I-ASIDE: Towards the Global Interpretability of Image Model Robustness through the Lens of Axiomatic Spectral Importance Decomposition

Anonymous authors
Paper under double-blind review

Abstract

Robust decisions leverage a high proportion of robust features. Natural images have spectral non-uniformity and the majority of spectral energy is concentrated on low-frequency components. A change with an infinitesimal amount of energy on the high-frequency components can rewrite the raw features dominated by high-frequency components. Image models are parameterized general non-linear signal filters. The spectral structures of the model responses to inputs determines the fragility of the learned feature representations. The spectral importance decomposition of models can thus reflect model robustness in response to feature perturbations. To this end, we formulate the spectral importance decomposition problem, and, present **I**mage **A**xiomatic **S**pectral **I**mportance **D**ecomposition **E**xplanation (**I-ASIDE**) – a model-agnostic global interpretability method – to quantify model global robustness and understand how models respond to perturbations. We theoretically show that **I-ASIDE** decomposes the mutual information between feature representations and labels onto spectrum. Our approach provides a unique insight into interpreting model global robustness from the perspective of information theory and enables a considerable number of applications in research, from understanding model robustness, to studying learning dynamics, to assessing label noise, to investigating adversarial vulnerability, to studying out-of-distribution robustness, etc. We showcase multiple applications to endorse these claims.

1 Introduction

Global interpretability summarizes the decision dynamics of neural networks *en masse* in contrast to instance-wise local interpretability (Lipton, 2018; Zhang et al., 2021). Local interpretability for image models has achieved great success (Sundararajan et al., 2017; Smilkov et al., 2017; Linardatos et al., 2020; Selvaraju et al., 2017; Arrieta et al., 2020; Zhou et al., 2016; Ribeiro et al., 2016; Lundberg & Lee, 2017; Lakkaraju et al., 2019; Guidotti et al., 2018; Bach et al., 2015; Montavon et al., 2019; Shrikumar et al., 2017), yet quantifiable global interpretability remains virtually unexplored. A brief literature review regarding global interpretability for image models is provided in Section 2.

The global robustness (Hendrycks & Dietterich, 2019; Bai et al., 2021; Goodfellow et al., 2014; Silva & Najafirad, 2020) of models reflects an intrinsic property of models and delineates a crucial aspect towards interpretability and trustworthiness. To this end, we present **I-ASIDE**[1][2], a model-agnostic method, to quantify the global robustness of image models and understand how models respond to perturbations. Unlike prior works, **I-ASIDE** directly quantifies model global robustness and enables a considerable number of applications in deep learning research across multiple domains (See Section 4). We also theoreticaly show that **I-ASIDE** *de facto* decomposes the mutual information between feature representations and labels. This can provide profound insights into understanding how features are represented inside black box models. We will provide five application showcases to demonstrate the potential applications of **I-ASIDE** framework.

[1]Anonymized reproducibility: `https://anonymous.4open.science/r/IASIDE_reproducibility-F8BC/`.
[2]The full showcase experiments have consumed 25000+ core hours on supercomputing center with multiple GPUs to complete.

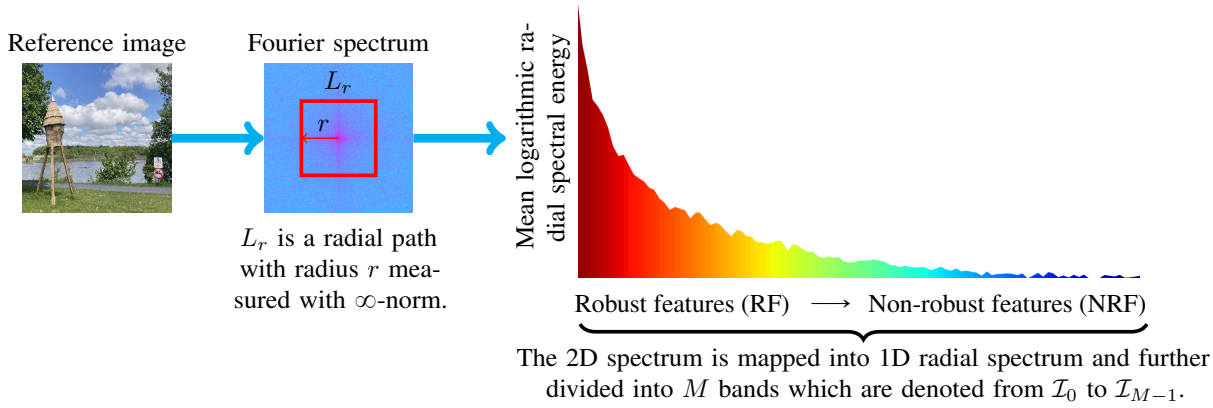

Figure 1: This diagram illustrates the correlation between the power-law like radial spectral energy density of natural images and feature robustness. According to the Parseval identity theorem, small spatial perturbations targeting on high-frequency components can rewrite the features dominated by high-frequency components.

Feature representations are the representations of input features that models learn to represent relevant information while neglect irrelevant information (Bengio et al., 2013; Goodfellow et al., 2016). The feature representation robustness largely delimits decision robustness afterwards. Robust decisions leverage more robust feature representations while non-robust decisions leverage less robust representations. The correlation between the spectral energy density distributions of input features and the spectral structures of the responses of models to inputs plays a crucial aspect in feature representation robustness. Consequently, the signal spectrum can index the robustness of features: Low-frequency indicates robust features while high-frequency indicates non-robust features.

Natural images have spectral non-uniformity and the majority of spectral energy concentrates on low-frequency components (See Figure 1 and Figure 3(a)). According to the Parseval identity theorem – the energy of the perturbations in the spatial domain and the spectral domain is conservative and equivalent, the small spatial perturbations targeting on high-frequency components can rewrite those feature representations dominated by high-frequency components. Figure 3 shows four example spectrum from adversarial perturbations to out-of-distribution (OOD) perturbations. The small spatial perturbations in both adversarial noises and OOD noises have considerable energy concentrating on high-frequency components and can have significant impacts on the features dominated by high-frequent components.

Ilyas et al. use the terms 'robust feature (RF)' and 'non-robust feature (NRF)' to theoretically analyze the adversarial robustness problem from the perspective of feature representations and argue that the presence of non-robust features can incur model robustness issues (Ilyas et al., 2019; Tsipras et al., 2018). We adopt the use of their terms and step forward to rigorously discuss the robust decision problem in a broader scope beyond adversarial robustness.

Measuring the ratio of the robust features in decisions can help to interpret model inference dynamics globally from the perspective of feature representations. Unfortunately, discriminating features between being robust and being non-robust is difficult. We have noticed the correlation between the power-law like spectral structures of natural images and feature robustness. Feature robustness can thus be indexed by radial spectrum (See Figure 1). We refer to the radial spectrum as the 'robustness spectrum' hereafter. The decompositions of the mutual information of the feature representations of inputs and labels with respect to the robustness spectrum can summarize model decision robustness, and quantitatively reflect how neural networks respond to signals in the frequency-domain.

The spectral decompositions are a unique solution satisfying a set of desirable axioms. We formulate this problem by using Shapley value theory (Roth, 1988; Aumann & Shapley, 2015). Our approach provides

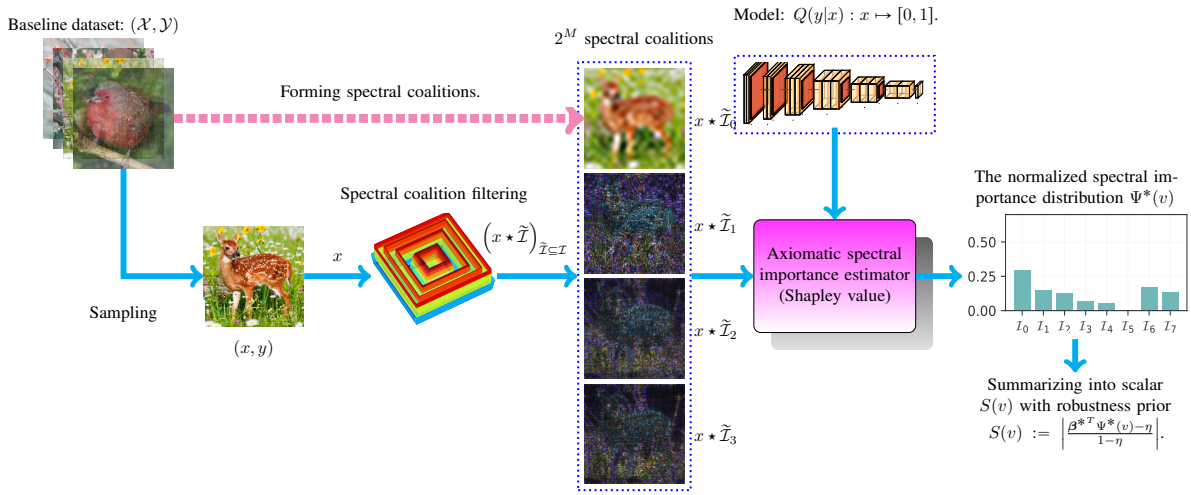

Figure 2: This diagram illustrates the overview of **I-ASIDE** with some baseline dataset $(\mathcal{X}, \mathcal{Y})$ (where $\mathcal{X}$ denotes images and $\mathcal{Y}$ denotes labels) and some classifier $Q(y|x) : x \mapsto [0, 1]$ which outputs a normalized probability distribution for $y \in \mathcal{Y}$. The images are sampled from the baseline dataset $(\mathcal{X}, \mathcal{Y})$ and fed into a spectral coalition filter to form 2^M spectral coalitions. The spectral coalitions are then used to compute the spectral importance distribution $\Psi^*(v)$. The normalized spectral importance distribution $\Psi^*(v)$ can be summarized into a robustness score $S(v)$.

a unique insight into interpreting the global decision robustness of image models and enables multiple applications in research such as:

- Quantifying model global robustness by examining spectral importance distributions (See the experiments regarding spectral importance distributions in Figure 4, summarized numerical comparison and t-SNE projection in Figure 7);
- Understanding the learning behaviours of image models on the datasets with supervision noise (See the experiments in Figure 8);
- Investigating the learning dynamics of models from the perspective of the evolution of feature representation robustness in optimization (See the experiments in Figure 9);
- Providing an insight into understanding the adversarial vulnerability of models by examining spectral importance distributions (See the experiments in Figure 10);
- Studying model out-of-distribution (OOD) robustness when input domains shift (See the experiments in Figure 11). **I-ASIDE** can help to answer model robustness concerns from the perspective of feature representations in frequency-domain.

