The Geometry of Stability: A Cohomological View on Preference Cycles and Algorithmic Robustness

Anonymous authors
Paper under double-blind review

Abstract

Algorithmic stability—the robustness of predictions to training data perturbations—is fundamental to reliable machine learning. While methods like bagging, regularization, and inflated operators improve stability, they appear as disconnected techniques. We propose a unified mathematical framework demonstrating that algorithmic instability often arises from fundamental inconsistencies in local data preferences, mathematically analogous to Condorcet cycles in social choice theory. We formalize these inconsistencies as cohomological obstructions ($H^1 \neq 0$), leveraging established connections between social choice theory and algebraic topology. This framework reveals bagging as a strategy for obstruction prevention (smoothing the preference landscape) and inflated operators as a strategy for obstruction resolution (target space enlargement). Furthermore, we derive a novel technique from this framework, obstruction-aware regularization, which directly enforces mathematical consistency. We provide direct empirical validation for our claims. First, we demonstrate that engineered Condorcet cycles induce high instability in standard methods, which is resolved by inflated operators. Second, using Hodge decomposition, we confirm that bagging significantly reduces the magnitude of cohomological obstructions. Third, we show that our proposed obstruction-aware regularization successfully reduces mathematical inconsistencies and yields substantial improvements across multiple metrics of algorithmic stability.

1 Introduction

The generalization performance of machine learning models is deeply connected to their stability—the ability to produce consistent predictions when trained on slightly different datasets (Bousquet & Elisseeff, 2002). This has motivated substantial research into stability enhancement methods, including the canonical approach of bagging (Breiman, 1996; Soloff et al., 2024), regularization or regularization-based ensembling (Bousquet & Elisseeff, 2002; Chen et al., 2023), and more recently, "inflated" operators (Adrian et al., 2024; Liang et al., 2025).

While effective, these methods appear disconnected. Bagging relies on statistical variance reduction, regularization on optimization constraints, and inflation on enlarging output spaces to handle ambiguity. This fragmentation raises a fundamental question: *Is there a unifying mathematical structure underlying these diverse approaches to stability?*

We argue the answer is yes, provided by **algebraic topology**, specifically the tools of **cohomology** as applied to the geometry of preferences. We propose that a significant source of instability arises when different subsets of the data, or different regions of the data space, induce inconsistent local preferences over the hypothesis space.

Example 1.1 (Running Example: The Ranking Cycle)**.** *Consider selecting the best model among three candidates $\{A, B, C\}$. Suppose one subset of the data strongly suggests A is better than B ($A \succ B$). A second subset suggests $B \succ C$, and a third suggests $C \succ A$. This creates a cycle: $A \succ B \succ C \succ A$. No single global ranking can satisfy all local preferences. An algorithm forced to choose a single winner will be highly unstable in this scenario; small perturbations in the data composition can drastically change which preference dominates, flipping the global choice.*

This scenario is the well-known Condorcet Paradox from social choice theory. We demonstrate that these paradoxes, which are rigorously characterized by cohomological obstructions (Chichilnisky, 1980; Baryshnikov, 1993), are the root cause of many instabilities in machine learning selection and ranking tasks.

Our framework formalizes the stability problem as one of preference aggregation under metric constraints. We show that:

1. Inconsistent local preferences create **cohomological obstructions** ($H^1 \neq 0$) that prevent the existence of a stable, accurate, single-valued solution.
2. Bagging and regularization are methods for **obstruction prevention**. They smooth the preference landscape, thereby breaking the cycles (reducing the H^1 component).
3. Inflated operators implement **obstruction resolution via systematic target space enlargement**. When a single consensus cannot be found, they return the minimal set encoding the ambiguity (e.g., returning $\{A, B, C\}$ in Example 1.1), which is a stable solution in the enlarged space.

This framework provides mathematical unification, revealing deep connections between algorithmic stability, social choice theory, and algebraic topology, and suggesting principled avenues for designing stable algorithms.

2 Related Work

Our work sits at the intersection of algorithmic stability, topological data analysis, and social choice theory.

2.1 Foundations of Algorithmic Stability

The foundational work of Bousquet & Elisseeff (2002) established the link between algorithmic stability and generalization. The standard paradigm often addresses stability through statistical variance reduction, canonically exemplified by bagging (Breiman, 1996), whose assumption-free stability properties have been recently formalized (Soloff et al., 2024).

More recently, research has focused on the instability inherent in selection procedures like the argmax. Inflated operators (Adrian et al., 2024; Liang et al., 2025) address this by returning a stable set of near-optimal solutions, motivated largely by the metric sensitivity of the loss function under near-ties. Our work provides a unifying theoretical framework that reinterprets these methods as strategies for handling topological obstructions in the underlying data preferences.

2.2 Topological Data Analysis in Machine Learning

The application of Topological Data Analysis (TDA) to machine learning is a growing field. Much of this work focuses on the geometry of the *parameter space* or the loss landscape. For instance, persistent homology and Morse theory have been used to identify obstructions to optimization, such as the structure of local minima and saddle points (Barannikov et al., 2022).

