

Angular Visual Hardness

Beidi Chen¹ Weiyang Liu² Zhiding Yu³ Jan Kautz³ Anshumali Shrivastava¹
Animesh Garg^{3 4 5} Anima Anandkumar^{3 6}

Abstract

Recent convolutional neural networks (CNNs) have led to impressive performance but often suffer from poor calibration. They tend to be overconfident, with the model confidence not always reflecting the underlying true ambiguity and hardness. In this paper, we propose angular visual hardness (AVH), a score given by the normalized angular distance between the sample feature embedding and the target classifier to measure sample hardness. We validate this score with an in-depth and extensive scientific study, and observe that CNN models with the highest accuracy also have the best AVH scores. This agrees with an earlier finding that state-of-art models improve on the classification of harder examples. We observe that the training dynamics of AVH is vastly different compared to the training loss. Specifically, AVH quickly reaches a plateau for all samples even though the training loss keeps improving. This suggests the need for designing better loss functions that can target harder examples more effectively. We also find that AVH has a statistically significant correlation with human visual hardness. Finally, we demonstrate the benefit of AVH to a variety of applications such as self-training for domain adaptation and domain generalization.

1. Introduction

Convolutional neural networks (CNNs) have achieved great progress on many computer vision tasks such as image classification (He et al., 2016; Krizhevsky et al., 2012), face recognition (Sun et al., 2014; Liu et al., 2017b; 2018a), and scene understanding (Zhou et al., 2014; Long et al., 2015a). On certain large-scale benchmarks such as ImageNet, C-




Figure 1. Example images that confuse humans. Top row: images with degradation. Bottom row: images with semantic ambiguity.

NNs have even surpassed human-level performance (Deng et al., 2009). Despite the notable progress, CNNs are still far from matching human-level visual recognition in terms of robustness (Goodfellow et al., 2014; Wang et al., 2018c), adaptability (Finn et al., 2017) and few-shot generalizability (Hariharan & Girshick, 2017; Liu et al., 2019), and could suffer from various biases. For example, ImageNet-trained CNNs are reported to be biased towards textures, and these biases may result in CNNs being overconfident, or prone to domain gaps and adversarial attacks (Geirhos et al., 2019).

Softmax score has been widely used as a confidence measure for CNNs but it tends to give over-confident output (Guo et al., 2017; Li & Hoiem, 2018). To fix this issue, one line of work considers confidence calibration from a Bayesian point of view (Springenberg et al., 2016; Lakshminarayanan et al., 2017). Most of these methods tend to focus on the calibration and rescaling of model confidence by matching expected error or ensemble. But how much they are correlated with human confidence is yet to be thoroughly studied. On the other hand, several recent works (Liu et al., 2016; 2017c; 2018b) conjecture that softmax feature embeddings tend to naturally decouple into norms and angular distances that are related to intra-class confidence and inter-class semantic difference. Though inspiring, the conjecture lacks thorough investigation and we make surprising observations partially contradicting to the conjecture on intra-class confidence. This motivates us to conduct rigorous studies for reliable and semantics-related confidence measure.

Human vision is considered much more robust than current CNNs, but this does not mean humans cannot be confused.

¹Rice University ²Georgia Institute of Technology ³NVIDIA
⁴University of Toronto ⁵Vector Institute, Toronto ⁶Caltech. Correspondence to: Beidi Chen <beidi.chen@rice.edu>, Weiyang Liu <wylu@gatech.edu>, Zhiding Yu <zhidingy@nvidia.com>.




Figure 2. Visualization of embeddings on MNIST by setting their dimensions to 2 in a CNN.

Many images appear ambiguous or hard for humans due to various image degradation factors such as lighting conditions, occlusions, visual distortions, etc. or due to semantic ambiguity in not understanding the label category, as shown in Figure 1. It is therefore natural to consider such human ambiguity or visual hardness on images as the gold-standard for confidence measures. However, explicitly encoding human visual hardness in a supervised manner is generally not feasible, since hardness scores can be highly subjective and difficult to obtain. Fortunately, a surrogate for human visual hardness was recently made available on the ImageNet validation set (Recht et al., 2019). This is based on **Human Selection Frequency (HSF)** - the average number of times an image gets picked by a crowd of annotators from a pool belonging to certain specified category. We adopt HSF as a surrogate for human visual hardness in this paper to validate our proposed angular hardness measure in CNNs.

Contribution: Angular Visual Hardness (AVH). Given a CNN, we propose a novel score function for measuring sample hardness. It is the normalized angular distance between the image feature embedding and the weights of the target category (See Figure 2 as a toy example). The normalization takes into account the angular distances to other categories.

We make observations on the dynamic evolution of AVH scores during ImageNet training. We find that AVH plateaus early in training even though the training (cross-entropy) loss keeps decreasing. This is due to the nature of parameterization in softmax loss, of which the minimization goes in two directions: either by aligning the angles between feature embeddings and classifiers or by increasing the norms of feature embeddings. We observe two phases popping up during training: (1) Phase 1, where the softmax improvement is primarily due to angular alignment, and later, (2) Phase 2, where the improvement is primarily due to significant increase in feature-embedding norms.

The above findings suggest that the AVH can be a robust universal measure of hardness since angular scores are mostly frozen early in training. In addition, they suggest the need to design better loss functions over softmax loss that can improve performance on hard examples and focus on optimiz-

ing angles, e.g., (Liu et al., 2017b; Deng et al., 2019; Wang et al., 2018b;a). We verify that better models tend to have better average AVH scores, which validates the argument in (Recht et al., 2019) that improving on hard examples is the key to improved generalization. We show that AVH has a statistically significant stronger correlation with human selection frequency than widely used confidence measures such as softmax score and embedding norm across several CNN models. This makes AVH a potential proxy of human perceived hardness when such information is not available.

Finally, we empirically show the superiority of AVH with its application to self-training for unsupervised domain adaptation and domain generalization. With AVH being an improved confidence measure, our proposed self-training framework renders considerably improved pseudo-label selection and category estimation, leading to state-of-the-art results with significant performance gain over baselines. Our proposed new loss function based on AVH also shows drastic improvement for the task of domain generalization.

2. Related Work

Example hardness measures. An automatic detection of examples that are hard for human vision has numerous applications. (Recht et al., 2019) showed that state-of-the-art models perform better on hard examples. This implies that in order to improve generalization, the models need to improve accuracy on hard examples. This can be achieved through various learning algorithms such as curriculum learning (Bengio et al., 2009) and self-paced learning (Kumar et al., 2010) where being able to detect hard examples is crucial. Measuring sample confidence is also important in partially-supervised problems such as semi-supervised learning (Zhu; Zhou et al., 2012), unsupervised domain adaptation (Chen et al., 2011) and weakly-supervised learning (Tang et al., 2017) due to their under-constrained nature. Sample hardness can also be used to identify implicit distribution imbalance in datasets to ensure fairness and remove societal biases (Buolamwini & Gebru, 2018).

Angular distance in neural networks. (Zhang et al., 2018) uses deep features to quantify the semantic difference between images, indicating that deep features contain the most crucial semantic information. It empirically shows that the angular distance between feature maps in deep neural networks is very consistent with the human in distinguishing the semantic difference. (Liu et al., 2017c) proposes a hyperspherical neural network that constrains the parameters of neurons on a unit hypersphere and uses angular similarity to replace the inner product similarity. (Liu et al., 2018b) proposes to decouple the inner product as norm and angle, arguing that norms correspond to intra-class variation, and angles corresponds to inter-class semantic difference. However, this work does not perform in-depth studies to prove this conjecture. Recent research (Liu et al., 2018a; Lin et al.,




Figure 3. Toy example of two overlapping Gaussian distributions (classes) on a unit sphere. Left: samples from the distributions as input to a multi layer perceptron (MLP). Middle: AVH heat map produced by MLP, where samples in lighter colors (higher hardness) are mostly overlapping hard examples. Right: ℓ_2 -norm heat map, where certain non-overlapping samples also have higher values.

2020; Liu et al., 2020) comes up with an angle-based hyperspherical energy to characterize the neuron diversity and improve generalization by minimizing this energy.

Deep model calibration. Confidence calibration aims to predict probability estimates representative of the true correctness likelihood (Guo et al., 2017). It is well-known that the deep neural networks tend to be mis-calibrated and there has been a rich set of literature trying to solve this problem (Kumar et al., 2018; Guo et al., 2017). While establishing correlation between model confidence and prediction correctness, the connection to human confidence has not been widely studied from a training dynamics perspective.

Uncertainty estimation. In uncertainty estimation, two types of uncertainties are often considered: (1) *Aleatoric* uncertainty which captures noise inherent in the observations; (2) *Epistemic* uncertainty which accounts for uncertainty in the model due to limited data (Der Kiureghian & Ditlevsen, 2009). The latter is widely modeled by Bayesian inference (Kendall & Gal, 2017) and its approximation with dropout (Gal & Ghahramani, 2016; Gal et al., 2017), but often at the cost of additional computation. The fact that AVH correlates well with Human Selection Frequency indicates its underlying connection to aleatoric uncertainty. This makes it suitable for tasks such as self-training. Yet unlike Bayesian inference, AVH can be naturally computed during regular softmax training, making it convenient to obtain with only one-time training and a drop-in uncertainty measure for most existing neural networks.

3. Discoveries in CNN Training Dynamics

Notation. Denote \mathbb{S}^n as the unit n -sphere, i.e., $\mathbb{S}^n = \{\mathbf{x} \in \mathbb{R}^{n+1} \mid \|\mathbf{x}\|_2 = 1\}$. Below by $\mathcal{A}(\cdot, \cdot)$, we denote the angular distance between two points on \mathbb{S}^n , i.e., $\mathcal{A}(\mathbf{u}, \mathbf{v}) = \arccos(\frac{\langle \mathbf{u}, \mathbf{v} \rangle}{\|\mathbf{u}\| \|\mathbf{v}\|})$. Let \mathbf{x} be the feature embedding input for the last layer of the classifiers in the pretrained CNNs (e.g., FC-1000 in VGG-19). Let \mathcal{C} be the number of classes for a classification task. Denote $\mathbf{W} = \{\mathbf{w}_i \mid 0 < i \leq \mathcal{C}\}$ as the set of weights for all \mathcal{C} classes in the final layer of the classifier.

Definition 1 (Model Confidence). We define *Model Confidence* on a single sample as the probability score of the true objective class output by the CNN models, $\frac{e^{\mathbf{w}_y \mathbf{x}}}{\sum_{i=1}^{\mathcal{C}} e^{\mathbf{w}_i \mathbf{x}}}$.

Definition 2 (Human Selection Frequency). We define one way to measure human visual hardness on pictures as *Human Selection Frequency (HSF)*. Quantitatively, given m number of human workers in a labeling process described in (Recht et al., 2019), if b out of m label a picture as a particular class and that class is the target class of that picture in the final dataset, then HSF is defined as $\frac{b}{m}$.

