
Stability Guarantees for Feature Attributions with Multiplicative Smoothing

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Explanation methods for machine learning models tend to not provide any formal
2 guarantees and may not reflect the underlying decision-making process. In this
3 work, we analyze stability as a property for reliable feature attribution methods.
4 We prove that a relaxed variant of stability is guaranteed if the model is sufficiently
5 Lipschitz with respect to the masking of features. To achieve such a model, we
6 develop a smoothing method called Multiplicative Smoothing (MuS). We show
7 that MuS overcomes theoretical limitations of standard smoothing techniques and
8 can be integrated with any classifier and feature attribution method. We evaluate
9 MuS on vision and language models with a variety of feature attribution methods,
10 such as LIME and SHAP, and demonstrate that MuS endows feature attributions
11 with non-trivial stability guarantees.

12 1 Introduction

13 Modern machine learning models are incredibly powerful at challenging prediction tasks but notori-
14 ously black-box in their decision-making. One can therefore achieve impressive performance without
15 fully understanding *why*. In settings like like medical diagnosis [1, 2] and legal analysis [3, 4], where
16 accurate and well-justified decisions are important, such power without proof is insufficient. In order
17 to fully wield the power of such models while ensuring reliability and trust, a user needs accurate and
18 insightful *explanations* of model behavior.

19 One popular family of explanation methods is *feature attributions* [5, 6, 7, 8]. Given a model and
20 input, a feature attribution method generates a score for each input feature that denotes its importance
21 to the overall prediction. For instance consider Figure 1, in which the Vision Transformer [9] classifier
22 predicts the full image (left) as “Goldfish”. We then use a feature attribution method like SHAP [7]
23 to score each feature and select the top-25%, for which the masked image (middle) is consistently
24 predicted as “Goldfish”. However, additionally including a single patch of features (right) alters
25 the prediction confidence so much that it now yields “Axolotl”. This suggests that the explanation
26 is brittle [10], as small changes easily cause it to now induce some other class. In this paper we
27 study how to overcome such behavior by analyzing the *stability* of an explanation: we consider an
28 explanation to be stable if once the explanatory features are included, the addition of more features
29 does not change the prediction.

30 Stability implies that the selected features are enough to explain the prediction [11, 12, 13] and that this
31 selection maintains strong explanatory power even in the presence of additional information [10, 14].
32 Similar properties are studied in literature and identified as useful for interpretability [15], and we
33 emphasize that our main focus is on analyzing and achieving provable guarantees. Stability guarantees
34 in particular are useful as they allow one to accurately predict how model behavior varies with the
35 explanation. Given a stable explanation, one can include more features, e.g. adding context, while
36 maintaining confidence in the consistency of the underlying explanatory power. Crucially, we observe

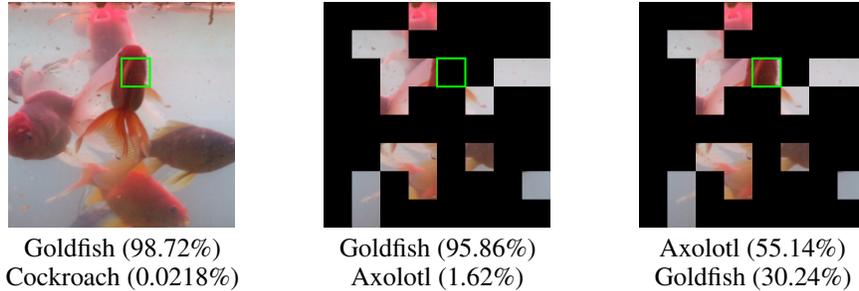


Figure 1: Classification by VisionTransformer [9] on an attribution generated by SHAP [7] with top-25% selection. A single 28×28 pixel patch of difference between the two attributions (marked green) significantly affects prediction confidence and results in a classification flip.

37 that such guarantees only make sense when jointly considering the model and explanation method:
 38 the explanation method necessarily depends on the model to yield an explanation, and stability is
 39 then evaluated with respect to the model.

40 Thus far, existing work on feature attributions with formal guarantees face challenges with com-
 41 putational tractability and explanatory utility. While some methods take an axiomatic approach
 42 [8, 16], others use metrics that appear reasonable but may not reliably reflect useful model behavior,
 43 a common and known limitation [17]. Such explanations have been criticized as at best a plausible
 44 guess, and at worst completely misleading [18] about model behavior.

45 In this paper we study how to construct explainable models with provable stability guarantees. We
 46 jointly consider the classification model and explanation method, and present a formalization for
 47 studying such properties that we call *explainable models*. We focus on *binary feature attributions* [19]
 48 wherein each feature is either marked as explanatory (1) or not explanatory (0). We present a
 49 method to solve this problem, which is inspired by techniques from adversarial robustness, in
 50 particular randomized smoothing [20, 21]. Our method can take *any* off-the-shelf classifier and
 51 feature attribution method to efficiently yield an explainable model that satisfies provable stability
 52 guarantees. In summary, our contributions are as follows:

- 53 • We formalize stability as a key property for binary feature attributions and study this in the
 54 framework of explainable models. We prove that relaxed variants of stability are guaranteed
 55 if the model is sufficiently Lipschitz with respect to the masking of features.
- 56 • To achieve the sufficient Lipschitz conditions, we develop a smoothing method called
 57 Multiplicative Smoothing (MuS). We show that MuS achieves strong smoothness conditions,
 58 overcomes key theoretical and practical limitations of standard smoothing techniques, and
 59 can be integrated with any classifier and feature attribution method.
- 60 • We evaluate MuS on vision and language models along with different feature attribution
 61 methods. We demonstrate that MuS-smoothed explainable models achieve strong stability
 62 guarantees at a small cost to accuracy.

63 2 Overview

64 We observe that formal guarantees for explanations must take into account both the model and
 65 explanation method, and for this we present in Section 2.1 a pairing that we call *explainable models*.
 66 This formulation allows us to then describe the desired stability properties in Section 2.2. We show
 67 in Section 2.3 that classifiers with sufficient Lipschitz smoothness with respect to feature masking
 68 allows us to yield provable guarantees of stability. Finally in Section 2.4 we show how to adapt
 69 existing feature attribution methods into our explainable model framework.

70 2.1 Explainable Models

71 We first present explainable models as a formalism for rigorously studying explanations. Let $\mathcal{X} = \mathbb{R}^n$
 72 be the space of inputs, a classifier $f : \mathcal{X} \rightarrow [0, 1]^m$ maps inputs $x \in \mathcal{X}$ to m logits (class probabilities)

73 that sum to 1, where the class of $f(x) \in [0, 1]^m$ is taken to be the largest coordinate. Similarly,
 74 an explanation method $\varphi : \mathcal{X} \rightarrow \{0, 1\}^n$ maps an input $x \in \mathcal{X}$ to an explanation $\varphi(x) \in \{0, 1\}^n$
 75 that indicates which features are considered explanatory for the prediction $f(x)$. In particular, we
 76 may pick and adapt φ from among a selection of existing feature attribution methods like LIME [6],
 77 SHAP [7], and many others [5, 8, 22, 23, 24], wherein φ may be thought of as a top- k feature selector.
 78 Note that the selection of input features necessarily depends on the explanation method executing or
 79 analyzing the model, and so it makes sense to jointly study the model and explanation method: given
 80 a classifier f and explanation method φ , we call the pairing $\langle f, \varphi \rangle$ an *explainable model*. Given some
 81 $x \in \mathcal{X}$, the explainable model $\langle f, \varphi \rangle$ maps x to both a prediction and explanation. We show this in
 82 Figure 2, where $\langle f, \varphi \rangle(x) \in [0, 1]^m \times \{0, 1\}^n$ pairs the class probabilities and the feature attribution.

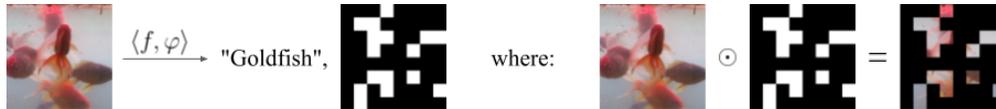


Figure 2: An explainable model $\langle f, \varphi \rangle$ outputs both a classification and a feature attribution. The feature attribution is a binary-valued mask (white 1, black 0) that can be applied over the original input. Here f is Vision Transformer [9] and φ is SHAP [7] with top-25% feature selection.

83 For an input $x \in \mathcal{X}$, we will evaluate the quality of the binary feature attribution $\varphi(x)$ through its
 84 masking on x . That is, we will study the behavior of f on the masked input $x \odot \varphi(x) \in \mathcal{X}$, where \odot
 85 is the element-wise vector product. To do this, we define a notion of *prediction equivalence*: for two
 86 $x, x' \in \mathcal{X}$, we write $f(x) \cong f(x')$ to mean that $f(x)$ and $f(x')$ yield the same class. This allows us
 87 to formalize the intuition that an explanation $\varphi(x)$ should recover the prediction of x under f .

88 **Definition 2.1.** *The explainable model $\langle f, \varphi \rangle$ is consistent at x if $f(x) \cong f(x \odot \varphi(x))$.*

89 Evaluating f on $x \odot \varphi(x)$ this way lets us apply the model as-is and therefore avoids the challenge
 90 of constructing a surrogate model that is accurate to the original [25]. Moreover, this approach is
 91 reasonable, especially in domains like vision — where one intuitively expects that a masked image
 92 retaining only the important features should induce the intended prediction. Indeed, architectures like
 93 Vision Transformer [9] can maintain high accuracy with only a fraction of the image present [26].

94 Particularly, we would like for $\langle f, \varphi \rangle$ to generate explanations that are stable and concise (i.e. sparse).
 95 The former is our central guarantee, and is ensured through smoothing. The latter implies that $\varphi(x)$
 96 has few ones entries, and is a desirable property since a good explanation should not contain too much
 97 redundant information. However, sparsity is a difficult property to enforce, as this is contingent on
 98 the model having high accuracy with respect to heavily masked inputs. For sparsity we present a
 99 simple heuristic in Section 2.4 and evaluate its effectiveness in Section 4.

100 2.2 Stability Properties of Explainable Models

101 Given an explainable model $\langle f, \varphi \rangle$ and some $x \in \mathcal{X}$, stability means that the prediction does not
 102 change even if one adds more explanatory features to $\varphi(x)$. For instance, the model-explanation pair
 103 in Figure 1 is *not* stable, as the inclusion of a single feature group (patch) changes the prediction. To
 104 formalize this notion of stability, we first introduce a partial ordering: for $\alpha, \alpha' \in \{0, 1\}^n$, we write
 105 $\alpha \succeq \alpha'$ iff $\alpha_i \geq \alpha'_i$ for all $i = 1, \dots, n$. That is, $\alpha \succeq \alpha'$ iff α includes all the features selected by α' .

106 **Definition 2.2.** *The explainable model $\langle f, \varphi \rangle$ is stable at x if $f(x \odot \alpha) \cong f(x \odot \varphi(x))$ for all
 107 $\alpha \succeq \varphi(x)$.*

108 Note that the constant explanation $\varphi(x) = \mathbf{1}$, the vector of ones, makes $\langle f, \varphi \rangle$ trivially stable at every
 109 $x \in \mathcal{X}$, though this is not a concise explanation. Additionally, stability at x implies consistency at x .

110 Unfortunately, stability is a difficult property to enforce in general, as it requires that f satisfy a
 111 monotone-like behavior with respect to feature inclusion — which is especially challenging for
 112 complex models like neural networks. Checking stability without additional assumptions on f is
 113 also hard: if $k = \|\varphi(x)\|_1$ is the number of ones in $\varphi(x)$, then there are 2^{n-k} possible $\alpha \succeq \varphi(x)$ to
 114 check. This large space of possible $\alpha \succeq \varphi(x)$ motivates us to instead examine *relaxations* of stability.
 115 We introduce lower and upper-relaxations of stability below.

116 **Definition 2.3.** *The explainable model $\langle f, \varphi \rangle$ is incrementally stable at x with radius r if $f(x \odot \alpha) \cong$
 117 $f(x \odot \varphi(x))$ for all $\alpha \succeq \varphi(x)$ where $\|\alpha - \varphi(x)\|_1 \leq r$.*

118 Incremental stability is the lower-relaxation since it considers the case where the mask α has only
 119 a few features more than $\varphi(x)$. For instance, if one can provably add up to r features to a masked
 120 $x \odot \varphi(x)$ without altering the prediction, then $\langle f, \varphi \rangle$ would be incremental stable at x with radius r .
 121 We next introduce the upper-relaxation that we call decremental stability.

122 **Definition 2.4.** *The explainable model $\langle f, \varphi \rangle$ is decrementally stable at x with radius r if $f(x \odot \alpha) \cong$
 123 $f(x \odot \varphi(x))$ for all $\alpha \succeq \varphi(x)$ where $\|\mathbf{1} - \alpha\|_1 \leq r$.*

124 Decremental stability is a subtractive form of stability, in contrast to the additive nature of incremental
 125 stability. Particularly, decremental stability considers the case where α has much more features than
 126 $\varphi(x)$. If one can provably remove up to r features from the full x without altering the prediction,
 127 then $\langle f, \varphi \rangle$ is decrementally stable at x with radius r . Note also that decremental stability necessarily
 128 entails consistency of $\langle f, \varphi \rangle$, but for simplicity of definitions we do not enforce this for incremental
 129 stability. Furthermore, observe that for a sufficiently large radius of $r = \lceil (n - \|\varphi(x)\|_1)/2 \rceil$,
 130 incremental and decremental stability together imply stability.

131 **Remark 2.5.** *Similar notions to the above have been proposed in literature, and we refer to [15] for
 132 an extensive survey. In particular for [15], consistency is akin to preservation and stability is similar
 133 to continuity, except we are concerned with adding features. In this regard, incremental stability is
 134 most similar to incremental addition and decremental stability to incremental deletion.*

135 2.3 Lipschitz Smoothness Entails Stability Guarantees

136 If $f : \mathcal{X} \rightarrow [0, 1]^m$ is Lipschitz with respect to the masking of features, then we can guarantee relaxed
 137 stability properties for the explainable model $\langle f, \varphi \rangle$. In particular, we require for all $x \in \mathcal{X}$ that
 138 $f(x \odot \alpha)$ is Lipschitz with respect to the mask $\alpha \in \{0, 1\}^n$. This then allows us to establish our
 139 main result (Theorem 3.3), which we preview below in Remark 2.6.