It is also notable that the potential applications are not restricted to above fields. Other research fields such as data augmentation and self-supervised learning can also use **I-ASIDE** as a device to understand the dynamics of the learned feature representations of models. For example, in data augmentation research, **I-ASIDE** can interpret the training dynamics with ablation experiments by understanding how hyper parameters and augmentation tricks can affect the robustness of the learned features.

2 Related work

We summarize related works from four categories to provide research context in visual models: (1) Global interpretability, (2) model robustness, (3) frequency-domain research and (4) mutual information in deep learning.

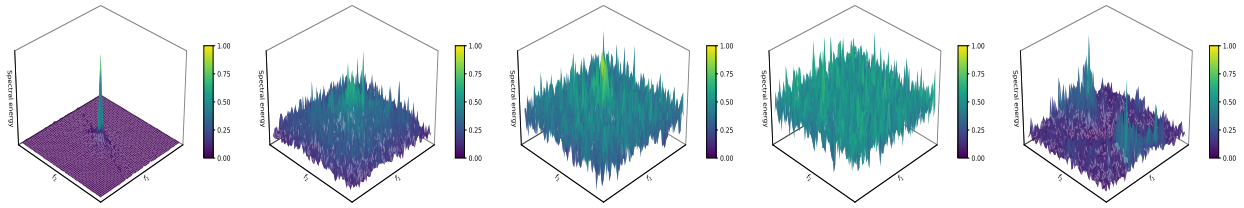

(a) Reference (b) FGSM (c) PGD (d) White noise (e) Gaussian-blur

Figure 3: This shows the spectrum examples from four perturbation sources: (b) Adversarial FGSM perturbation with $eps = 0.1$, (c) adversarial PGD perturbation with $eps = 0.1$, (d) OOD white noise with $\sigma = 0.1$ and (e) OOD blurring with Gaussian kernel 3×3. The (a) is the reference spectrum of the reference image (See Figure 1). The adversarial samples are obtained by using a *resnet50* pre-trained on *ImageNet*. The perturbation noises are computed by $\Delta\boldsymbol{x} := \boldsymbol{x}^* - \boldsymbol{x}$ where $\boldsymbol{x}$ is some clean sample and $\boldsymbol{x}^*$ is the corresponding perturbed sample. We apply Fourier analysis on the min-max normalized $\Delta\boldsymbol{x}$ and show the results. The normalization does not affect the results from Fourier analysis since Fourier operator is a linear operator. The examples show that several usual adversarial perturbations and OOD perturbations have considerable spectral energy concentrating on high-frequency components. The perturbations targeting on high-frequency components can rewrite the features dominated high-frequency components.

Global interpretability: Global interpretability summarizes the decision behaviours of models and provides a holistic view. Feature visualization with neuron or class activation maximization can show the ideal inputs for specific neurons or classes by optimizing inputs, and help to understand the learned feature representations of models (Olah et al., 2017; Nguyen et al., 2019; Zeiler et al., 2010; Simonyan et al., 2013; Nguyen et al., 2016a;b). Yet, neural networks are surjective maps and the results often yield 'surrealistic' inputs which are not interpretable. Network dissection attempts to establish the connection between the functions of the units (e.g. a set of channels or a set of layers) in convolutional neural networks and some concepts – e.g. eyes or ears (Bau et al., 2017). Concept-based approach measures the activations of networks with respect to concepts, and interpret networks at concept level (Kim et al., 2018; Ghorbani et al., 2019; Koh et al., 2020; Chen et al., 2020). Global feature importance analysis measures how important the features contribute to predictions (Altmann et al., 2010; Greenwell et al., 2018; Lundberg & Lee, 2017; Ribeiro et al., 2016; Simonyan et al., 2013; Sundararajan et al., 2017). For example, Lundberg & Lee present an additive importance measure which unifies numerous methods in literature, and proposes SAGE by applying Shapley value theory to interpret feature importance. Yet, the implementations of these works are largely restricted to examining feature importance in spatial domain. Interpreting the global feature importance in spatial domain has inherited limits. For example, the affine transforms on images can change considerable pixel values. Thus the interpretations in spatial domain often suffer from these limits and provide very limited insights regarding robustness. Our work differs from these prior works in that we interpret global model robustness in frequency domain.

Model robustness: Szegedy et al. notice the robustness problem of deep models arising from adversarial perturbations. Thereafter, empirical observations demonstrate that boosting model robustness with adversarial training is at the cost of the performance degradation on model standard accuracy for visual tasks (Goodfellow et al., 2014). Later theoretical analyses show standard accuracy is at odds with model robustness (Zhang et al., 2019; Tsipras et al., 2018). Ilyas et al. argue features can be distinguished by their brittleness into robust features and non-robust features, and show that adversarial robustness relates to non-robust features. Our research is based on the above insights.

Frequency-domain research: Neural networks are non-linear parameterized signal processing filters. Investigating how neural networks respond to inputs in the frequency-domain can provide a unique insight into understanding its functions. Rahaman et al. approximate ReLU networks with piece-wise continuous linear functions in order to perform Fourier analysis. Their results suggest neural networks have a 'spectral bias' on smooth hypotheses (Raghu et al., 2017; Montufar et al., 2014; Rahaman et al., 2019). Xu et al. investigate the learning dynamics of neural networks in frequency-domain in their work 'F-Principle' (Xu et al.,

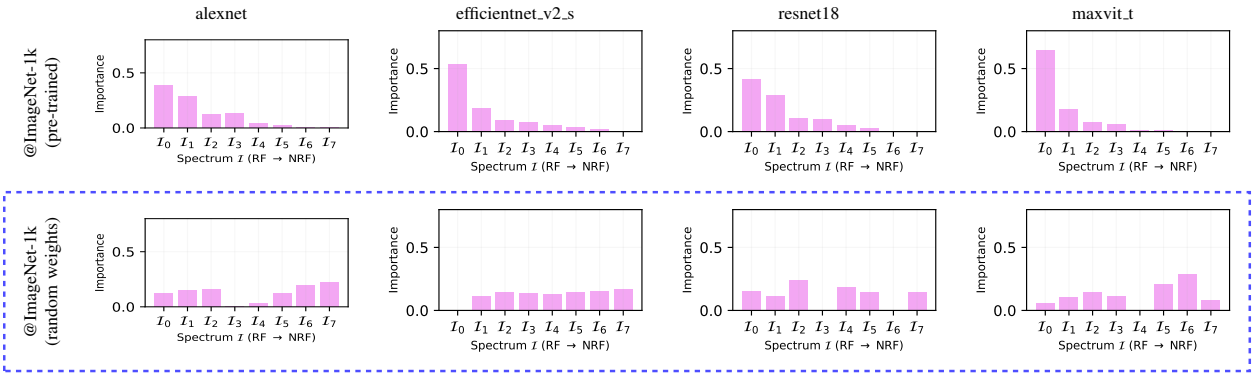

Figure 4: We showcase the spectral importance distributions from multiple models pre-trained on *ImageNet-1k*. We also include the models with random weights as a control marked by blue box. We have noticed that: (1) The spectral importance of the robustness of feature representations of trained models is non-uniform, and, (2) the spectral importance of the robustness of the feature representations of un-trained models is uniform.

2019a;b). Their work suggests that the learning behaviors of neural networks are spectral non-uniformity: Neural networks fit low-frequency components first, then high-frequency components later. Tsuzuku & Sato affirm convolutional neural networks have spectral non-uniformity regarding Fourier bases (Tsuzuku & Sato, 2019). Wang et al. conducts an empirical study of the connection between supervision signals and feature representations from the perspective of frequency (Wang et al., 2020). Kolek et al. propose 'CartoonX' and attempt to interpret the instance-wise decisions of image classifiers in wavelet domain (Kolek et al., 2022). Both **I-ASIDE** and CartoonX are the attempts to interpret the decision behaviors of models in frequency-domain. Our work differs from CartoonX in that: (1) We use Fourier transform instead of wavelet transform because the feature robustness to perturbations correlates with the Fourier spectrum of inputs, and (2) we assess the global decision behaviors over a population other than instance-wise.

Mutual information in deep learning: The optimization objectives of a variety of works from supervised learning and un-supervised learning can be unified under the view of maximizing variational mutual information (Poole et al., 2019; Nowozin et al., 2016; Kingma & Welling, 2013). We use some theoretical results in information theory to justify and characterize the choice of the characteristic function in **I-ASIDE**. We summarize the major results in terms of the theoretical advances and applications of the information theory in neural networks. Tishby et al. introduce information bottleneck theory for interpreting the learning dynamics of neural networks (Tishby et al., 2000; Saxe et al., 2019). They claim that neural networks are learnable compression algorithms which extract relevant information while neglect irrelevant information for inferences (Saxe et al., 2019). Nowozin et al. theoretically show that enormous discriminators in Generative Adversarial Networks are variational divergence estimators of measuring the statistical distances between training sample distributions and generated sample distributions in the min-max game (Nowozin et al., 2016). They then propose the f-GAN framework (Nowozin et al., 2016) to unify a variety of GAN frameworks. For example, when the $g_f(v)$ is set to $g_f(v) := -\log(1 + e^v)$ and the corresponding conjugate function is set to $f^*(t) := -\log(1 - e^{-t})$, the optimization objective of f-GAN gives the optimization objective of vanilla GAN which measures the Jensen-Shannon divergences between training sample distributions and generated sample distributions (Nowozin et al., 2016; Goodfellow et al., 2020). Veličković et al. propose Infomax loss to supervise representation learning by maximizing the variational mutual information between hidden embeddings and inputs (Veličković et al., 2018). The Variational Autoencoder (VAE) uses a variational mutual information bound of the mutual information between encoding representations and inputs – the evidence lower bound (ELBO) – to supervise the representation learning of encoders (Kingma & Welling, 2013). Barber & Agakov have shown a variational lower bound between inputs and labels (Barber & Agakov, 2004). Qin et al. use the result and show that the training objective of the classifiers trained with usual cross-entropy is equivalent to estimating a variational bound of the mutual information between inputs and labels (Qin et al., 2021). Donsker & Varadhan show a tight variational mutual information representation of Kullback-Leibler divergence between two distributions (Donsker-Varadhan Representation