3.1. Proposal and Intuition

Definition 3 (Angular Visual Hardness). The AVH score, for any (\mathbf{x}, y) , is defined as:

$$\text{AVH}(\mathbf{x}) = \frac{\mathcal{A}(\mathbf{x}, \mathbf{w}_y)}{\sum_{i=1}^{\mathcal{C}} \mathcal{A}(\mathbf{x}, \mathbf{w}_i)}, \quad (1)$$

which \mathbf{w}_y represents the weights of the target class.

Theoretical Foundations of AVH. There are theoretical supports of AVH from both machine learning and vision science perspectives. On the machine learning side, we have briefly discussed above that AVH is directly related to the angle between feature embedding and the classifier weight of ground truth class. (Soudry et al., 2018) theoretically shows that the logit of ground truth class must diverge to infinity in order to minimize cross-entropy loss to zero under gradient descent. Assuming input feature embeddings have fixed unit-norm, the norm of classifier weight grows to infinity. Similar result is also shown in (Wei & Ma, 2019) where generalization error of a linear classifier is controlled by the output margins normalized with classifier norm. Although the above analyses make certain assumptions, they indicate that norm is a less calibrated variable towards measuring properties of model/data compared to angle. This conclusion is comprehensively validated by our experiments in Section 3. On the vision science side, there have been wide studies showing that human vision is highly adapted for extracting structural information (Zhang et al., 2018; Wang et al., 2004), while the angular distance in AVH is precisely

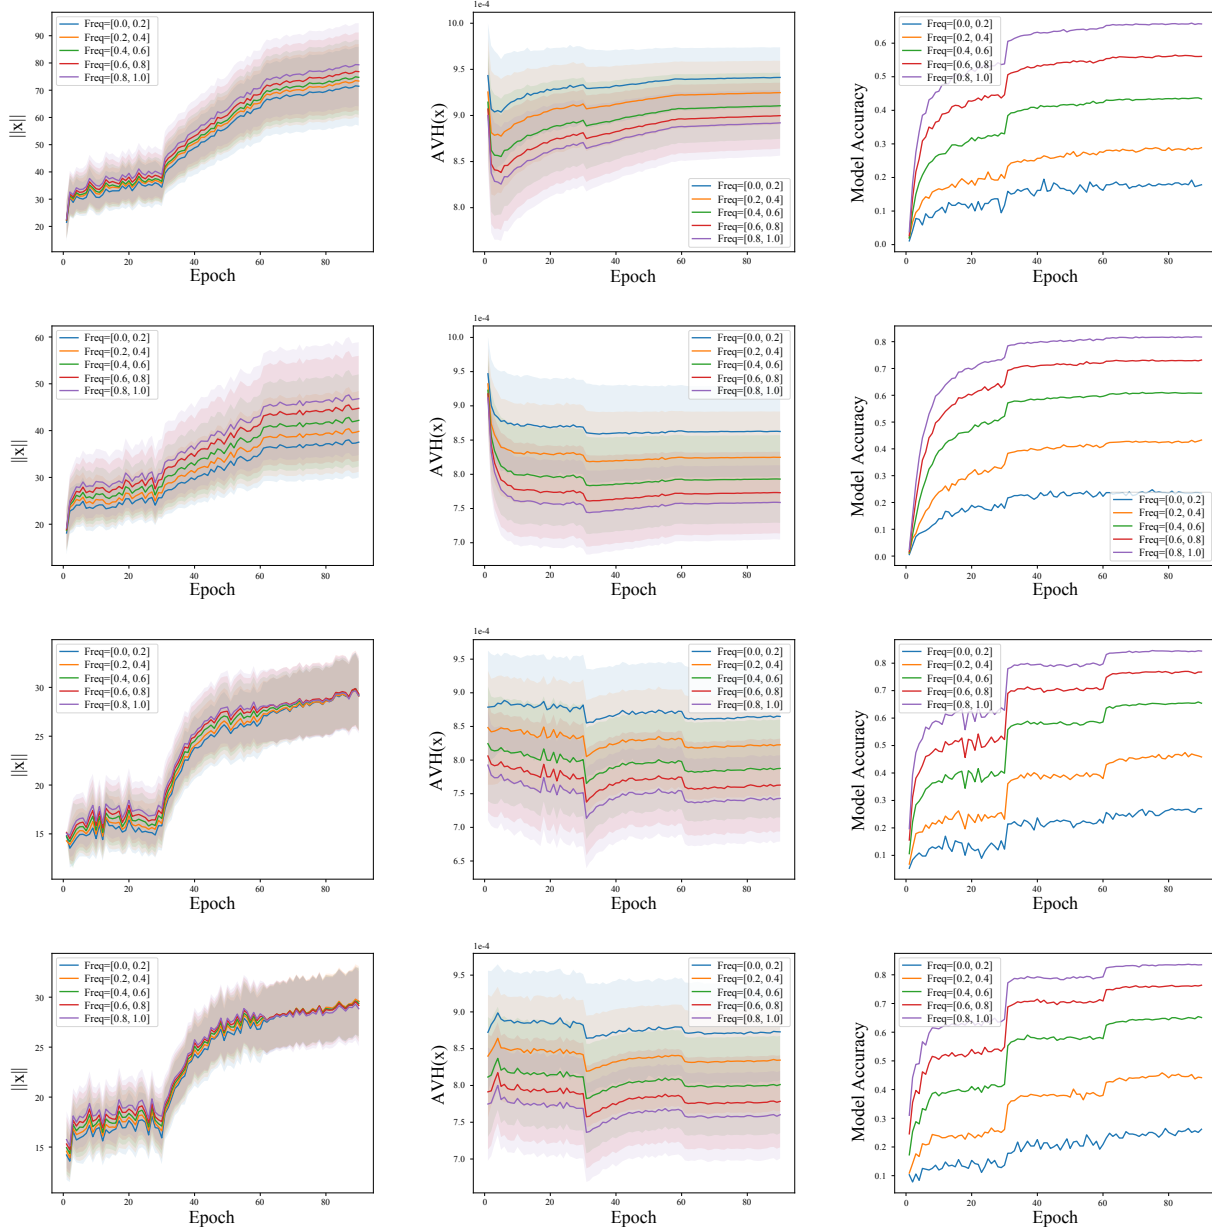


Figure 4. Averaged training dynamics across different Human Selection Frequency levels on ImageNet validation set. Columns from left to right: number of epochs vs. average ℓ_2 norm, number of epochs vs. average AVH score, and number of epochs vs. model accuracy. Rows from top to bottom: dynamics corresponding to AlexNet, VGG-19, ResNet-50, and DenseNet-121. Shadows in the figures of the first two columns denote the corresponding standard deviations.

good at capturing such information (Liu et al., 2018b). This also justifies our angular based design as an inductive bias towards measuring human visual hardness.

The AVH score is inspired by the observation from Figure 2 as well as (Liu et al., 2018b) that samples from each class concentrate in a convex cone in the embedding space along with some interesting theoretical results that are discussed above. Naturally, we conjecture AVH, a measure with angle or margin information, could be the useful component of

softmax score indicating the input sample hardness. We also perform a simulation providing visual intuition of how AVH instead of feature embedding norms corresponds to visually hard examples on two Gaussians in Figure 3 (simulation details and analyses in Appendix A).

3.2. Observations and Conjecture

Setup. We aim to observe the complete training dynamics of models that are trained from scratch on ImageNet instead of the pretrained models. Therefore, we follow the

standard training process of AlexNet (Krizhevsky et al., 2012), VGG-19 (Simonyan & Zisserman, 2014), ResNet-50 (He et al., 2016) and DenseNet-121 (Huang et al., 2017). For consistency, we train all models for 90 epochs and decay the initial learning rate by a factor of 10 every 30 epochs. The initial learning rate for AlexNet and VGG-19 is 0.01 and for DenseNet-121 and ResNet-50 is 0.1. We split all the validation images into 5 bins, $[0.0, 0.2]$, $[0.2, 0.4]$, $[0.4, 0.6]$, $[0.6, 0.8]$, $[0.8, 1.0]$, based on their HSF respectively. In Appendix B, we further provide experimental results on different datasets, such as MNIST, CIFAR10/100, and degraded ImageNet with different contrast or noise level, to better validate our proposal. For all the figures in this section, epoch starts from 1.

Optimization algorithms are used to update weights and biases, *i.e.*, the internal parameters of a model to improve the training loss. Both the angles between the feature embedding and classifiers, and the L_2 norm of the embedding can influence the loss. While it is well-known that the training loss or accuracy keeps improving but it is not obvious what would be the dynamics of the angles and norms separately during training. we design the experiments to observe the training dynamics of various network architectures.

Observation 1: The norm of feature embeddings keeps increasing during training.

The first column of Figure 4 presents the dynamics of averaged $\|x\|_2$ on validation samples with the same range of HSF over 90 epochs of training. Different figures also cover different network architectures. Note that we are using the validation data for dynamics observation and therefore have never fed them into the model. The average $\|x\|_2$ increases with a small initial slope but it suddenly climbs after 30 epochs when the first learning rate decay happens. The accuracy curve is very similar to that of the average $\|x\|_2$. These observations are consistent in all models and compatible with (Soudry et al., 2018) although it is more about the norm of the classifier weights. More interestingly, we find that neural networks with shortcuts (*e.g.*, ResNets and DenseNets) tend to make the norm of the images with different HSF the same, while neural networks without shortcuts (*e.g.*, AlexNet and VGG) tend to keep the gap of norm among the images with different human visual hardness.

Observation 2: AVH hits a plateau very early even when the accuracy or loss is still improving.

Middle row of Figure 4 exhibits the change of average AVH for validation samples in 90 epochs of training on three models. The average AVH for AlexNet and VGG-19 decreases sharply at the beginning and then starts to bounce back a little bit before converging. However, the dynamics of the average AVH for DenseNet-121 and ResNet-50 are different. They both decrease slightly and then quickly hits

a plateau in all three learning rate decay stages. But the common observation is that they all stop improving even when $\|x\|_2$ and model accuracy are increasing. AVH is more important than $\|x\|_2$ in the sense that it is the key factor deciding which class the input sample is classified to. However, optimizing the norm under the current softmax cross-entropy loss would be easier, which cause the plateau of angles for easy examples. However, the plateau for the hard examples can be caused by the limitation of the model itself and we show a simple illustration in Appendix C. It shows the necessity of designing loss functions that focus on optimizing angles.

Observation 3: AVH’s correlation with Human Selection Frequency consistently holds across models throughout the training process.

In Figure 4, we average over validation samples in five HSF bins or five degradation level bins separately, and then compute the average embedding norm, AVH and model accuracies. We can observe that for $\|x\|_2$, the gaps between the samples with different human visual hardness are not obvious in ResNet and DenseNet, while they are quite obvious in AlexNet and VGG. However, for AVH, such AVH gaps are very significant and consistent across every network architecture during the entire training process. Interestingly, even if the network is far from being converged, such AVH gaps are still consistent across different HSF. Also the norm gaps are also consistent. The intuition behind this could be that the angles for hard examples are much harder to decrease and probably never in the region for correct classification. Therefore the corresponding norms would not increase otherwise hurting the loss. It validates that AVH is a consistent and robust measure for visual hardness (and even generalization).

