 - C54) are used throughout the study to benchmark selection across different models.

Code	Criterion	Systemic Importance
C1	Std Dev Monthly Usage	33.76
C2	Total Combined Variation Count	33.50
C3	Warehouse Transfers In Units	32.93
C4	Van Warehouses Without Stock Count	32.56
C5	Avg Monthly Usage	32.48
C6	Adjustment Transaction Count	32.38
C7	Direct To Warehouse Purchase Order Units	32.36
C8	Total Combined Positive Variation Units	32.19
C9	Total Combined Negative Variation Units	31.73
C10	Job Issue Build Units	31.70
C11	Warehouse Transfers In Transaction Count	31.69
C12	Physical Warehouses With Stock Count	31.49
C13	Total Combined Nett Variation	31.44
C14	Purchase Order Units	31.40
C15	Warehouse Transfers In Value	31.23
C16	Direct To Warehouse Purchase Order Transaction Count	31.20
C17	Job Issue Incident Units	30.81
C18	Total Warehouse Variation Count	30.75
C19	Warehouses Without Stock Count	30.63
C20	On BOM Qty	30.61

(*continued on next page*)

(*continued*)

Code	Criterion	Systemic Importance
C21	Job Issue Build Transaction Count	30.53
C22	Job Issue Incident Transaction Count	30.30
C23	BOM Site Count	30.26
C24	Total Negative Warehouse Variation Units	29.99
C25	Total Qty On Hand	29.55
C26	Direct To Warehouse Purchase Order Value	29.53
C27	Total Stock Value	29.45
C28	Job Issue Build Value	29.33
C29	Job Issue Incident Value	29.31
C30	Purchase Order Value	29.18
C31	Total Positive Warehouse Variation Units	29.08
C32	Adjustment Value	27.87
C33	Total Positive Van Variation Units	27.86
C34	Total Nett Warehouse Variation	27.67
C35	Consumption Units	27.65
C36	Total Van Variation Count	27.55
C37	Consumption Value	27.53
C38	Adjustment Units	27.46
C39	Total Nett Van Variation	27.45
C40	Purchase Order Transaction Count	27.16
C41	Consumption Transaction Count	27.07
C42	Physical Warehouses Without Stock Count	26.99
C43	Total Negative Van Variation Units	25.66
C44	CoV Demand	25.62
C45	Van Warehouses With Stock Count	25.44
C46	Warehouses With Stock Count	20.48
C47	Direct To Site Purchase Order Units	18.19
C48	Direct To Site Purchase Order Value	17.52
C49	Last Cost	16.62
C50	Direct To Site Purchase Order Transaction Count	12.81
C51	Direct To Van Purchase Order Value	9.05
C52	Direct To Van Purchase Order Units	8.59
C53	Direct To Van Purchase Order Transaction Count	7.53
C54	Buy Price	7.17

Appendix B

Comparative feature selection matrix. A universal comparison of criteria selected by T-SAV, RFE, PCA, and MI (1 = Selected, 0 = Pruned). The matrix lists all 54 candidate criteria, and the column totals reconcile with the selection counts reported in Table 3 (T-SAV 23, RFE 5, PCA 22, MI 26). The matrix highlights the divergent selection profiles across the dataset.

Code	Criterion Name	Dimension	T-SAV	RFE	PCA	MI
C1	Std Dev Monthly Usage	Variation	1	0	0	1
C2	Total Combined Variation Count	Variation	0	0	0	1
C3	Warehouse Transfers In Units	Volume	1	0	0	1
C4	Van Warehouses Without Stock Count	Frequency	1	1	0	1
C5	Avg Monthly Usage	Volume	1	0	0	1
C6	Adjustment Transaction Count	Frequency	1	0	0	1
C7	Direct To Warehouse Purchase Order Units	Volume	1	0	0	0
C8	Total Combined Positive Variation Units	Variation	1	0	0	1
C9	Total Combined Negative Variation Units	Variation	1	0	0	0
C10	Job Issue Build Units	Volume	0	0	0	0
C11	Warehouse Transfers In Transaction Count	Frequency	1	0	0	1
C12	Physical Warehouses With Stock Count	Frequency	0	0	1	1
C13	Total Combined Nett Variation	Variation	1	0	0	1
C14	Purchase Order Units	Volume	1	0	0	0
C15	Warehouse Transfers In Value	Value	1	1	1	1
C16	Direct To Warehouse Purchase Order Transaction Count	Frequency	0	0	0	0
C17	Job Issue Incident Units	Volume	0	0	1	0
C18	Total Warehouse Variation Count	Frequency	1	0	1	0
C19	Warehouses Without Stock Count	Frequency	1	0	1	1
C20	On BOM Qty	Volume	0	0	0	0
C21	Job Issue Build Transaction Count	Frequency	0	0	0	0
C22	Job Issue Incident Transaction Count	Frequency	0	0	0	1
C23	BOM Site Count	Frequency	0	0	0	0
C24	Total Negative Warehouse Variation Units	Variation	1	0	0	0
C25	Total Qty On Hand	Volume	1	0	1	1
C26	Direct To Warehouse Purchase Order Value	Value	0	0	0	0
C27	Total Stock Value	Value	1	0	1	1
C28	Job Issue Build Value	Value	0	0	0	0