Crucially, our work analyzes the geometry of the *preference space*—how data induces rankings over the hypothesis space. The obstructions we identify are not features of the optimization landscape, but rather fundamental inconsistencies in the problem definition itself (conflicting data preferences), which exist independently of the optimization procedure.

2.3 Cohomology and Hodge Theory in ML

The mathematical tools we employ have precursors in related ML domains, though their application and interpretation differ significantly from ours.

Fairness. The foundational theory of topological social choice (Chichilnisky, 1980; Baryshnikov, 1993) has been applied to diagnose impossibility results in fairness. Prior work demonstrates that simultaneously satisfying disparate fairness criteria can be impossible due to Condorcet-like cycles across different metrics. Our work generalizes this insight from fairness to the foundational property of algorithmic stability.

Rank Aggregation. Combinatorial Hodge theory was introduced to the problem of rank aggregation by Jiang et al. (2008). They utilized the Hodge decomposition to extract a global ranking (the gradient component) from pairwise comparison data and quantify the remaining inconsistencies (the cyclical components). While this provides a powerful tool for analyzing ranking data, our work makes the crucial connection between this mathematical structure and *algorithmic stability*. We reinterpret the cyclical component as the H^1 obstruction that *causes* instability, thereby providing a unifying framework for understanding stability-enhancing methods like bagging and inflated operators.

3 Stability as Constrained Preference Aggregation

We formalize the learning problem and stability constraints, analyzing how local data preferences can conflict.

3.1 The Learning Setup and Stability

Let $\mathcal{Z}$ be the data space, and $\mathcal{D}$ the set of finite training datasets. We can view $\mathcal{D}$ as a graph where edges connect adjacent datasets (differing by one point). Let $\mathcal{H}$ be the hypothesis space, equipped with a metric $d_{\mathcal{H}}$. Let $L : \mathcal{H} \times \mathcal{D} \to \mathbb{R}$ be the loss function. A learning algorithm $\mathcal{A} : \mathcal{D} \to \mathcal{H}$ aims to select $\mathcal{A}(S) \approx \arg\min_{h\in\mathcal{H}} L(h, S)$.

Definition 3.1 (Algorithmic Stability). *An algorithm $\mathcal{A}$ is ϵ-stable if for any two adjacent datasets $S, S' \in \mathcal{D}$:*

$$d_{\mathcal{H}}(\mathcal{A}(S), \mathcal{A}(S')) \leq \epsilon.$$

3.2 Local Preferences and the Accuracy-Stability Tradeoff

The loss function induces preferences over the hypothesis space for each dataset.

Definition 3.2 (Induced Preference Relation). *For a dataset S, we say h_1 is preferred over h_2 (denoted $h_1 \succ_S h_2$) if $L(h_1, S) < L(h_2, S)$.*

To capture the set of viable hypotheses, we define the set of near-optimal solutions.

Definition 3.3 (Near-Optimal Hypothesis Set). *For accuracy tolerance $\delta \geq 0$, the δ-optimal hypothesis set for dataset S is:*

$$O_\delta(S) := \{h \in \mathcal{H} \mid L(h, S) \leq \min_{h'\in\mathcal{H}} L(h', S) + \delta\}.$$

The fundamental challenge is finding an algorithm that is simultaneously accurate and stable.

Definition 3.4 (The Stable Selection Problem). *Given $\epsilon, \delta \geq 0$, the Stable Selection Problem is to find a map $\mathcal{A} : \mathcal{D} \to \mathcal{H}$ such that for all $S, S' \in \mathcal{D}$:*

1. *$\mathcal{A}(S) \in O_\delta(S)$ (Accuracy)*
2. *If S, S' are adjacent, $d_{\mathcal{H}}(\mathcal{A}(S), \mathcal{A}(S')) \leq \epsilon$ (Stability)*

To ensure the problem is well-posed, we assume standard regularity conditions.

Assumption 3.5. *The dataset space $\mathcal{D}$ is connected.*

4 Cohomological Obstructions to Stability

We leverage the connection between social choice theory and algebraic topology to rigorously characterize when inconsistent local preferences prevent stable aggregation. This connection uses cohomology to detect the presence of Condorcet cycles (Chichilnisky, 1980; Baryshnikov, 1993).

4.1 The Geometry of Preferences

We analyze the structure of preferences induced by the data. We focus on a finite subset of relevant hypotheses $\mathcal{H}_k = \{h_1, \ldots, h_k\}$ (e.g., models in an ensemble or candidates in model selection). We view the collection of datasets $\mathcal{D}$ (or individual data points in an ensemble context) as "voters."

Definition 4.1 (Preference Profile)**.** *A preference profile P is a collection of preference relations $\{\succ_S\}_{S\in\mathcal{D}}$ over $\mathcal{H}_k$.*

An aggregation rule seeks a global consensus ranking $\succ_G$ that best represents the profile P.

4.2 Cohomological Interpretation of Cycles

We briefly introduce the necessary concepts; a detailed mathematical treatment is provided in Appendix A. The key insight from topological social choice theory is that the impossibility of fair aggregation stems from the topology of the *space of preferences.*

We encode preferences algebraically as cochains on a simplicial complex representing the hypothesis space. A 0-cochain represents a utility function (or loss), and a 1-cochain represents pairwise comparisons. The coboundary operator d^0 transforms utility into induced comparisons. Coboundaries represent perfectly consistent (acyclic) preferences derivable from a global ranking.