Observation 4: AVH is an indicator of a model’s generalization ability.

From Figure 4, we observe that better models (*i.e.*, higher accuracy) have lower average AVH throughout the training process and also across samples under different human visual hardness. For instance, Alexnet is the weakest model, and its overall average AVH and average AVH on each of the five bins are higher than those of the other three models. In addition, we have found that when testing Hypothesis 3 for better models, their AVH correlations with HSF are significantly stronger than correlations of Model Confidence. The above observations are aligned with earlier observations from (Recht et al., 2019) that better models also tend to generalize better on samples across different levels of human visual hardness. In addition, AVH is potentially a better measure for the generalization of a pretrained model. As shown in (Liu et al., 2017b), the norms of feature embeddings are often related to training data priors such as data imbalance and class granularity (Krizhevsky et al., 2012).




Figure 5. The left one presents HSF v.s. $\mathcal{AVH}(\mathbf{x})$, which we can see strong correlation. The second plot presents the correlation between HSF and Model Confidence with ResNet-50. It is not surprising that the density is highest on the right corner. The third one presents HSF v.s. $\|\mathbf{x}\|_2$. There are no obvious correlation between them. Note that different color indicates the density of samples in that bin.

Table 1. Spearman’s rank correlation coefficients between HSF and AVH, Model Confidence and L2 Norm of the Embedding in ResNet-50 for different visual hardness bin of samples. Note that here we show the absolute value of the coefficient which represents the strength of the correlation. For example, $[0, 0.2]$ denotes the samples that have HSF from 0 to 0.2.

	z-score	Total Coef	$[0, 0.2]$	$[0.2, 0.4]$	$[0.4, 0.6]$	$[0.6, 0.8]$	$[0.8, 1.0]$
Number of Samples	-	29987	837	2732	6541	11066	8811
AVH	0.377	0.36	0.228	0.125	0.124	0.103	0.094
Model Confidence	0.337	0.325	0.192	0.122	0.102	0.078	0.056
$\ \mathbf{x}\ _2$	-	0.0017	0.0013	0.0007	0.0005	0.0004	0.0003

However, when extracting features from unseen classes that do not exist in the training set, such training data prior is often undesired. Since AVH does not consider such feature embedding norms, it potentially presents a better measure towards the open set generalization of a deep network.

Conjecture on training dynamics of CNNs. From Figure 4 and observations above, we conjecture that the training of CNN has two phases. 1) At the beginning of the training, the softmax cross-entropy loss will first optimize the angles among different classes while the norm will fluctuate and increase very slowly. We argue that it is because changing the norm will not decrease the loss when the angles are not separated enough for correct classification. As a result, the angles get optimized firstly. 2) As the training continues, the angles become more stable and change very slowly while the norm increases rapidly. On the one hand, for easy examples, it is because when the angles get decreased enough for correct classification, the softmax cross-entropy loss can be well minimized by purely increasing the norm. On the other hand, for hard examples, the plateau is caused by that the CNN is unable to decrease the angle to correctly classify examples and thereby also unable to increase the norms (because it may otherwise increase the loss).

4. Connections to Human Visual Hardness

From Section 3.2, we conjecture that AVH has a stronger correlation with Human Selection Frequency - a reflection of human visual hardness that is related to aleatoric uncertainty. In order to validate this claim, we design statistical testings for the connections between Model Confidence,

AVH, $\|\mathbf{x}\|_2$ and HSF. Studying the precise connection or gap between human visual hardness and model uncertainty is usually prohibitive because it is laborious to collect such highly subjective human annotations. In addition, these annotations are application or dataset specific, which significantly reduces the scalability of uncertainty estimation models that are directly supervised by them. This makes yet another motivation to this work since AVH is naturally obtained for free without any confidence supervision. In our case, we only leverage such human annotated visual hardness measure for correlation testing. In this section, We first provide four hypothesis and test them accordingly.

Hypothesis 1. *AVH has a correlation with Human Selection Frequency.*

Outcome: Null Hypothesis Rejected

We use the pre-trained network model to extract the feature embedding \mathbf{x} from each validation sample and also provide the class weights \mathbf{w} to compute $\mathcal{AVH}(\mathbf{x})$. Note that we linearly scale the range of $\mathcal{AVH}(\mathbf{x})$ to $[0, 1]$. Table 1 shows the overall consistent and stronger correlation between $\mathcal{AVH}(\mathbf{x})$ and HSF (p-value < 0.001 rejects the null hypothesis). From the coefficients shown in different bins of sample hardness, we can see that the harder the sample, the weaker the correlation. Also Note that we validate the results across different CNN architectures and found that better models tend to have higher coefficients.

The plot on the left in Figure 5 indicates the strong correlation between $\mathcal{AVH}(\mathbf{x})$ and HSF on validation images. One intuition behind this correlation is that the class weights W might correspond to human perceived semantics for each

Table 2. This table presents the Spearman’s rank, Pearson, and Kendall’s Tau correlation coefficients between Human Selection Frequency and AVH, Model Confidence on ResNet-50, along with significance testings between coefficient pairs. Note that having p-value < 0.05 indicates that the result is statistically significant.

Type	Coef with AVH	Coef with Model Confidence	Z_{avh}	Z_{mc}	Z value	p-value
Spearman’s rank	0.360	0.325	0.377	0.337	4.85	< .00001
Pearson	0.385	0.341	0.406	0.355	6.2	< .00001
Kendall’s Tau	0.257	0.231	0.263	0.235	3.38	.0003

category and thereby $\mathcal{AVH}(\mathbf{x})$ corresponds to human’s semantic categorization of an image. In order to test if the strong correlation holds for all models, we perform the same set of experiments on different backbones, including AlexNet, VGG-19 and DenseNet-121.

Hypothesis 2. *Model Confidence has a correlation with Human Selection Frequency.*

Outcome: Null Hypothesis Rejected

An interesting observation in (Recht et al., 2019) shows that HSF has strong influence on the Model Confidence. Specifically, examples with low HSF tends to have relatively low Model Confidence. Naturally we examine if the correlation between Model Confidence and HSF is strong. Specifically, all ImageNet validation images are evaluated by the pre-trained models. The corresponding output is simply the Model Confidence on each image. From Table 1, we can first see that it is clear that because p-value is < 0.001, Model Confidence does have a strong correlation with HSF. However, the correlation coefficient for Model Confidence and HSF is consistently lower than that of AVH and HSF.

The middle plot in Figure 5 presents a two-dimensional histogram for the correlation visualization. The x-axis represents HSF, and the y-axis represents Model Confidence. Each bin exhibits the number of images which lie in the corresponding range. We can observe the high density at the right corner, which means the majority of the images have both high human and model accuracy. However, there is a considerable amount of density on the range of medium human accuracy but either extremely low or high model accuracy. One may question that the difference of the correlation coefficient is not large, thereby we also run statistical testing on the significance of the gap, naturally our next step is to test if the difference is significant.

Hypothesis 3. *AVH has a stronger correlation to Human Selection Frequency than Model Confidence.*

Outcome: Null Hypothesis Rejected

There are three steps for testing if two correlation coefficient are significantly different. First step is applying Fisher Z-Transformation to both coefficients. The Fisher Z-Transformation is a way to transform the sampling distribution of the correlation coefficient so that it becomes normally distributed. Therefore, we apply fisher transformation for each correlation coefficient: Z score for coefficient

of AVH becomes 0.377 and that of Model Confidence becomes 0.337. The second step is to compute the Z value of two Z scores. Then we determined the Z value to be 4.85 from the two above-mentioned Z scores and sample sizes. The last step is that find out the p-value according to the Z table. According to Z table, p-value is 0.00001. Therefore, we reject the null hypothesis and conclude that AVH has statistically significant stronger correlation with HSF than Model Confidence. In later Section 5.1, we also empirically show that such stronger correlation brings cumulative advantages in some applications.

In Table 2, besides the Spearman correlation coefficient, we also show the coefficients of Pearson and Kendall Tau. In addition, in Appendix D, we run the same tests on four different architectures to test whether the same conclusion holds for different models. Our conclusion is that for all the considered models, AVH correlates significantly stronger than Model Confidence, and the correlation is even stronger for better models. This indicates that besides what we have shown in Section 3, AVH is also better aligned with human visual hardness which is related to aleatoric uncertainty.

Hypothesis 4. $\|\mathbf{x}\|_2$ *has a correlation with Human Selection Frequency.*

Outcome: Failure to Reject Null Hypothesis

(Liu et al., 2018b) conjectures that $\|\mathbf{x}\|_2$ accounts for intra-class Human/Model Confidence. Particularly, if the norm is larger, the prediction from the model is also more confident, to some extent. Therefore, we conduct similar experiments like previous section to demonstrate the correlation between $\|\mathbf{x}\|_2$ and HSF. Initially, we compute the $\|\mathbf{x}\|_2$ for every validation sample for all models. Then we normalize $\|\mathbf{x}\|_2$ within each class. Table 1 presents the results for the correlation test. We omit the results for p-value in the table and report here that they are all much higher than 0.05, indicating there is no correlation between $\|\mathbf{x}\|_2$ and HSF. The right plot in Figure 5 uses a two-dimensional histogram to show the correlation for all the validation images. Given that the norm has been normalized with each class, naturally, there is notable density when the norm is 0 or 1. Except for that, there is no obvious correlation between $\|\mathbf{x}\|_2$ and HSF.

We also provide a detailed discussion on the difference between AVH and Model Confidence in Appendix E.

5. Applications

5.1. AVH for Self-training and Domain Adaptation

Unsupervised domain adaptation (Ben-David et al., 2010) presents an important transfer learning problem and deep self-training (Lee, 2013) recently emerged as a powerful framework to this problem (Saito et al., 2017a; Shu et al., 2018; Zou et al., 2018; 2019). Here we show the application of AVH as an improved confidence measure in self-training that could significantly benefit domain adaptation.

Dataset: We conduct experiments on the VisDA-17 (Peng et al., 2017) dataset which is a widely used major benchmark for domain adaptation in image classification. The dataset contains a total number of 152,409 2D synthetic images from 12 categories in the source training set, and 55,400 real images from MS-COCO (Lin et al., 2014) with the same set of categories as the target domain validation set. We follow the protocol of previous works to train a source model with the synthetic training set, and report the model performance on target validation set upon adaptation.

Baseline: We use class-balanced self-training (CBST) (Zou et al., 2018) as a state-of-the-art self-training baseline. We also compare our model with confidence regularized self-training (CRST)¹ (Zou et al., 2019), a more recent framework improved over CBST with network prediction/pseudo-label regularized with smoothness. Specifically, our work follows the exact implementation of CBST/CRST.