(continued on next page)

(*continued*)

Code	Criterion Name	Dimension	T-SAV	RFE	PCA	MI
C29	Job Issue Incident Value	Value	0	0	1	1
C30	Purchase Order Value	Value	0	0	0	0
C31	Total Positive Warehouse Variation Units	Variation	1	0	1	0
C32	Adjustment Value	Value	0	0	1	1
C33	Total Positive Van Variation Units	Variation	0	0	0	1
C34	Total Nett Warehouse Variation	Variation	1	0	1	0
C35	Consumption Units	Volume	0	1	0	1
C36	Total Van Variation Count	Variation	0	0	0	1
C37	Consumption Value	Value	1	0	0	0
C38	Adjustment Units	Volume	0	0	1	0
C39	Total Nett Van Variation	Variation	1	0	0	1
C40	Purchase Order Transaction Count	Frequency	0	0	1	0
C41	Consumption Transaction Count	Frequency	0	1	0	1
C42	Physical Warehouses Without Stock Count	Frequency	1	0	1	1
C43	Total Negative Van Variation Units	Variation	0	0	1	1
C44	CoV Demand	Variation	0	0	1	0
C45	Van Warehouses With Stock Count	Frequency	1	0	0	1
C46	Warehouses With Stock Count	Frequency	0	0	1	1
C47	Direct To Site Purchase Order Units	Volume	0	0	1	0
C48	Direct To Site Purchase Order Value	Value	0	0	1	0
C49	Last Cost	Value	0	0	1	0
C50	Direct To Site Purchase Order Transaction Count	Frequency	0	0	1	0
C51	Direct To Van Purchase Order Value	Value	0	0	0	0
C52	Direct To Van Purchase Order Units	Volume	0	0	0	0
C53	Direct To Van Purchase Order Transaction Count	Frequency	0	0	0	0
C54	Buy Price	Value	0	1	1	0