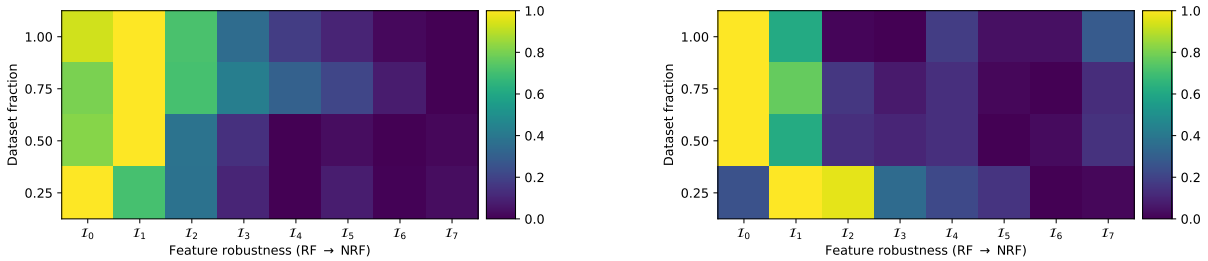

(a) *resnet18 @CIFAR10*

(b) *resnet18 @CIFAR100*

Figure 5: This demonstrates how the spectral importance distributions of feature representations change with respect to the training set sizes. We choose *resnet18* on *CIFAR10* and *CIFAR100*. The preliminary results show that there does not exist a unique pattern regarding how spectral importance distributions with respect to training set sizes.

Theorem): $D_{KL}(P||Q) = \sup_{T:\Omega\mapsto\mathbb{R}} \left(\mathbb{E}_{P}[T] - \log\left(\mathbb{E}_{Q}[e^{T}] \right) \right)$ (Donsker & Varadhan, 1983) where Ω denotes some compact domain on $\mathbb{R}^d$ and P and Q are distributions over Ω. Belghazi et al. deploy the Donsker-Varadhan Representation Theorem as a mutual information estimator to estimate the mutual information between two random variables by training (Belghazi et al., 2018).

3 Axiomatic spectral importance decomposition

We rigorously formulate the spectral importance decomposition problem as a fairness utility distribution problem in the context of coalition game. We directly deploy the Shapley value theory to derive the solution. We also justify and characterize the choice of the characteristic function by showing that **I-ASIDE** decomposes the mutual information between feature representations and labels onto spectrum. Figure 2 shows the overview of **I-ASIDE**.

3.1 Spectral coalition game

The probability predictions of an image classifier can be viewed as the worth (utility) scores in a coalition game with spectral players. The spectral players are the radial spectral regions in such a spectral coalition game (See Figure 1). The coalition game is defined by some image classifier. Spectral players *cooperate*, *interact*, and *contribute* to the decisions of the classifiers with various bargaining powers. The worth division should satisfy a set of axioms: *efficiency*, *symmetry*, *linearity* and *dummy player* (Roth, 1988; Hart, 1989; Winter, 2002). Shapley value theory is the unique solution to such a division.

Let $(\mathcal{X}, \mathcal{Y})$ be some image baseline dataset where $\mathcal{X}$ denotes images and $\mathcal{Y}$ denotes labels. Let $Q(y|x) : x \mapsto [0, 1]$ be some image classifier which predicts the discrete category probability distribution of some x for $y \in \mathcal{Y}$. Let $[0, 1]$ be the normalized 1-dimensional radial spectrum (See Figure 1). The radial spectrum is partitioned into M equispaced regions. Each region is a ℓ_∞ ball with various radius. Each partition then is a '*spectral player*'. The i-th spectral player is denoted as $\mathcal{I}_i$ (where $i \in [M] \stackrel{\text{def}}{=} \{0, 1, \cdots, M-1\}$). The M spectral players constitute a player set $\mathcal{I} := \{\mathcal{I}_i\}_{i=0}^{M-1}$. The consideration of using ℓ_∞ ball is because: (1) It is easier to handle the distant pixels at implementation and (2) the pixel density in the region remains constant compared with ℓ_p $(1 < p < \infty)$. Let $v : 2^{\mathcal{I}} \mapsto \mathbb{R}$ be some characteristic function which sends $2^{\mathcal{I}}$ to the field $\mathbb{R}$ such that $v(\varnothing) = 0$. The map v is given as the form of the statistical expectation of some $Q(y|x)$ over some baseline dataset $(\mathcal{X}, \mathcal{Y})$. We will show that this choice can be justified in that the characteristic function v measures a variational lower bound (Poole et al., 2019; Barber & Agakov, 2004) of the mutual information between images and labels in Section 3.4. We define a coalition game $(\mathcal{I}, v)$. Let $\psi_i(\mathcal{I}, v)$ be the spectral importance component for $\mathcal{I}_i$. Let $\Psi(v) := (\psi_i)_{i\in[M]}$ be the collection of $\psi_i(\mathcal{I}, v)$.

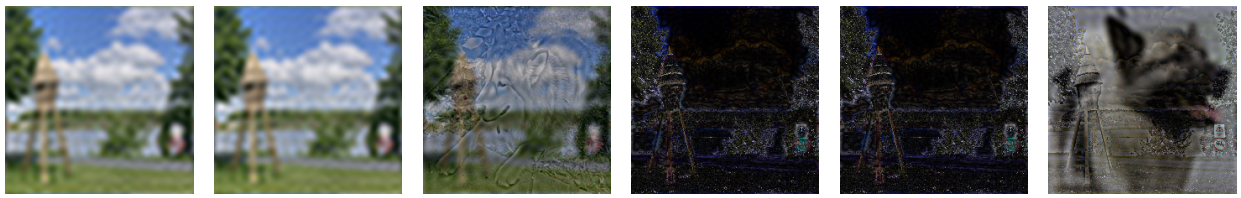

(a) Zeroing (b) Complex Gaussian (c) Replacement (d) Zeroing (e) Complex Gaussian (f) Replacement

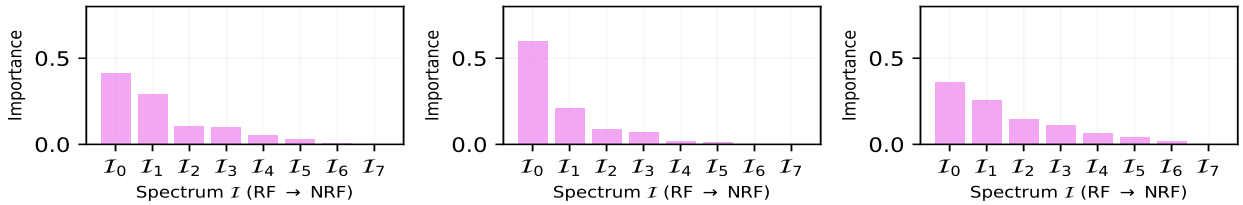

(g) *resnet18* w/ Zeroing (h) *resnet18* w/ Complex Gaussian (i) *resnet18* w/ Replacement

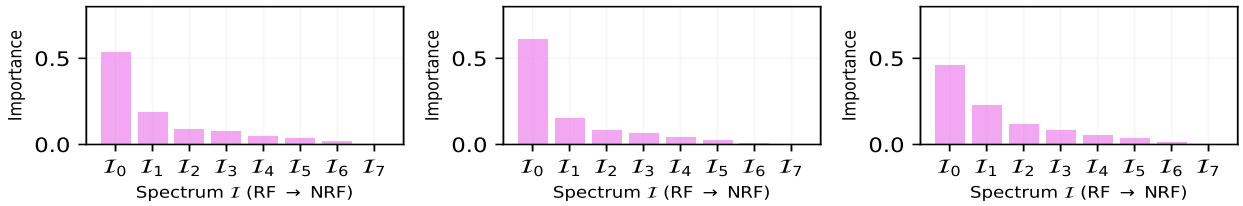

(j) *efficientnet_v2_s* w/ Zeroing (k) *efficientnet_v2_s* w/ Complex Gaussian (l) *efficientnet_v2_s* w/ Replacement

Figure 6: This shows the filtering examples with three masking (absence assignment) strategies: (1) Assigning the spectral absences with constant zeros (Zeroing), (2) assigning the spevtral absences with Gaussian noise (Complex Gaussian) and (3) randomly sampling spectral components from the same image datasets (Replacement). The standard complex Gaussian distribution is given by: $\mathcal{N}(0, \frac{1}{2}) + i\mathcal{N}(0, \frac{1}{2})$. The figures (a), (b) and (c) show the coalition filtering results with the spectral coalition: $\{\mathcal{I}_0\}$. The figures (d), (e) and (f) show the coalition filtering results with the spectral coalition: $\{\mathcal{I}_1, \mathcal{I}_2, \mathcal{I}_3, \mathcal{I}_4, \mathcal{I}_5, \mathcal{I}_6, \mathcal{I}_7\}$. The figures (g) to (l) show the examples of the measured spectral importance distributions of a *resnet18* and a *efficientnet_v2_s* (both are pre-trained on *ImageNet*) with the three assignment strategies.

3.2 Axioms and decomposition

***Efficiency* axiom**: This axiom states that the sum of the divisions of all players must be summed to the worth of player set (the grand coalition): $\sum_{i=0}^{M-1} \psi_i(\mathcal{I}, v) = v(\mathcal{I})$.