Specifically, given the labeled source domain training set $\mathbf{x}_s \in \mathbf{X}_S$ and the unlabeled target domain data $\mathbf{x}_t \in \mathbf{X}_T$, with known source labels $\mathbf{y}_s = (y_s^{(1)}, \dots, y_s^{(K)}) \in \mathbf{Y}_S$ and unknown target labels $\hat{\mathbf{y}}_t = (\hat{y}_t^{(1)}, \dots, \hat{y}_t^{(K)}) \in \hat{\mathbf{Y}}_T$ from K classes, CBST performs joint network learning and pseudo-label estimation by treating pseudo-labels as discrete learnable latent variables with the following loss:

$$\begin{aligned} \min_{\mathbf{w}, \hat{\mathbf{Y}}_T} \mathcal{L}_{CB}(\mathbf{w}, \hat{\mathbf{Y}}) = & - \sum_{s \in S} \sum_{k=1}^K y_s^{(k)} \log p(k|\mathbf{x}_s; \mathbf{w}) \\ & - \sum_{t \in T} \sum_{k=1}^K \hat{y}_t^{(k)} \log \frac{p(k|\mathbf{x}_t; \mathbf{w})}{\lambda_k} \text{ s.t. } \hat{\mathbf{y}}_t \in \mathbf{E}^K \cup \{\mathbf{0}\}, \forall t \end{aligned} \quad (2)$$

where the feasible set of pseudo-labels is the union of $\{\mathbf{0}\}$ and the K dimensional one-hot vector space \mathbf{E}^K , and \mathbf{w} and $p(k|\mathbf{x}; \mathbf{w})$ represent the network weights and the classifier’s softmax probability for class k , respectively. In addition, λ_k serves as a class-balancing parameter controlling the pseudo-label selection of class k , and is determined by the softmax confidence ranked at portion p (in descending order) among samples predicted to class k . Therefore, only one parameter p is used to determine all λ_k ’s. The optimization problem in (2) can be solved via minimizing with respect to \mathbf{w} and

$\hat{\mathbf{Y}}$ alternatively, and the solver of $\hat{\mathbf{Y}}$ can be written as:

$$\hat{y}_t^{(k)*} = \begin{cases} 1, & \text{if } k = \arg \max_c \left\{ \frac{p(c|\mathbf{x}_t; \mathbf{w})}{\lambda_c} \right\} \text{ and } p(k|\mathbf{x}_t; \mathbf{w}) > \lambda_k \\ 0, & \text{otherwise} \end{cases}$$

The optimization with respect to \mathbf{w} is simply network re-training with source labels and estimated pseudo-labels. The complete self-training process involves alternative repeat of network re-training and pseudo-label estimation.

CBST+AVH: We seek to improve the pseudo-label solver with better confidence measure from AVH. We propose the following definition of angular visual confidence (AVC) to represent the predicted probability of class c :

$$\mathcal{AVC}(c|\mathbf{x}; \mathbf{w}) = \frac{\pi - \mathcal{A}(\mathbf{x}, \mathbf{w}_c)}{\sum_{k=1}^K (\pi - \mathcal{A}(\mathbf{x}, \mathbf{w}_k))}, \quad (3)$$

and pseudo-label estimation in CBST+AVH is defined as:

$$\hat{y}_t^{(k)*} = \begin{cases} 1, & \text{if } k = \arg \max_c \left\{ \frac{p(c|\mathbf{x}_t; \mathbf{w})}{\lambda_c} \right\} \\ & \text{and } \mathcal{AVC}(k|\mathbf{x}_t; \mathbf{w}) > \beta_k \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

where $p(k|\mathbf{x}_t; \mathbf{w})$ is the softmax output of \mathbf{x}_t . λ_k and β_k are determined respectively by referring to $p(k|\mathbf{x}_t; \mathbf{w})$ and $\mathcal{AVC}(k|\mathbf{x}_t; \mathbf{w})$ ranked at a particular portion among samples predicted to class k , following the same definition of λ_k in CBST. In addition, network re-training in CBST+AVH follows the softmax self-training loss in (2).

One could see that AVH changes the self-training behavior by having improved pseudo-label selection in (4) in terms of $\mathcal{AVC}(k|\mathbf{x}_t; \mathbf{w}) > \beta_k$. Specifically, the condition determines which samples are not ignored during self-training based on AVC. With the improved confidence measure that better resembles human visual hardness, this aspect is likely to influence the final performance of self-training.

Experimental Results: We present the results of the proposed method in Table 3, and also show its performance with respect to different self-training epochs in Figure 6. One could see that CBST+AVH outperforms both CBST and CRST by a very significant margin. We would like to emphasize that this is a very compelling result under “apples to apples” comparison with the same source model, implementation and hyper-parameters.

Analysis: A major challenge of self-training is the amplification of error due to misclassified pseudo-labels. Therefore, traditional self-training methods such as CBST often use model confidence as the measure to select confidently labeled examples. The hope is that higher confidence potentially implies lower error rate. While this generally proves useful, the model tends to focus on the “less informative” samples, whereas ignoring the “more informative”, harder ones near classifier boundaries that could be essential for learning a better classifier. More details are in Appendix F.

¹We consider MRKLD+LRENT which is reported to be the highest one in (Zou et al., 2019).

Table 3. Class-wise and mean classification accuracies on VisDA-17.

Method	Aero	Bike	Bus	Car	Horse	Knife	Motor	Person	Plant	Skateboard	Train	Truck	Mean
Source (Saito et al., 2018)	55.1	53.3	61.9	59.1	80.6	17.9	79.7	31.2	81.0	26.5	73.5	8.5	52.4
MMD (Long et al., 2015b)	87.1	63.0	76.5	42.0	90.3	42.9	85.9	53.1	49.7	36.3	85.8	20.7	61.1
DANN (Ganin et al., 2016)	81.9	77.7	82.8	44.3	81.2	29.5	65.1	28.6	51.9	54.6	82.8	7.8	57.4
ENT (Grandvalet & Bengio, 2005)	80.3	75.5	75.8	48.3	77.9	27.3	69.7	40.2	46.5	46.6	79.3	16.0	57.0
MCD (Saito et al., 2017b)	87.0	60.9	83.7	64.0	88.9	79.6	84.7	76.9	88.6	40.3	83.0	25.8	71.9
ADR (Saito et al., 2018)	87.8	79.5	83.7	65.3	92.3	61.8	88.9	73.2	87.8	60.0	85.5	32.3	74.8
Source (Zou et al., 2019)	68.7	36.7	61.3	70.4	67.9	5.9	82.6	25.5	75.6	29.4	83.8	10.9	51.6
CBST (Zou et al., 2019)	87.2	78.8	56.5	55.4	85.1	79.2	83.8	77.7	82.8	88.8	69.0	72.0	76.4
CRST (Zou et al., 2019)	88.0	79.2	61.0	60.0	87.5	81.4	86.3	78.8	85.6	86.6	73.9	68.8	78.1
Proposed	93.3	80.2	78.9	60.9	88.4	89.7	88.9	79.6	89.5	86.8	81.5	60.0	81.5




Figure 6. Adaptation accuracy vs. epoch for different comparing methods on VisDA-17.

Table 4. Statistics of the examples selected by CBST+AVH and CBST/CRST.

Method	TP Rate	AVH (avg)	Model Confidence	Norm $\ x\ $
CBST+AVH	0.844	0.118	0.961	20.84
CBST/CRST	0.848	0.117	0.976	21.28

An advantage we observe from AVH is that the improved calibration leads to more frequent sampling of harder samples, whereas the pseudo-label classification on these hard samples generally outperforms softmax results. Table 4 shows the statistics of examples selected with AVH and model confidence respectively at the beginning of the training process. The true positive rate (TP Rate) for CBST+AVH remains similar to CBST/CRST, indicating AVH overall is not introducing additional noise compare to model confidence. On the other hand, it is observed that the average model confidence of AVH selected samples is lower, indicating there are more selected hard samples that are closer to the decision boundary. It is also observed that the average sample norm by AVH is also lower, confirming the influence of sample norms on ultimate model confidence.

5.2. AVH-based Loss for Domain Generalization

The problem of domain generalization (DG) is to learn from multiple training domains, and extract a domain-agnostic model that can then be applied to an unseen domain. Since we have no assumption on how the unseen domain looks like, the generalization on the unseen domains will mostly depend on the generalizability of the neural network. We

Table 5. Domain generalization accuracy (%) on PACS dataset.

Method	Painting	Cartoon	Photo	Sketch	Avg
AlexNet (Li et al., 2017)	62.86	66.97	89.50	57.51	69.21
MLDG (Li et al., 2018)	66.23	66.88	88.00	58.96	70.01
MetaReg (Balaji et al., 2018)	69.82	70.35	91.07	59.26	72.62
Feature-critic (Li et al., 2019)	64.89	71.72	89.94	61.85	72.10
Baseline CNN-10	66.46	67.88	89.70	51.72	68.94
CNN-10 + AVH	72.02	66.42	90.12	61.26	72.46

use the challenging PACS dataset (Li et al., 2017) which consists of Art painting, Cartoon, Photo and Sketch domains. For each domain, we leave it out as the test set and train our models on rest of the three domains.

Specifically, we train a 10-layer plain CNN with the following AVH-based loss (additional details in Appendix F):

$$\mathcal{L}_{AVH} = \sum_i \frac{\exp(s(\pi - \mathcal{A}(\mathbf{x}_i, \mathbf{w}_{y_i})))}{\sum_{k=1}^K \exp(s(\pi - \mathcal{A}(\mathbf{x}_i, \mathbf{w}_k)))}, \quad (5)$$

where s is hyperparameter that adjusts the scale of the output logits and implicitly controls the optimization difficulty. This hyperparameter is typically set by cross-validation. Experimental results are reported in Table 5. With the proposed new loss which directly has an AVH-based design, a simple CNN is outperforming baseline and recent methods that are based on more complex models. In fact, similar learning objectives have also been shown useful in image recognition (Liu et al., 2017c) and face recognition (Wang et al., 2017; Ranjan et al., 2017), indicating that AVH is generally effective to improve generalization in various tasks.

6. Concluding Remarks

We propose a novel measure for CNN models known as Angular Visual Hardness. Our comprehensive empirical studies show that AVH can serve as an indicator of generalization abilities of neural networks, and improving SOTA accuracy entails improving accuracy on hard examples. AVH also has a significantly stronger correlation with Human Selection Frequency. We empirically show the advantage of AVH over Model Confidence in self-training for domain adaptation task and loss function for domain generalization task. AVH can be useful in other applications such as deep metric learning, fairness, knowledge transfer, etc. and we plan to investigate them in the future (discussions in Appendix G).

Acknowledgements

Work done during internship at NVIDIA. We would like to thank Shiyu Liang, Yue Zhu and Yang Zou for the valuable discussions that enlighten our research. We are also grateful to the anonymous reviewers for their constructive comments that significantly helped to improve our paper. Weiyang Liu is partially supported by Baidu scholarship and NVIDIA GPU grant. This work was supported by NSF-1652131, NSF-BIGDATA 1838177, AFOSR-YIPFA9550-18-1-0152, Amazon Research Award, and ONR BRC grant for Randomized Numerical Linear Algebra.

