

Explainable Adversarial Attacks on Coarse-to-Fine Classifiers

Akram Heidarizadeh

*Dept. of Electrical and Computer Engineering
University of Central Florida
Orlando, FL, USA
akram.heidarizadeh@ucf.edu*

Connor Hatfield

*Dept. of Electrical and Computer Engineering
University of Central Florida
Orlando, FL, USA
connorhat123@gmail.com*

Lorenzo Lazzarotto

*School of Technology
Pontificia Universidade Católica
do Rio Grande do Sul
Porto Alegre, Brazil
l.lazzarotto@edu.pucrs.br*

HanQin Cai

*Dept. of Statistics and Data Science
Dept. of Computer Science
University of Central Florida
Orlando, FL, USA
hqcai@ucf.edu*

George Atia

*Dept. of Electrical and Computer Engineering
Dept. of Computer Science
University of Central Florida
Orlando, FL, USA
george.atia@ucf.edu*

Abstract—Traditional adversarial attacks typically aim to alter the predicted labels of input images by generating perturbations that are imperceptible to the human eye. However, these approaches often lack explainability. Moreover, most existing work on adversarial attacks focuses on single-stage classifiers, but multi-stage classifiers are largely unexplored. In this paper, we introduce instance-based adversarial attacks for multi-stage classifiers, leveraging Layer-wise Relevance Propagation (LRP), which assigns relevance scores to pixels based on their influence on classification outcomes. Our approach generates explainable adversarial perturbations by utilizing LRP to identify and target key features critical for both coarse and fine-grained classifications. Unlike conventional attacks, our method not only induces misclassification but also enhances the interpretability of the model’s behavior across classification stages, as demonstrated by experimental results.

Index Terms—Adversarial attacks, Explainability, Hierarchical classifiers, Layer-wise relevance propagation.

I. INTRODUCTION

The remarkable success of neural networks across diverse domains, from image recognition to natural language processing, has led to their widespread adoption in critical applications such as autonomous driving [1], [2], healthcare [3], and security systems [4]. Despite these advancements, deep neural networks (DNNs) remain vulnerable to adversarial attacks, where imperceptible perturbations to input data can lead to incorrect predictions [5], [6], [7], [8].

Traditional attack methods, such as the Fast Gradient Sign Method (FGSM) [9] and Projected Gradient Descent (PGD) [10], have been designed to generate small perturbations within an ℓ_1 ball that alter the model’s prediction while remaining imperceptible to human observers. While these attacks expose the vulnerability of neural networks, they provide limited insight into the reasoning behind the model’s

decision changes, thus lacking explainability. This limitation also applies to hierarchical classifiers, which make decisions across multiple stages. While a few works, such as [11] and [12], have explored traditional attacks in hierarchical settings, to our knowledge, no research has addressed explainable attacks in this context.

Explainability and interpretability have gained significant attention in recent years [13]. Methods like Layer-wise Relevance Propagation (LRP) [14], [15], GradCAM [16], LIME [17] and SHAP [18] provide various mechanisms for explanations by identifying the key features or regions of an input that most influence a model’s output.

However, research on explainable adversarial attacks remains limited, even for single-stage classifiers, with only a few studies such as [19], [20] and [21], which focus on single-stage universal attacks. The exploration of such methods for coarse-to-fine (C2F) classifiers, where understanding the decision-making process across multiple stages is crucial, remains largely unexplored.

In this paper, we study explainable adversarial attacks on C2F classifiers by introducing an approach that leverages LRP to guide the generation of interpretable adversarial perturbations. Unlike traditional attacks that solely aim to change the output label, our method targets key features identified at both the coarse and fine classification stages, as highlighted by LRP-generated heatmaps. By focusing on perturbing these critical features, we generate attacks that not only fool the model but also provide insights into the features the model relies on at each stage of the classification process.

We evaluate the effectiveness of our approach through experiments on a C2F classifier architecture, using benchmark datasets like ImageNet [22]. Our results show that the generated perturbations successfully mislead the models at both stages of classification, while highlighting a key trade-off between explainability and perceptibility. Compared to the other

methods, we allow for greater perturbation, while keeping it imperceptible and achieving enhanced explainability.