***Symmetry* axiom**: Let $\widetilde{\mathcal{I}} \in 2^{\mathcal{I}}$ be some spectral player coalition. Let $\mathcal{I}_i, \mathcal{I}_j \notin \widetilde{\mathcal{I}}$ be two spectral players. The condition $v(\widetilde{\mathcal{I}} \cup \{\mathcal{I}_i\}) = v(\widetilde{\mathcal{I}} \cup \{\mathcal{I}_j\})$ implies that the decomposed contributions for $\mathcal{I}_a$ and $\mathcal{I}_b$ must satisfy the '*equal treatment of equals*' principle such that: $\psi_i(\mathcal{I}, v) = \psi_j(\mathcal{I}, v)$.

***Linearity* axiom**: Let u and v be two characteristic functions. Let $(\mathcal{I}, u)$ and $(\mathcal{I}, v)$ be two coalition games. Let $(u + v)(\widetilde{\mathcal{I}}) := u(\widetilde{\mathcal{I}}) + v(\widetilde{\mathcal{I}})$ where $\widetilde{\mathcal{I}} \in 2^{\mathcal{I}}$. The divisions of the new coalition game $(\mathcal{I}, u + v)$ should satisfy: $\psi_i(\mathcal{I}, u + v) = \psi_i(\mathcal{I}, u) + \psi_i(\mathcal{I}, v)$.

***Dummy player* axiom**: A dummy player (null player) $\mathcal{I}^*$ is the player who has no contribution such that: $V_{\{\mathcal{I}^*\}}(\mu) = 0$ and $V_{\widetilde{\mathcal{I}} \cup \{\mathcal{I}^*\}}(\mu) = V_{\widetilde{\mathcal{I}}}(\mu)$ for $\forall \mathcal{I}^* \notin \widetilde{\mathcal{I}} \subseteq \mathcal{I}$.

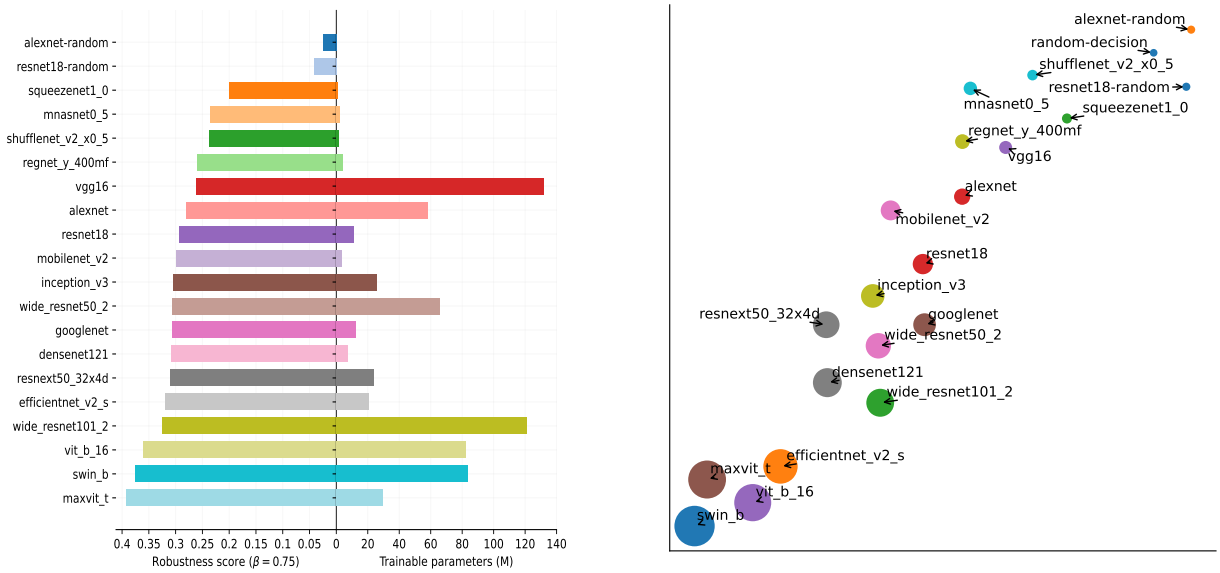

(a) Spectral importance score comparison

(b) Spectral importance distribution projection

Figure 7: This is an application showcase (A1) in research with two experiments. The left bidirectional chart showcases to model spectral importance scores with respect to trainable parameters. The right figure showcases to visualize model robustness by projecting the spectral importance distributions with t-SNE. The shape sizes correspond to robustness (the larger is the better). The experiments are performed on dataset *ImageNet*. The results also correlate with the experiments in Figure 10 regarding adversarial perturbations. In our experiments, the numbers of the trainable parameters of models do not play a crucial role on model robustness.

We directly deploy Shapley value theory to decompose the predictions of models amongst spectral players. The decompositions are uniquely given by:

$$\psi_i(\mathcal{I}, v) = \sum_{\tilde{\mathcal{I}} \subseteq \mathcal{I} \setminus \mathcal{I}_i} \frac{1}{M} \binom{M-1}{|\tilde{\mathcal{I}}|}^{-1} \left\{ v(\tilde{\mathcal{I}} \cup \{\mathcal{I}_i\}) - v(\tilde{\mathcal{I}}) \right\} \tag{1}$$

where $\frac{1}{M}\binom{M-1}{|\tilde{\mathcal{I}}|}^{-1}$ gives the probability of sampling a spectral player $\mathcal{I}_i$ and a spectral player coalition $\tilde{\mathcal{I}}$.

3.3 Spectral coalition filtering

We must implement spectral coalition sampling. We implement the spectral coalition with ideal planar comb filtering. Let $\mathscr{F}$ be the Discrete Fourier transform (DFT) operator and $\mathscr{F}^{-1}$ be inverse DFT (IDFT) operator (Tan & Jiang, 2018). We take some spectral coalition $\tilde{\mathcal{I}}$ and crop the frequency components not present in $\tilde{\mathcal{I}}$ by using a channel-wise plane comb filter on frequency-domain. Let $\mathbb{H}_{\tilde{\mathcal{I}}} = (\mathbb{H}_{\tilde{\mathcal{I}}}(m,n))_{(m,n)\in[M]\times[N]}$ be the filter transfer function for some spectral coalition $\tilde{\mathcal{I}}$ such that:

$$\mathbb{H}_{\tilde{\mathcal{I}}}(m,n) = \begin{cases} 1, & \text{if the frequency point (m,n) in } \tilde{\mathcal{I}}, \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

where $M \times N$ is the dimension of image. We define an operator '$\star$' to represent the filtering by:

$$\tilde{x}_{\tilde{\mathcal{I}}} = x \star \tilde{\mathcal{I}} = \mathscr{F}^{-1}\left[\underbrace{\mathscr{F}(x) \odot \mathbb{H}_{\tilde{\mathcal{I}}}}_{\text{Spectral presence}} + \underbrace{\boldsymbol{b} \odot (\mathbb{1} - \mathbb{H}_{\tilde{\mathcal{I}}})}_{\text{Spectral absence}} \right] \tag{3}$$

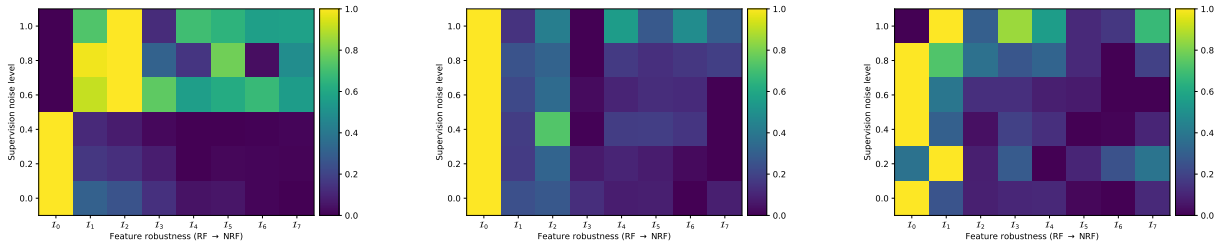

(a) *googlenet @Caltech101* (b) *resnet18* (c) *vit*

Figure 8: This is an application showcase (A2) demonstrating how models respond to the label noise of a human-annotated dataset in training. We create the noisy-label datasets from clean human-annotated *Caltech101* by randomly re-assigning a proportion of labels. We vary the proportions (label noise levels) from 0 to 1 with stride 0.1. We train *googlenet*, *resnet18* and *vit* on the derived noisy-label datasets with 70 epochs, learning rate 0.0025 and SGD optimizer. The image sizes are set to 64×64, 64×64 and 224×224 respectively. The results are visualized in 2D heat maps. We have observed the existence of a 'spectral drift' effect correlated with the levels of supervision signal noise in our preliminary showcase experiments. The models trained with lower levels of supervision signal noise tend to rely on the features dominated by low-frequency components, while the models trained with higher levels of supervision noise do not exhibit such a spectral preference. It is noteworthy that the 'spectral drift' effect is not a simple monotonic correlation with the levels of supervision noise. This phenomenon requires further investigation with additional experiments. Models exhibit disparate learning dynamics on human-annotated datasets and random-annotated datasets. As the theoretical analysis shown in Section 3.4, **I-ASIDE** is a quantification method for measuring the spectral importance of the variational mutual information between feature representations and labels. Intuitively, we surmise this '*spectral drift*' effect originates from that the supervision signals come from humans and the relevant features perceived by humans largely concentrate on low-frequency components due to the biological limit of our vision system. The training objectives for usual supervised classifiers are equivalent to estimate a variational lower bound of the mutual information between images and labels (See Section 3.4). Yet, it is not in the scope of this research.

where '$\odot$' denotes Hadamard product, $\mathbb{1}$ denotes an all-ones matrix and $\boldsymbol{b}$ represents the assignments of the absences of spectral players. At our implementation, we empirically set $\boldsymbol{b} = \mathbf{0}$.