References

- Balaji, Y., Sankaranarayanan, S., and Chellappa, R. Metareg: Towards domain generalization using meta-regularization. In *NeurIPS*, 2018.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- Bengio, Y., Louradour, J., Collobert, R., and Weston, J. Curriculum learning. In *ICML*, 2009.
- Berardino, A., Ball, J., Laparra, V., and Simoncelli, E. P. Eigen-distortions of hierarchical representations, 2017.
- Buolamwini, J. and Gebru, T. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on Fairness, Accountability and Transparency*, pp. 77–91, 2018.
- Chen, M., Weinberger, K. Q., and Blitzer, J. Co-training for domain adaptation. In *Advances in neural information processing systems*, pp. 2456–2464, 2011.
- Dekel, R. Human perception in computer vision, 2017.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- Deng, J., Guo, J., Xue, N., and Zafeiriou, S. Arcface: Additive angular margin loss for deep face recognition. In *CVPR*, 2019.
- Der Kiureghian, A. and Ditlevsen, O. Aleatory or epistemic? does it matter? *Structural safety*, 31(2):105–112, 2009.
- Dodge, S. and Karam, L. A study and comparison of human and deep learning recognition performance under visual distortions. *2017 26th International Conference on Computer Communication and Networks (ICCCN)*, Jul 2017. doi: 10.1109/icccn.2017.8038465.
- Fellbaum, C. Wordnet and wordnets. 2005.
- Finn, C., Abbeel, P., and Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *ICML*, 2017.
- Gal, Y. and Ghahramani, Z. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *ICML*, 2016.
- Gal, Y., Hron, J., and Kendall, A. Concrete dropout. In *NeurIPS*, 2017.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *JMLR*, 2016.
- Geirhos, R., Temme, C. R., Rauber, J., Schütt, H. H., Bethge, M., and Wichmann, F. A. Generalisation in humans and deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 7538–7550, 2018.
- Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., and Brendel, W. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. 2019.
- Goodfellow, I. J., Shlens, J., and Szegedy, C. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *NeurIPS*, 2005.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.
- Hadsell, R., Chopra, S., and LeCun, Y. Dimensionality reduction by learning an invariant mapping. In *CVPR*, 2006.
- Hariharan, B. and Girshick, R. Low-shot visual recognition by shrinking and hallucinating features. In *ICCV*, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hinton, G., Vinyals, O., and Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

- Kendall, A. and Gal, Y. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017.
- Kheradpisheh, S. R., Ghodrati, M., Ganjtabesh, M., and Masquelier, T. Deep networks can resemble human feed-forward vision in invariant object recognition. *Scientific Reports*, 6(1), Sep 2016. ISSN 2045-2322. doi: 10.1038/srep32672.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Kumar, A., Sarawagi, S., and Jain, U. Trainable calibration measures for neural networks from kernel mean embeddings. In *International Conference on Machine Learning*, pp. 2810–2819, 2018.
- Kumar, M. P., Packer, B., and Koller, D. Self-paced learning for latent variable models. In *NeurIPS*, 2010.
- Lakshminarayanan, B., Pritzel, A., and Blundell, C. Simple and scalable predictive uncertainty estimation using deep ensembles. In Guyon, I., Luxburg, U. V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30*, pp. 6402–6413. Curran Associates, Inc., 2017.
- Lee, D.-H. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *ICML Workshop on Challenges in Representation Learning*, 2013.
- Li, D., Yang, Y., Song, Y.-Z., and Hospedales, T. M. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pp. 5542–5550, 2017.
- Li, Y., Yang, Y., Zhou, W., and Hospedales, T. M. Feature-critic networks for heterogeneous domain generalization. *arXiv preprint arXiv:1901.11448*, 2019.
- Li, Z. and Hoiem, D. Reducing over-confident errors outside the known distribution. *arXiv preprint arXiv:1804.03166*, 2018.
- Li, Z., Xu, Z., Ramamoorthi, R., Sunkavalli, K., and Chandraker, M. Learning to reconstruct shape and spatially-varying reflectance from a single image. In *SIGGRAPH Asia 2018 Technical Papers*, pp. 269. ACM, 2018.
- Lin, R., Liu, W., Liu, Z., Feng, C., Yu, Z., Rehg, J. M., Xiong, L., and Song, L. Regularizing neural networks via minimizing hyperspherical energy. In *CVPR*, 2020.
- Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., and Zitnick, C. L. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Lindsay, P. H. and Norman, D. A. *Human information processing: An introduction to psychology*. Academic press, 2013.
- Liu, W., Wen, Y., Yu, Z., and Yang, M. Large-margin softmax loss for convolutional neural networks. In *ICML*, 2016.
- Liu, W., Dai, B., Humayun, A., Tay, C., Yu, C., Smith, L. B., Rehg, J. M., and Song, L. Iterative machine teaching. In *ICML*, 2017a.
- Liu, W., Wen, Y., Yu, Z., Li, M., Raj, B., and Song, L. Sphereface: Deep hypersphere embedding for face recognition. In *CVPR*, 2017b.
- Liu, W., Zhang, Y.-M., Li, X., Yu, Z., Dai, B., Zhao, T., and Song, L. Deep hyperspherical learning. In *NeurIPS*, 2017c.
- Liu, W., Lin, R., Liu, Z., Liu, L., Yu, Z., Dai, B., and Song, L. Learning towards minimum hyperspherical energy. In *NeurIPS*, 2018a.
- Liu, W., Liu, Z., Yu, Z., Dai, B., Lin, R., Wang, Y., Rehg, J. M., and Song, L. Decoupled networks. In *CVPR*, 2018b.
- Liu, W., Liu, Z., Rehg, J. M., and Song, L. Neural similarity learning. In *NeurIPS*, 2019.
- Liu, W., Lin, R., Liu, Z., Rehg, J. M., Xiong, L., Weller, A., and Song, L. Orthogonal over-parameterized training. *arXiv preprint arXiv:2004.04690*, 2020.
- Long, J., Shelhamer, E., and Darrell, T. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015a.
- Long, M., Cao, Y., Wang, J., and Jordan, M. I. Learning transferable features with deep adaptation networks. In *ICML*, 2015b.
- Martin Cichy, R., Khosla, A., Pantazis, D., and Oliva, A. Dynamics of scene representations in the human brain revealed by magnetoencephalography and deep neural networks. *NeuroImage*, 153:346358, Jun 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2016.03.063.
- Oh Song, H., Xiang, Y., Jegelka, S., and Savarese, S. Deep metric learning via lifted structured feature embedding. In *CVPR*, 2016.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.

- Pramod, R. T. and Arun, S. P. Do computational models differ systematically from human object perception? *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun 2016. doi: 10.1109/cvpr.2016.177.
- Ranjan, R., Castillo, C. D., and Chellappa, R. L2-constrained softmax loss for discriminative face verification. *arXiv preprint arXiv:1703.09507*, 2017.
- Recht, B., Roelofs, R., Schmidt, L., and Shankar, V. Do imagenet classifiers generalize to imagenet? *arXiv preprint arXiv:1902.10811*, 2019.
- Saito, K., Ushiku, Y., and Harada, T. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, 2017a.
- Saito, K., Watanabe, K., Ushiku, Y., and Harada, T. Maximum classifier discrepancy for unsupervised domain adaptation. 2017b.
- Saito, K., Ushiku, Y., Harada, T., and Saenko, K. Adversarial dropout regularization. In *ICLR*, 2018.
- Schroff, F., Kalenichenko, D., and Philbin, J. Facenet: A unified embedding for face recognition and clustering. In *CVPR*, 2015.
- Shu, R., Bui, H. H., Narui, H., and Ermon, S. A dirt-t approach to unsupervised domain adaptation. 2018.
- Simonyan, K. and Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Sohn, K. Improved deep metric learning with multi-class n-pair loss objective. In *NeurIPS*, 2016.
- Soudry, D., Hoffer, E., Nacson, M. S., Gunasekar, S., and Srebro, N. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.
- Springenberg, J. T., Klein, A., Falkner, S., and Hutter, F. Bayesian optimization with robust bayesian neural networks. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 29*, pp. 4134–4142. Curran Associates, Inc., 2016.
- Sun, Y., Wang, X., and Tang, X. Deep learning face representation from predicting 10,000 classes. In *CVPR*, 2014.
- Tang, P., Wang, X., Bai, X., and Liu, W. Multiple instance detection network with online instance classifier refinement. In *CVPR*, 2017.
- Wang, F., Xiang, X., Cheng, J., and Yuille, A. L. Normface: L2 hypersphere embedding for face verification. In *ACM-MM*, 2017.
- Wang, F., Liu, W., Liu, H., and Cheng, J. Additive margin softmax for face verification. *arXiv preprint arXiv:1801.05599*, 2018a.
- Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., and Liu, W. Cosface: Large margin cosine loss for deep face recognition. In *CVPR*, 2018b.
- Wang, Y., Liu, W., Ma, X., Bailey, J., Zha, H., Song, L., and Xia, S.-T. Iterative learning with open-set noisy labels. In *CVPR*, 2018c.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Wei, C. and Ma, T. Improved sample complexities for deep networks and robust classification via an all-layer margin. *arXiv preprint arXiv:1910.04284*, 2019.
- Wu, C.-Y., Manmatha, R., Smola, A. J., and Krahenbuhl, P. Sampling matters in deep embedding learning. In *ICCV*, 2017.
- Yi, D., Lei, Z., Liao, S., and Li, S. Z. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.
- Zhou, B., Lapedriza, A., Xiao, J., Torralba, A., and Oliva, A. Learning deep features for scene recognition using places database. In *NeurIPS*, 2014.
- Zhou, Y., Kantarcioglu, M., and Thuraisingham, B. Self-training with selection-by-rejection. In *2012 IEEE 12th international conference on data mining*, pp. 795–803. IEEE, 2012.
- Zhu, X. Semi-supervised learning tutorial.
- Zou, Y., Yu, Z., Kumar, B., and Wang, J. Domain adaptation for semantic segmentation via class-balanced self-training. *arXiv preprint arXiv:1810.07911*, 2018.
- Zou, Y., Yu, Z., Liu, X., Kumar, B., and Wang, J. Confidence regularized self-training. *arXiv preprint arXiv:1908.09822*, 2019.

Appendix: Angular Visual Hardness

A. Simulation Details

Gaussian Simulation Plot: We generate 2000 3-d random vectors from two multivariate normal distribution (1000 for each) and normalize to unit norm, shown in red and green color on the left plot in Figure 7. Then these data points are passed as the inputs to a simple multi layer perceptron classification model with one 3×2 hidden layer. Upon convergence, we compute the AVH scores for each data point. The middle image shows the visualization of AVH scores for all data points, with lighter color representing higher AVH scores. It is obvious that AVH scores for points lying on the intersection of two clusters are higher, which agrees with the intuition that those are hard examples. We also compute the ℓ_2 norm of the feature embeddings shown in the right plot. One can see there is less correlation with hard examples.