Theorem 4.2 (Cohomological Obstruction Principle (Adapted))**.** *A preference profile P contains a Condorcet cycle if and only if the configuration of local preferences leads to a structural inconsistency. This inconsistency is mathematically characterized by the existence of a non-trivial class in the first cohomology group H^1 of the underlying topological space of preferences.*

Proof. This theorem is a restatement of foundational results in topological social choice theory. The proof relies on showing that the space of non-cyclic preference profiles is contractible (topologically trivial). The presence of a cycle introduces a "hole" in this space, which is detected by cohomology. See Appendix A for a detailed adaptation of this framework. □

Example 4.3 (The Condorcet Cocycle)**.** *The profile $(A \succ B, B \succ C, C \succ A)$ from Example 1.1 is the canonical generator of H^1. It represents a fundamental structural inconsistency that prevents the existence of a single global ranking satisfying all local preferences simultaneously.*

4.3 Mapping Stability to Cohomological Obstructions

We now connect this theory back to the Stable Selection Problem.

Proposition 4.4 (Instability Induced by Obstructions)**.** *If the preference profile induced by the data contains a structural inconsistency (a Condorcet cycle characterized by a non-trivial H^1 class), then for sufficiently small accuracy tolerance δ and stability parameter ϵ, the Stable Selection Problem is unsolvable.*

Proof. Let the non-trivial H^1 class be represented by a cycle of preferences over hypotheses $\{h_1, \ldots, h_k\}$, such that $h_1 \succ h_2 \succ \cdots \succ h_k \succ h_1$. This implies the existence of subsets of the dataset space, $\mathcal{D}_i \subseteq \mathcal{D}$, where h_i is strongly preferred over h_{i+1} (indices taken modulo k).

Formally, there exists a separation constant $\gamma > 0$ such that for $S \in \mathcal{D}_i$, $L(h_i, S) < L(h_{i+1}, S) - \gamma$.

We choose the accuracy tolerance $\delta < \gamma/2$. By the definition of $O_\delta(S)$, if $S \in \mathcal{D}_i$, then $h_{i+1} \notin O_\delta(S)$. This means the optimal sets in different regions are disjoint with respect to the hypotheses in the cycle.

Now consider the stability constraint ϵ. To satisfy the accuracy constraint, an algorithm $\mathcal{A}$ must select $\mathcal{A}(S) \in O_\delta(S)$. The structural inconsistency (the cycle) forces the algorithm to choose between distinct hypotheses in different regions to maintain accuracy.

If the stability constraint ϵ is too tight relative to the separation of the hypotheses and the geometry of $\mathcal{D}$, the algorithm cannot transition between the preferred hypotheses h_i while remaining stable.

Consider the simplified case where $\mathcal{H}$ has the discrete metric ($d(h_i, h_j) = 1$ if $i \neq j$) and $\epsilon < 1$. Stability implies $\mathcal{A}(S) = \mathcal{A}(S')$ for adjacent S, S'. Since $\mathcal{D}$ is connected (Assumption 3.5), this forces a single global choice h^*. However, the cycle implies that for any h^*, there exists a region $\mathcal{D}_i$ where h^* is strongly dispreferred. For small δ, $h^* \notin O_\delta(S)$ for $S \in \mathcal{D}_i$. Thus, the Stable Selection Problem is unsolvable. □

When $H^1 \neq 0$, no single-valued algorithm can simultaneously satisfy the inconsistent preferences and the metric constraints of stability.

5 A Unified View on Stability Methods

Our cohomological framework provides a unified perspective on how different algorithmic stability methods address these obstructions. They fall into two main categories: Obstruction Prevention (modifying the input preferences) and Obstruction Resolution (modifying the output space).

5.1 Obstruction Prevention: Smoothing the Landscape

Methods like bagging and regularization aim to prevent obstructions by smoothing the preference landscape, effectively breaking the cycles. To analyze this rigorously, we must consider the *ordinal* preferences (rankings) induced by the losses, rather than the cardinal losses themselves. This distinction is crucial because cardinal aggregation (e.g., averaging losses) always yields a consistent gradient flow and thus cannot detect Condorcet cycles (see Appendix A.4.2). We aggregate these rankings (e.g., using Pairwise Majority Vote) to construct the preference cochain C_P, which we analyze using the Hodge decomposition.

5.1.1 Bagging: Statistical Cycle Breaking

Bagging utilizes bootstrap samples $\{S_i\}$ from S. This induces a smoothed loss landscape, which in turn leads to more consistent ordinal preferences across the data. While its effectiveness is statistically grounded in variance reduction (Soloff et al., 2024), our framework provides a structural interpretation.