Spectral player absence assignment: In the implementation of the coalition filtering, the masking (assignment) strategies of the absences of the spectral players in coalition filtering can change the distributions of filtered images. There exist multiple choices for the assignments of the absences of spectral layers in coalition filtering design: (1) Assigning to constant zeros (Zeroing), (2) assigning to complex Gaussian noise (Complex Gaussian) and (3) assigning to the corresponding frequency components randomly sampled from other images at the same dataset (Replacement). When we adopt 'Complex Gaussian' strategy, the equation (3) changes to: $\boldsymbol{b} = \mathcal{N}(\boldsymbol{\mu}, \frac{\boldsymbol{\sigma}^2}{2}) + i\mathcal{N}(\boldsymbol{\mu}, \frac{\boldsymbol{\sigma}^2}{2})$ (where $\mathcal{N}(\boldsymbol{\mu}, \frac{\boldsymbol{\sigma}^2}{2})$ is some multivariate Gaussian distribution with the same dimension as input image $\boldsymbol{x}$). When we adopt 'Replacement' strategy, the equation (3) changes to: $\boldsymbol{b} = \mathscr{F}(x^*)$ (where $x^* \sim \mathcal{X}$ is a randomly sampled image from some set $\mathcal{X}$). At our implementation, we simply choose 'zeroing' as our filtering strategy: $\boldsymbol{b} = \mathbf{0}$. In Figure 6, we show the filtered image examples by using above three strategies and also show the examples of measured spectral importance distributions. In our empirical observations, the three strategies have rather similar performance. In this research, we do not unfold the discussions regarding the masking strategy choices. We will investigate the masking strategy choices in our future work.

3.4 Characteristic function design and characterization

Let $Q(y|x) : x \mapsto [0, 1]$ be some image classifier. For a sample x, $Q(y|x)$ sends the x to a discrete category probability distribution for $y \in \mathcal{Y}$. Let $Q(\mathcal{X})$ denote the representations of image set $\mathcal{X}$ with given classifier

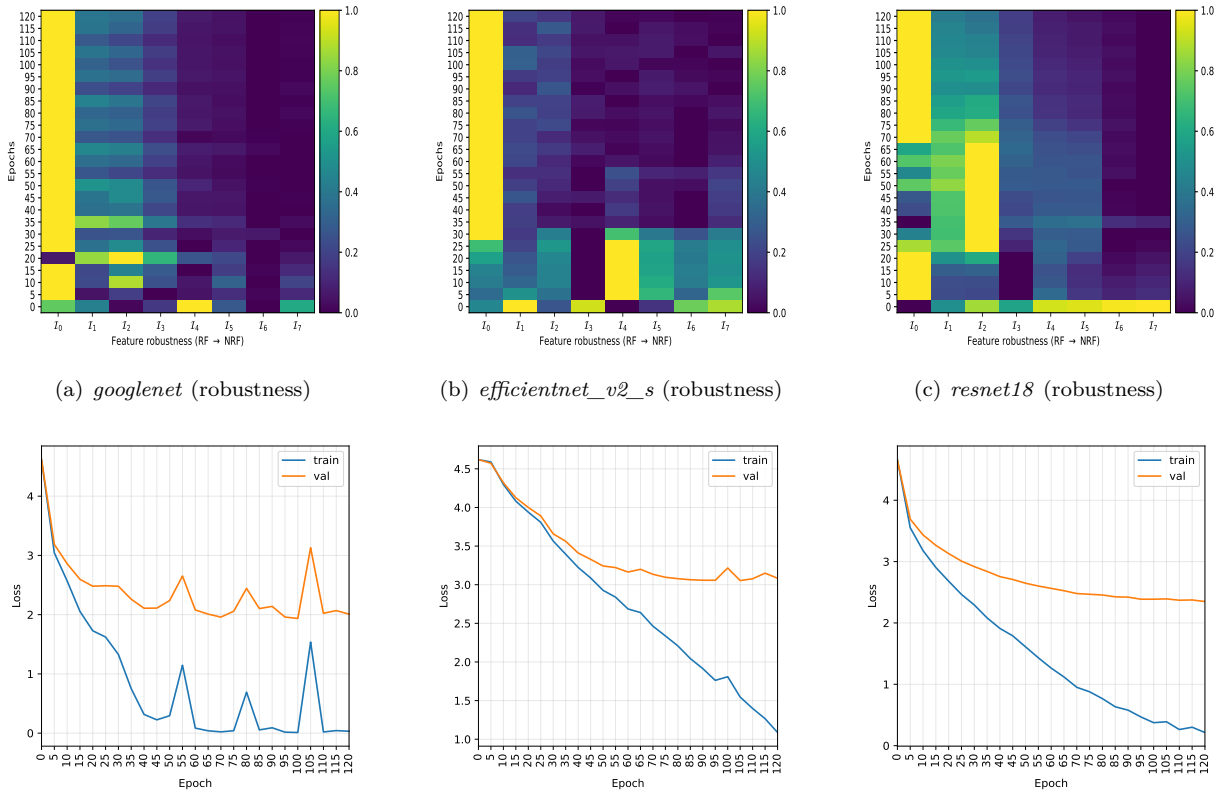

(a) *googlenet* (robustness) (b) *efficientnet_v2_s* (robustness) (c) *resnet18* (robustness)

(d) *googlenet* (loss) (e) *efficientnet_v2_s* (loss) (f) *resnet18* (loss)

Figure 9: This is an application showcase (A3) in research demonstrating how the robustness of feature representations of models evolves with respect to training epochs. We train '*googlenet*', '*efficientnet_v2_s*' and '*resnet18*' with 120 epochs, SGD optimizer and learning rate 0.0025. The training dataset is *Caltech*101. The spectral importance distributions are measured for every 5 epochs using **I-ASIDE**. We present the results using 2D heat maps with the corresponding training and validation loss curves. The first row is the training dynamics of feature robustness. The second row is the corresponding training and validation loss curve. The results suggest that models have two major learning stages: (1) Feature exploration stage and (2) feature exploitation stage. At the feature exploration stage (1), both training loss and validation loss drop synchronously. For example, the epochs 1 to 35 for *googlenet*, the epochs 1 to 40 for *efficientnet_v2_s* and the epochs 1 to 70 for *resnet18*. This stage explores the features which are more general and the peaks of spectral distributions move towards low-frequency components. At the exploitation (2), models continue to refine the learned feature representations for further lower training loss. However, overfitting emerges at this stage. This stage uses more features dominated by low-frequency components. The results echo with the claim in a prior literature (Wang et al., 2020): High-frequency components are essential for the generalization capability of image models (Wang et al., 2020).

Q. We design the characteristic function as:

$$v(\widetilde{\mathcal{I}}) := \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \left\{ \log Q(y|x \star \widetilde{\mathcal{I}}) - \log Q(y|x \star \varnothing) \right\} \tag{4}$$

which satisfies $v(\varnothing) = 0$ if the absences of spectral players are assigned to zeros in coalition filtering (See Section 3.3).

We justify such a choice by showing that the characteristic function measures the mutual information between the representations $Q(\mathcal{X})$ and labels $\mathcal{Y}$. Image classifiers trained with cross-entropy loss estimate a variational lower bound of the mutual information between images and labels (Qin et al., 2021; Barber & Agakov, 2004).

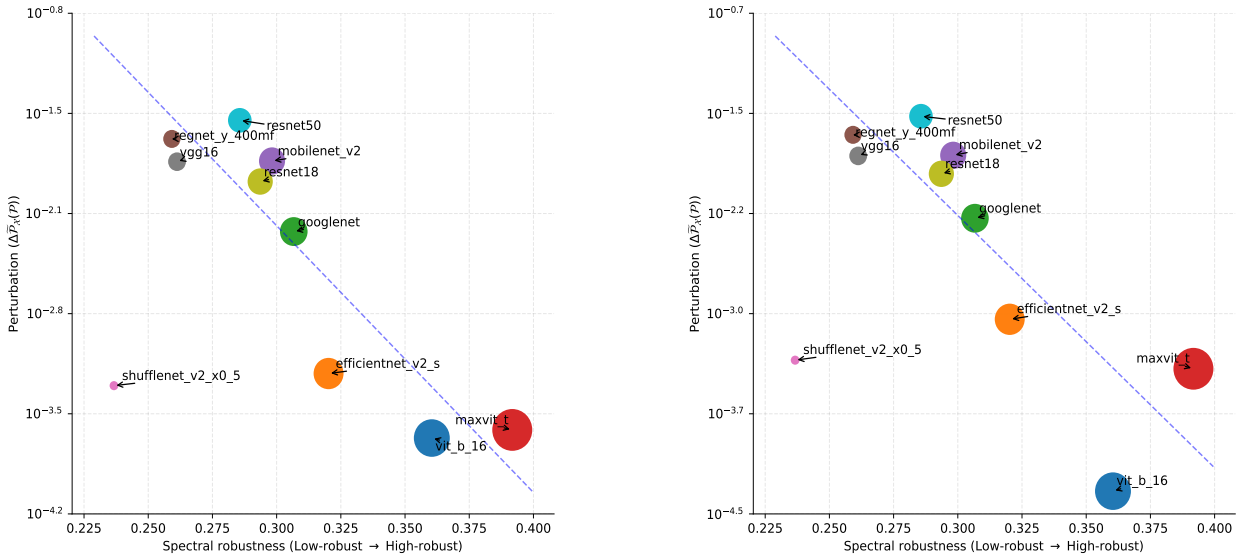

(a) FGSM ($eps = 0.1$)

(b) PGD ($eps = 0.1$)

Figure 10: This is an application showcase (A4) in research demonstrating that there is a correlation between adversarial perturbations and the spectral importance scores using **I-ASIDE**. The adversarial perturbations are measured by the statistical expectations of prediction probability variations (See Section 4.4). We choose multiple pretrained models (on *ImageNet*) and perform un-targeted FGSM/PGD attacking with $eps = 0.1$. The adversarial samples carry a high proportion of non-robust features (Ilyas et al., 2019) (See Figure 3). For most models, adversarial perturbations are negatively proportional to spectral importance scores. The circle sizes are proportional to spectral importance scores. It is also notable there are some outliers (e.g. *shufflenet*), and this implies the perturbation-based robustness research can merely capture an aspect of the holistic robustness of models.