Figure 7. Toy example of two overlapping Gaussian distributions (classes) on a unit sphere. Left: samples from the distributions as input to a multi layer perceptron (MLP). Middle: AVH heat map produced by MLP, where samples in lighter colors (higher hardness) are mostly overlapping hard examples. Right: ℓ_2 -norm heat map, where certain non-overlapping samples also have higher values.

MNIST Simulation Plot: We train MNIST with a very simple CNN model which the dimension of the embedding (right before the classifier) is 2. Figure 8 shows the visualization of those 2D embeddings.




Figure 8. Visualization of embeddings on MNIST by setting their dimensions to 2 in a CNN.

B. Additional Results of Training Dynamics

B.1. Additional Results on ImageNet

Model Confidence: Figure 9 shows the training dynamics of the model confidence corresponding to AlexNet, VGG-19, and ResNet-50.




Figure 9. Number of epochs vs. Model Confidence. Results from left to right correspond to AlexNet, VGG-19 and ResNet-50.

Averaged training dynamics: In Figure 10, we plot the average embedding norm, AVH and model accuracies for AlexNet, VGG-19, ResNet-50 and DenseNet-121 over the validation samples.

Image degradation: Because CNNs and humans can achieve similar accuracy on large-scale benchmark dataset such as ImageNet, a number of works have investigated similarities and differences between CNNs and human vision (Martin Cichy et al., 2017; Kheradpisheh et al., 2016; Dodge & Karam, 2017; Dekel, 2017; Pramod & Arun, 2016; Berardino et al., 2017). Since human annotation data is relatively hard to obtain, researchers have proposed an alternative measure of visual hardness on images based on image degradation (Lindsay & Norman, 2013). This involves adding noise or changing image properties such as contrast, blurriness, and brightness. (Geirhos et al., 2018) employed psychological studies to validate the degradation method as a way to measure human visual hardness. It should be noted that the artificial visual hardness introduced by degradation is a different concept from the natural visual hardness. The hardness based on degradation only reflects the hardness of a single original image with various of transformations, while natural visual hardness based on the ambiguity of human perception across a distribution of natural images. In the following additional experiments, we also consider different level of degradation as the surrogate of human visual hardness besides Human Selection Frequency.

Definition 4 (Image Degradation Level). We define another way to measure human visual hardness on pictures as Image Degradation Level. We consider two degradation methods in this paper, decreasing contrast and adding noise. Quantitatively, Image Degradation Level for decreasing contrast is directly the contrast level. Image Degradation Level for adding noise is the amount of pixel-wise additive uniform noise.

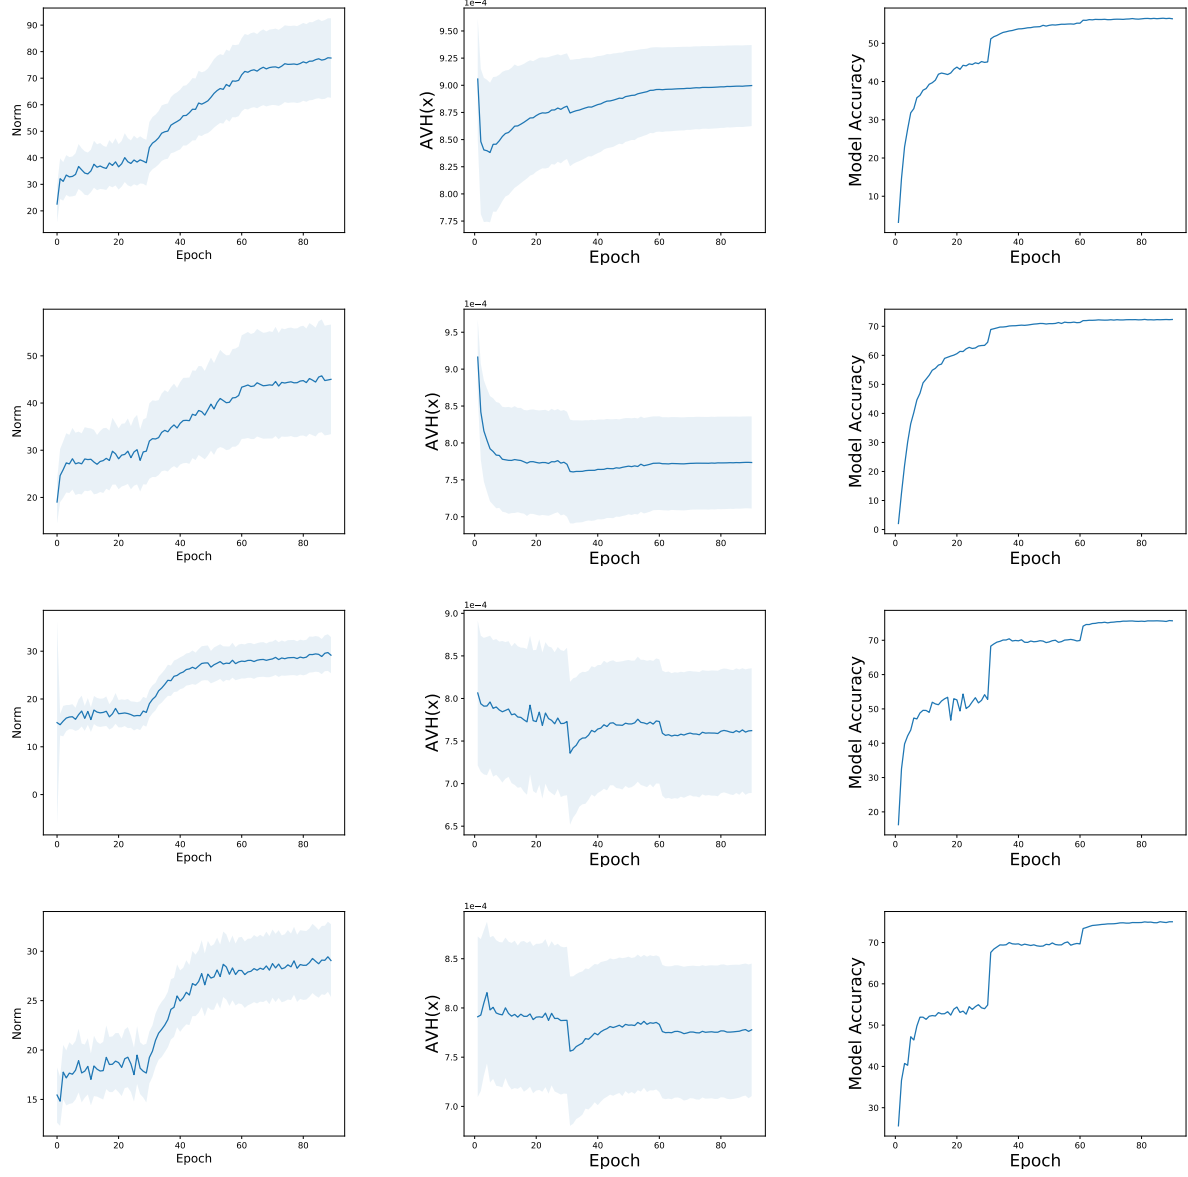


Figure 10. Averaged training dynamics on ImageNet validation set. Columns from left to right: number of epochs vs. average ℓ_2 norm, number of epochs vs. average AVH score, and number of epochs vs. model accuracy. Rows from top to bottom: dynamics corresponding to AlexNet, VGG-19, ResNet-50, and DenseNet-121.

Dynamics across noise degradation levels: In Figure 11, we illustrate the averaged training dynamics on the ImageNet validation set across five image noise degradation levels - [0.4, 0.3, 0.2, 0.1, 0.0].