Proposition 5.1 (Bagging Reduces Cohomological Obstructions)**.** *Bagging transforms the original preference profile P induced by an unstable base learner $\mathcal{A}$ into a smoother profile P_{bag} induced by the bagged learner $\mathcal{A}_{bag}$. This operation systematically reduces the magnitude of the cohomological obstruction in the aggregated ordinal preferences.*

Proof. We aim to show that bagging reduces the magnitude of the cohomological obstruction, i.e., $\mathbb{E}[\|C_{\text{cycle}}^{\text{bag}}\|] \leq \mathbb{E}[\|C_{\text{cycle}}\|]$. We utilize the Hodge decomposition (Appendix A.4) on the aggregated ordinal preference cochain C_P:

$$C_P = C_{\text{gradient}} + C_{\text{cycle}}.$$

C_{cycle} represents the harmonic component (H^1 obstruction), quantifying the deviation from a globally consistent ranking.

The key mechanism of bagging is variance reduction through averaging over bootstrap samples (Breiman, 1996). We analyze how this statistical stabilization impacts the structure of the ordinal preferences.

Let $V_{ij}(S) = L(h_j, S) - L(h_i, S)$ be the cardinal preference score between hypotheses h_i and h_j induced by dataset S. The corresponding ordinal preference is $O_{ij}(S) = \text{sign}(V_{ij}(S))$. The aggregated ordinal cochain (via Pairwise Majority Vote, see Appendix A.4.2) is approximately $C_P(i, j) = \mathbb{E}_S[O_{ij}(S)]$.

An unstable base learner $\mathcal{A}$ exhibits high variance in $V_{ij}(S)$. This instability implies that the rankings between hypotheses are sensitive to the specific dataset S. Consequently, the ordinal preferences $O_{ij}(S)$ frequently flip signs across the distribution of datasets (frequent rank reversals).

This high frequency of rank reversals leads to localized inconsistencies. Consider a triplet of hypotheses (h_i, h_j, h_k). Instability increases the probability that the local preferences form a cycle (e.g., $O_{ij}(S) = 1, O_{jk}(S) = 1, O_{ki}(S) = 1$). When aggregated over S, these local inconsistencies contribute to the global cyclical component C_{cycle}.

The bagged learner $\mathcal{A}_{\text{bag}}$ smooths the cardinal preferences, resulting in $V_{ij}^{\text{bag}}(S)$ with significantly reduced variance compared to $V_{ij}(S)$ (Soloff et al., 2024). This stabilization makes the ordinal preferences $O_{ij}^{\text{bag}}(S)$ more robust and less likely to fluctuate across the dataset distribution.

As the preferences become more stable, the occurrence of localized inconsistencies (local cycles) decreases. The preference profile becomes more aligned with a global consensus ranking. Mathematically, this means the resulting aggregated cochain $C_{P_{\text{bag}}}$ is closer to the space of coboundaries, $\text{im}(d^0)$ (the gradient subspace).

By the orthogonality of the Hodge decomposition, a cochain that is closer to the gradient subspace must have a smaller projection onto the orthogonal harmonic subspace (the cyclical component). Therefore, the stabilization provided by bagging systematically reduces the magnitude of the H^1 obstruction. □

5.1.2 Regularization: Enforcing Consensus

Regularization-based ensemble stability adds penalty terms that discourage disagreement, e.g., $\mathcal{L}_{\text{glue}} = \sum_{i,j} \|h_i(x) - h_j(x)\|^2$. This explicitly enforces consensus among the ensemble members (the "voters"). By minimizing $\mathcal{L}_{\text{glue}}$, the optimization procedure forces the local preferences to align, directly minimizing the inconsistencies that contribute to the cyclical component C_{cycle}.

5.2 Obstruction Resolution: Target Space Enlargement

When obstructions are inherent ($H^1 \neq 0$), a stable single-valued solution does not exist. The mathematical solution is to change the target space.

5.2.1 The Mathematical Principle

In social choice theory, when a single consensus winner does not exist due to cycles, the solution is to identify a consensus *set*.

Definition 5.2 (Top Trading Cycle (TTC) Set)**.** *The Top Trading Cycle set (related to the Smith or Schwartz sets) is the minimal set of hypotheses $H^* \subseteq \mathcal{H}_k$ such that hypotheses in H^* dominate those outside H^*.*

When a Condorcet cycle exists, the TTC set contains all hypotheses involved in the cycle. This set represents the minimal enlargement of the output required to resolve the ambiguity.

5.2.2 Inflated Operators as Obstruction Resolution

Inflated operators, rigorously analyzed by Adrian et al. (2024) and Liang et al. (2025), implement this principle by enlarging the target space from $\mathcal{H}$ to $\mathcal{P}(\mathcal{H})$.

Definition 5.3 (Inflated Argmax (Adrian et al., 2024))**.** $\mathit{sargmax}_\epsilon(w) := \{j \mid w_j \geq \max_k w_k - \epsilon\}$.

Proposition 5.4 (Cohomological Significance of Inflation)**.** *Inflated operators resolve cohomological obstructions by systematic target space enlargement. The set-valued output of the inflated operator approximates the Top Trading Cycle set associated with the underlying preference cycle.*

Proof. Consider a profile P with a cohomological obstruction, represented by a cycle $h_1 \succ \cdots \succ h_k \succ h_1$. As shown in Proposition 4.4, this leads to the failure of the single-valued Stable Selection Problem.

The Top Trading Cycle (TTC) set associated with this cycle is $H^* = \{h_1, \ldots, h_k\}$.