II. BACKGROUND

A. Coarse-to-Fine Model Formulation

Consider a pre-trained hierarchical classifier structured in a C2F hierarchy, where an initial coarse-level classifier provides a broad classification, which is then further refined by subsequent finer classifiers.

Each data point $x \in X \subseteq \mathbb{R}^N$ is assigned a coarse label $i \in [M]$, where M denotes the total number of coarse labels and $[M] := \{1, 2, \dots, M\}$. Additionally, there is a fine label $l \in [M_i]$, where M_i represents the number of fine classes associated with the i -th coarse label. Let $C : \mathbb{R}^N \rightarrow [M]$ be the coarse classifier function that assigns x to a coarse class. For classifier C , we assume the existence of M discriminant functionals, $C_i(x) : \mathbb{R}^N \rightarrow \mathbb{R}$ for $i \in [M]$, which are used for coarse classification such that

$$C(x) = \arg \max_{i \in [M]} C_i(x). \quad (1)$$

For each coarse label $i \in [M]$, let F^i represent the i -th fine classifier function. Similar to the formulation in (1), we define $F_j^i(x) : \mathbb{R}^N \rightarrow \mathbb{R}$ for $j \in [M_i]$ as the discriminant functions used to determine the finer class of coarse label i , such that

$$F^i(x) = \arg \max_{j \in [M_i]} F_j^i(x). \quad (2)$$

B. Layer-wise Relevance Propagation

Layer-wise Relevance Propagation (LRP) is a technique to explain the decisions made by neural networks, determining the contribution of each parts of the input data to the final decision [14]. To interpret the network's prediction for a specific class c , we propagate relevance scores R from the output layer L back to the input layer, using the activations and network weights. For the output layer, relevance is defined by:

$$R_i^L = \delta_{i,c}, \quad (3)$$

where $\delta_{i,c}$ is the Kronecker delta, which selects the relevance for class c by setting $R_i^L = 1$ when $i = c$ and $R_i^L = 0$ otherwise. The relevance scores are then propagated back through all layers except the first using the z+ rule [15]:

$$R_i^l = \sum_j \frac{a_i^l (W^l)_{ij}^+}{\sum_k a_k^l (W^l)_{kj}^+} R_j^{l+1}, \quad (4)$$

where $(W^l)^+$ denotes the positive weights of the l -th layer and a^l is the activation vector of the l -th layer. Finally, the relevance scores at the input layer are calculated using the z β rule [15]:

$$R_i^0 = \sum_j \frac{a_i^0 W_{ij}^0 - l_i (W^0)_{ij}^+ - h_i (W^0)_{ij}^-}{\sum_k (a_i^0 W_{kj}^0 - l_i (W^0)_{kj}^+ - h_i (W^0)_{kj}^-)} R_j^1, \quad (5)$$

where l_i and h_i are the lower and upper bounds of the input domain, respectively. For simplicity, we henceforth use the notation $LRP_G(x; c)$ to indicate the relevance scores at the input layer of a classifier G , for an input image x and label c .

III. LRP ATTACK FORMULATION

Unlike the methods in [6], [10], and [23] that directly use the gradients of the DNN's outputs and inputs to generate perturbations, the attacks we propose herein target the heatmaps produced by the LRP method. By assuming that the LRP interpretation indicates the DNN's attention, our algorithm is designed to create perturbations by disrupting this attention.

A. Fooling the Coarse Level

In this attack, the goal is to generate imperceptible additive perturbations to fool the coarse-level classification, satisfying the requirement $C(x + \eta) \neq C(x)$. To formalize this, let $r_{\text{org}} = C(x)$ and $r_{\text{adv}} = C(x + \eta)$ represent the original and adversarial coarse labels, respectively, where these labels correspond to the coarse categories with the highest prediction probability for the input image x and the perturbed image $x + \eta$. The attacker selects r_{adv} as

$$r_{\text{adv}} = \arg \max_{i \in [M] \setminus r_{\text{org}}} C_i(x), \quad (6)$$

the coarse label with the second-highest probability after r_{org} .