140 **Remark 2.6** (Sketch of main result). *Consider an explainable model $\langle f, \varphi \rangle$ where for all $x \in \mathcal{X}$ the
 141 function $g(x, \alpha) = f(x \odot \alpha)$ is λ -Lipschitz in $\alpha \in \{0, 1\}^n$ with respect to the ℓ^1 norm. Then at any
 142 x , the radius of incremental stability r_{inc} and radius of decremental stability r_{dec} are respectively:*

$$r_{\text{inc}} = \frac{g_A(x, \varphi(x)) - g_B(x, \varphi(x))}{2\lambda}, \quad r_{\text{dec}} = \frac{g_A(x, \mathbf{1}) - g_B(x, \mathbf{1})}{2\lambda},$$

143 where $g_A - g_B$ is referred to as the logit gap, with g_A, g_B the first and second-largest logits:

$$k^* = \operatorname{argmax}_{1 \leq k \leq m} g_k(x, \alpha), \quad g_A(x, \alpha) = g_{k^*}(x, \alpha), \quad g_B(x, \alpha) = \max_{i \neq k^*} g_i(x, \alpha). \quad (1)$$

144 Observe that Lipschitz smoothness is in fact a stronger assumption than necessary, as besides
 145 $\alpha \succeq \varphi(x)$ it also imposes guarantees on $\alpha \preceq \varphi(x)$. Nevertheless, Lipschitz smoothness is one of the
 146 few classes of properties that can be guaranteed and analyzed at scale on arbitrary models [21, 27].

147 2.4 Adapting Existing Feature Attribution Methods

148 Most existing feature attribution methods assign a real-valued score to feature importance, rather
 149 than a binary value. We therefore need to convert this to a binary-valued method for use with a
 150 stable explainable model. Let $\psi : \mathcal{X} \rightarrow \mathbb{R}^n$ be such a continuous-valued method like LIME [6] or
 151 SHAP [7], and fix some desired incremental stability radius r_{inc} and decremental stability radius r_{dec} .
 152 Given some $x \in \mathcal{X}$ a simple construction for binary $\varphi(x) \in \{0, 1\}^n$ is described next.

153 **Remark 2.7** (Iterative construction of $\varphi(x)$). *Consider any $x \in \mathcal{X}$ and let ρ be an index ordering on
 154 $\psi(x)$ from high-to-low (i.e. largest logit first). Initialize $\alpha = \mathbf{0}$, and for each $i \in \rho$: assign $\alpha_i \leftarrow 1$
 155 then check whether $\langle f, \varphi : x \mapsto \alpha \rangle$ is now consistent, incrementally stable with radius r_{inc} , and
 156 decrementally stable with radius r_{dec} . If so then terminate with $\varphi(x) = \alpha$; otherwise continue.*

157 Note that the above method of constructing $\varphi(x)$ does not impose sparsity guarantees in the way that
 158 we may guarantee stability through Lipschitz smoothness. Instead, the ordering from a continuous-
 159 valued $\psi(x)$ serves as a greedy heuristic for constructing $\varphi(x)$. We show in Section 4 that some
 160 feature attributions (e.g. SHAP [7]) tend to yield sparser selections on average compared to others
 161 (e.g. Vanilla Gradient Saliency [5]).

162 **3 Multiplicative Smoothing for Lipschitz Constants**

163 In this section we present our main technical contribution in Multiplicative Smoothing (MuS).
 164 The goal is to transform an arbitrary base classifier $h : \mathcal{X} \rightarrow [0, 1]^m$ into a smoothed classifier
 165 $f : \mathcal{X} \rightarrow [0, 1]^m$ that is Lipschitz with respect to the masking of features. This then allows one to
 166 appropriately couple an explanation method φ with f in order to form an explainable model $\langle f, \varphi \rangle$
 167 with provable stability guarantees.

168 We give an overview of our MuS in Section 3.1, where we illustrate a principal motivation for its
 169 development is because standard smoothing techniques may violate an property that we call *masking*
 170 *equivalence*. We present the Lipschitz constant of the smoothed classifier f in Section 3.2 and show
 171 how this is used to certify stability. Finally we give an efficient computation of MuS in Section 3.3,
 172 allowing us to exactly evaluate f at a low sample complexity.

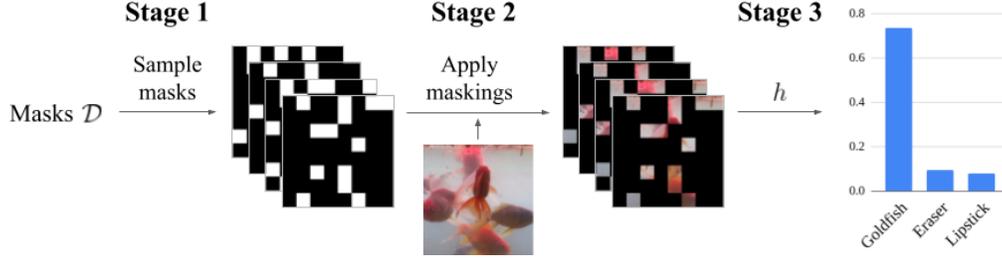


Figure 3: Evaluating $f(x)$ is done in three stages. **(Stage 1)** Generate N samples of binary masks $s^{(1)}, \dots, s^{(N)} \in \{0, 1\}^n$, where each coordinate is Bernoulli with parameter λ (here $\lambda = 1/4$). **(Stage 2)** Apply each mask on the input to yield $x \odot s^{(i)}$ for $i = 1, \dots, N$. **(Stage 3)** Average over $h(x \odot s^{(i)})$ to compute $f(x)$, and note that the predicted class is given by a weighted average.

173 **3.1 Technical Overview of MuS**

174 Our key insight is that randomly dropping (i.e. zeroing) features attains the desired smoothness. In
 175 particular, we uniformly drop features with probability $1 - \lambda$ by sampling binary masks $s \in \{0, 1\}^n$
 176 from some distribution \mathcal{D} where each coordinate is distributed as $\Pr[s_i = 1] = \lambda$. Then define f as:

$$f(x) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot s), \quad \text{such that } s_i \sim \mathcal{B}(\lambda) \text{ for } i = 1, \dots, n \quad (2)$$

177 where $\mathcal{B}(\lambda)$ is the Bernoulli distribution with parameter $\lambda \in [0, 1]$. We give an overview of evaluating
 178 $f(x)$ in Figure 3. Importantly, our main results (Theorem 3.2, Theorem 3.3) hold provided \mathcal{D}
 179 is coordinate-wise Bernoulli with λ , and so we avoid restricting ourselves to any one particular choice
 180 until necessary. However, it will be easy to intuit the exposition with $\mathcal{D} = \mathcal{B}^n(\lambda)$, the coordinate-wise
 181 i.i.d. Bernoulli distribution with λ .

182 We can equivalently parametrize f using the mapping $g(x, \alpha) = f(x \odot \alpha)$, where it follows that:

$$g(x, \alpha) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot \tilde{\alpha}), \quad \tilde{\alpha} = \alpha \odot s. \quad (3)$$

183 Note that one could have alternatively first defined g and then f due to the identity $g(x, \mathbf{1}) = f(x)$.
 184 We require that the relationship between f and g follows an identity that we call *masking equivalence*:

$$g(x \odot \alpha, \mathbf{1}) = f(x \odot \alpha) = g(x, \alpha), \quad \text{for all } x \in \mathcal{X} \text{ and } \alpha \in \{0, 1\}^n. \quad (4)$$

185 This follows by definition of g , and the relevance to stability is this: if masking equivalence holds,
 186 then we can rewrite stability properties involving f in terms of g 's second parameter as follows:

$$f(x \odot \alpha) = g(x, \alpha) \cong g(x, \varphi(x)) = f(x \odot \varphi(x)), \quad \text{for all } \alpha \succeq \varphi(x), \quad (\text{c.f. Definition 2.2})$$

187 where incremental and decremental stability may be analogously defined. This translation is useful,
 188 as we will prove that g is λ -Lipschitz in its second parameter (Theorem 3.2), which then allows us to
 189 establish the desired stability properties (Theorem 3.3).

190 Observe that we have not given an exact construction for \mathcal{D} , since many choices are in fact valid.
 191 Rather, so long as each coordinate of $s \sim \mathcal{D}$ obeys $s_i \sim \mathcal{B}(\lambda)$ then the Lipschitz properties for g

192 follow. The implication here is that although simple distributions like $\mathcal{B}^n(\lambda)$ suffices for \mathcal{D} , they may
 193 not be sample efficient. We show in Section 3.3 how to exploit a structured statistical dependence in
 194 order to reduce the sample complexity of computing MuS.

195 Importantly, we are motivated to develop MuS because standard smoothing techniques, namely
 196 additive smoothing [20, 21], may fail to satisfy masking equivalence. Additive smoothing is by far
 197 the most popular smoothing technique, and differs from our scheme (3) in how noise is applied,
 198 where let \mathcal{D}_{add} and $\mathcal{D}_{\text{mult}}$ be any two distributions on \mathbb{R}^n :

$$g(x, \alpha) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot \tilde{\alpha}), \quad \tilde{\alpha} = \begin{cases} \alpha + s, & s \sim \mathcal{D}_{\text{add}}, & \text{additive noise} \\ \alpha \odot s, & s \sim \mathcal{D}_{\text{mult}}, & \text{multiplicative noise} \end{cases}$$

199 Particularly, additive smoothing has counterexamples to masking equivalence.

200 **Proposition 3.1.** *There exists $h : \mathcal{X} \rightarrow [0, 1]$ and distribution \mathcal{D} , where for*

$$g^+(x, \alpha) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot \tilde{\alpha}), \quad \tilde{\alpha} = \alpha + s,$$

201 *we have $g^+(x, \alpha) \neq g^+(x \odot \alpha, \mathbf{1})$ for some $x \in \mathcal{X}$ and $\alpha \in \{0, 1\}^n$.*

202 *Proof.* Observe that it suffices to have h, x, α such that $h(x \odot (\alpha + s)) > h((x \odot \alpha) \odot (\mathbf{1} + s))$ for
 203 a non-empty set of $s \in \mathbb{R}^n$. Let \mathcal{D} be a distribution on these s , then:

$$g^+(x, \alpha) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot (\alpha + s)) > \mathbb{E}_{s \sim \mathcal{D}} h((x \odot \alpha) \odot (\mathbf{1} + s)) = g^+(x \odot \alpha, \mathbf{1})$$

204 □

205 Intuitively, this occurs because additive smoothing primarily applies noise by perturbing feature
 206 values, rather than completely masking them. As such, there might be “information leakage” when
 207 non-explanatory bits of α are changed into non-zero values. This then causes each sample of $h(x \odot \tilde{\alpha})$
 208 within $g(x, \alpha)$ to observe more features of x than it would have been able to otherwise.

209 3.2 Certifying Stability Properties with Lipschitz Classifiers

210 Our core technical result is in showing that f as defined in (2) is Lipschitz to the masking of features.
 211 We present MuS in terms of g , where it is parametric with respect to the distribution \mathcal{D} : so long as \mathcal{D}
 212 satisfies a coordinate-wise Bernoulli condition, then it is usable with MuS.

213 **Theorem 3.2 (MuS).** *Let \mathcal{D} be any distribution on $\{0, 1\}^n$ where each coordinates of $s \sim \mathcal{D}$ is
 214 distributed as $s_i \sim \mathcal{B}(\lambda)$. Consider any $h : \mathcal{X} \rightarrow [0, 1]$ and define $g : \mathcal{X} \times \{0, 1\}^n \rightarrow [0, 1]$ as*

$$g(x, \alpha) = \mathbb{E}_{s \sim \mathcal{D}} h(x \odot \tilde{\alpha}), \quad \tilde{\alpha} = \alpha \odot s.$$

215 *Then the function $g(x, \cdot) : \{0, 1\}^n \rightarrow [0, 1]$ is λ -Lipschitz in the ℓ^1 norm for all $x \in \mathcal{X}$.*

216 The strength of this result is in its weak assumptions. First, the theorem applies to any model h and
 217 input $x \in \mathcal{X}$. It further suffices that each coordinate is distributed as $s_i \sim \mathcal{B}(\lambda)$, and we emphasize
 218 that statistical independence between different s_i, s_j is *not assumed*. This allows us to construct
 219 \mathcal{D} with structured dependence in Section 3.3, such that we may exactly and efficiently evaluate
 220 $g(x, \alpha)$ at a sample complexity of $N \ll 2^n$. A low sample complexity is important for making MuS
 221 practically usable, as otherwise one must settle for of the expected value subject to probabilistic
 222 guarantees. For instance, simpler distributions like $\mathcal{B}^n(\lambda)$ do in fact satisfy the requirements of
 223 Theorem 3.2 — but costs 2^n samples because of coordinate-wise independence.

224 Whatever choice of \mathcal{D} , one can guarantee stability so long as g is Lipschitz in its second argument.

225 **Theorem 3.3 (Stability).** *Consider any $h : \mathcal{X} \rightarrow [0, 1]^m$ with coordinates h_1, \dots, h_m . Fix $\lambda \in [0, 1]$
 226 and let g_1, \dots, g_m be the respectively smoothed coordinates as in Theorem 3.2, using which we
 227 analogously define $g : \mathcal{X} \times \{0, 1\}^n \rightarrow [0, 1]^m$. Also define $f(x) = g(x, \mathbf{1})$. Then for any explanation
 228 method φ and input $x \in \mathcal{X}$, the explainable model $\langle f, \varphi \rangle$ is incrementally stable with radius r_{inc} and
 229 decrementally stable with radius r_{dec} :*

$$r_{\text{inc}} = \frac{g_A(x, \varphi(x)) - g_B(x, \varphi(x))}{2\lambda}, \quad r_{\text{dec}} = \frac{g_A(x, \mathbf{1}) - g_B(x, \mathbf{1})}{2\lambda},$$

230 *where g_A, g_B are the first and second largest logits of g as in (1).*

231 Note that it is only in the case where the radius ≥ 1 do non-trivial stability guarantees exist. Because
 232 each g_k has range in $[0, 1]$, this means that a Lipschitz constant of $\lambda \leq 1/2$ is necessary to attain at
 233 least one radius of stability. We present in Appendix A.2 some extensions to MuS that allows one to
 234 achieve higher coverage of features.