The optimization objective of some image classifier with usual cross-entropy loss in learning is to maximize the likelihood such that:

$$Q = \arg \sup_{\forall\ Q^*} \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q^*(y|x) \tag{5}$$

in which:

$$\sup_{\forall\ Q^*} \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q^*(y|x) \leqslant \mathbb{I}(\mathcal{X};\mathcal{Y}) - H(\mathcal{Y}) \tag{6}$$

where $H(\mathcal{Y})$ is the differential entropy of $\mathcal{Y}$ and $\mathbb{I}(\mathcal{X};\mathcal{Y})$ is the mutual information between images and labels (Qin et al., 2021; Poole et al., 2019; Barber & Agakov, 2004). We will provide a brief proof below. Hence, the equation 6 gives a variational lower bound of the mutual information between images and labels. We denote such a bound as:

$$\mathbb{I}_{BA}(Q(\mathcal{X});\mathcal{Y}) \stackrel{\text{def}}{=} \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q(y|x) + H(\mathcal{Y}). \tag{7}$$

From this point of view, the characteristic function measures the mutual information between representations and labels, and can also be written as:

$$v(\widetilde{\mathcal{I}}) := \mathbb{I}_{BA}(Q(\mathcal{X} \star \widetilde{\mathcal{I}}),\mathcal{Y}) - \mathbb{I}_{BA}(Q(\mathcal{X} \star \varnothing),\mathcal{Y}). \tag{8}$$

Proof. Barber & Agakov shows a tight variational lower bound of mutual information (Barber & Agakov, 2004). Qin et al. also demonstrates that classifiers trained with cross-entropy loss estimate the mutual

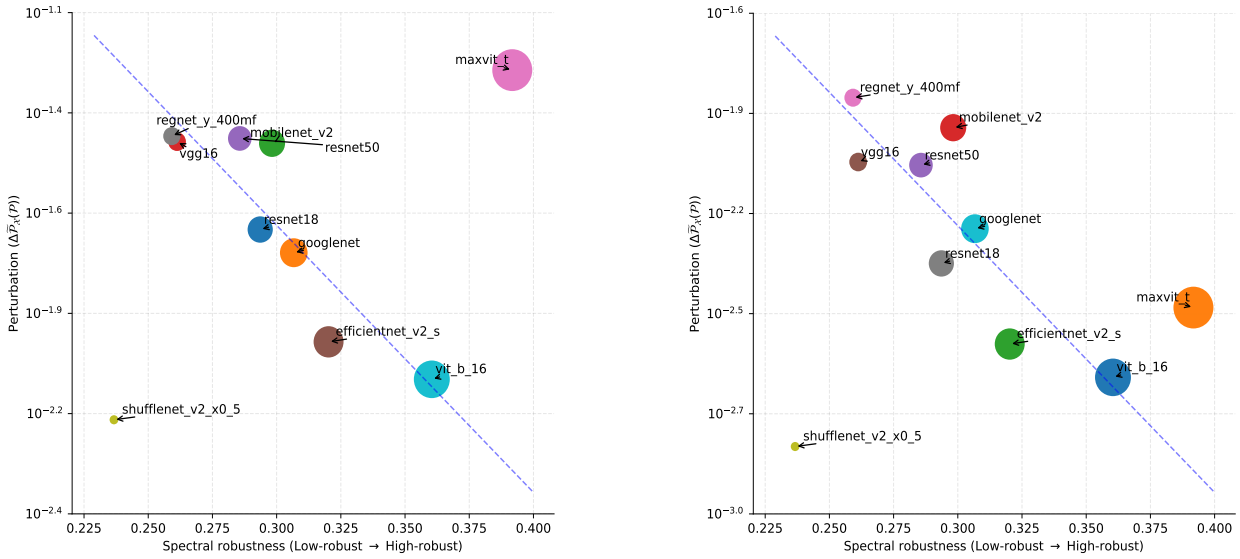

(a) White noise ($\sigma = 0.1$)

(b) Gaussian-blur (3×3)

Figure 11: This is an application showcase (A5) in research demonstrating that there is a correlation between out-of-distribution (OOD) perturbations and the spectral importance scores using **I-ASIDE**. The OOD perturbations are measured by the statistical expectations of prediction probability variations (See Section 4.5). We choose multiple pretrained models (on *ImageNet*) and evaluate the perturbations with two OOD transforms: Gaussian ($\sigma = 0.1$) and Gaussian-blur ($kernel = 3 \times 3$). The results show that the OOD perturbations are negatively proportional to spectral importance scores. The circle sizes are proportional to spectral importance scores. It is also notable there are some outliers (e.g. *maxvit_t*), and this implies the perturbation-based robustness research can merely capture an aspect of the holistic robustness of models.

information between inputs and labels (Qin et al., 2021). We provide a brief proof for equation 6 and equation 7. The proof technique can be found in multiple literature (Poole et al., 2019; Barber & Agakov, 2004; Qin et al., 2021).

Let $(\mathcal{X}, \mathcal{Y})$ be some dataset. Assume $\mathcal{X} \times \mathcal{Y}$ is compact in $\mathbb{R}^d$ and all probability measures on which are absolutely continuous in terms of Lebesgue measure. Let $P(y|x)$ and $P(y)$ be the prior conditional distribution and the prior label distribution on the dataset. The mutual information between $\mathcal{X}$ and $\mathcal{Y}$ is bounded above the mutual information between the representations of $\mathcal{X}$ (denoted as $Q(\mathcal{X})$) and the labels for the image classifier $Q(y|x)$. Let $Q^*(y|x) \in \{T|T(y|x) : x \mapsto (0,1]\}$ be an arbitrary probability measure map. Starting from the definition of mutual information:

$$\mathbb{I}(\mathcal{X};\mathcal{Y}) = \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log \frac{P(y|x)}{P(y)} \geqslant \sup_{Q^*} \left\{ \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log \frac{P(y|x)}{P(y)} \frac{Q^*(y|x)}{Q^*(y|x)} \right\} \tag{9}$$

$$= \sup_{Q^*} \left\{ \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q^*(y|x) - \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log P(y) + \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log \frac{P(y|x)}{Q^*(y|x)} \right\} \tag{10}$$

$$\geqslant \sup_{Q^*} \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q^*(y|x) + H(\mathcal{Y}) \tag{11}$$

$$= \sup_{Q^*} \mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log Q^*(y|x) + H(\mathcal{Y}) \stackrel{\text{def}}{=} \mathbb{I}_{BA}(Q(\mathcal{X});\mathcal{Y}) \tag{12}$$

where $\mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} \log \frac{P(y|x)}{Q^*(y|x)} = \mathop{\mathbb{E}}_{x\sim\mathcal{X}} \{KL\left[P(y|x)||Q^*(y|x)\right]\} \geqslant 0$ since $P(x,y) = P(y|x)P(x)$ and the non-negativity property of Kullback–Leibler divergence (Gibbs' inequality) (Poole et al., 2019; Barber & Agakov, 2004; Qin et al., 2021; Kullback, 1997). The bound is tight *if and only if*: $Q^* = P$. Readers can refer the similar proof technique and the definition to the equation (2) in literature (Poole et al., 2019) and the equation (1) and the equation (3) in literature (Barber & Agakov, 2004). □

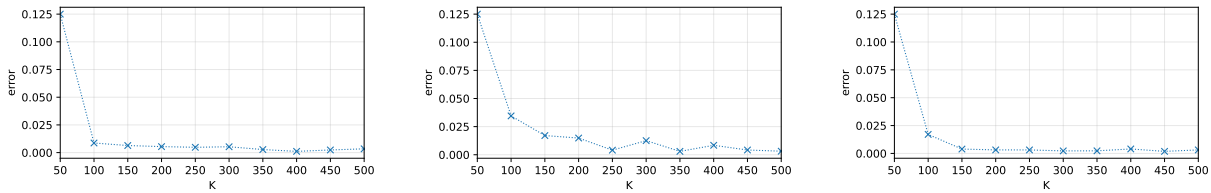

(a) *resnet18 @CIFAR10* (b) *resnet18 @CIFAR100* (c) *resnet18 @ImageNet*

Figure 12: This shows how the relative estimation errors converge with respect to the numbers of samples K. The errors are measured by: $\frac{1}{M}||\Psi^{(i+1)}(v) - \Psi^{(i)}(v)||_1$ where $\Psi^{(i)}(v)$ denotes the i-th measured spectral importance distribution with respect to characteristic function v. The experiments are conducted on *CIFAR10*, *CIFAR100* and *ImageNet* with *resnet18*.

3.5 Summarizing spectral importance distribution

We next discuss how to summarize the distributions into real scalars. Let $S(v) : \Psi(v) \mapsto [0, 1]$ be a functional on field $\mathbb{R}$ which summarizes the derived spectral importance distributions into real scalars. There are several choices to summarize according to our needs. We are interested in feature representation robustness. We summarize the scores by summing with a feature representation robustness prior series as weights: The learned feature representations dominated by low-frequency components are more robust than those feature representations dominated by high-frequency components.