To achieve the objective of redirecting the coarse classifier's attention from r_{org} to r_{adv} , we define a loss function based on the ℓ_p -norm of the positive and negative relevance scores produced by LRP with respect to these labels. The loss function is designed to decrease all positive relevance scores and increase all negative relevance scores of the heatmap $LRP_C(x + \eta; r_{\text{org}})$, while simultaneously increasing all positive relevance scores and decreasing all negative relevance scores of $LRP_C(x + \eta; r_{\text{adv}})$. Thus, we define the loss function for the LRP Coarse-Level Attack (LRPC):

$$\mathcal{L}_C = \|LRP_C(x + \eta; r_{\text{org}})^+\|_p - \|LRP_C(x + \eta; r_{\text{adv}})^+\|_p - \|LRP_C(x + \eta; r_{\text{org}})^-\|_p + \|LRP_C(x + \eta; r_{\text{adv}})^-\|_p. \quad (7)$$

Here, we set $p = 1$, calculating the ℓ_1 -norm of the heatmaps.

B. Fooling the Fine Level

In this scenario, the goal is to create perturbations that alter the finer classification while keeping the coarse prediction unchanged. Specifically, we aim to ensure that $C(x + \eta) = C(x) = r_{\text{org}}$, while $F^{r_{\text{org}}}(x + \eta) \neq F^{r_{\text{org}}}(x)$. The coarse label is first determined and used as a prerequisite to identify the corresponding fine label. Similar to the coarse level attack, the attacker selects f_{adv} as:

$$f_{\text{adv}} = \arg \max_{j \in [M_{r_{\text{org}}}] \setminus f_{\text{org}}} F_j^{r_{\text{org}}}(x), \quad (8)$$

where $f_{\text{org}} := F^{r_{\text{org}}}(x)$ and f_{adv} is the fine label with the second-highest probability.

We apply the same approach at the fine level by defining the loss function similarly, based on the ℓ_1 -norm of the positive relevances produced by LRP for the r_{org} -th fine classifier, $F^{r_{\text{org}}}$, with respect to the fine-level labels. Thus, the loss function for the LRP Fine-Level Attack (LRPF) is defined as:

$$\mathcal{L}_F = \|LRP_{F^{r_{\text{org}}}}(x + \eta; f_{\text{org}})^+\|_p - \|LRP_{F^{r_{\text{org}}}}(x + \eta; f_{\text{adv}})^+\|_p - \|LRP_{F^{r_{\text{org}}}}(x + \eta; f_{\text{org}})^-\|_p + \|LRP_{F^{r_{\text{org}}}}(x + \eta; f_{\text{adv}})^-\|_p. \quad (9)$$