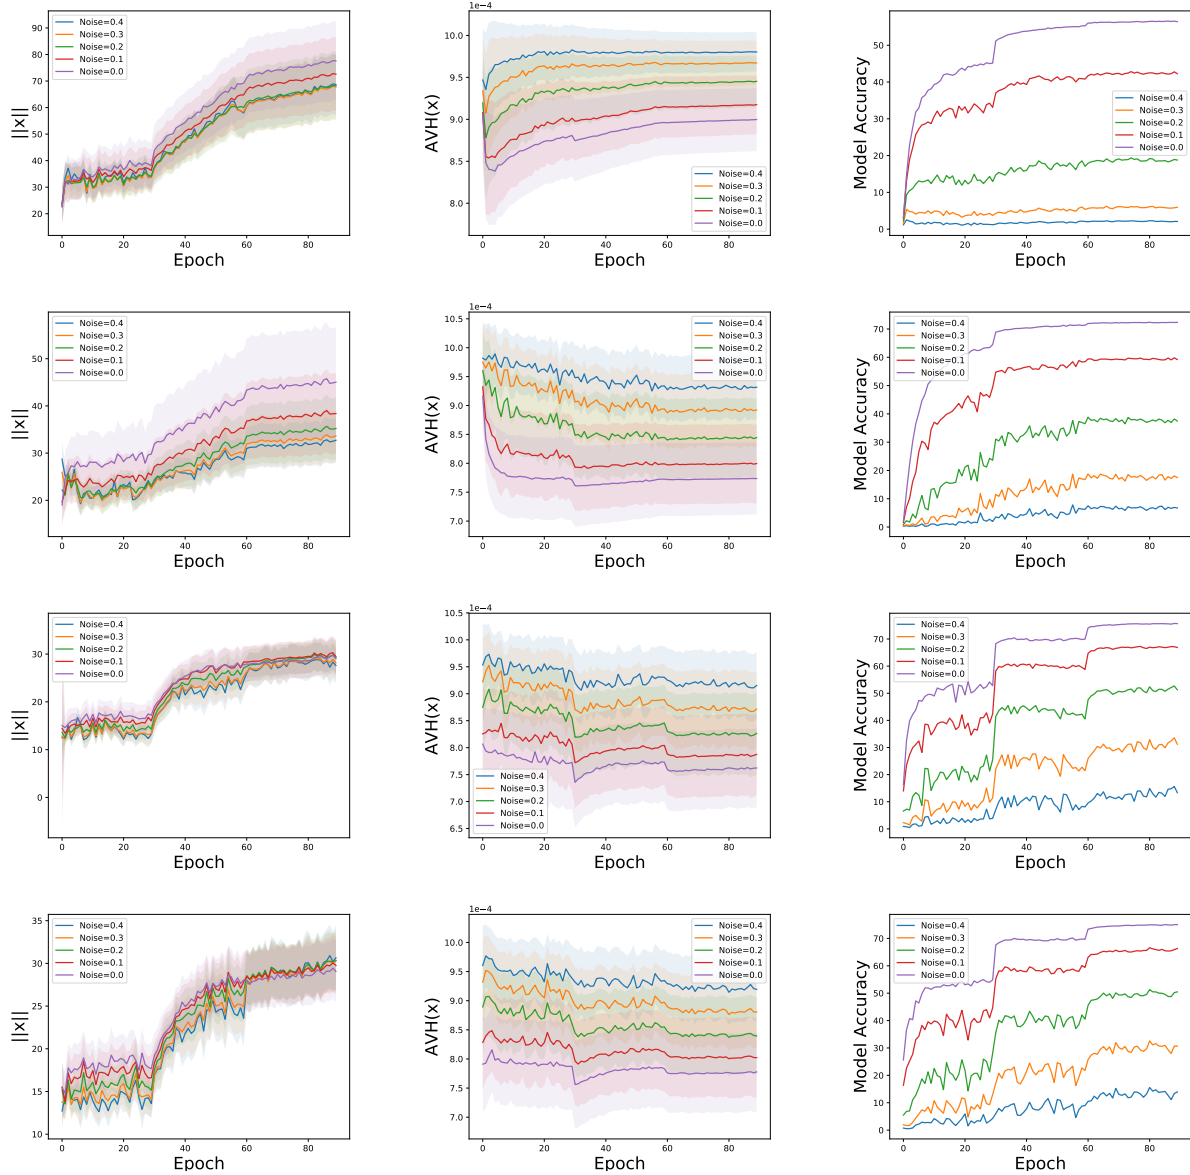


Figure 11. Averaged training dynamics across different noise degradation levels. Columns from left to right: number of epochs vs. average ℓ_2 norm, number of epochs vs. average AVH score, and number of epochs vs. model accuracy. Rows from top to bottom: dynamics corresponding to AlexNet, VGG-19, ResNet-50, and DenseNet-121.

Dynamics across contrast degradation levels: In Figure 12, we illustrate the averaged training dynamics on the ImageNet validation set across five image contrast degradation levels - [0.1, 0.2, 0.3, 0.6, 1.0].

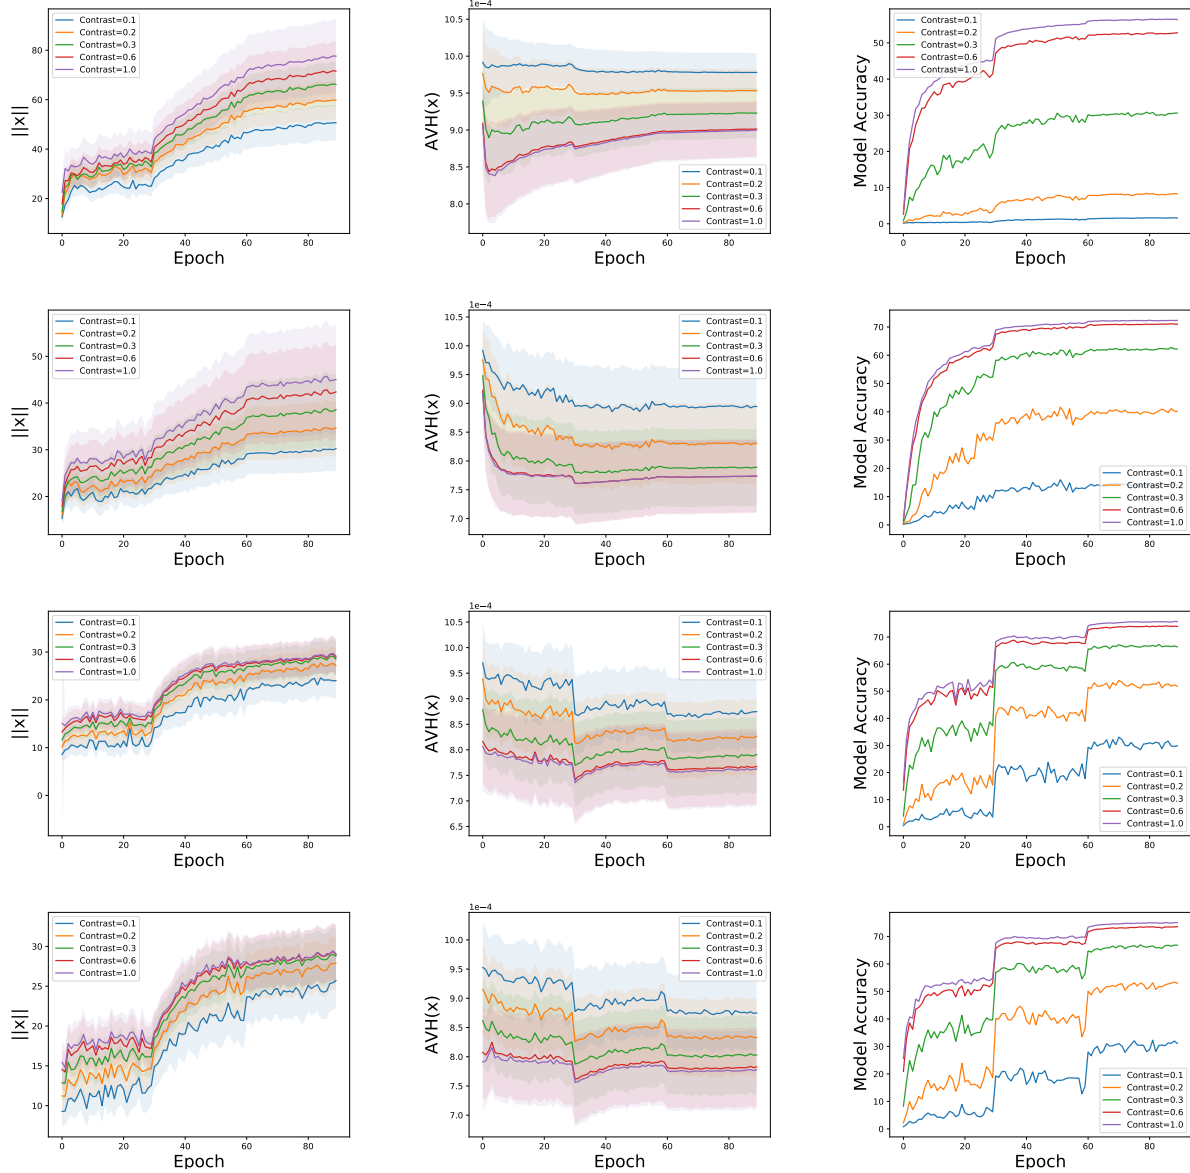


Figure 12. Averaged training dynamics across different contrast degradation levels. Columns from left to right: number of epochs vs. average ℓ_2 norm, number of epochs vs. average AVH score, and number of epochs vs. model accuracy. Rows from top to bottom: dynamics corresponding to AlexNet, VGG-19, ResNet-50, and DenseNet-121.

One can see that the observations from Section 3 in the main paper also hold on this set of experiments.

B.2. Additional Results on CIFAR-10, CIFAR-100 and MNIST

Figure 13 and 14 show the dynamics of average ℓ_2 norm of the embeddings and average AVH(x) on CIFAR-10 and CIFAR-100 datasets respectively. We can observe the similar phenomena we have discussed in section 3. It further supports our theoretical foundation from (Soudry et al., 2018) that gradient descent converges to the same direction as maximum margin solutions irrelevant to the ℓ_2 norm of classifier weights or feature embeddings.

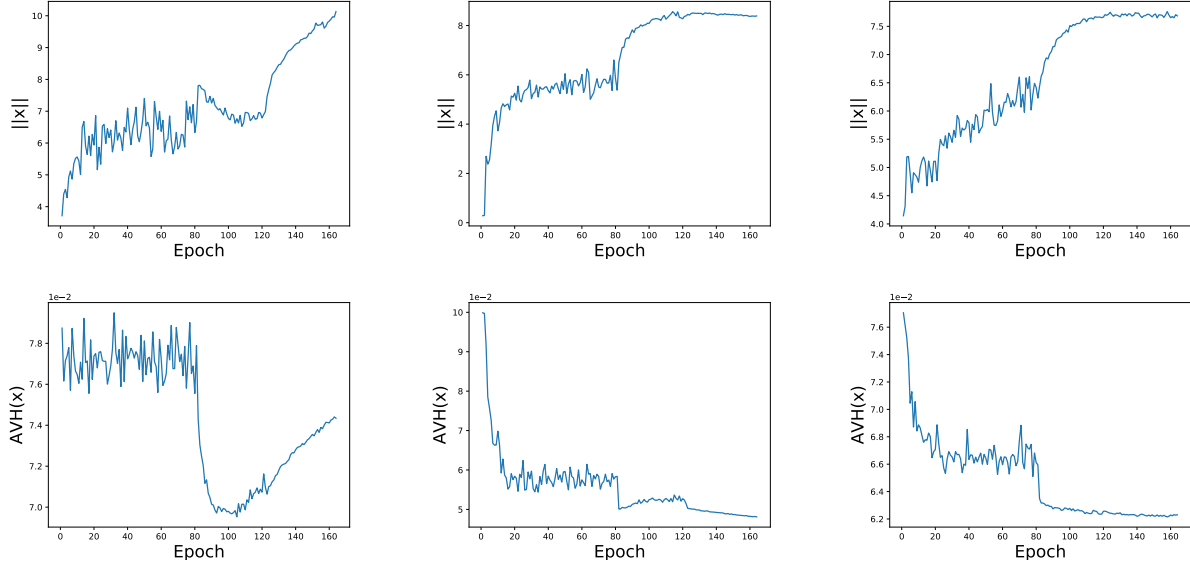


Figure 13. The top three plots show the number of Epochs v.s. Average ℓ_2 norm across CIFAR-10 validation samples. The bottom three plots represent number of Epochs v.s. Average AVH(x). From left to right, we use AlexNet, VGG-19 and ResNet-50.