Intuitively, we can rank the importance of the M spectral players as: $\widetilde{\boldsymbol{\beta}} := (0, -1, -2, \cdots, -(M-1))$. We use an exponential transform with base β to convert such a ranking vector $\widetilde{\boldsymbol{\beta}}$ into a non-negative series: $\boldsymbol{\beta} = (\beta^0, \beta^1, \cdots, \beta^{M-1})$ where $0 < \beta < 1$. The weighting vector $\boldsymbol{\beta}$ preserves the ranking information but the weights vary according to the choice of β. Set $\Psi^*(v) = \frac{\Psi(v) - \min \Psi(v)}{||\Psi(v) - \min \Psi(v)||_1}$ for normalizing $\Psi(v)$. The normalized $\Psi^*(v)$ can be summarized by the inner product with the prior weighting vector $\boldsymbol{\beta}$: $\boldsymbol{\beta}^T \Psi^*(v)$. We further take into account the random decision as a baseline by the same model with randomly initialized weights:

$$\left| \boldsymbol{\beta}^T \Psi^*(v) - \underbrace{\mathop{\mathbb{E}}_{\widetilde{v}}[\boldsymbol{\beta}^T \Psi^*(\widetilde{v})]}_{\text{baseline}} \right| \tag{13}$$

where $\widetilde{v}$ denotes the characteristic function of some model with randomly initialized weights and $\mathop{\mathbb{E}}_{\widetilde{v}}[\boldsymbol{\beta}^T \Psi^*(\widetilde{v})]$ denotes the statistical expectation of the measured distributions when the model over all random weights. We set:

$$\mathop{\mathbb{E}}_{\widetilde{v}}[\boldsymbol{\beta}^T \Psi^*(\widetilde{v})] \approx \boldsymbol{\beta}^T \frac{\mathbb{1}}{M}. \tag{14}$$

The result can further be normalized into $[0, 1]$ by:

$$S(v) := \frac{\left|\boldsymbol{\beta}^T \Psi^*(v) - \boldsymbol{\beta}^T \frac{\mathbb{1}}{M}\right|}{\sup \left|\boldsymbol{\beta}^T \Psi^*(v) - \boldsymbol{\beta}^T \frac{\mathbb{1}}{M}\right|} = \left| \frac{\boldsymbol{\beta}^{*T} \Psi^*(v) - \eta}{1 - \eta} \right| \tag{15}$$

where $\beta \in (0, 1)$, $\boldsymbol{\beta}^* = \frac{\boldsymbol{\beta}}{||\boldsymbol{\beta}||_2}$ and $\eta = \frac{1}{M} \frac{||\boldsymbol{\beta}||_1}{||\boldsymbol{\beta}||_2}$. The deduction of this formula is provided in supplementary material.

3.6 Estimation of error upper bound

Suffering from the limit of computational resources, the statistical expectation is evaluated by using Monte Carlo sampling. We analyze the error bound by taking K samples. Let:

$$\Delta v(\tilde{\mathcal{I}}, \mathcal{I}_i) := v(\tilde{\mathcal{I}} \cup \{\mathcal{I}_i\}) - v(\tilde{\mathcal{I}}) \tag{16}$$

and

$$\Delta v(\mathcal{I}_i) := \left(\Delta v(\tilde{\mathcal{I}}, \mathcal{I}_i)\right)_{\tilde{\mathcal{I}} \subseteq \mathcal{I}} \tag{17}$$

and

$$W := \left(\frac{1}{M} \binom{M-1}{|\tilde{\mathcal{I}}|}^{-1} \right)_{\tilde{\mathcal{I}} \subseteq \mathcal{I}}. \tag{18}$$

Hence:

$$\psi_i(\mathcal{I}, v) = W^T \Delta v(\mathcal{I}_i) \tag{19}$$

where $||W||_1 \equiv 1$ since W is a probability distribution. Let ψ_i^*, $\Delta v^*(\mathcal{I}_i)$ and $\Delta v^*(\tilde{\mathcal{I}}, \mathcal{I}_i)$ be estimations with K samples using Monte Carlo sampling. The error bound with ℓ_1 norm is given by:

$$\epsilon \stackrel{\text{def}}{=} \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} ||\psi_i^*(\mathcal{I}, v) - \psi_i(\mathcal{I}, v)||_1 = \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} ||W^T \Delta v^*(\mathcal{I}_i) - W^T \Delta v(\mathcal{I}_i)||_1 \tag{20}$$

$$\leqslant \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} ||W||_1 \cdot ||\Delta v^*(\mathcal{I}_i) - \Delta v(\mathcal{I}_i)||_\infty \qquad \left(\text{Hölder's inequality}\right) \tag{21}$$

$$= \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} || \sum_{\tilde{\mathcal{I}} \subseteq \mathcal{I} \setminus \mathcal{I}_i} \left(\Delta v^*(\tilde{\mathcal{I}}, \mathcal{I}_i) - \Delta v(\tilde{\mathcal{I}}, \mathcal{I}_i)\right) ||_\infty \tag{22}$$

$$\leqslant \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} 2^{M-1} \sup_{\tilde{\mathcal{I}}} ||\Delta v^*(\tilde{\mathcal{I}}, \mathcal{I}_i) - \Delta v(\tilde{\mathcal{I}}, \mathcal{I}_i)||_\infty \tag{23}$$

$$= \sup_i \mathop{\mathbb{E}}_{(\mathcal{X},\mathcal{Y})} 2^{M-1} \sup_{\tilde{\mathcal{I}}} ||\Delta v^*(\tilde{\mathcal{I}}, \mathcal{I}_i) - \Delta v(\tilde{\mathcal{I}}, \mathcal{I}_i)||_1 \tag{24}$$

$$\leqslant 2^{M-1} \cdot \left\{ \frac{Var(\Delta v^*)}{K} \right\}^{\frac{1}{2}} \tag{25}$$

where $Var(\Delta v^*)$ gives the upper bound of the variance of $\Delta v^*(\tilde{\mathcal{I}}, \mathcal{I}_i)$. Figure 12 shows the experiments of the estimation errors of a *resnet18* model on *CIFAR10*, *CIFAR100* and *ImageNet* datasets. In our experiments, we empirically choose $K = 200$ due to this error upper bound.

4 Applications

We showcase multiple applications to demonstrate the potential uses of **I-ASIDE**[3]: (A1) Quantifying model global robustness by examining spectral importance distributions, (A2) understanding the learning behaviours of image models on the datasets with supervision noise, (A3) investigating the learning dynamics of models from the perspective of the evolution of feature representation robustness in optimization, (A4) understanding the adversarial vulnerability of models by examining spectral importance distributions and (A5) studying image model out-of-distribution (OOD) robustness.

4.1 Showcase A1: Quantifying model global robustness

We measure the robustness by examining the spectral importance distributions or the spectral importance scores.

Spectral importance distribution: The experiments in Figure 4 showcase the measured spectral distributions of multiple models with multiple datasets. The results show that the spectral importance of the learned feature representations of trained models is non-uniform. Furthermore, models trained on larger training datasets exhibit higher robustness.

Numerical comparison: The experiment (a) in Figure 7 showcases the application of numerically comparing model robustness. The results correlate with the adversarial perturbation experiments in Figure 10 in which

[3]The core code implementation is provided in supplementary material.

we measure the prediction probability variations between clean samples and adversarial samples for given labels with un-targeted FGSM/PGD attacking (Szegedy et al., 2013; Moosavi-Dezfooli et al., 2016; Goodfellow et al., 2014; Madry et al., 2017). The result implies that the numbers of trainable parameters seem not to play a crucial aspect on model robustness.

Visualizing by projection: The experiment (b) in Figure 7 showcases the projection of spectral importance distributions. The projection is performed by using t-SNE (Hinton & Roweis, 2002).

4.2 Showcase A2: Understanding how models respond to supervision noise

The experiments in Figure 8 showcase the investigation of how models respond to various label noise levels in the training process. We conduct the experiment using the *Caltech101* dataset. We create noisy-label datasets from the clean *Caltech101* by randomly re-assigning a proportion (referred to as'label noise level' or 'supervision noise level') of clean labels. We vary the label noise level from 0 to 1 with stride 0.1. The results show that the models trained with higher label noise levels tend to use a higher proportion of non-robust features.

4.3 Showcase A3: Investigating how feature representations are learned in training

The experiments in Figure 9 showcase how the feature representations learned by models evolve with respect to epochs in training process. The learning dynamics of models exhibit two stages. In the first stage, models readily learn more robust feature representations. In the second stage, models learn more non-robust feature representations to achieve higher accuracy.

4.4 Showcase A4: Understanding adversarial vulnerability

Adversarial samples carry a high proportion of non-robust features (Su et al., 2018; Ilyas et al., 2019; Bai et al., 2021; Tsipras et al., 2018). Intuitively, the models with high spectral importance scores are less likely susceptible to adversarial perturbations. In Figure 10, we show how FGSM/PGD (Yuan et al., 2019; Madry et al., 2017; Akhtar & Mian, 2018; Chakraborty et al., 2018) adversarial perturbations correlate with the spectral importance scores by **I-ASIDE**. The full results are provided in supplementary material. We measure the prediction probability variations between the prediction probabilities of clean samples and adversarial samples. We use the 'smooth' version of Softmax function with temperature $\mathcal{T}$ to convert logits into probabilities (Hinton et al., 2015; Goodfellow et al., 2016; Jang et al., 2016).

Experimental method: Let I be some clean image and I^* be the corresponding adversarial sample. Let $\mathcal{P}(y|x) : x \mapsto [0, 1]$ be some image classifier which outputs the probability with respect to some category y. We define the adversarial perturbation $\Delta\widetilde{\mathcal{P}}$ as $\mathop{\mathbb{E}}_{x,y\sim(\mathcal{X},\mathcal{Y})} |\mathcal{P}(y|x) - \mathcal{P}(y|x^*)|$ over some dataset $(\mathcal{X}, \mathcal{Y})$ where (x, y) denotes an image and the corresponding label and x^* denotes adversarial sample for x. We plot the prediction perturbations against spectral importance scores.

The results in Figure 10 shows that the adversarial perturbations are negatively correlated with spectral importance scores. The outliers remind us: Spectral importance distributions can merely reflect one aspect of the robustness of models.

4.5 Showcase A5: Studying model out-of-distribution (OOD) robustness

We also showcase how the spectral importance scores by using **I-ASIDE** correlates with the OOD perturbations to further justify our methodology. The measurement method is the same as in Section 4.4. We choose two OOD perturbations: Gaussian ($\sigma = 0.1$) and blurring with Gaussian kernel ($kernel = 3 \times 3$). The preliminary results show that there exists a correlation between OOD perturbations and spectral importance scores.

5 Limitations

I-ASIDE provides a unique insight into the interpretability of image models from the feature robustness on spectral perspective. Yet, our methodology has several limitations: (1) Spectral perspective can merely reflect one aspect of the holistic view of model robustness, (2) the masking strategies in coalition filtering have yet to be investigated in future, (3) the choice of β in summarizing model robustness scores and (4) the resolution of the measured spectral distributions for image models.