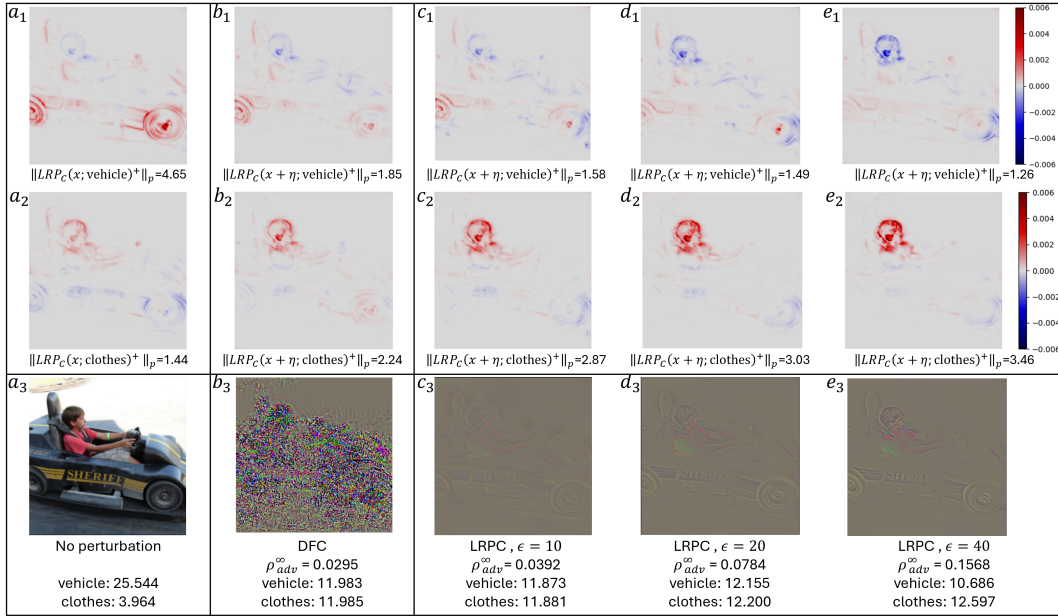


Fig. 1. LRP visualizations before and after LRPC and DFC attacks. (a₁) LRP of the original coarse class (r_{org} : “vehicle”) before the attack. (a₂) LRP of the adversarial coarse class (r_{adv} : “clothes”) before the attack. (a₃) Benign image. (c₁, d₁, e₁) LRP of r_{org} after LRPC attack for $\epsilon = 10, 20, 40$, compared to (b₁) for DFC. (c₂, d₂, e₂) LRP of r_{adv} after LRPC attack for $\epsilon = 10, 20, 40$, compared to (b₂) for DFC. Perturbations generated with LRPC ($\epsilon = 10, 20, 40$) are shown in (c₃, d₃, e₃), and for DFC in (b₃). LRP norms and prediction scores are displayed below the respective cases.

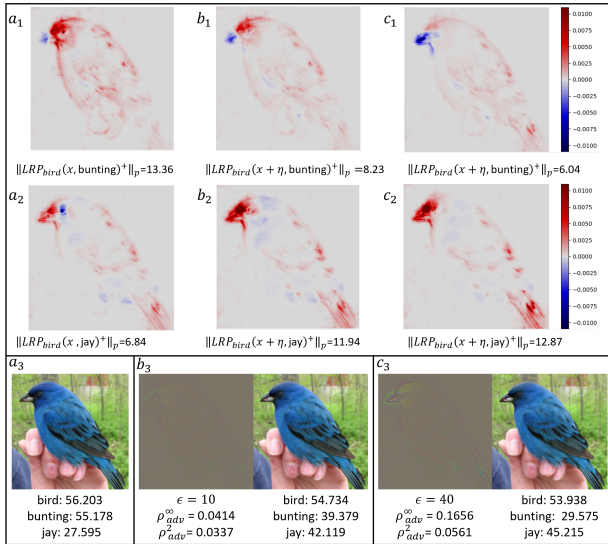


Fig. 2. Example of LRP visualization of a₁) the original fine class (f_{org}) and a₂) the adversarial fine class (f_{adv}) before attack. a₃) Benign image. b₁, c₁) LRP of f_{org} after LRPF attack with $\epsilon = 10, 40$. b₂, c₂) LRP of f_{adv} after LRPF attack with $\epsilon = 10, 40$. b₃, c₃) Perturbation with LRPF for $\epsilon = 10, 40$ and perturbed image. The prediction scores for each case are displayed below the respective images. r_{org} : “bird”, f_{org} : “bunting” and f_{adv} : “jay”.

Algorithm 1 describes our approach. The perturbation is initialized to zero at the beginning of the attack. At each iteration, the perturbation η is clipped by $\min(\epsilon, \eta)$, where ϵ is the maximum permissible l_∞ -norm. η is then added to the benign images to create perturbed images, which are clipped again to maintain pixel values within the range $[0, 255]$. Both

benign and perturbed images are fed into the pre-trained model to obtain the original and adversarial labels. The adversarial labels could change with each iteration, as the goal is to generate small perturbations that progressively move the perturbed image toward the nearest labels at each step. The gradient-descent algorithm is used to minimize the loss function as defined in (7) and (9) and optimize the perturbation η via the loss gradients $\frac{\partial \mathcal{L}_C}{\partial \eta}$ and $\frac{\partial \mathcal{L}_F}{\partial \eta}$ for LRPC and LRPF, respectively. Depending on the attack type, the optimization process terminates either when the coarse label changes (LRPC) or when the fine label changes while maintaining coarse label consistency (LRPF), ensuring effective perturbation generation.