Now consider the application of the inflated operator (e.g., returning the near-optimal set $O_\delta(S)$). If the cycle is balanced, the losses of the hypotheses in the cycle will be close across the relevant region of the dataset space: $|L(h_i, S) - L(h_j, S)| \leq \epsilon'$ for $h_i, h_j \in H^*$.

If we choose the tolerance parameter $\delta \geq \epsilon'$, the near-optimal set $O_\delta(S)$ will contain H^*. If hypotheses outside the cycle have significantly higher loss, then $O_\delta(S) \approx H^*$.

The inflated operator thus identifies the set of hypotheses involved in the structural ambiguity characterized by H^1. The stability of this approach relies on the established result (Adrian et al., 2024) that for nearby datasets S, S' (under bounded loss variations), their near-optimal sets have a non-empty intersection: $O_\delta(S) \cap O_\delta(S') \neq \emptyset$. This ensures the set-valued output remains consistent under perturbation, resolving the obstruction by changing the requirements on the output. □

6 Obstruction-Aware Ensembling: A Framework Inspired by Cohomological Insights

Our cohomological perspective suggests principled approaches to designing stable algorithms by systematically targeting these mathematical structures.

6.1 Cohomologically-Inspired Regularization

Motivated by the mathematical structure of disagreement cycles, we propose regularization terms that target higher-order inconsistencies, moving beyond pairwise agreement. Consider an ensemble with models $\{h_i\}$ and learnable alignment maps $\{\phi_{ji}\}$ between their representation spaces. Global consistency requires the alignment maps to satisfy the cocycle condition: $\phi_{ki} = \phi_{kj} \circ \phi_{ji}$.

Definition 6.1 (Obstruction-Aware Ensemble Loss)**.** *A cohomologically-inspired loss includes a term enforcing higher-order consistency:*

$$\mathcal{L}_{cocycle} = \sum_{i,j,k} \|\phi_{ki} - \phi_{kj} \circ \phi_{ji}\|_F^2$$

The total loss is $\mathcal{L}_{total} = \mathcal{L}_{task} + \lambda_1 \mathcal{L}_{glue} + \lambda_2 \mathcal{L}_{cocycle}$.

Minimizing $\mathcal{L}_{\text{cocycle}}$ encourages global consistency in the geometric relationships between models, directly targeting the structures identified by the H^1 framework in the representation space.

Connection to Cycle Consistency. The mathematical form of $\mathcal{L}_{\text{cocycle}}$ is analogous to the "Cycle-Consistency Loss" widely used in unsupervised learning domains like image translation and unsupervised machine translation Grover et al. (2020). In those domains, cycle consistency serves as an engineering heuristic to constrain ill-posed, unsupervised problems. Our work re-contextualizes this mechanism: we derive the cocycle constraint from a theoretical framework, arguing that violations of this consistency are manifestations of a mathematical obstruction with direct consequences for stability. This elevates the mechanism from a heuristic to a principled, theoretically-grounded tool for improving robustness.

6.2 Adaptive Target Space Enlargement

Inspired by the success of inflated operators, we propose adaptive mechanisms that automatically transition between single-valued outputs when consensus exists (suggesting $H^1 = 0$) and set-valued outputs when obstructions may be present ($H^1 \neq 0$), based on measures of ensemble disagreement.

7 Computational Experiments

We conduct three experiments to validate our theoretical framework (full code provided as supplementary material). The first two provide direct evidence for the core claims connecting instability to cohomological obstructions and the mechanisms of existing methods. The third validates the cohomologically-inspired regularization.

7.1 Experiment 1: Instability Induced by Engineered Condorcet Cycles

This experiment validates Proposition 4.4 (obstructions cause instability) and Proposition 5.4 (inflation resolves it).

Setup: We engineered a scenario with three hypotheses (A, B, C) where the underlying data distribution is a mixture of three regions (Z_1, Z_2, Z_3) inducing a balanced Condorcet cycle. Z_1 prefers $A \succ B \succ C$; Z_2 prefers $B \succ C \succ A$; Z_3 prefers $C \succ A \succ B$. We simulate the process of drawing a dataset S (by sampling

Table 1: Stability under an engineered Condorcet cycle ($H^1 \neq 0$).

Method	Stability (Agreement/Intersection)	Avg. Output Size
Standard Argmax	0.4034	1.00
Inflated Argmax	1.0000	3.00

equally from the regions) and an adjacent dataset S' (by perturbing the sampling weights by a small amount, equivalent to changing one data point). We ran 5000 trials.

Metrics: We measure the stability of the standard ‘argmax‘ (Agreement Rate: the probability that the winner is the same for S and S') and the ‘inflated argmax‘ ($\epsilon = 0.01$) using Intersection Consistency (the probability that the output sets for S and S' have a non-empty intersection).

Results: As shown in Table 1, the standard ‘argmax‘ exhibits low stability (0.4034). In contrast, the ‘inflated argmax‘ achieves perfect stability (1.0000). It successfully identifies the ambiguity inherent in the cycle, consistently returning the consensus set $\{A, B, C\}$. This confirms that cohomological obstructions induce instability in single-valued algorithms, which is effectively resolved by target space enlargement.