235 3.3 Exploiting Structured Dependency

236 We now present $\mathcal{L}_{qv}(\lambda)$, a distribution on $\{0, 1\}^n$ that allows for efficient and exact evaluation of a
 237 MuS-smoothed classifier. Our construction is an adaption of [27] from uniform to Bernoulli noise,
 238 where the primary insight is that one can parametrize n -dimensional noise using a single dimension
 239 via structured coordinate-wise dependence. In particular, we use a *seed vector* v , where with an
 240 integer *quantization parameter* $q > 1$ there will only exist q distinct choices of $s \sim \mathcal{L}_{qv}(\lambda)$. All
 241 the while, we still enforce that any such s is coordinate-wise Bernoulli with $s_i \sim \mathcal{B}(\lambda)$. Thus for a
 242 sufficiently small quantization parameter (i.e. $q \ll 2^n$) we may tractably enumerate through all q
 243 possible choices of s and thereby evaluate a MuS-smoothed model with only q samples.

244 **Proposition 3.4.** Fix integer $q > 1$ and consider any vector $v \in \{0, 1/q, \dots, (q-1)/q\}^n$ and scalar
 245 $\lambda \in \{1/q, \dots, q/q\}$. Define $s \sim \mathcal{L}_{qv}(\lambda)$ to be a random vector in $\{0, 1\}^n$ with coordinates given by

$$s_i = \mathbb{I}[t_i \leq \lambda], \quad t_i = v_i + s_{\text{base}} \bmod 1, \quad s_{\text{base}} \sim \mathcal{U}(\{1/q, \dots, q/q\}) - 1/(2q).$$

246 Then there are q distinct values of s and each coordinate is distributed as $s_i \sim \mathcal{B}(\lambda)$.

247 *Proof.* First, observe that each of the q distinct values of s_{base} defines a unique value of s , since
 248 we have assumed v and λ to be fixed. Next, observe that each t_i has q unique values uniformly
 249 distributed as $t_i \sim \mathcal{U}(1/q, \dots, q/q) - 1/(2q)$. Because $\lambda \in \{1/q, \dots, q/q\}$ we therefore have
 250 $\Pr[t_i \leq \lambda] = \lambda$, which implies that $s_i \sim \mathcal{B}(\lambda)$. \square

251 The seed vector v is the source of our structured coordinate-wise dependence and the one-dimensional
 252 source of randomness s_{base} is used to generate the n -dimensional s . Such $s \sim \mathcal{L}_{qv}(\lambda)$ then satisfies
 253 the conditions for use in MuS (Theorem 3.2), and this noise allows for an exact evaluation of the
 254 smoothed classifier in q samples. We have found $q = 64$ to be sufficient in practice and values as
 255 low as $q = 16$ to also yield good performance. We remark that one drawback is that one may get an
 256 unlucky seed v , but we have not yet observed this in our experiments.

257 4 Empirical Evaluations

258 We evaluate the quality of MuS on different classification models and explanation methods as they
 259 relate to stability guarantees. To that end, we perform the following experiments.

260 **(E1) How good are the stability guarantees?** There exists a natural measure of quality for stability
 261 guarantees over a dataset: what radii are achieved, and at what frequency. We investigate how
 262 different combinations of models, explanation methods, and λ affect this measure.

263 **(E2) What is the cost of smoothing?** To increase the radius of a provable stability guarantee, we
 264 must decrease the Lipschitz constant λ . As λ decreases, however, more features are dropped during
 265 the smoothing process. This experiment investigates this stability-accuracy trade-off.

266 **(E3) Which explanation method is best?** We evaluate which existing feature attribution methods
 267 are amenable to strong stability guarantees. We examine LIME [6], SHAP [7], Vanilla Gradient
 268 Saliency (VGrad) [5], and Integrated Gradients (IGrad) [8], with focus on the size of the explanation.

269 **(Experimental Setup)** We use on two vision models (Vision Transformer [9] and ResNet50 [28])
 270 and one language model (RoBERTa [29]). For the vision dataset we use ImageNet1K [30] and
 271 for the language dataset we use TweetEval [31] sentiment analysis. We use *feature grouping* from
 272 Appendix A.2.1 on ImageNet1K to reduce the $3 \times 224 \times 224$ dimensional input into $n = 64$ superpixel
 273 patches. We report stability radii r in terms of fraction of features, i.e. r/n . In all our experiments
 274 we use the quantized noise as in Section 3.3 with $q = 64$ unless specified otherwise. We refer to
 275 Appendix B for training details and the comprehensive experiments.

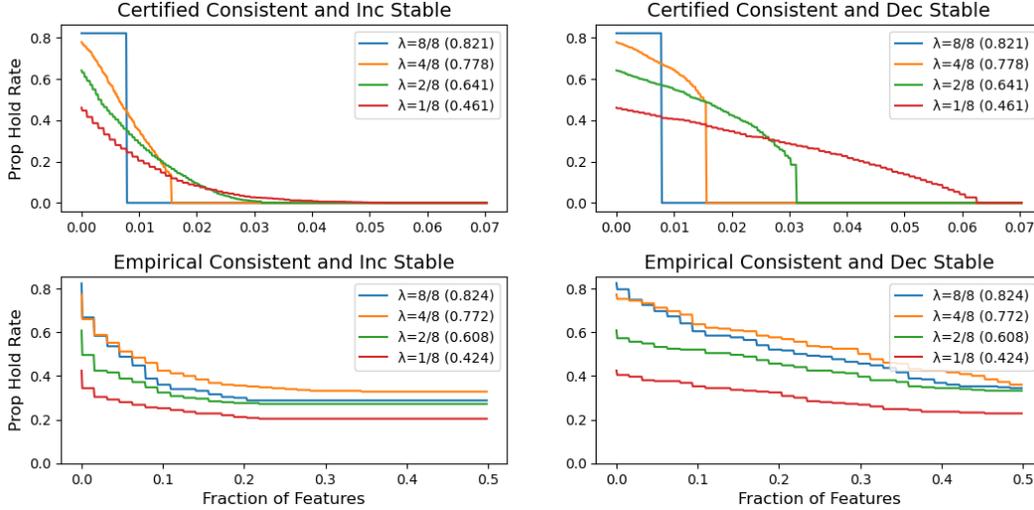


Figure 4: Rate of consistency and incremental (decremental) stability up to radius r vs. fraction of feature coverage r/n . Left: certified $N_{\text{cert}} = 2000$; Right: empirical $N_{\text{emp}} = 250$ with $q = 16$.

276 4.1 (E1) Quality of Stability Guarantees

277 We study how much radius of consistent and incremental (resp. decremental) stability is achieved,
 278 and how often. We take an explainable model $\langle f, \varphi \rangle$ where f is Vision Transformer and φ is SHAP
 279 with top-25% feature selection. We plot the rate at which a property holds (e.g. consistent and
 280 incrementally stable with radius r) as a function of radius (expressed as a fraction of features r/n).

281 We show our results in Figure 4, where on the left we have the certified guarantees for $N_{\text{cert}} = 2000$
 282 samples from ImageNet1K; on the right we have the empirical radii for $N_{\text{emp}} = 250$ samples
 283 obtained by applying a standard box attack [32] strategy with $q = 16$. We observe for the certified
 284 results that the decremental stability have larger radii than the incremental stability. This is reasonable
 285 since the base classifier sees much more of the input when analyzing decremental stability, and can
 286 therefore be more confident on average, i.e. achieve a larger logit gap. Moreover, our empirical radii
 287 often cover up to half of the input, which suggests that our certified analysis is quite conservative.

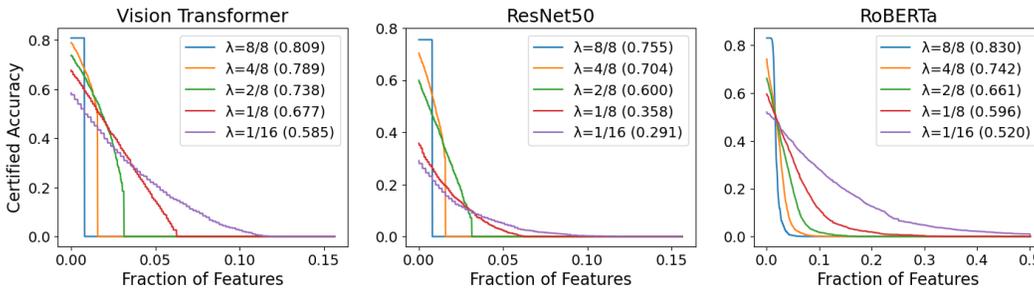


Figure 5: Certified accuracy vs. decremental stability radius. $N = 2000$.

288 4.2 (E2) Stability-Accuracy Trade-Offs

289 We next investigate how smoothing impacts the classifier accuracy. As λ decreases due to more
 290 smoothing, the base classifier sees increasingly zeroed out features — which should hurt accuracy.
 291 We took $N = 2000$ samples for each classifier on their respective datasets and plotted the certified
 292 accuracy vs. radius of decremental stability.

293 We show the results in Figure 5, where as expected the clean accuracy (in parentheses) decreases with
 294 λ . This accuracy drop is especially pronounced for ResNet50, and we suspect that the transformer

295 architecture of Vision Transformer and RoBERTa make them more resilient to the randomized
 296 masking of features. Nevertheless, this experiment demonstrates that large models, especially
 297 transformers, can tolerate non-trivial noise from MuS while maintaining high accuracy.

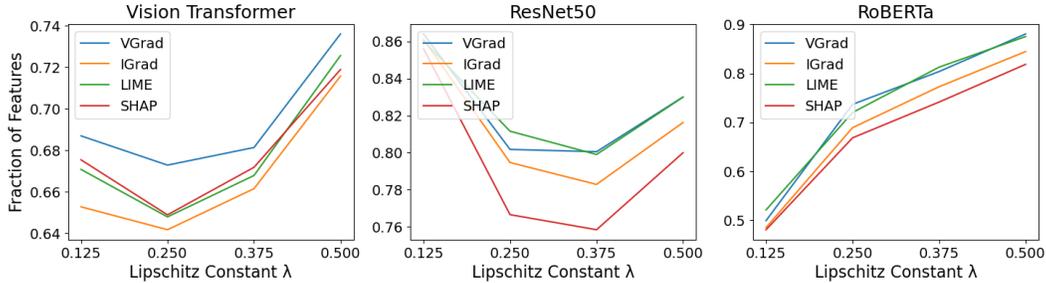


Figure 6: Average k/n vs. λ , where $k = \|\varphi(x)\|_1$ is the number of features for $\langle f, \varphi \rangle$ to be consistent, incrementally stable with radius 1, and decrementally stable with radius 1. $N = 250$.

298 4.3 (E3) Which Explanation Method to Pick?

299 Finally, we explore which feature attribution method is best suited to stability guarantees of explain-
 300 able model $\langle f, \varphi \rangle$. All four methods $\psi \in \{\text{LIME}, \text{SHAP}, \text{VGrad}, \text{IGrad}\}$ are continuous-valued, for
 301 which we samples $N = 250$ inputs from each model’s respective dataset. For each input x we use
 302 the feature importance ranking generated by $\psi(x)$ to iteratively build $\varphi(x)$ in a greedy manner like in
 303 Section 2.4. For some x , let $k_x = \varphi(x)/n$ be the number fraction of features needed for $\langle f, \varphi \rangle$
 304 to be consistent, incrementally stable with radius 1, and decrementally stable with radius 1. We then
 305 plot the average k_x for different methods at $\lambda \in \{1/8, \dots, 4/8\}$ in Figure 6, where note that SHAP
 306 tends to require fewer features to achieve the desired properties, while VGrad tends to require more.
 307 However, we do not believe these to be decisive results, as many curves are relatively close, especially
 308 for Vision Transformer and ResNet50.

309 5 Related Works

310 For extensive surveys on explainability methods see [15, 19, 33, 34, 35, 36, 37]. Notable feature
 311 attribution methods include Vanilla Gradient Saliency [5], SmoothGrad [22], Integrated Gradients [8],
 312 Grad-CAM [38], Occlusion [39], LIME [6], SHAP [7], and their variants. Of these, Shapley
 313 valued [16] based methods [7, 23, 24] are rooted in axiomatic principles, as are Integrated Gradients [8,
 314 40]. The work of [41] finds confidence intervals over attribution scores. A study of common
 315 feature attribution methods is done in [42]. Similar to our approach is [43], which studies binary-
 316 valued classifiers and presents an algorithm with succinctness and probabilistic precision guarantees.
 317 Different metrics for evaluating feature attributions are studied in [15, 17, 44, 45, 46, 47, 48, 49, 50].
 318 Whether an attribution correctly identifies relevant features is a well-known issue [51, 52]. Many
 319 methods are also susceptible to adversarial attacks [53, 54]. As a negative result, [55] shows that
 320 feature attributions have provably poor performance on sufficiently rich model classes. Related to
 321 feature attributions are *data attributions* [56, 57, 58], which assigns values to training data points.

322 6 Conclusion

323 We study provable stability guarantees for binary feature attribution methods through the framework
 324 of explainable models. A selection of features is stable if the additional inclusion of other features do
 325 not alter its explanatory power. We show that if the classifier is Lipschitz with respect to the masking
 326 of features, then one can guarantee relaxed variants of stability. To achieve this Lipschitz condition
 327 we develop a smoothing method called Multiplicative Smoothing (MuS). This method is parametric
 328 to the choice of noise distribution, allowing us to construct and exploit distributions with structured
 329 dependence for exact and efficient evaluation. We evaluate MuS on vision and language models, and
 330 demonstrate that MuS yields strong stability guarantees at only a small cost to accuracy.