IV. EXPERIMENTAL RESULTS

We evaluate our attack on the C2F architecture using ImageNet [22] and compare it to other attacks, such as PGD [10] and DeepFool [6]. Since these two methods were designed for single-stage classifiers, we adapted them for the C2F setting by applying separate attacks at both the coarse and fine levels. We refer to these adapted versions as DFC and DFF for DeepFool’s coarse and fine attacks, and PGDC and PGDF for PGD’s coarse and fine attacks, where perturbations in DeepFool are measured using the l_2 -norm, and PGD is implemented with $p = \infty$. We calculate the average perceptibility of the attack as:

$$\rho_{\text{adv}}^p(f) = \frac{1}{|D|} \sum_{x \in D} \frac{\|\eta\|_p}{\|x\|_p}, \quad (10)$$

where D is the dataset, and η is the adversarial perturbation corresponding to input x . Additionally, we analyze the fooling ratio, defined as the proportion of images whose labels are changed by the attack relative to the total number of images.

Algorithm 1 LRP-based Attack for Coarse-to-Fine Classifiers

```

1: Input: Pre-trained coarse model  $C$  and fine models  $F_i$ ,
   input image  $x$ , ground truth labels  $(y, z)$ , max iterations
    $K$ , step size  $lr$ , clip parameter  $\epsilon$ , attack= $\{\text{LRPC, LRPF}\}$ 
2: Initialize perturbation  $\eta \leftarrow \mathbf{0}$ ,  $k \leftarrow 0$ 
3: Compute labels  $r_{\text{org}}, r_{\text{adv}}, f_{\text{org}}$  and  $f_{\text{adv}}$ 
4: while  $k < K$  do
5:   Calculate LRPs for both labels
6:   Compute Loss  $\mathcal{L}$  according to (7) or (9)
7:   Update perturbation  $\eta \leftarrow \eta - lr \cdot \frac{\partial \mathcal{L}}{\partial \eta}$ 
8:    $\eta \leftarrow \text{clip}(\eta; -\epsilon, \epsilon)$ 
9:   Update adv label for  $x + \eta$ 
10:   $k \leftarrow k + 1$ 
11:  if attack==LRPC then
12:    if  $r_{\text{org}} \neq y$  then
13:      Increment fooling count and break loop
14:    end if
15:  else if attack==LRPF then
16:    if  $f_{\text{org}} \neq z$  and  $r_{\text{org}} = y$  then
17:      Increment fooling count and break loop
18:    end if
19:  end if
20: end while
21: Output: perturbation  $\eta$ 

```

Coarse-to-fine classification framework. We have introduced a C2F classifier with a total of $M = 8$ coarse labels for the hierarchical classification of the ImageNet dataset. The classification occurs in two stages: a coarse classifier C assigns input images to one of eight broad categories: {fish, bird, reptile, clothes, food, vehicle, electrical device, dog}, which are further classified by separate fine-level classifiers F^i within each coarse category into fine-grained labels.

Dataset. Our dataset is sourced from the ImageNet dataset, which contains 1000 distinct labels. However, only 393 of these labels were used in our model, distributed across the 8 coarse categories. We randomly selected 80% of the training set from the 393 classes of the ILSVRC2012 train dataset for training the classifiers, with the remaining 20% used for validation. We evaluate our attack on the VGG-16 network [24], which has an accuracy of 96% for the coarse classifier and 78%-90% for each of the fine classifiers.