Regarding the limitation (1), for example, carefully crafted malicious adversarial perturbations on low-frequency components can fool neural networks (Luo et al., 2022; Liu et al., 2023; Maiya et al., 2021). This further implies the complexity of this research topic. Regarding the limitation (2), we have preliminarily discussed several masking strategies (See Figure 6 and Section 3.3). We need further investigation regarding how to choose the optimal masking strategy. Regarding the limitation (3), providing a better summarizing function ($S : \mathbb{R}^M \mapsto [0, 1]$) needs further investigation in future work. Regarding the limitation (4), as the spectral energy follows power law alike distribution, measuring the expectations of spectral contributions on decisions on high-frequency components poses a challenge. As such, **I-ASIDE** does not use sampling based approach on spectral coalitions to avoid inaccurate results and thus suffers from the computation cost by $\mathcal{O}(2^M)$. Fortunately, we do not need high spectral resolution to analyze model robustness problem.

6 Conclusions

This research is motivated by the link between feature robustness and the non-uniformity of the spectral structures of natural images, in tandem with the insight that neural networks are parameterized non-linear signal filters. **I-ASIDE** investigates how the 'neural signal filter' responds to signals on spectrum. We formulate the global feature importance as a decomposition problem on spectrum using coalition game theory. We theoretically show that our approach decomposes the mutual information between feature representations and labels onto spectrum. Our work provides a unique insight into deep learning research and enables a considerable number of applications as we have demonstrated. We have conducted a large number of experiments to support our claims.

References

Naveed Akhtar and Ajmal Mian. Threat of adversarial attacks on deep learning in computer vision: A survey. *Ieee Access*, 6:14410–14430, 2018.

André Altmann, Laura Toloşi, Oliver Sander, and Thomas Lengauer. Permutation importance: a corrected feature importance measure. *Bioinformatics*, 26(10):1340–1347, 2010.

Alejandro Barredo Arrieta, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bennetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information fusion*, 58:82–115, 2020.

Robert J Aumann and Lloyd S Shapley. *Values of non-atomic games*. Princeton University Press, 2015.

Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015.

Tao Bai, Jinqi Luo, Jun Zhao, Bihan Wen, and Qian Wang. Recent advances in adversarial training for adversarial robustness. *arXiv preprint arXiv:2102.01356*, 2021.

David Barber and Felix Agakov. The im algorithm: a variational approach to information maximization. *Advances in neural information processing systems*, 16(320):201, 2004.

David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 6541–6549, 2017.

Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and R Devon Hjelm. Mine: mutual information neural estimation. *arXiv preprint arXiv:1801.04062*, 2018.

Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. Adversarial attacks and defences: A survey. *arXiv preprint arXiv:1810.00069*, 2018.

Zhi Chen, Yijie Bei, and Cynthia Rudin. Concept whitening for interpretable image recognition. *Nature Machine Intelligence*, 2(12):772–782, 2020.

Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain markov process expectations for large time. iv. *Communications on pure and applied mathematics*, 36(2):183–212, 1983.

Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. Towards automatic concept-based explanations. *Advances in Neural Information Processing Systems*, 32, 2019.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144, 2020.

Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.

Brandon M Greenwell, Bradley C Boehmke, and Andrew J McCarthy. A simple and effective model-based variable importance measure. *arXiv preprint arXiv:1805.04755*, 2018.

Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Dino Pedreschi, Franco Turini, and Fosca Giannotti. Local rule-based explanations of black box decision systems. *arXiv preprint arXiv:1805.10820*, 2018.

Sergiu Hart. Shapley value. In *Game theory*, pp. 210–216. Springer, 1989.

Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019.

Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.

Geoffrey E Hinton and Sam Roweis. Stochastic neighbor embedding. *Advances in neural information processing systems*, 15, 2002.

Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. *Advances in neural information processing systems*, 32, 2019.

Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.

Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*, pp. 2668–2677. PMLR, 2018.

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

Pang Wei Koh, Thao Nguyen, Yew Siang Tang, Stephen Mussmann, Emma Pierson, Been Kim, and Percy Liang. Concept bottleneck models. In *International Conference on Machine Learning*, pp. 5338–5348. PMLR, 2020.

Stefan Kolek, Duc Anh Nguyen, Ron Levie, Joan Bruna, and Gitta Kutyniok. Cartoon explanations of image classifiers. In *European Conference on Computer Vision*, pp. 443–458. Springer, 2022.

Solomon Kullback. *Information theory and statistics*. Courier Corporation, 1997.

Himabindu Lakkaraju, Ece Kamar, Rich Caruana, and Jure Leskovec. Faithful and customizable explanations of black box models. In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pp. 131–138, 2019.

Pantelis Linardatos, Vasilis Papastefanopoulos, and Sotiris Kotsiantis. Explainable ai: A review of machine learning interpretability methods. *Entropy*, 23(1):18, 2020.

Zachary C Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31–57, 2018.

Jiyuan Liu, Bingyi Lu, Mingkang Xiong, Tao Zhang, and Huilin Xiong. Low frequency sparse adversarial attack. *Computers & Security*, 132:103379, 2023.

Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30, 2017.

Cheng Luo, Qinliang Lin, Weicheng Xie, Bizhu Wu, Jinheng Xie, and Linlin Shen. Frequency-driven imperceptible adversarial attack on semantic similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15315–15324, 2022.

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.

Shishira R Maiya, Max Ehrlich, Vatsal Agarwal, Ser-Nam Lim, Tom Goldstein, and Abhinav Shrivastava. A frequency perspective of adversarial robustness. *arXiv preprint arXiv:2111.00861*, 2021.

Grégoire Montavon, Alexander Binder, Sebastian Lapuschkin, Wojciech Samek, and Klaus-Robert Müller. Layer-wise relevance propagation: an overview. *Explainable AI: interpreting, explaining and visualizing deep learning*, pp. 193–209, 2019.

Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the number of linear regions of deep neural networks. *Advances in neural information processing systems*, 27, 2014.

Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard. Deepfool: a simple and accurate method to fool deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2574–2582, 2016.

Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. *Advances in neural information processing systems*, 29, 2016a.

Anh Nguyen, Jason Yosinski, and Jeff Clune. Multifaceted feature visualization: Uncovering the different types of features learned by each neuron in deep neural networks. *arXiv preprint arXiv:1602.03616*, 2016b.

Anh Nguyen, Jason Yosinski, and Jeff Clune. Understanding neural networks via feature visualization: A survey. *Explainable AI: interpreting, explaining and visualizing deep learning*, pp. 55–76, 2019.

Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-gan: Training generative neural samplers using variational divergence minimization. *Advances in neural information processing systems*, 29, 2016.

Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature visualization. *Distill*, 2(11):e7, 2017.

Ben Poole, Sherjil Ozair, Aaron Van Den Oord, Alex Alemi, and George Tucker. On variational bounds of mutual information. In *International Conference on Machine Learning*, pp. 5171–5180. PMLR, 2019.

Zhenyue Qin, Dongwoo Kim, and Tom Gedeon. Neural network classifier as mutual information evaluator. *arXiv preprint arXiv:2106.10471*, 2021.

Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *international conference on machine learning*, pp. 2847–2854. PMLR, 2017.

Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. On the spectral bias of neural networks. In *International Conference on Machine Learning*, pp. 5301–5310. PMLR, 2019.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pp. 1135–1144, 2016.

Alvin E Roth. *The Shapley value: essays in honor of Lloyd S. Shapley*. Cambridge University Press, 1988.

Andrew M Saxe, Yamini Bansal, Joel Dapello, Madhu Advani, Artemy Kolchinsky, Brendan D Tracey, and David D Cox. On the information bottleneck theory of deep learning. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124020, 2019.

Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pp. 618–626, 2017.

Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *International conference on machine learning*, pp. 3145–3153. PMLR, 2017.

Samuel Henrique Silva and Peyman Najafirad. Opportunities and challenges in deep learning adversarial robustness: A survey. *arXiv preprint arXiv:2007.00753*, 2020.

Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013.

Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. *arXiv preprint arXiv:1706.03825*, 2017.

Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy?–a comprehensive study on the robustness of 18 deep image classification models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 631–648, 2018.

Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *International conference on machine learning*, pp. 3319–3328. PMLR, 2017.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks. *arXiv preprint arXiv:1312.6199*, 2013.

Lizhe Tan and Jean Jiang. *Digital signal processing: fundamentals and applications*. Academic Press, 2018.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. *arXiv preprint arXiv:1805.12152*, 2018.

Yusuke Tsuzuku and Issei Sato. On the structural sensitivity of deep convolutional networks to the directions of fourier basis functions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 51–60, 2019.

Petar Veličković, William Fedus, William L Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep graph infomax. *arXiv preprint arXiv:1809.10341*, 2018.

Haohan Wang, Xindi Wu, Zeyi Huang, and Eric P Xing. High-frequency component helps explain the generalization of convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8684–8694, 2020.

Eyal Winter. The shapley value. *Handbook of game theory with economic applications*, 3:2025–2054, 2002.

Zhi-Qin John Xu, Yaoyu Zhang, Tao Luo, Yanyang Xiao, and Zheng Ma. Frequency principle: Fourier analysis sheds light on deep neural networks. *arXiv preprint arXiv:1901.06523*, 2019a.

Zhi-Qin John Xu, Yaoyu Zhang, and Yanyang Xiao. Training behavior of deep neural network in frequency domain. In *International Conference on Neural Information Processing*, pp. 264–274. Springer, 2019b.

Xiaoyong Yuan, Pan He, Qile Zhu, and Xiaolin Li. Adversarial examples: Attacks and defenses for deep learning. *IEEE transactions on neural networks and learning systems*, 30(9):2805–2824, 2019.

Matthew D Zeiler, Dilip Krishnan, Graham W Taylor, and Rob Fergus. Deconvolutional networks. In *2010 IEEE Computer Society Conference on computer vision and pattern recognition*, pp. 2528–2535. IEEE, 2010.

Hongyang Zhang, Yaodong Yu, Jiantao Jiao, Eric Xing, Laurent El Ghaoui, and Michael Jordan. Theoretically principled trade-off between robustness and accuracy. In *International conference on machine learning*, pp. 7472–7482. PMLR, 2019.

Yu Zhang, Peter Tiňo, Aleš Leonardis, and Ke Tang. A survey on neural network interpretability. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2021.

Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2921–2929, 2016.