331 **References**

- 332 [1] Mauricio Reyes, Raphael Meier, Sérgio Pereira, Carlos A Silva, Fried-Michael Dahlweid,
333 Hendrik von Tengg-Kobligk, Ronald M Summers, and Roland Wiest. On the interpretability
334 of artificial intelligence in radiology: challenges and opportunities. *Radiology: artificial*
335 *intelligence*, 2(3):e190043, 2020.
- 336 [2] Erico Tjoa and Cuntai Guan. A survey on explainable artificial intelligence (xai): Toward
337 medical xai. *IEEE transactions on neural networks and learning systems*, 32(11):4793–4813,
338 2020.
- 339 [3] Sandra Wachter, Brent Mittelstadt, and Chris Russell. Counterfactual explanations without
340 opening the black box: Automated decisions and the gdpr. *Harv. JL & Tech.*, 31:841, 2017.
- 341 [4] Adrien Bibal, Michael Lognoul, Alexandre De Streel, and Benoît Frénay. Legal requirements
342 on explainability in machine learning. *Artificial Intelligence and Law*, 29:149–169, 2021.
- 343 [5] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks:
344 Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*,
345 2013.
- 346 [6] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining
347 the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international*
348 *conference on knowledge discovery and data mining*, pages 1135–1144, 2016.
- 349 [7] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions.
350 *Advances in neural information processing systems*, 30, 2017.
- 351 [8] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In
352 *International conference on machine learning*, pages 3319–3328. PMLR, 2017.
- 353 [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai,
354 Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al.
355 An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint*
356 *arXiv:2010.11929*, 2020.
- 357 [10] Amirata Ghorbani, Abubakar Abid, and James Zou. Interpretation of neural networks is fragile.
358 In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3681–3688,
359 2019.
- 360 [11] Gavin Brown. A new perspective for information theoretic feature selection. In *Artificial*
361 *intelligence and statistics*, pages 49–56. PMLR, 2009.
- 362 [12] Jianbo Chen, Le Song, Martin Wainwright, and Michael Jordan. Learning to explain: An
363 information-theoretic perspective on model interpretation. In *International Conference on*
364 *Machine Learning*, pages 883–892. PMLR, 2018.
- 365 [13] Xiaoping Li, Yadi Wang, and Rubén Ruiz. A survey on sparse learning models for feature
366 selection. *IEEE transactions on cybernetics*, 52(3):1642–1660, 2020.
- 367 [14] Akhilan Boopathy, Sijia Liu, Gaoyuan Zhang, Cynthia Liu, Pin-Yu Chen, Shiyu Chang, and
368 Luca Daniel. Proper network interpretability helps adversarial robustness in classification. In
369 *International Conference on Machine Learning*, pages 1014–1023. PMLR, 2020.
- 370 [15] Meike Nauta, Jan Trienes, Shreyasi Pathak, Elisa Nguyen, Michelle Peters, Yasmin Schmitt,
371 Jörg Schlotterer, Maurice van Keulen, and Christin Seifert. From anecdotal evidence to
372 quantitative evaluation methods: A systematic review on evaluating explainable ai. *arXiv*
373 *preprint arXiv:2201.08164*, 2022.
- 374 [16] L.S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, pages
375 307–317, 1953.
- 376 [17] Yilun Zhou, Serena Booth, Marco Tulio Ribeiro, and Julie Shah. Do feature attribution methods
377 correctly attribute features? In *Proceedings of the AAAI Conference on Artificial Intelligence*,
378 volume 36, pages 9623–9633, 2022.

- 379 [18] Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable nlp systems: How should we
380 define and evaluate faithfulness? In *Proceedings of the 58th Annual Meeting of the Association
381 for Computational Linguistics*, pages 4198–4205, 2020.
- 382 [19] Jundong Li, Kewei Cheng, Suhang Wang, Fred Morstatter, Robert P Trevino, Jiliang Tang, and
383 Huan Liu. Feature selection: A data perspective. *ACM computing surveys (CSUR)*, 50(6):1–45,
384 2017.
- 385 [20] Jeremy Cohen, Elan Rosenfeld, and Zico Kolter. Certified adversarial robustness via randomized
386 smoothing. In *International Conference on Machine Learning*, pages 1310–1320. PMLR, 2019.
- 387 [21] Greg Yang, Tony Duan, J Edward Hu, Hadi Salman, Ilya Razenshteyn, and Jerry Li. Randomized
388 smoothing of all shapes and sizes. In *International Conference on Machine Learning*, pages
389 10693–10705. PMLR, 2020.
- 390 [22] Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smooth-
391 grad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.
- 392 [23] Mukund Sundararajan and Amir Najmi. The many shapley values for model explanation. In
393 *International conference on machine learning*, pages 9269–9278. PMLR, 2020.
- 394 [24] Yongchan Kwon and James Zou. Weightedshap: analyzing and improving shapley based feature
395 attributions. *arXiv preprint arXiv:2209.13429*, 2022.
- 396 [25] Reza Alizadeh, Janet K Allen, and Farrokh Mistree. Managing computational complexity using
397 surrogate models: a critical review. *Research in Engineering Design*, 31:275–298, 2020.
- 398 [26] Hadi Salman, Saachi Jain, Eric Wong, and Aleksander Madry. Certified patch robustness via
399 smoothed vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision
400 and Pattern Recognition*, pages 15137–15147, 2022.
- 401 [27] Alexander J Levine and Soheil Feizi. Improved, deterministic smoothing for l₁ certified
402 robustness. In *International Conference on Machine Learning*, pages 6254–6264. PMLR, 2021.
- 403 [28] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
404 recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*,
405 pages 770–778, 2016.
- 406 [29] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike
407 Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining
408 approach. *arXiv preprint arXiv:1907.11692*, 2019.
- 409 [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng
410 Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual
411 recognition challenge. *International journal of computer vision*, 115:211–252, 2015.
- 412 [31] Francesco Barbieri, Jose Camacho-Collados, Leonardo Neves, and Luis Espinosa-Anke. Tweet-
413 eval: Unified benchmark and comparative evaluation for tweet classification. *arXiv preprint
414 arXiv:2010.12421*, 2020.
- 415 [32] Pin-Yu Chen, Huan Zhang, Yash Sharma, Jinfeng Yi, and Cho-Jui Hsieh. Zoo: Zeroth order
416 optimization based black-box attacks to deep neural networks without training substitute models.
417 In *Proceedings of the 10th ACM workshop on artificial intelligence and security*, pages 15–26,
418 2017.
- 419 [33] Nadia Burkart and Marco F Huber. A survey on the explainability of supervised machine
420 learning. *Journal of Artificial Intelligence Research*, 70:245–317, 2021.
- 421 [34] Jianlong Zhou, Amir H Gandomi, Fang Chen, and Andreas Holzinger. Evaluating the quality of
422 machine learning explanations: A survey on methods and metrics. *Electronics*, 10(5):593, 2021.
- 423 [35] Qing Lyu, Marianna Apidianaki, and Chris Callison-Burch. Towards faithful model explanation
424 in nlp: A survey. *arXiv preprint arXiv:2209.11326*, 2022.

- 425 [36] Xuhong Li, Haoyi Xiong, Xingjian Li, Xuanyu Wu, Xiao Zhang, Ji Liu, Jiang Bian, and Dejing
426 Dou. Interpretable deep learning: Interpretation, interpretability, trustworthiness, and beyond.
427 *Knowledge and Information Systems*, 64(12):3197–3234, 2022.
- 428 [37] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von
429 Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the
430 opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- 431 [38] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi
432 Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based
433 localization. In *Proceedings of the IEEE international conference on computer vision*, pages
434 618–626, 2017.
- 435 [39] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In
436 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September*
437 *6-12, 2014, Proceedings, Part I 13*, pages 818–833. Springer, 2014.
- 438 [40] Daniel D Lundstrom, Tianjian Huang, and Meisam Razaviyayn. A rigorous study of integrated
439 gradients method and extensions to internal neuron attributions. In *International Conference on*
440 *Machine Learning*, pages 14485–14508. PMLR, 2022.
- 441 [41] Dylan Slack, Anna Hilgard, Sameer Singh, and Himabindu Lakkaraju. Reliable post hoc
442 explanations: Modeling uncertainty in explainability. *Advances in neural information processing*
443 *systems*, 34:9391–9404, 2021.
- 444 [42] Tessa Han, Suraj Srinivas, and Himabindu Lakkaraju. Which explanation should i choose?
445 a function approximation perspective to characterizing post hoc explanations. *arXiv preprint*
446 *arXiv:2206.01254*, 2022.
- 447 [43] Guy Blanc, Jane Lange, and Li-Yang Tan. Provably efficient, succinct, and precise explanations.
448 *Advances in Neural Information Processing Systems*, 34:6129–6141, 2021.
- 449 [44] Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim.
450 Sanity checks for saliency maps. *Advances in neural information processing systems*, 31, 2018.
- 451 [45] Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. A benchmark for in-
452 terpretability methods in deep neural networks. *Advances in neural information processing*
453 *systems*, 32, 2019.
- 454 [46] Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard
455 Socher, and Byron C Wallace. Eraser: A benchmark to evaluate rationalized nlp models. *arXiv*
456 *preprint arXiv:1911.03429*, 2019.
- 457 [47] Jasmijn Bastings, Sebastian Ebert, Polina Zablotskaia, Anders Sandholm, and Katja Filippova.
458 " will you find these shortcuts?" a protocol for evaluating the faithfulness of input salience
459 methods for text classification. *arXiv preprint arXiv:2111.07367*, 2021.
- 460 [48] Yao Rong, Tobias Leemann, Vadim Borisov, Gjergji Kasneci, and Enkelejda Kasneci. A consis-
461 tent and efficient evaluation strategy for attribution methods. *arXiv preprint arXiv:2202.00449*,
462 2022.
- 463 [49] Yilun Zhou and Julie Shah. The solvability of interpretability evaluation metrics. *arXiv preprint*
464 *arXiv:2205.08696*, 2022.
- 465 [50] Julius Adebayo, Michael Muelly, Harold Abelson, and Been Kim. Post hoc explanations may
466 be ineffective for detecting unknown spurious correlation. In *International Conference on*
467 *Learning Representations*, 2022.
- 468 [51] Mengjiao Yang and Been Kim. Benchmarking attribution methods with relative feature impor-
469 tance. *arXiv preprint arXiv:1907.09701*, 2019.
- 470 [52] Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven
471 Dähne, Dumitru Erhan, and Been Kim. The (un) reliability of saliency methods. In *Explainable*
472 *AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 267–280. Springer, 2019.

- 473 [53] Dylan Slack, Anna Hilgard, Himabindu Lakkaraju, and Sameer Singh. Counterfactual expla-
474 nations can be manipulated. *Advances in neural information processing systems*, 34:62–75,
475 2021.
- 476 [54] Adam Ivankay, Ivan Girardi, Chiara Marchiori, and Pascal Frossard. Fooling explanations in
477 text classifiers. *arXiv preprint arXiv:2206.03178*, 2022.
- 478 [55] Blair Bilodeau, Natasha Jaques, Pang Wei Koh, and Been Kim. Impossibility theorems for
479 feature attribution. *arXiv preprint arXiv:2212.11870*, 2022.
- 480 [56] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions.
481 In *International conference on machine learning*, pages 1885–1894. PMLR, 2017.
- 482 [57] Andrew Ilyas, Sung Min Park, Logan Engstrom, Guillaume Leclerc, and Aleksander Madry.
483 Datamodels: Predicting predictions from training data. *arXiv preprint arXiv:2202.00622*, 2022.
- 484 [58] Sung Min Park, Kristian Georgiev, Andrew Ilyas, Guillaume Leclerc, and Aleksander Madry.
485 Trak: Attributing model behavior at scale. *arXiv preprint arXiv:2303.14186*, 2023.
- 486 [59] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint*
487 *arXiv:1412.6980*, 2014.