7.2 Experiment 2: Bagging Reduces Obstructions via Hodge Decomposition

This experiment validates Proposition 5.1, demonstrating that bagging reduces the magnitude of the cohomological obstruction.

Setup: We used a complex synthetic dataset (5 classes, high noise) designed to be ambiguous. We compared an unstable base learner $\mathcal{A}$ (a deep, unpruned Decision Tree) with its stabilized bagged counterpart $\mathcal{A}_{\text{bag}}$ (20 estimators). Over 50 trials, we trained $K = 20$ instances of the base learner on different small subsamples ($N = 80$) of the training data. For the bagged learner, we trained a single instance on the full training data and analyzed its internal estimators. This setup maximizes the variance of the base learner, making obstructions more likely.

Methodology: We treat the independent test data points as "voters" and the K generated hypotheses as "candidates." We constructed the Ordinal Preference Cochain C_P using Pairwise Majority Vote based on 0-1 loss (see Appendix A.4.2). We then applied the Hodge decomposition $C_P = C_{\text{gradient}} + C_{\text{cycle}}$. The L2 norm $\|C_{\text{cycle}}\|$ quantifies the magnitude of the H^1 obstruction.

Table 2: Experiment 2: Hodge decomposition of induced preferences.

Algorithm	Mean cyclical norm $\|C_{\text{cycle}}\|$
Base Learner (Unstable)	2.066×10^{-16}
Bagged Learner (Stable)	1.003×10^{-16}

Results: As shown in Table 2, the unstable Base Learner exhibited a small cyclical norm. We note that the absolute magnitudes for both learners are near machine precision (10^{-16}), suggesting that in this specific experimental realization, structural obstructions were not the dominant source of instability compared to statistical variance. Despite the small absolute magnitude, the key validation of our framework lies in the relative comparison: the Bagged Learner significantly reduced the cyclical component compared to the Base Learner ($p < 10^{-19}$). This empirically confirms the mechanism described in Proposition 5.1: the statistical stabilization of bagging translates directly into a structural smoothing of the preference landscape by reducing the H^1 obstruction.

7.3 Experiment 3: Validation of Obstruction-Aware Regularization

We validate the cohomologically-inspired regularization $\mathcal{L}_{\text{cocycle}}$ proposed in Section 6 by demonstrating a direct correspondence between enforced mathematical consistency and improved algorithmic stability.

Setup: We constructed a system of three interconnected 2D representation spaces with learnable linear transformations ϕ_{ji}. We engineered conflicting ground truth transformations (using rotations of 60°, 60°, and -30° for $\phi_{10}, \phi_{21}, \phi_{20}$ respectively). This results in a 150° "error on the loop" ($\phi_{21} \circ \phi_{10} \neq \phi_{20}$), simulating an engineered structural inconsistency (a non-trivial H^1 obstruction) in the representation space. We compared standard training (using only pairwise alignment loss) with framework-guided training (adding $\mathcal{L}_{\text{cocycle}}$ with weight λ_2).

Methodology: To robustly assess the correspondence between mathematical consistency and stability, we performed a comprehensive parameter sweep across various levels of data noise (0.01 to 0.5) and regularization strengths (λ_2 from 1.0 to 20.0), running 8 independent trials per configuration. We measured the Cocycle Defect (the mathematical inconsistency) and several stability metrics based on the variance across trials: Parameter Variance (consistency of learned weights), Prediction Variance (consistency of outputs), Functional Variance (stability of the behavior of composed maps), and Transformation Consistency (average distance between learned transformations across runs).

Table 3: Experiment 3: Effect of obstruction-aware regularization (Best correspondence configuration: Medium noise, $\lambda_2 = 5.0$).

Metric	Standard training	Framework-guided	Improvement
Mathematical Consistency			
Cocycle defect	2.738	0.110	24.9× reduction
Algorithmic Stability			
Parameter variance	0.000418	0.000177	2.4× improvement
Prediction variance	0.000159	0.000088	1.8× improvement
Transformation consistency	0.028042	0.017872	1.6× improvement
Functional variance	0.000085	0.000056	1.5× improvement

Results: The parameter sweep revealed a strong correspondence between obstruction reduction and stability improvement under moderate conditions. The configuration demonstrating the best overall correspondence (Noise=0.1, $\lambda_2 = 5.0$) is detailed in Table 3. Enforcing the mathematical constraints significantly reduced the cocycle defect (24.9×). Crucially, this correlated strongly with improvements across all four key stability metrics, ranging from 1.5× (Functional Variance) to 2.4× (Parameter Variance). This strongly supports the central claim of this work: targeting the mathematical structures identified by obstruction theory provides a principled and effective mechanism for improving algorithmic stability.

8 Limitations and Future Directions

While our framework provides a unifying perspective on algorithmic stability, it has several limitations that suggest promising directions for future research.

8.1 Limitations

Continuous Hypothesis Spaces: The framework of topological social choice theory and Hodge decomposition is primarily developed for a finite set of alternatives $\mathcal{H}_k$. While we argue that this captures the essential dynamics of instability arising from conflicts between distinct local optima (e.g., in model selection or analyzing ensemble members), machine learning often involves continuous hypothesis spaces $\mathcal{H}$ (e.g., neural network parameters). Extending this framework to handle the geometry of preferences over infinite-dimensional spaces remains a significant challenge.