References

[1] S. Axsäter, in: C. Price (Ed.), Inventory Control, 2015, Springer, Cham, 2015, p. 271, https://doi.org/10.1007/978-3-319-15729-0.
[2] E.A. Silver, D.F. Pyke, D.J. Thomas, Inventory and production management in supply chains, CRC press, 2016, https://doi.org/10.1201/9781315374406.
[3] J.M.R. Andaur, G.A. Ruz, M. Goycoolea, Predicting out-of-stock using machine learning: an application in a retail packaged foods manufacturing company, Electronics 10 (22) (2021) 2787.
[4] F.M. Theunissen, C.N. Bezuidenhout, S. Alam, Exploring the shortcomings in formal criteria selection for multicriteria decision making based inventory classification models: a systematic review and future directions, Int. J. Prod. Res. (2024) 1–21, https://doi.org/10.1080/00207543.2024.2320680.
[5] M. Kügemann, H. Polatidis, Methodological framework to select evaluation criteria for multi-criteria decision analysis of road transportation fuels and vehicles, Energies 15 (14) (2022) 5267, https://doi.org/10.3390/en15145267.
[6] J. Rezaei, N. Salimi, Optimal ABC inventory classification using interval programming, Int. J. Syst. Sci. 46 (11) (2015) 1944–1952.
[7] K. Lee, K.T. Carlberg, Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders, J. Comput. Phys. 404 (2020) 108973.
[8] M. Harandi, M. Salzmann, R. Hartley, Dimensionality reduction on SPD manifolds: the emergence of geometry-aware methods, IEEE Trans. Pattern Anal. Mach. Intell. 40 (1) (2017) 48–62.
[9] I.T. Jollife, J. Cadima, Principal component analysis: a review and recent developments, Philos. Trans. R. Soc. A Math. Phys. Eng. Sci. 374 (2065) (2016), https://doi.org/10.1098/rsta.2015.0202.
[10] S. Curry, I. Lee, S. Ma, N. Serban, Global sensitivity analysis via a statistical tolerance approach, Eur. J. Oper. Res. 296 (1) (2022) 44–59.
[11] B.F. Darst, K.C. Malecki, C.D. Engelman, Using recursive feature elimination in random forest to account for correlated variables in high dimensional data, BMC Genet. 19 (1) (2018) 65.
[12] F.M. Theunissen, S. Alam, A. Sajjad, An analytical framework for decision criteria validation in complex supply chains, Supply Chain Anal. (2025) 100169, https://doi.org/10.1016/j.sca.2025.100169.
[13] B. Roy, Multicriteria methodology for decision aiding, Kluwer Academic Publishers, 1996.
[14] F.Y. Partovi, J. Burton, Using the analytic hierarchy process for ABC analysis, Int. J. Oper. & Prod. Manag. (1993), https://doi.org/10.1108/01443579310043619.
[15] B.E. Flores, D.L. Olson, V.K. Dorai, Management of Multicriteria Inventory Classification, Math. Comput. Model. 16 (12) (1992) 71–82, https://doi.org/10.1016/0895-7177(92)90021-C.
[16] M. Keshavarz-Ghorabaee, E.K. Zavadskas, L. Olfat, Z. Turskis, Multi-criteria inventory classification using a new method of evaluation based on distance from average solution (EDAS), Informatica 26 (3) (2015) 435–451, https://doi.org/10.15388/Informatica.2015.57.
[17] S.A. Torabi, S.M. Hatefi, B.S. Pay, ABC inventory classification in the presence of both quantitative and qualitative criteria, Comput. Ind. Eng. 63 (2) (2012) 530–537.
[18] T.J. Van Kampen, R.D. Akkerman, Pieter, P. Van Donk, SKU classification: a literature review and conceptual framework, Int. J. Oper. Prod. Manag. 32 (7) (2012) 850–876, https://doi.org/10.1108/01443571211250112.
[19] P. Sharma, M.K. Mehlawat, P. Gupta, W. Ding, Integrating feature selection and fuzzy decision-making: a spherical triangular fuzzy number based framework for large-scale decision-making, Appl. Softw. Comput. 182 (2025), https://doi.org/10.1016/j.asoc.2025.113535.
[20] A. Petroni, M. Braglia, Vendor selection using principal component analysis, J. Supply Chain Manag. 36 (1) (2000) 63–69, https://doi.org/10.1111/j.1745-493X.2000.tb00078.x.
[21] Vergara J.R., Estévez P.A. A Review of Feature Selection Methods Based on Mutual Information. Neural Computing & Applications, vol 24 (1), pp 175-186, 20142015.
[22] U.M. Modibbo, M. Hassan, A. Ahmed, I. Ali, Multi-criteria decision analysis for pharmaceutical supplier selection problem using fuzzy TOPSIS, Manag. Decis. 60 (3) (2022) 806–836.
[23] S. Mishra, S.P. Singh, J. Johansen, Y. Cheng, S. Farooq, Evaluating indicators for international manufacturing network under circular economy, Manag. Decis. 57 (4) (2019) 811–839.
[24] A.Pd.A. Costa, R. Choren, D.Ad.M. Pereira, A.V. Terra, I.Pd.A. Costa, Cd.S. R. Junior, et al., Integrating multicriteria decision making and principal component analysis: a systematic literature review, Cogent Eng. 11 (1) (2024) 2374944.
[25] R.M. Wu, Z. Zhang, W. Yan, J. Fan, J. Gou, B. Liu, et al., A comparative analysis of the principal component analysis and entropy weight methods to establish the indexing measurement, PloS. One 17 (1) (2022) e0262261.
[26] E.C. Özelkan, L. Duckstein, Analysing water resources alternatives and handling criteria by multi criterion decision techniques, J. Environ. Manag. 48 (1) (1996) 69–96.
[27] A. Brint, A. Genovese, C. Piccolo, G.J. Taboada-Perez, Reducing data requirements when selecting key performance indicators for supply chain management: the case of a multinational automotive component manufacturer, Int. J. Prod. Econ. 233 (2021) 107967.
[28] M. Marttunen, F. Haag, V. Belton, J. Mustajoki, J. Lienert, Methods to inform the development of concise objectives hierarchies in multi-criteria decision analysis, Eur. J. Oper. Res. 277 (2) (2019) 604–620, https://doi.org/10.1016/j.ejor.2019.02.039.
[29] B. Rolf, A. Beier, I. Jackson, M. Müller, T. Reggelin, H. Stuckenschmidt, et al., A review on unsupervised learning algorithms and applications in supply chain management, Int. J. Prod. Res. 63 (5) (2025) 1933–1983.
[30] C. Wu, Y. Jia, D. Barnes, Criteria system construction for sustainable supplier selection: a product-category-oriented intelligent model, Appl. Softw. Comput. 145 (2023) 110586.
[31] M. Shahlaei, Descriptor selection methods in quantitative structure–activity relationship studies: a review study, Chem. Rev. 113 (10) (2013) 8093–8103.
[32] L. Hou, K. Lu, G. Bi, Predicting the credit risk of small and medium-sized enterprises in supply chain finance using machine learning algorithms, Manag. Decis. Econ. 45 (4) (2024) 2393–2414.
[33] M. Alkhodair, H. Alkhudhayr, Harnessing industry 4.0 for SMEs: advancing smart manufacturing and logistics for sustainable supply chains, Sustainability 17 (3) (2025) 813.