488 **A Proofs and Theoretical Discussions**

489 Here we present the proofs of our main results, as well as some extensions to MuS.

490 **A.1 Proofs of Main Results**

491 **A.1.1 Proof of Theorem 3.2**

492 *Proof.* By linearity we have:

$$g(x, \alpha) - g(x, \alpha') = \mathbb{E}_{s \sim \mathcal{D}} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}')], \quad \tilde{\alpha} = \alpha \odot s, \quad \tilde{\alpha}' = \alpha' \odot s,$$

493 so it suffices to analyze an arbitrary term by fixing some $s \sim \mathcal{D}$. Consider any $x \in \mathcal{X}$, let
 494 $\alpha, \alpha' \in \{0, 1\}^n$, and define $\delta = \alpha - \alpha'$. Observe that $\tilde{\alpha}_i \neq \tilde{\alpha}'_i$ exactly when $|\delta_i| = 1$ and $s_i = 1$.
 495 Since $s_i \sim \mathcal{B}(\lambda)$, we thus have $\Pr[\tilde{\alpha}_i \neq \tilde{\alpha}'_i] = \lambda|\delta_i|$, and applying the union bound:

$$\Pr_{s \sim \mathcal{D}} [\tilde{\alpha} \neq \tilde{\alpha}'] = \Pr_{s \sim \mathcal{D}} \left[\bigcup_{i=1}^n \tilde{\alpha}_i \neq \tilde{\alpha}'_i \right] \leq \sum_{i=1}^n \lambda|\delta_i| = \lambda \|\delta\|_1,$$

496 and so:

$$\begin{aligned} |g(x, \alpha) - g(x, \alpha')| &= \left| \mathbb{E}_{s \sim \mathcal{D}} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}')] \right| \\ &= \left| \Pr_{s \sim \mathcal{D}} [\tilde{\alpha} \neq \tilde{\alpha}'] \cdot \mathbb{E}_{s \sim \mathcal{D}} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}') \mid \tilde{\alpha} \neq \tilde{\alpha}'] \right. \\ &\quad \left. - \Pr_{s \sim \mathcal{D}} [\tilde{\alpha} = \tilde{\alpha}'] \cdot \mathbb{E}_{s \sim \mathcal{D}} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}') \mid \tilde{\alpha} = \tilde{\alpha}'] \right|. \end{aligned}$$

497 Note that $\mathbb{E} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}') \mid \tilde{\alpha} = \tilde{\alpha}'] = 0$, and so

$$\begin{aligned} |g(x, \alpha) - g(x, \alpha')| &= \Pr_{s \sim \mathcal{D}} [\tilde{\alpha} \neq \tilde{\alpha}'] \cdot \underbrace{\left| \mathbb{E}_{s \sim \mathcal{D}} [h(x \odot \tilde{\alpha}) - h(x \odot \tilde{\alpha}') \mid \tilde{\alpha} \neq \tilde{\alpha}'] \right|}_{\leq 1 \text{ because } h(\cdot) \in [0, 1]} \\ &\leq \Pr_{s \sim \mathcal{D}} [\tilde{\alpha} \neq \tilde{\alpha}'] \\ &\leq \lambda \|\delta\|_1. \end{aligned}$$

498 Thus, $g(x, \cdot)$ is λ -Lipschitz in the ℓ^1 norm. □

499 **A.1.2 Proof of Theorem 3.3**

500 *Proof.* We first show incremental stability. Consider any $x \in \mathcal{X}$, then by masking equivalence:

$$f(x \odot \varphi(x)) = g(x \odot \varphi(x), \mathbf{1}) = g(x, \varphi(x)),$$

501 and let g_A, g_B be the top two logits of g as defined in (1). By Theorem 3.2, both g_A, g_B are Lipschitz
 502 in their second parameter, and so for all $\alpha \in \{0, 1\}^n$:

$$\begin{aligned} \|g_A(x, \varphi(x)) - g_A(x, \alpha)\|_1 &\leq \lambda \|\varphi(x) - \alpha\|_1 \\ \|g_B(x, \varphi(x)) - g_B(x, \alpha)\|_1 &\leq \lambda \|\varphi(x) - \alpha\|_1 \end{aligned}$$

503 Observe that if α is sufficiently close to $\varphi(x)$, i.e.:

$$2\lambda \|\varphi(x) - \alpha\|_1 \leq g_A(x, \varphi(x)) - g_B(x, \varphi(x)),$$

504 then the top logit index of $g(x, \varphi(x))$ and $g(x, \alpha)$ are the same. This means that $g(x, \varphi(x)) \cong g(x, \alpha)$
 505 and thus $f(x \odot \varphi(x)) \cong f(x \odot \alpha)$, thus proving incremental stability with radius $d(x, \varphi(x))/(2\lambda)$.

506 The decremental stability case is similar, except we replace $\varphi(x)$ with $\mathbf{1}$. □

507 **A.2 Some Basic Extensions**

508 Below we present some extensions to MuS that help increase the fraction of the input to which we
 509 can guarantee stability.

510 **A.2.1 Feature Grouping**

511 We have so far assumed that $\mathcal{X} = \mathbb{R}^n$, but sometimes it may be desirable to group features together,
512 e.g. color channels of the same pixel. Our results also hold for more general $\mathcal{X} = \mathbb{R}^{d_1} \times \dots \times \mathbb{R}^{d_n}$,
513 where for such $x \in \mathcal{X}$ and $\alpha \in \mathbb{R}^n$ we lift \odot as

$$\odot : \mathcal{X} \times \mathbb{R}^n \rightarrow \mathcal{X}, \quad (x \odot \alpha)_i = x_i \cdot \mathbb{I}[\alpha_i = 1] \in \mathbb{R}^{d_i}.$$

514 All of our proofs are identical under this construction, with the exception of the dimensionalities of
515 terms like $(x \odot \alpha)$. An example of feature grouping is given in Figure 1.

516 **B All Experiments**

517 **Models, Datasets, and Explanation Methods** We evaluate on two vision models (Vision Trans-
518 former [9] and ResNet50 [28]) and one language model (RoBERTa [29]). For the vision dataset we
519 use ImageNet1K [30] and for the language dataset we use TweetEval [31] sentiment analysis. We
520 use four explanation methods in SHAP [7], LIME [6], Integrated Gradients (IGrad) [8], and Vanilla
521 Gradient Saliency (VGrad) [5]; where we take $\varphi(x)$ as the top- k weighted features.

522 **Training Details** We use Adam [59] as our optimizer with default parameters and learning rate
523 10^{-6} . For each $\lambda \in \{1/8, \dots, 8/8\}$ we fine-tuned each model for 1 epoch, which results in a total of
524 $8 \times 3 = 24$ models used in our experiments. To train with a particular λ : for each training input x we
525 generate two random maskings — one where λ of the features are zeroed and one where $\lambda/2$ of the
526 features are zeroed. This additional $\lambda/2$ zeroing is to account for the fact that inputs to a smoothed
527 model will be subject to masking by λ as well as $\varphi(x)$, where the scaling factor of $1/2$ is informed
528 by our prior experience about the size of a stable explanation.

529 **Miscellaneous Preprocessing** For images in ImageNet1K we use feature grouping (Section A.2.1)
530 to group the $3 \times 224 \times 224$ dimensional image into patches of size $3 \times 28 \times 28$, such that there
531 remains $n = 64$ feature groups. Each feature of a feature group then receives the same value of
532 noise during smoothing. We report radii of stability as a fraction of the feature groups covered. For
533 example, if at some input from ImageNet1K we get an incremental stability radius of r , then we
534 report $r/64$ as the fraction of features up to which we are guaranteed to be stable. This is especially
535 amenable to evaluating RoBERTa on TweetEval where inputs do not have uniform token lengths, i.e.
536 do not have uniform feature dimensions. In all of our experiments we use the quantized noise as in
537 Section 3.3 with a quantization parameter of $q = 64$, with the exception of Appendix B.2 where for
538 the box attack search we use $q = 16$.

539 Our experiments are organized as follows:

- 540 • (Section B.1) What is the quality of stability guarantees?
- 541 • (Section B.2) What is the theoretical vs empirical stability that can be guaranteed?
- 542 • (Section B.3) What are the stability-accuracy trade-offs?
- 543 • (Section B.4) Which explanation method is best?

544 **B.1 Quality of Stability Guarantees**

545 Here we study what radii of stability are certifiable, and how often these can be achieved with different
546 models and explanation methods. We therefore consider explainable models $\langle f, \varphi \rangle$ constructed
547 from base models $h \in \{\text{Vision Transformer, ResNet50, RoBERTa}\}$ and explanation methods $\varphi \in$
548 $\{\text{SHAP, LIME, IGrad, VGrad}\}$ with top- $k \in \{1/8, 2/8, 3/8\}$ feature selection. We take $N = 2000$
549 samples from each model’s respective datasets and compute the following value for each radius:

$$\text{value}(r) = \frac{\#\{x : \langle f, \varphi \rangle \text{ consistent and inc (dec) stable with radius } \leq r\}}{N}.$$

550 Plots of incremental stability are on the left; plots of decremental stability are on the right.

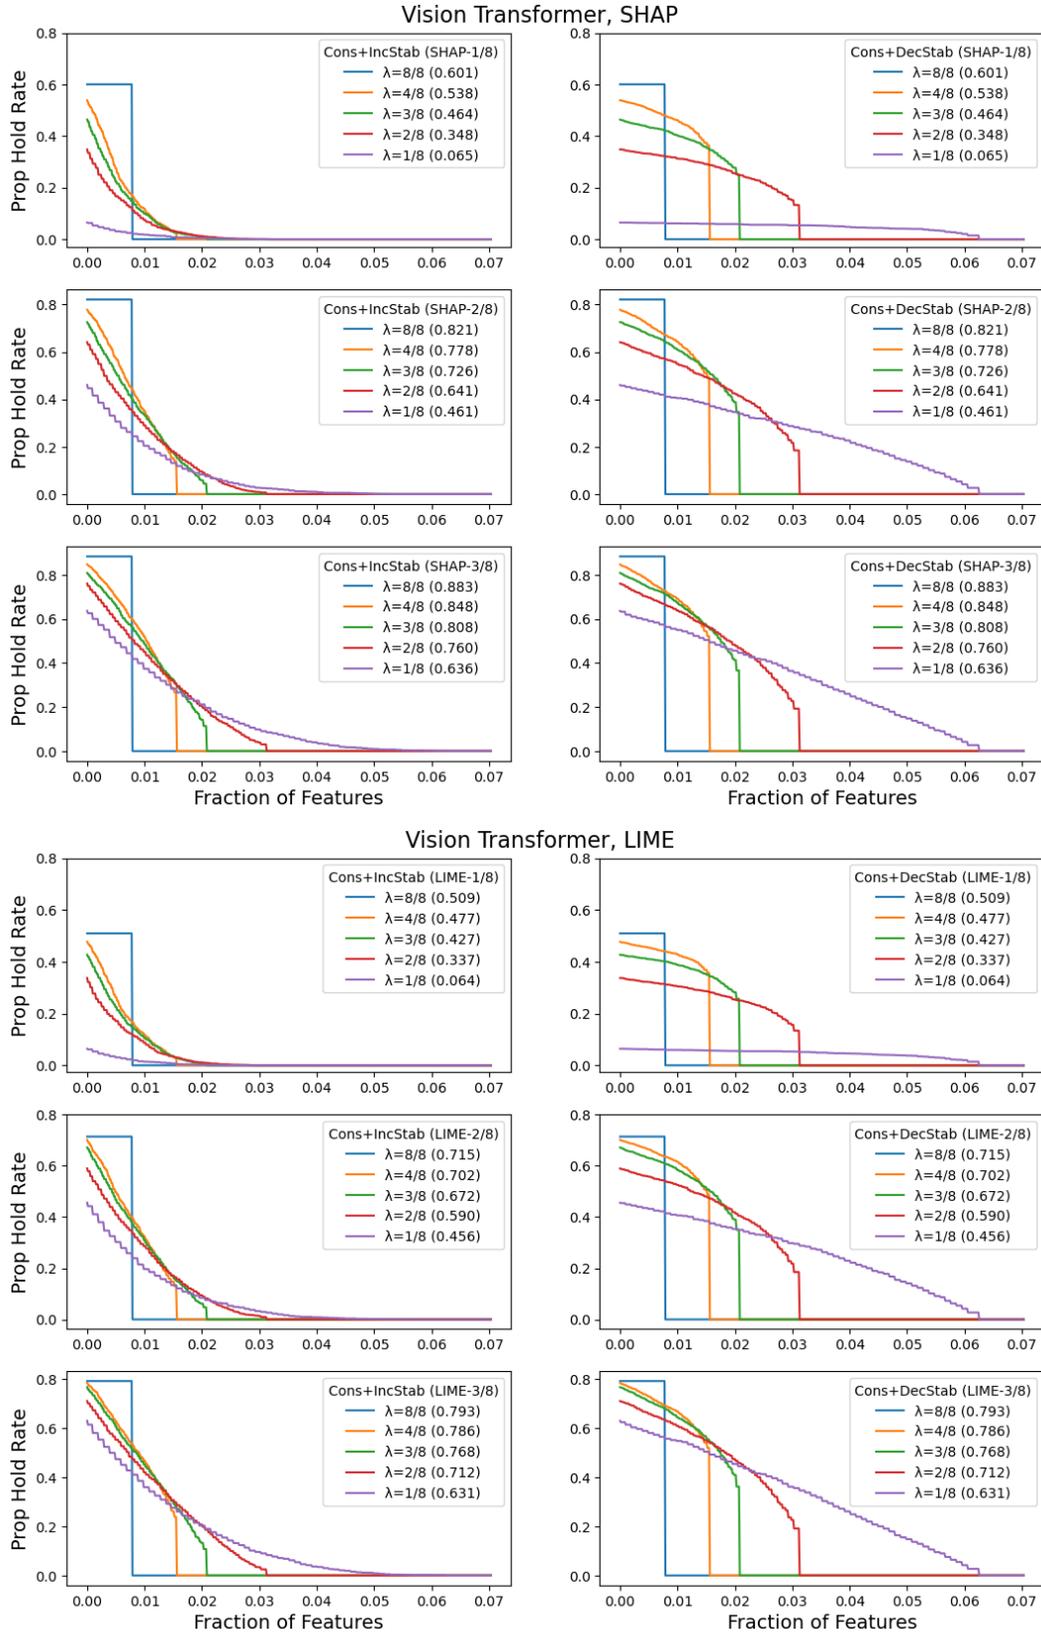


Figure 7: (Top) Vision Transformer with SHAP; (Bottom) Vision Transformer with LIME. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