Explainability-perceptibility tradeoff. Fig. 1 illustrates LRP results before and after the coarse attack for LRPC with varying perturbation levels. The heatmaps highlight regions contributing to coarse-level predictions, showing how the LRPC attack shifts the model’s attention from the car to the person inside, which leads to misclassification from “vehicle” to “clothes.” Both positive and negative relevance scores for the original and adversarial classes are manipulated, with reduced relevance for car pixels (c_1, d_1, e_1) and increased relevance for the person (c_2, d_2, e_2). As perturbation increases (via higher ϵ), explainability improves, revealing a tradeoff between explainability and perceptibility. Our results clarify this

relationship, both quantitatively (by illustrating the tradeoff between perceptibility and relevance scores), and qualitatively (by progressively emphasizing key regions while reducing attention on others as perturbation increases). In contrast, the DFC attack (b_1, b_2, b_3) scatters the perturbation, whereas our attack concentrates it on critical image regions. Fig. 2 demonstrates that our algorithm is also explainable at the finer level, although the features are very similar in this stage. In this example, the attack shifts the model’s attention from the eye and body of the “indigo” to areas typically associated with a “jay”, such as the beak and tail feathers, leading to misclassification.

Performance. We evaluate the performance of our attack for both coarse- and fine-level attack algorithms. Table I shows $\rho_{\text{adv}}^1, \rho_{\text{adv}}^2, \rho_{\text{adv}}^\infty$ and fooling ratio for LRPC, DFC, and PGDC, tested on the first 1000 RGB images of the validation set. LRPC achieves perceptibility comparable to PGDC in both ℓ_1 and ℓ_∞ -norms. Moreover, for $\epsilon = 10$, the ℓ_∞ -norm is on par with DFC, while the fooling rate remains high. Table II provides the fooling ratio and perceptibility values for the fine-level attack algorithms. LRPF achieves high fooling rates even with smaller perturbations, indicated by lower ϵ values. This demonstrates that LRPF can maintain competitive fooling rates while effectively controlling perceptibility. We emphasize that our approach prioritizes explainability rather than asserting quantitative superiority in metrics like fooling ratio or perceptibility.

V. CONCLUSION

We developed an approach for explainable adversarial attacks on coarse-to-fine classifiers by leveraging Layer-wise Relevance Propagation (LRP) to generate interpretable perturbations. Our method targets critical features identified at both classification stages, providing insights into the model’s decision-making process while successfully misleading the model and outperform traditional methods in providing clearer interpretations without compromising attack imperceptibility.

TABLE I
FOOLING RATIO AND PERCEPTIBILITY OF COARSE-LEVEL ATTACKS.

Algorithm	LRPC $\epsilon = 10$	LRPC $\epsilon = 20$	LRPC $\epsilon = 40$	DFC	PGDC
ρ_{adv}^2	0.0294	0.0323	0.0405	0.0045	0.0262
ρ_{adv}^1	0.0216	0.0174	0.0195	0.0031	0.0224
ρ_{adv}^∞	0.0399	0.0778	0.1557	0.0408	0.0101
Fooling(%)	87.1	92.5	99.3	100	100

TABLE II
FOOLING RATIO AND PERCEPTIBILITY OF FINE-LEVEL ATTACKS.

Algorithm	LRPF $\epsilon = 10$	LRPF $\epsilon = 20$	LRPF $\epsilon = 40$	DFD	PGDF
ρ_{adv}^2	0.0127	0.0145	0.0151	0.0020	0.0078
ρ_{adv}^1	0.0084	0.0079	0.0066	0.0013	0.0092
ρ_{adv}^∞	0.0241	0.0542	0.0819	0.0029	0.0035
Fooling(%)	98.7	100	100	100	95.7