[34] F.E. Bezerra, G.Cd Oliveira Neto, G.M. Cervi, R. Francesconi Mazetto, A.Md Faria, M. Vido, et al., Impacts of feature selection on predicting machine failures by machine learning algorithms, Appl. Sci. 14 (8) (2024) 3337.
[35] N.K. Chowdhury, M.A. Kabir, M.M. Rahman, S.M.S. Islam, Machine learning for detecting COVID-19 from cough sounds: an ensemble-based MCDM method, Comput. Biol. Med. 145 (2022) 105405.
[36] F.Z. Janane, T. Ouaderhman, H. Chamlal, A filter feature selection for high-dimensional data, J. Algorithms Comput. Technol. 17 (2023) 17483026231184171.
[37] M. Eskafi, M. Kowsari, A. Dastgheib, G.F. Ulfarsson, P. Taneja, R.I. Thorarinsdottir, Mutual information analysis of the factors influencing port throughput, Marit. Bus. Rev. 6 (2) (2021) 129–146.
[38] J. Mielniczuk, Information theoretic methods for variable selection—a review, Entropy 24 (8) (2022) 1079.
[39] X. Wang, Y. Zhou, Multi-label feature selection with conditional mutual information, Comput. Intell. Neurosci. 2022 (1) (2022) 9243893.
[40] N.N. Nashed, C. Lahoud, M.H. Abel, Feature selection effect on context-aware teacher-support systems, J. Multi-Criteria Decis. Anal. 32 (2) (2025), https://doi.org/10.1002/mcda.70014.
[41] N. Pourkhodabakhsh, M.M. Mamoudan, A. Bozorgi-Amiri, Effective machine learning, Meta-heuristic algorithms and multi-criteria decision making to minimizing human resource turnover, Appl. Intell. 53 (12) (2023) 16309–16331.
[42] R.L. Keeney, H. Raiffa, Decisions with multiple objectives: Preferences and value trade-offs, Cambridge university press, 1993.
[43] F. Chazal, B. Michel, An introduction to topological data analysis: fundamental and practical aspects for data scientists, Front. Artif. Intell. 4 (2021) 667963.
[44] Y. Skaf, R. Laubenbacher, Topological data analysis in biomedicine: a review, J. Biomed. Inform. 130 (2022) 104082.
[45] K.F. Mojdehi, B. Amiri, A. Haddadi, A novel hybrid Model for Credit Risk Assessment of Supply Chain Finance based on Topological Data Analysis and Graph Neural Network, IEEE Access. (2025), https://doi.org/10.1109/ACCESS.2025.3528373.
[46] K. Zhou, Z. Gong, G. Wei, R. Slowinski, Preference disaggregation analysis with criteria selection in a regularization framework, Omega 133 (2025) 103252, 2025.
[47] I. Gupta, A. Martinez, S. Correa, H. Wicaksono, A comparative assessment of causal machine learning and traditional methods for enhancing supply chain resiliency and efficiency in the automotive industry, Supply Chain Anal. 10 (2025), https://doi.org/10.1016/j.sca.2025.100116.
[48] Ž. Stević, S. Miškić, D. Vojinović, E. Huskanović, M. Stanković, D. Pamučar, Development of a model for evaluating the efficiency of transport companies: PCA–DEA–MCDM model, Axioms 11 (3) (2022) 140.