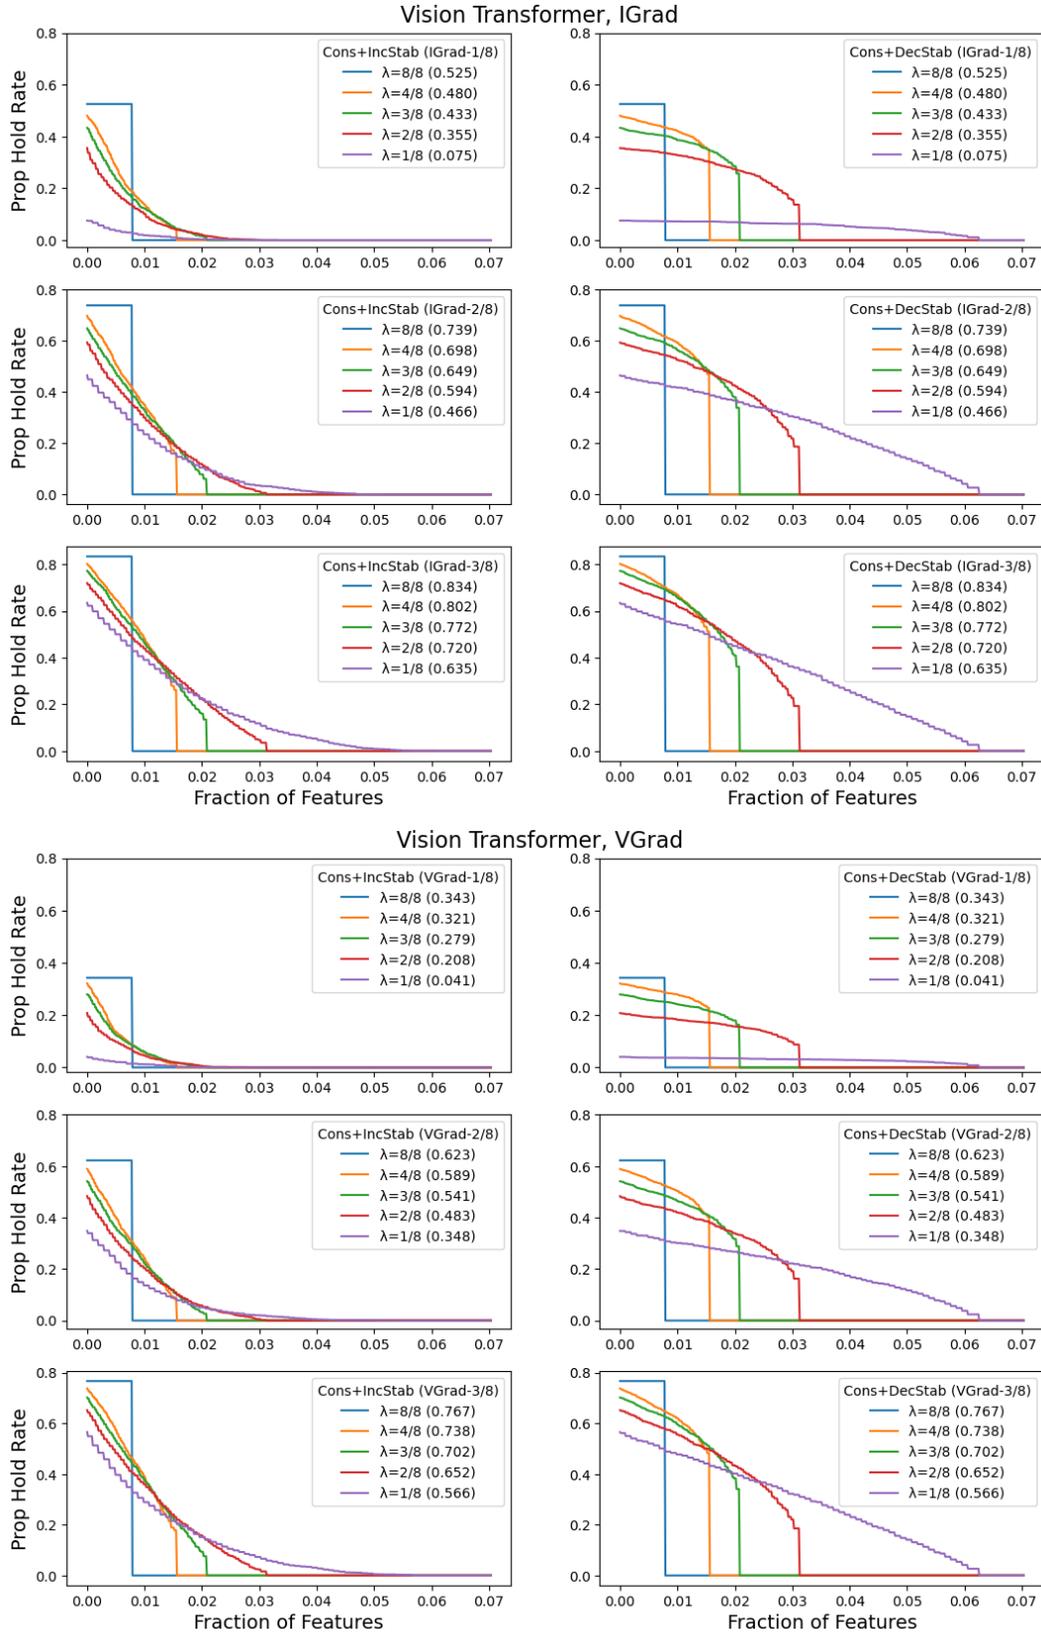


Figure 8: (Top) Vision Transformer with IGrad; (Bottom) Vision Transformer with VGrad. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

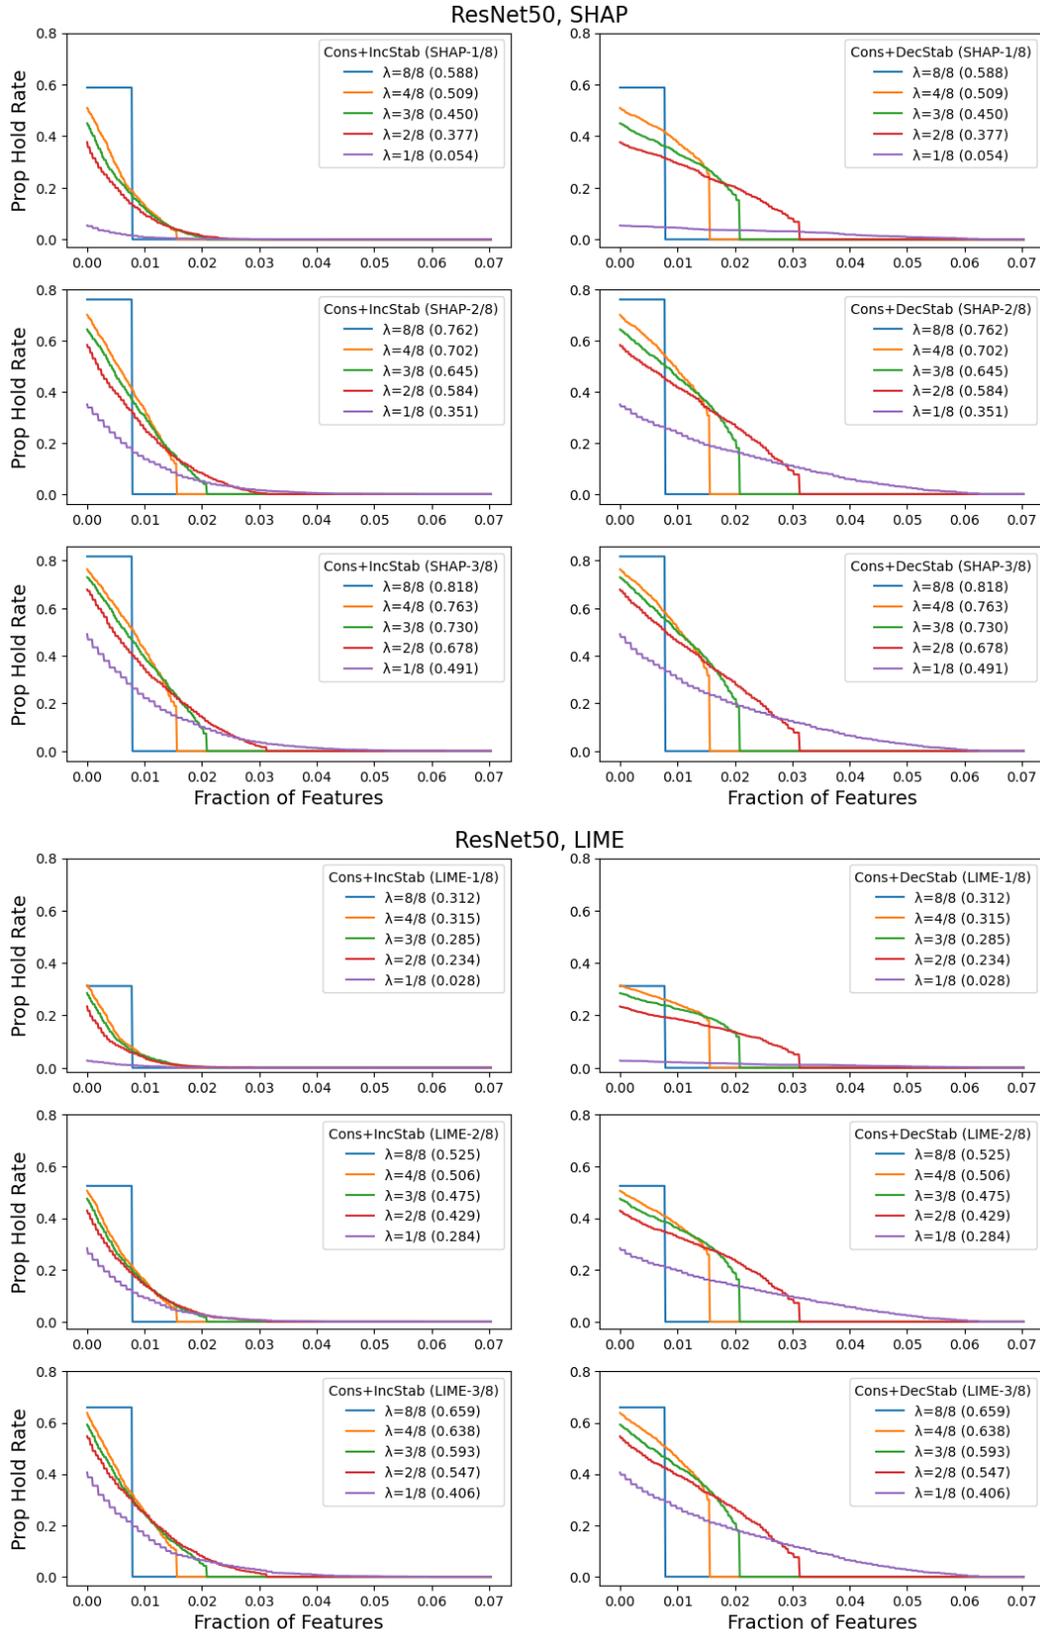


Figure 9: (Top) ResNet50 with SHAP; (Bottom) ResNet50 with LIME. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

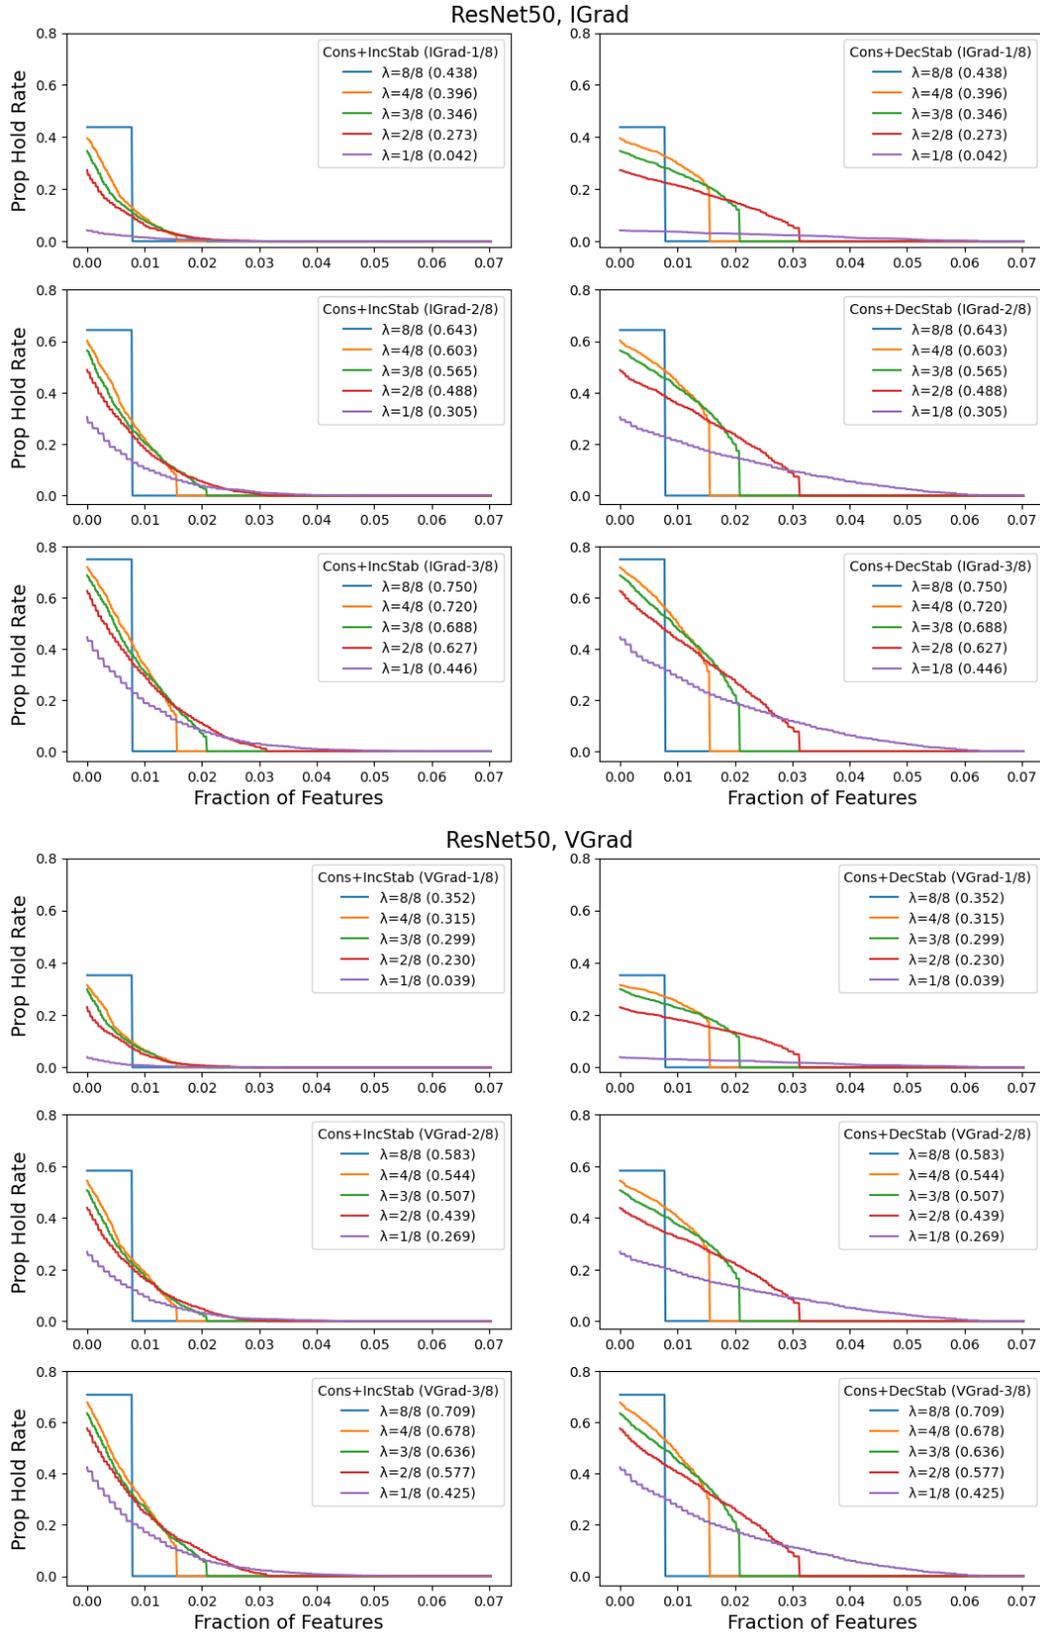


Figure 10: (Top) ResNet50 with IGrad; (Bottom) ResNet50 with VGrad. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