REFERENCES

- [1] M. Kozłowski, S. Racewicz, and S. Wierzbiński, “Image analysis in autonomous vehicles: A review of the latest AI solutions and their comparison,” 2024.
- [2] H. Maghsoumi, N. Masoumi, and B. N. Araabi, “RoADSaVe: A robust lane detection method based on validity borrowing from reliable lines,” *IEEE Sensors Journal*, vol. 23, no. 13, pp. 14 571–14 582, 2023.
- [3] S. A. Alowais, S. S. Alghamdi, N. Alsuhbany, T. Alqahtani, A. I. Alshaya, S. N. Almohareb, A. Aldairem, M. Alrashed, K. Bin Saleh, H. A. Badreldin *et al.*, “Revolutionizing healthcare: the role of artificial intelligence in clinical practice,” *BMC medical education*, vol. 23, no. 1, p. 689, 2023.
- [4] S. Silvestri, S. Islam, S. Papastergiou, C. Tzagkarakis, and M. Ciampi, “A machine learning approach for the nlp-based analysis of cyber threats and vulnerabilities of the healthcare ecosystem,” *Sensors*, vol. 23, no. 2, p. 651, 2023.
- [5] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, “Intriguing properties of neural networks,” *arXiv preprint arXiv:1312.6199*, 2013.
- [6] S.-M. Moosavi-Dezfooli, A. Fawzi, and P. Frossard, “Deepfool: a simple and accurate method to fool deep neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2574–2582.
- [7] H. Cai, Y. Lou, D. McKenzie, and W. Yin, “A zeroth-order block coordinate descent algorithm for huge-scale black-box optimization,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 1193–1203.
- [8] H. Cai, D. McKenzie, W. Yin, and Z. Zhang, “Zeroth-order regularized optimization (zoro): Approximately sparse gradients and adaptive sampling,” *SIAM Journal on Optimization*, vol. 32, no. 2, pp. 687–714, 2022.
- [9] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [10] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu, “Towards deep learning models resistant to adversarial attacks,” *arXiv preprint arXiv:1706.06083*, 2017.
- [11] I. R. Alkhouri and G. K. Atia, “Adversarial attacks on coarse-to-fine classifiers,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 2855–2859.
- [12] I. R. Alkhouri, A. Velasquez, and G. K. Atia, “Adversarial perturbation attacks on nested dichotomies classification systems,” in *IEEE 31st International Workshop on Machine Learning for Signal Processing (MLSP)*, 2021.
- [13] H. Baniecki and P. Biecek, “Adversarial attacks and defenses in explainable artificial intelligence: A survey,” *Information Fusion*, p. 102303, 2024.
- [14] S. Bach, A. Binder, G. Montavon, F. Klauschen, K.-R. Müller, and W. Samek, “On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation,” *PLoS one*, vol. 10, no. 7, p. e0130140, 2015.
- [15] G. Montavon, S. Lapuschkin, A. Binder, W. Samek, and K.-R. Müller, “Explaining nonlinear classification decisions with deep Taylor decomposition,” *Pattern Recognition*, vol. 65, pp. 211–222, 2017.
- [16] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 618–626.
- [17] M. T. Ribeiro, S. Singh, and C. 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*, 2016, pp. 1135–1144.
- [18] S. Lundberg, “A unified approach to interpreting model predictions,” *arXiv preprint arXiv:1705.07874*, 2017.
- [19] Z. Wang, X. Huang, J. Yang, and N. Kasabov, “Universal adversarial perturbation generated by attacking layer-wise relevance propagation,” in *IEEE 10th International Conference on Intelligent Systems (IS)*, 2020, pp. 431–436.
- [20] A. Ghorbani, A. Abid, and J. Zou, “Interpretation of neural networks is fragile,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 3681–3688.
- [21] A.-K. Dombrowski, M. Alber, C. Anders, M. Ackermann, K.-R. Müller, and P. Kessel, “Explanations can be manipulated and geometry is to blame,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [22] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, “ImageNet large scale visual recognition challenge,” *International Journal of Computer Vision*, vol. 115, pp. 211–252, 2015.
- [23] E. Wong, F. Schmidt, and Z. Kolter, “Wasserstein adversarial examples via projected sinkhorn iterations,” in *International Conference on Machine Learning*. PMLR, 2019, pp. 6808–6817.
- [24] K. Simonyan and A. Zisserman, “Very deep convolutional networks for large-scale image recognition,” *arXiv preprint arXiv:1409.1556*, 2014.