Scope of Instability Sources: Our framework focuses on instability arising from fundamental inconsistencies in data preferences (cohomological obstructions). It does not encompass all sources of instability in machine learning, such as those arising purely from optimization dynamics (e.g., sensitivity to initialization) or inherent non-smoothness in the model architecture. It is important to distinguish the geometry of the preference space (our focus) from the geometry of the loss landscape in the parameter space. Recent work has applied topological data analysis to the parameter space, using tools like the Morse complex and persistent homology (barcodes) to quantify obstructions to optimization and generalization (Barannikov et al., 2022). While these methods use related mathematical tools to identify "topological obstructions" they characterize the difficulty of navigating the optimization landscape (e.g., the "badness" of local minima) rather than the inconsistencies in the underlying data preferences. As seen in Experiment 7.2, instability can exist even when $H^1 = 0$ due to statistical variance in near-ties.

Computational Scalability: Calculating the Hodge decomposition requires operations (like the pseudoinverse of the Laplacian) that scale poorly when the number of relevant hypotheses K is very large.

8.2 Future Directions

Algorithmic Applications: Further development of obstruction-aware algorithms. This includes designing efficient methods to estimate when $H^1 \neq 0$ during training and adaptively deploying regularization (prevention) or target space enlargement (resolution) strategies based on the detected obstruction magnitude.

Theoretical Extensions: Investigating the precise mathematical relationship between the smoothing effect of bagging (variance reduction) and the reduction of the cyclical component in the Hodge decomposition. Furthermore, exploring the use of higher-order cohomology groups (H^q for $q \geq 2$) to identify more complex structural inconsistencies beyond pairwise cycles.

Connections to Fairness and Explainability: Preference cycles often underlie paradoxes in fairness and explainability, where different metrics or perspectives lead to conflicting conclusions. Extending this cohomological framework to analyze the stability and consistency of fairness metrics is a promising direction.

9 Conclusion

We have presented a unified framework for algorithmic stability based on algebraic topology and the geometry of preferences. By connecting instability to Condorcet cycles in social choice theory, we identify cohomological obstructions ($H^1 \neq 0$) as the fundamental mathematical structure preventing the stable aggregation of inconsistent local preferences. This framework rigorously unifies bagging, regularization, and inflated operators as strategies for obstruction prevention or resolution. We have provided direct empirical evidence demonstrating that these obstructions manifest in practice and that established stability methods successfully mitigate them, confirming the validity of the cohomological perspective.

References

Adrian, M., Soloff, J. A., and Willett, R. Stabilizing black-box model selection with the inflated argmax. *arXiv preprint arXiv:2410.18268*, 2024.

Barannikov, S., Voronkova, D., Trofimov, I., Korotin, A., Sotnikov, G., and Burnaev, E. Topological obstructions in neural networks' learning. *arXiv preprint arXiv:2012.15834v2*, 2022.

Baryshnikov, Y. M. Unifying impossibility theorems: A topological approach. *Advances in Applied Mathematics*, 14(4):404–415, 1993.

Bousquet, O. and Elisseeff, A. Stability and generalization. *Journal of Machine Learning Research*, 2(Mar):499–526, 2002.

Breiman, L. Bagging predictors. *Machine Learning*, 24(2):123–140, 1996.

Chen, Z., Yang, H., Quek, T., and Chong, K. F. Spectral co-distillation for personalized federated learning. In *Advances in Neural Information Processing Systems*, 36:8757–8773, 2023.

Chichilnisky, G. Social choice and the topology of spaces of preferences. *Advances in Mathematics*, 37(2):165–176, 1980.

Grover, A., Chute, C., Shu, R., Cao, Z., and Ermon, S. Alignflow: Cycle consistent learning from multiple domains via normalizing flows. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(4):4028–4035, 2020.

Jiang, X., Lim, L.-H., Yao, Y., and Ye, Y. Learning to rank with combinatorial Hodge theory. *arXiv preprint arXiv:0811.1067*, 2008.

Liang, R., Soloff, J. A., Barber, R. F., and Willett, R. Assumption-free stability for ranking problems. *arXiv preprint arXiv:2506.02257*, 2025.

Soloff, J. A., Barber, R. F., and Willett, R. Bagging provides assumption-free stability. *Journal of Machine Learning Research*, 25(131):1–35, 2024.

A Appendix: Mathematical Foundations: Cohomology of Preferences

This appendix provides the rigorous mathematical details underlying the connection between social choice theory, algebraic topology, and the cohomological obstructions discussed in the main text. We adapt the framework established by Chichilnisky (1980) and Baryshnikov (1993), and incorporate tools from Hodge theory for ranking.

A.1 The Geometry of Rankings and Preferences

Let $\mathcal{H}_k = \{h_1, \ldots, h_k\}$ be a finite set of hypotheses.