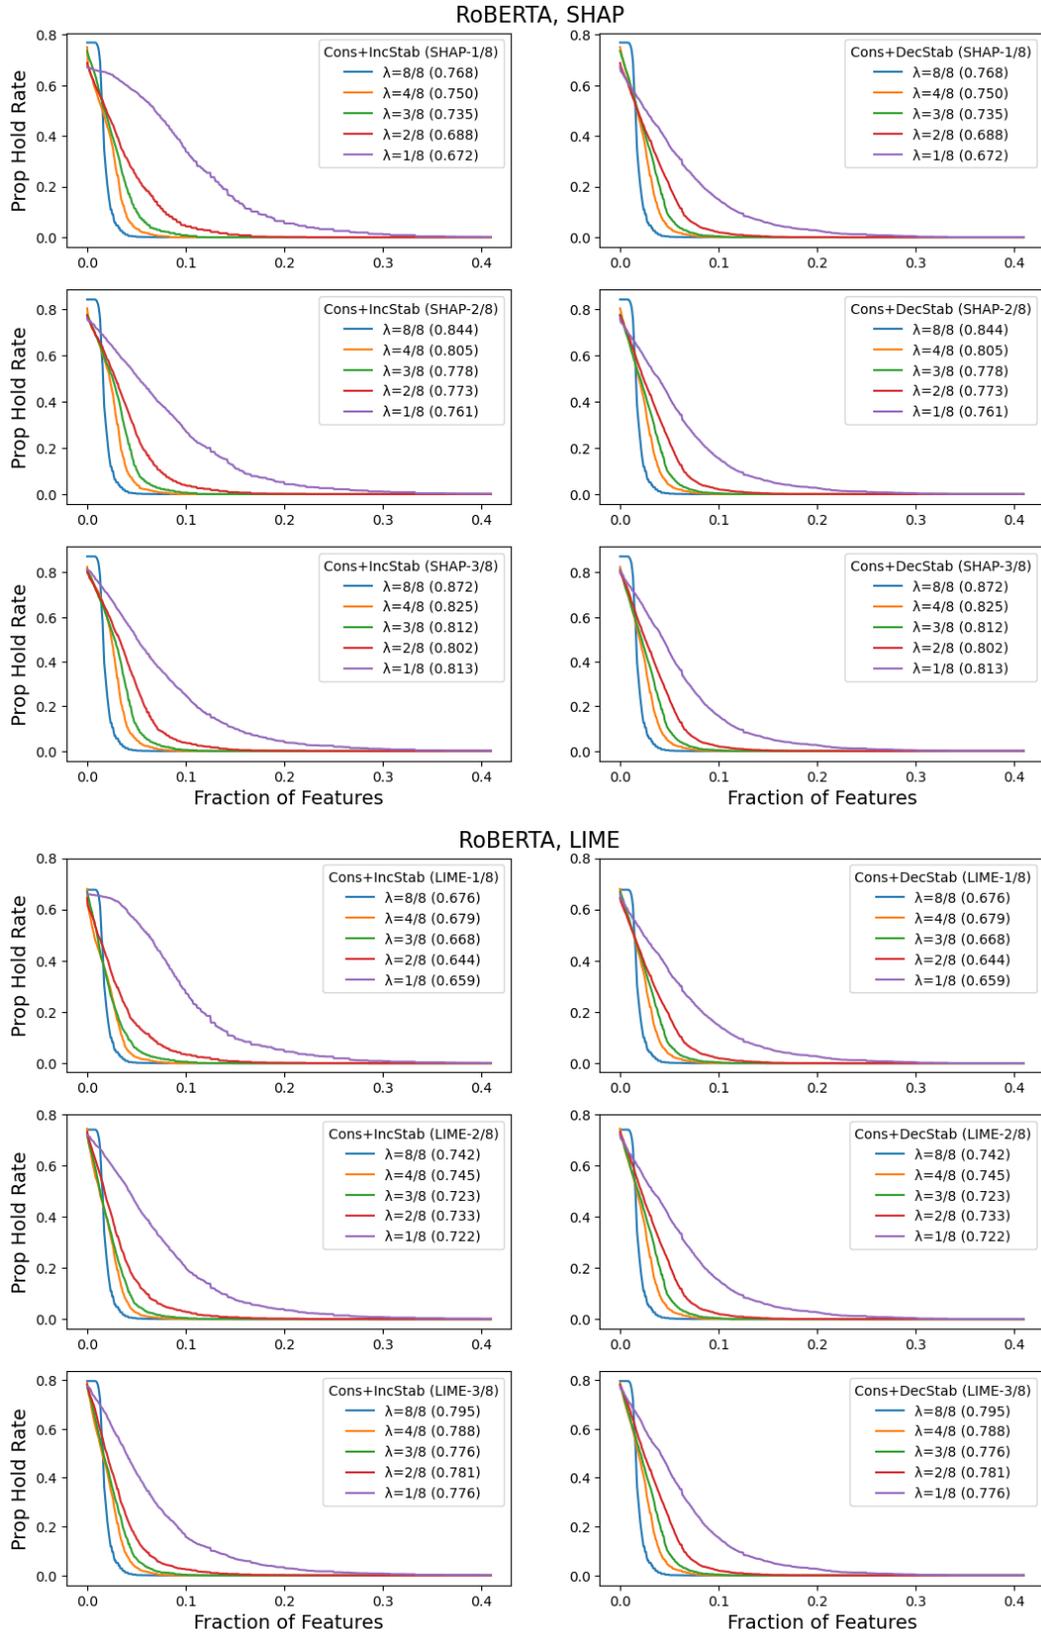


Figure 11: (Top) RoBERTa with SHAP; (Bottom) RoBERTa with LIME. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

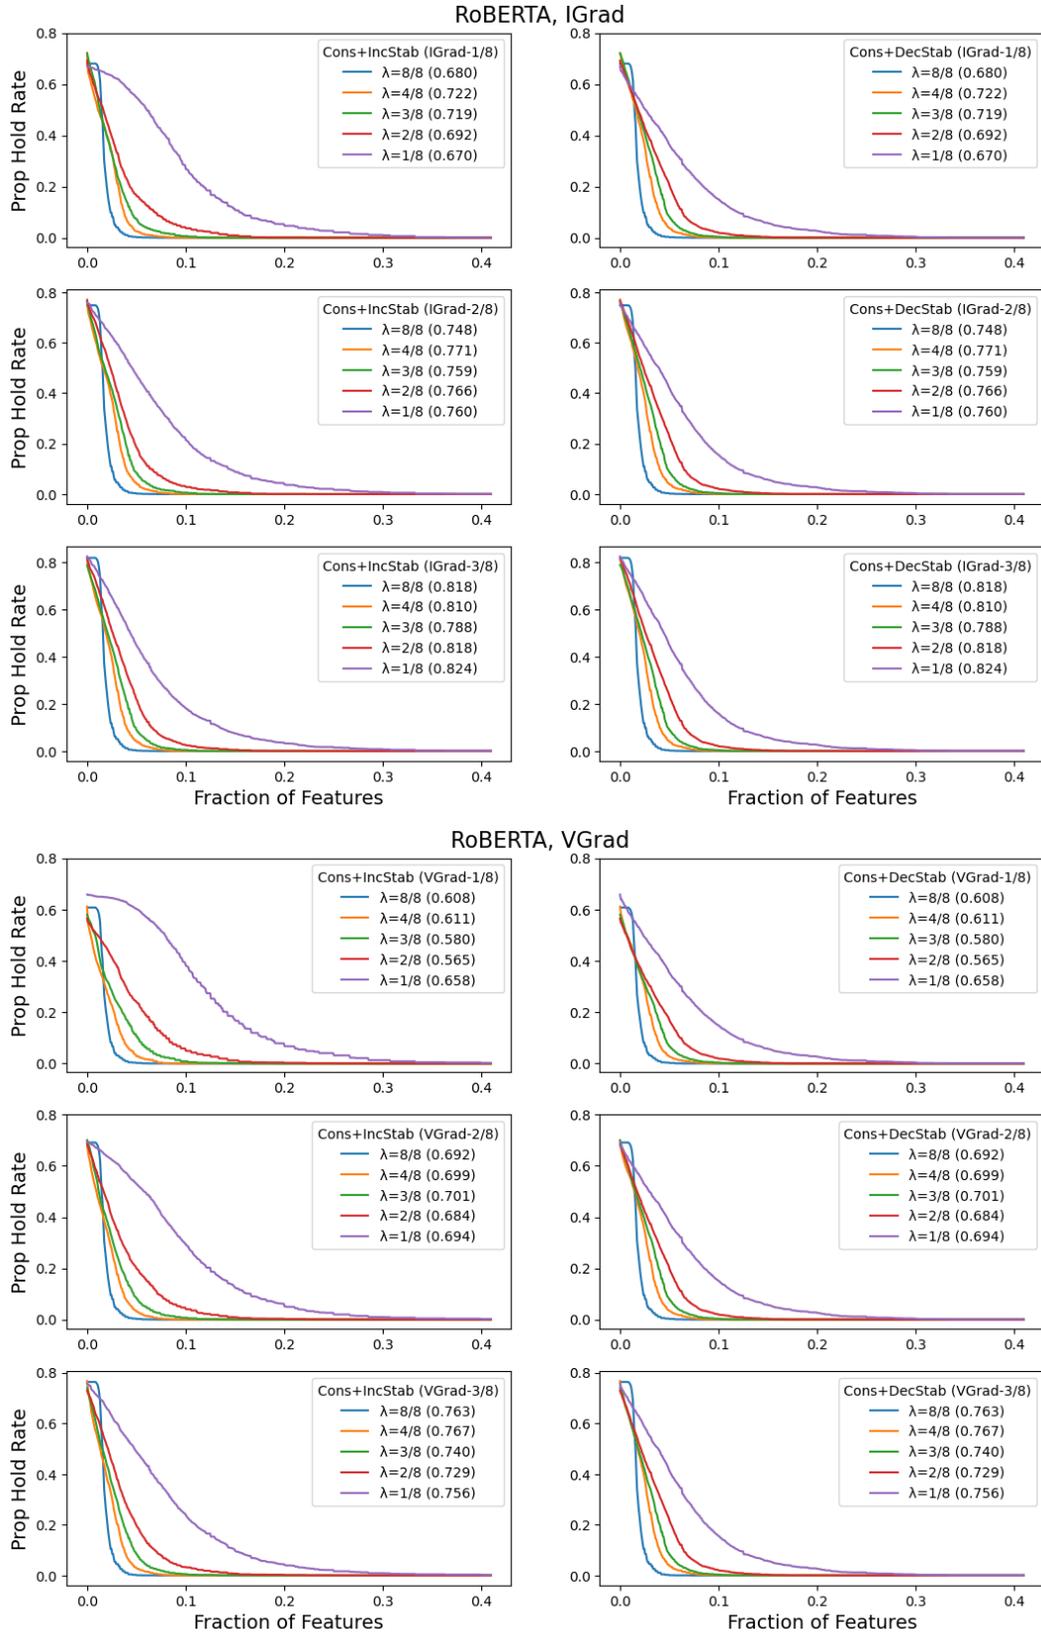


Figure 12: (Top) ROBERTa with IGrad; (Bottom) ROBERTa with VGrad. (Left) consistent and incrementally stable; (Right) consistent and decrementally stable.

551 **B.2 Theoretical vs Empirical**

552 We compare the certifiable theoretical stability guarantees with what is empirically attained via
 553 a standard box attack search [32]. This is an extension of Section 4.2, where we now show all
 554 models as evaluated with SHAP-top25%. The certified plots are identical from Appendix B.1. We
 555 take $N_{\text{cert}} = 2000$ samples for the certified plots, and $N_{\text{emp}} = 250$ for the empirical plots. This
 556 comparatively small selection of methods and data is because box attack is very time-intensive.

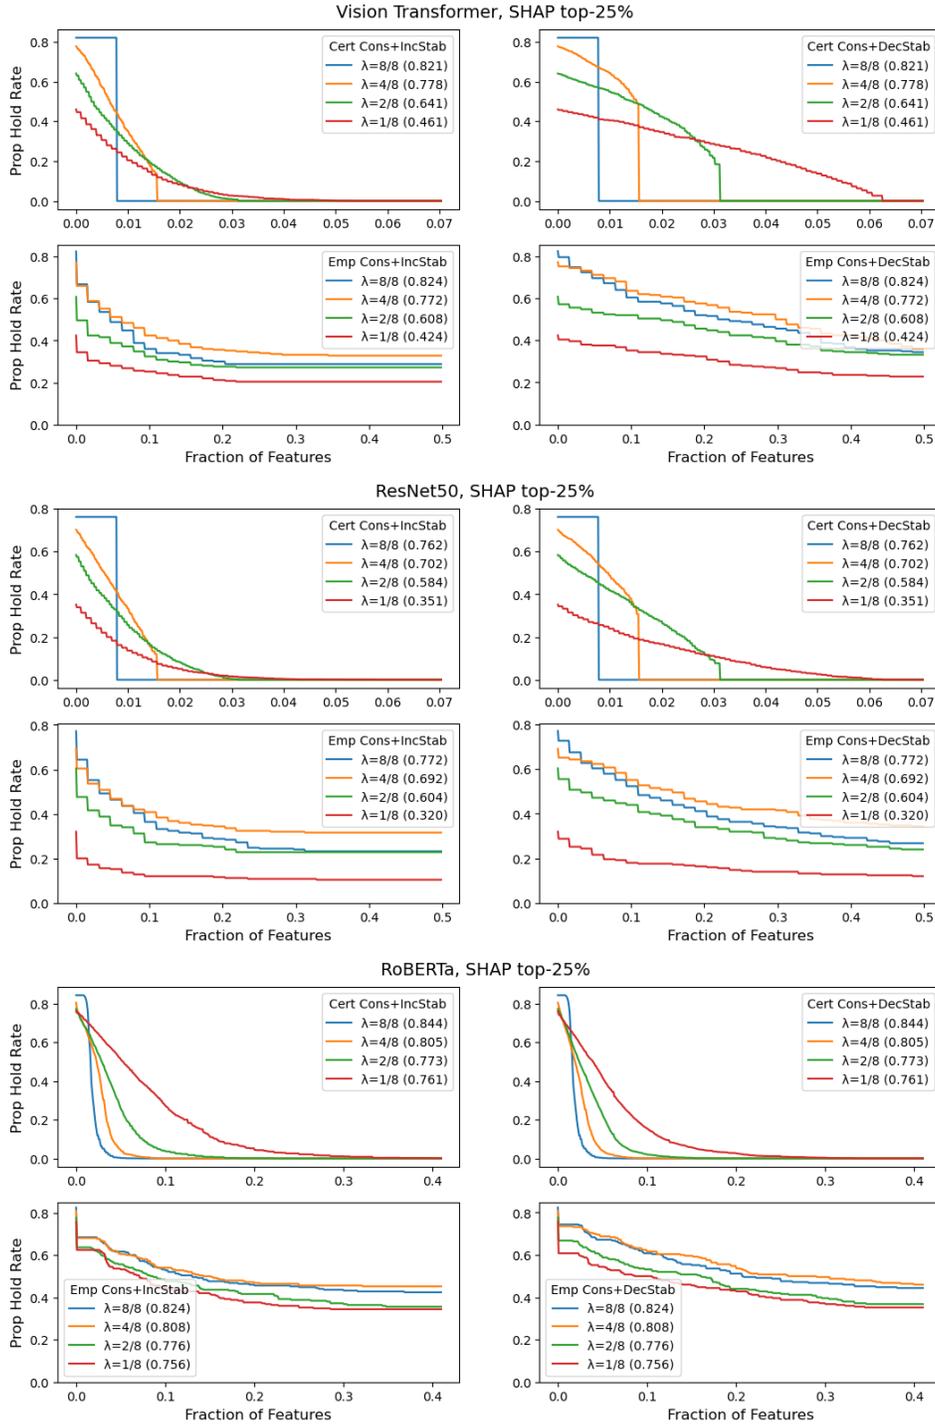


Figure 13: With SHAP top-25%: (Top) Vision Transformer; (Middle) ResNet50; (Bottom) RoBERTa.

557 **B.3 Stability-Accuracy Trade-Offs**

558 We study how the accuracy degrades with λ . We consider a smoothed model f constructed from a base
 559 classifier $h \in \{\text{Vision Transformer, ResNet50, RoBERTa}\}$ and vary $\lambda \in \{1/16, 1/8, 2/8, 4/8, 8/8\}$.
 560 We then take $N = 2000$ samples from each respective dataset and measure the accuracy of f at
 561 different radii. We use $f(x) \cong \text{true_label}$ to mean that f attained the correct prediction at $x \in \mathcal{X}$,
 562 and we plot the following value at each radius r :

$$\text{value}(r) = \frac{\#\{x : f(x) \cong \text{true_label} \text{ and dec stable with radius } \leq r\}}{N}$$

563 The overall accuracy with each λ is shown in the parentheses of each plot’s legend.

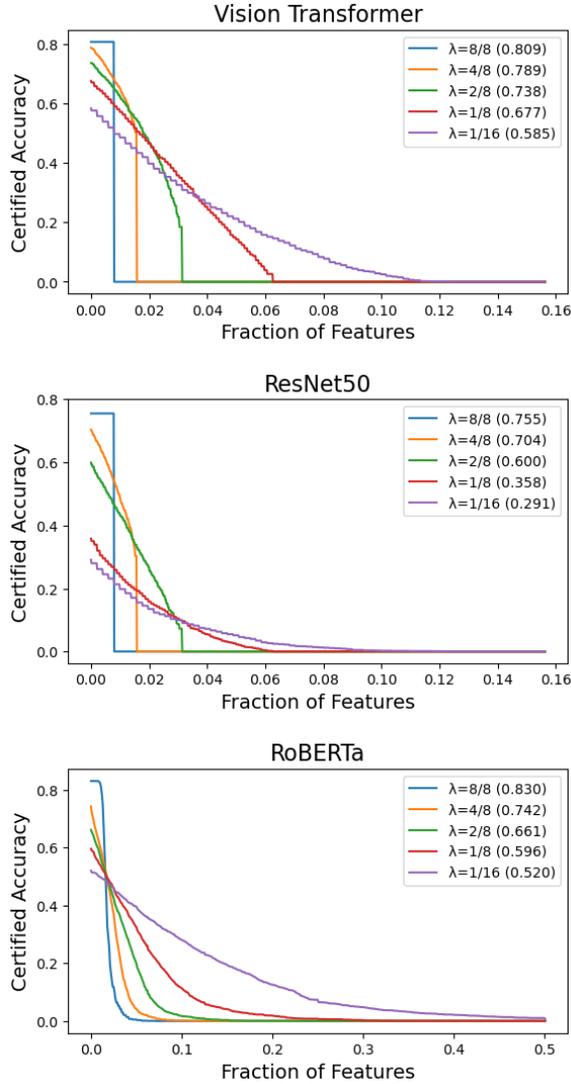


Figure 14: (Top) Vision Transformer; (Middle) ResNet50; (Bottom) RoBERTa.

564 **B.4 Which Explanation Method is the Best?**

565 We first investigate how many features are needed to yield consistent and non-trivially stable explanations, as done by the greedy selection algorithm in Section 2.4. For some $x \in \mathcal{X}$, let k_x denote the fraction of features that $\langle f, \varphi \rangle$ needs to be consistent, incrementally stable, and decrementally stable with radius 1. We vary $\lambda \in \{1/8, \dots, 4/8\}$, where recall $\lambda \leq 4/8$ is needed for non-trivial stability, and use $N = 250$ samples to plot the average k_x . This part is identical to Section 4.3.

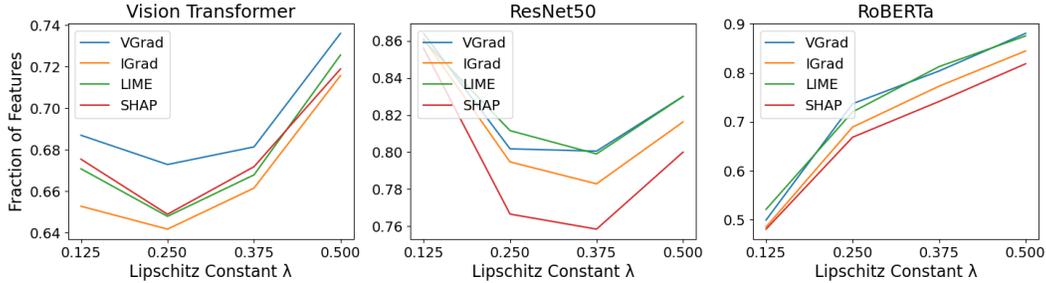


Figure 15: (Left) Vision Transformer; (Middle) ResNet50; (Right) RoBERTa.

570 We next investigate the ability of each method to predict features that lead to high accuracy. Let $f(x \odot \varphi(x)) \cong \text{true_label}$, mean that the masked input $x \odot \varphi(x)$ yields the correct prediction. We then plot this accuracy as we vary the top- $k \in \{1/8, 2/8, 3/8\}$ for different methods φ , and $\lambda \in \{1/8, \dots, 8/8\}$, using $N = 2000$ samples.

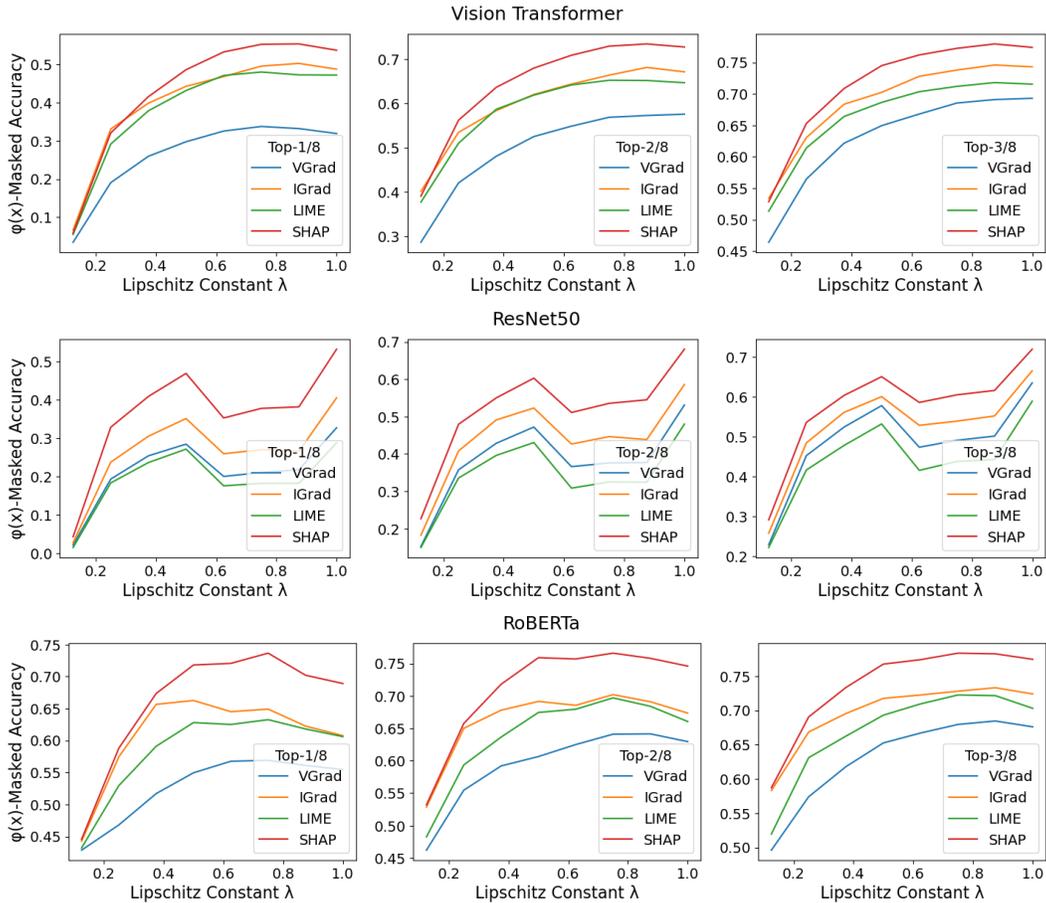


Figure 16: (Top) Vision Transformer; (Middle) ResNet50; (Bottom) RoBERTa.

574 B.5 Discussion

575 **Effect of Smoothing** We observe that smoothing can yield non-trivial stability guarantees, espe-
576 cially for Vision Transformer and RoBERTa, as evidenced in Appendix B.1. We see that smoothing
577 is least detrimental on these two transformer-based architectures, and most negatively impacts the
578 performance of ResNet50. We conjecture that although different training set-ups may improve
579 performance across every category, this still serves to illustrate the general trend.

580 **Theoretical vs Empirical** It is expected that the certifiable radii of stability is more conservative
581 than what is empirically observed. As mentioned in Section 3.2, for each λ there is a maximum
582 radius to which stability can be guaranteed, which is an inherent limitation of using logit gaps and
583 Lipschitz constants as the main theoretical technique. We emphasize that the notion of stability need
584 not be tied to smoothing, though we are currently not aware of other viable approaches.

585 **Why these Explanation Methods?** We chose SHAP, LIME, IGrad, and VGrad from among the
586 large variety of methods available primarily due to their popularity, and because we believe that
587 they are collectively representative of many techniques. In particular, we believe that LIME remains
588 representative baseline for surrogate model-based explanation methods. SHAP and IGrad are, to our
589 knowledge, the two most well-known families of axiomatic feature attribution methods. Finally, we
590 believe that VGrad is representative of a traditional gradient saliency-based approach.

591 **Which Explanation Method is the Best?** Based on our experiments in Appendix B.4 we see that
592 SHAP generally achieves higher accuracy using the same amount of top- k features as other methods.
593 On the other hand, VGrad tends to perform poorly. We remark that there is well-known critique
594 against the usefulness of saliency-based explanation methods [52].

595 C Miscellaneous

596 **Relevance to Other Explanation Methods** Our key theoretical contribution of MuS in Theorem 3.2
597 is a general-purpose smoothing method that is distinct from standard smoothing techniques, namely
598 additive smoothing. MuS is therefore applicable to other problem domains beyond what is studied in
599 this paper, and would be useful where Lipschitz constants with respect to maskings is desirable.

600 **Broader Impacts** Reliable explanations are necessary for making well-informed decisions, and are
601 increasingly important as machine learning models are integrated with fields like medicine, law, and
602 business — where the primary users may not be well-versed in the technical limitations of different
603 methods. Formal guarantees are therefore important for ensuring the predictability and reliability
604 of complex system, which then allows users to construct accurate mental models of interaction and
605 behavior. In this work we study a particular kind of guarantee known as stability, which is key to
606 feature attribution-based explanation methods.