Definition A.1 (Space of Preferences)**.** *A utility function $U : \mathcal{H}_k \to \mathbb{R}$ induces preferences. The space of utility functions is $\mathcal{U} = \mathbb{R}^k$. A preference profile is a collection of utility functions $\{U_S\}_{S \in \mathcal{D}}$. Alternatively, we can consider the space of strict linear orders $\mathcal{P}$.*

A.2 The Source of Obstructions: Topological Social Choice

The key insight of topological social choice is to analyze the topology of the *space of profiles*. Let $\mathcal{P}$ be the space of admissible preferences. A profile with N voters is in $\mathcal{P}^N$. An aggregation rule is a function $F : \mathcal{P}^N \to \mathcal{P}$.

A.2.1 Topological Interpretation (Chichilnisky's Approach)

Chichilnisky (1980) showed that the impossibility of fair aggregation (satisfying certain axioms like Arrow's conditions) stems from the topological properties of $\mathcal{P}$.

Theorem A.2 (Chichilnisky, 1980)**.** *If a continuous aggregation rule $F : \mathcal{P}^N \to \mathcal{P}$ exists satisfying certain desirable axioms (e.g., anonymity, unanimity), then $\mathcal{P}$ must be contractible.*

A space is contractible if it has no topological "holes" (i.e., its homology and cohomology groups are trivial).

A.2.2 Cohomology of Preference Spaces

We analyze the topology of the space of strict linear orders $\mathcal{S}_k$. This space is known to be non-contractible for $k \geq 3$.

The Condorcet cycle ($A \succ B, B \succ C, C \succ A$) represents a loop in the space of preferences. This loop cannot be contracted to a point while maintaining the structure of the preferences.

Lemma A.3. *The existence of a non-contractible loop in the space of preferences corresponds to a non-trivial element in the first cohomology group $H^1(\mathcal{S}_k)$.*

This is the rigorous mathematical foundation for Theorem 4.2. The obstruction arises from the structural inconsistency of the profile configuration itself.

A.3 Connecting Topology to Stability

The connection to algorithmic stability is established by recognizing that an unstable algorithm is one that is sensitive to the configuration of the input profile. When the data induces a preference profile that contains a Condorcet cycle (a non-trivial H^1 element), the landscape is structurally inconsistent. As shown in Proposition 4.4, satisfying both accuracy (respecting the local preferences) and stability (metric constraints on the output) becomes impossible.

A.4 Algebraic Tools: Cochains and Hodge Decomposition

While the fundamental obstruction lies in the topology of the preference space, we can analyze the properties of a specific profile using algebraic tools on the hypothesis complex. This is used in Section 5.1 and Experiment 7.2.

A.4.1 Cochains and the Coboundary Operator

Let K be the 1-skeleton (the complete graph) on the vertices $\mathcal{H}_k$.

- $C^0(K;\mathbb{R})$: 0-cochains (functions assigning values to vertices, i.e., utility functions).
- $C^1(K;\mathbb{R})$: 1-cochains (functions assigning values to edges, i.e., pairwise comparisons or flows).

The coboundary operator $d^0 : C^0 \to C^1$ maps a potential function $f \in C^0$ to a flow: $(d^0 f)(h_i, h_j) = f(h_j) - f(h_i)$. This represents a perfectly consistent ranking.

A.4.2 Preference Cochain Construction: Cardinal vs. Ordinal

We define the preference cochain $C_P \in C^1(K, \mathbb{R})$. The construction method is critical.

Cardinal Aggregation: If we aggregate the losses directly (e.g., average loss difference):

$$C_P^{\text{card}}(h_i, h_j) = \int_{\mathcal{D}} (L(h_j, S) - L(h_i, S)) d\mu(S).$$

This construction always results in a coboundary, as it is the gradient of the aggregate loss function $L_{\text{agg}}(h) = \int L(h, S) d\mu(S)$. Therefore, C_P^{card} cannot detect Condorcet cycles.

Ordinal Aggregation (Pairwise Majority Vote): To detect cohomological obstructions relevant to social choice theory, we must aggregate ordinal preferences. Let R_S be the ranking induced by $L(\cdot, S)$.

$$C_P^{\text{ord}}(h_i, h_j) = \int_{\mathcal{D}} \text{sign}(R_S(h_j) - R_S(h_i)) d\mu(S).$$

In practice (Experiment 7.2), this is calculated as: (Proportion of voters preferring h_j over h_i) - (Proportion preferring h_i over h_j). This construction allows for C_P to be non-coboundary.

A.4.3 Hodge Decomposition

We use the Hodge decomposition theorem to analyze flows (cochains) on the graph K.

Theorem A.4 (Hodge Decomposition on Graphs)**.** *The space of 1-cochains C^1 can be orthogonally decomposed as:*

$$C^1 = im(d^0) \oplus \ker(\delta^1)$$

where δ^1 is the adjoint of d^0 (the divergence operator).

Here, $\text{im}(d^0)$ is the space of coboundaries (gradient flows, C_{gradient}). $\ker(\delta^1)$ is the space of harmonic cochains (cyclical flows, C_{cycle}). This harmonic component corresponds to the cohomological obstruction identified in the topological framework.

This theorem allows us to rigorously decompose a preference profile C_P into its consistent part (C_{gradient}) and its cyclical obstruction (C_{cycle}), formalizing the analysis of how bagging reduces the cyclical component in Proposition 5.1.