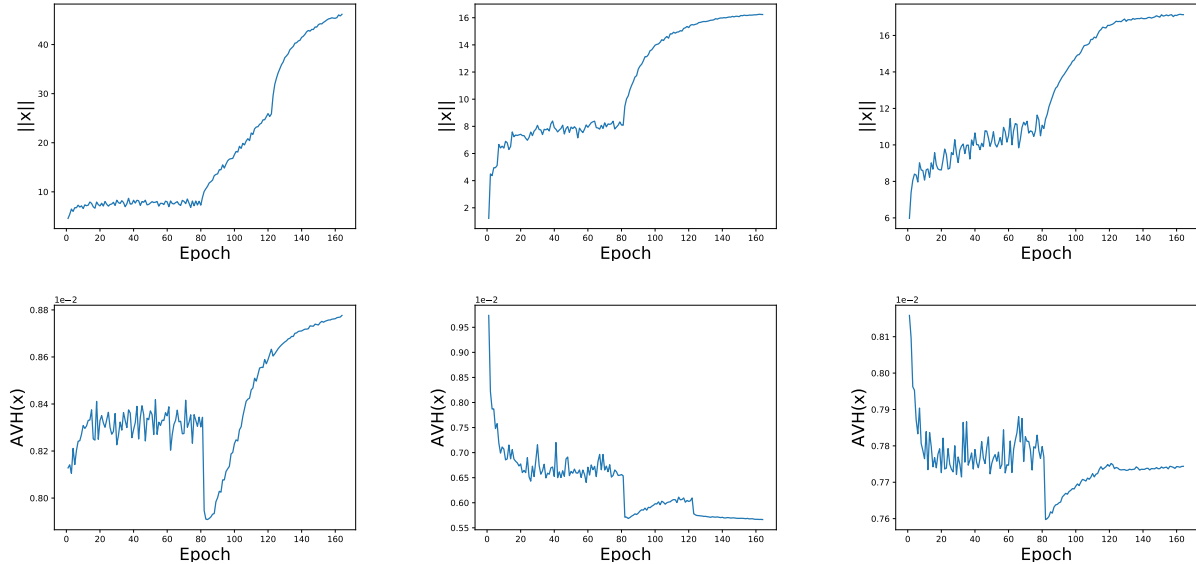


Figure 14. The top three plots show the number of Epochs v.s. Average ℓ_2 norm across CIFAR-100 validation samples. The bottom three plots represent number of Epochs v.s. Average AVH(x). From left to right, we use AlexNet, VGG-19 and ResNet-50.

Figure 15 illustrates how the average norm of the feature embedding and AVH between feature and class embedding for testing samples vary in 60 iterations during the training process on MNIST. The average norm increases with a large initial slope but it flattens slightly after 10 iterations. On the other hand, the average angle decreases sharply at the beginning and then becomes almost flat after 10 iterations.




Figure 15. Average ℓ_2 norm and angle of the embedding across all testing samples v.s. iteration number.

C. Additional Discussions for Observations in Training Dynamics

Observation 2 in section 3 describes that AVH hits a plateau very early even when the accuracy or loss is still improving. AVH is more important than $\|\mathbf{x}\|_2$ in the sense that it is the key factor deciding which class the input sample is classified to.

However, optimizing the norm under the current softmax cross-entropy loss would be easier for easy examples. Let us consider a simple binary classification case where the softmax score for class 1 is

$$\frac{\exp(\mathbf{w}_1 \mathbf{x})}{\sum_i \exp(\mathbf{w}_i \mathbf{x})} = \frac{\exp(\|\mathbf{w}_1\| \|\mathbf{x}\| \cos(\theta_{\mathbf{w}_1, \mathbf{x}}))}{\sum_i \exp(\|\mathbf{w}_i\| \|\mathbf{x}\| \cos(\theta_{\mathbf{w}_i, \mathbf{x}}))} \quad (6)$$

where \mathbf{w}_i is the classifier weights of class i , \mathbf{x} is the input deep feature and $\theta_{\mathbf{w}_i, \mathbf{x}}$ is the angle between \mathbf{w}_i and \mathbf{x} . To simplify, we assume the norm of \mathbf{w}_1 and \mathbf{w}_2 are the same, and then the classification result is based on the angle now. For easy examples, during early stage of the training, $\theta_{\mathbf{w}_1, \mathbf{x}}$ quickly becomes smaller than $\theta_{\mathbf{w}_2, \mathbf{x}}$ and the network will classify the sample \mathbf{x} as class 1. However, in order to further minimize the cross-entropy loss after making $\theta_{\mathbf{w}_1, \mathbf{x}}$ smaller than $\theta_{\mathbf{w}_2, \mathbf{x}}$, the network has a trivial solution: increasing the feature norm $\|\mathbf{x}\|$ instead of further minimizing the $\theta_{\mathbf{w}_1, \mathbf{x}}$. It is obviously a much more difficult task to minimize $\theta_{\mathbf{w}_1, \mathbf{x}}$ rather than increasing $\|\mathbf{x}\|$. Therefore, the network will tend to increase the feature norm $\|\mathbf{x}\|$ to minimize the cross-entropy loss, which is equivalent to maximizing the Model Confidence in class 1. In fact, this also matches our empirical observation that the feature norm keeps increasing during training. Moreover, this also matches our empirical result that AVH easily gets saturated while Model Confidence can keep improving. For hard examples, after some time of training, the feature norms are unavoidable also increasing (although slower than those of easy examples). We can see from equation 6 that when $\|\mathbf{x}\|$ is very large and $\cos(\theta_{\mathbf{w}_i, \mathbf{x}})$ is very small, improving the angle becomes much harder because for a bit improvement on these examples, the model needs to sacrifice a lot for those easy ones. For the case of value of $\cos(\theta_{\mathbf{w}_i, \mathbf{x}})$ is around the decision boundary, a little change to AVH can cause a lot improvement on loss and accuracy and thereby we can still observe the change of accuracy and loss while AVH plateaus. More details about why $\|\mathbf{x}\|$ might be harmful in the training process is in Appendix E.

D. Additional Experiments for Connections to Human Visual Hardness

D.1. Additional Results for Correlation Testings

In order to run rigorous correlation testings, besides computing the Spearman coefficient, we provide additional results on Pearson and Kendall Tau correlation coefficients. Moreover, we show results for all four architectures, AlexNet, VGG-19, ResNet-50 and DenseNet-121 in Table 6, 7, 8 and 9 respectively to support our claims in section 4.

Table 6. This table presents the Spearman’s rank correlation coefficients between Human Selection Frequency and AVH, Model Confidence on AlexNet. Note that we show the absolute value of the coefficient which represents the strength of the correlation. Z value is computed by Z scores of both coefficients. p-value < 0.05 indicates that the result is statistically significant.

Type	Coef with AVH	Coef with Model Confidence	Z_{avh}	Z_{mc}	Z value	p-value
Spearman’s rank	0.339	0.325	0.352	0.337	1.92	0.027
Pearson	0.324	0.31	0.336	0.320	1.90	0.028
Kendall’s Tau	0.244	0.23	0.249	0.234	1.81	0.035

Table 7. This table presents the Spearman’s rank correlation coefficients between Human Selection Frequency and AVH, Model Confidence on VGG-19. Note that we show the absolute value of the coefficient which represents the strength of the correlation. Z value is computed by Z scores of both coefficients. p-value < 0.05 indicates that the result is statistically significant.

	Coef with AVH	Coef with Model Confidence	Z_{avh}	Z_{mc}	Z value	p-value
Spearman’s rank	0.349	0.335	0.364	0.348	1.94	0.026
Pearson	0.358	0.343	0.374	0.357	2.09	0.018
Kendall’s Tau	0.244	0.229	0.249	0.233	1.94	0.026

Table 8. This table presents the Spearman’s rank correlation coefficient between Human Selection Frequency and AVH, Model Confidence on ResNet-50. Note that we show the absolute value of the coefficient which represents the strength of the correlation. Z value is computed by Z scores of both coefficients. p-value < 0.05 indicates that the result is statistically significant.

	Coef with AVH	Coef with Model Confidence	Z_{avh}	Z_{mc}	Z value	p-value
Spearman’s rank	0.360	0.325	0.377	0.337	4.85	< .00001
Pearson	0.385	0.341	0.406	0.355	6.2	< .00001
Kendall’s Tau	0.257	0.231	0.263	0.235	3.38	.0003

Table 9. This table presents the Spearman’s rank correlation coefficients between Human Selection Frequency and AVH, Model Confidence in DenseNet-121. Note that we show the absolute value of the coefficient which represents the strength of the correlation. Z value is computed by Z scores of both coefficients. p-value < 0.05 indicates that the result is statistically significant.

	Coef with AVH	Coef with Model Confidence	Z_{avh}	Z_{mc}	Z value	p-value
Spearman’s	0.367	0.329	0.4059	0.355	6.2	< .00001
Pearson	0.390	0.347	0.412	0.362	6.09	< .00001
Kendall’s Tau	0.262	0.234	0.268	0.238	3.65	.0001

D.2. Additional Plots for Hypothesis Testings

Additional plots for Section 4: Figure 16 presents the correlation between Human Selection Frequency and AVH using AlexNet, VGG-19 and DenseNet-121.




Figure 16. The three plots present the correlation between Human Selection Frequency and AVH using AlexNet, VGG-19 and DenseNet121.

Correlation between AVH and image degradation: In order to test if the results in Figure 5 from the main paper also hold on proxies other than human visual hardness (image degradation level), we perform the similar experiments but on the augmented ImageNet validation set. Figure 17 shows the correlation between $\mathcal{AVH}(\mathbf{x})$ and different noise degradation levels, while the plots in Figure 18 shows the correlation between $\mathcal{AVH}(\mathbf{x})$ and different contrast degradation levels. Along with Figure 16, these results all indicate that $\mathcal{AVH}(\mathbf{x})$ is a reliable measure of Human Visual Hardness.




Figure 17. Correlation between noise degradation levels and AVH scores on AlexNet, VGG-19, ResNet-50 and DenseNet-121. Note that the larger the noise level is, the harder a human can recognize the image.




Figure 18. Correlation between contrast degradation levels and AVH scores on AlexNet, VGG-19, ResNet-50 and DenseNet-121. Note that the larger the contrast Level is, the easier a human can recognize the image.

Additional plots for Hypothesis 4: We further verify if presenting all samples across 1000 different classes affects the visualization of the correlation. According to WordNet (Fellbaum, 2005) hierarchy, we map the original 1000 fine-grained classes to 45 higher hierarchical classes. Figure 19 exhibits the relationship between Human Selection Frequency and $\|x\|_2$ for three representative higher classes containing 58, 7, 1 fine-grained classes respectively. Noted that there is still not any visible direct proportion between these two variables across all plots.




Figure 19. ℓ_2 norm of the embedding vs. Human Selection Frequency under different class granularities (according to WordNet hierarchy). From left to right, there are 58, 7, 1 classes respectively. Human Selection Frequency is therefore computed based on the new class granularity.

E. Additional discussions on the Difference between AVH and Model Confidence

The difference between AVH and Model Confidence lies in the feature norm and its role during training. To illustrate the difference, we consider a simple binary classification case where the softmax score (i.e., Model Confidence) for class 1 is

$$\frac{\exp(\mathbf{w}_1 \mathbf{x})}{\sum_i \exp(\mathbf{w}_i \mathbf{x})} = \frac{\exp(\|\mathbf{w}_1\| \|\mathbf{x}\| \cos(\theta_{\mathbf{w}_1, \mathbf{x}}))}{\sum_i \exp(\|\mathbf{w}_i\| \|\mathbf{x}\| \cos(\theta_{\mathbf{w}_i, \mathbf{x}}))}$$

where \mathbf{w}_i is the classifier weights of class i , \mathbf{x} is the input deep feature and $\theta_{\mathbf{w}_i, \mathbf{x}}$ is the angle between \mathbf{w}_i and \mathbf{x} . To simplify, we assume the norm of \mathbf{w}_1 and \mathbf{w}_2 are the same, and then the classification result is based on the angle now. Once $\theta_{\mathbf{w}_1, \mathbf{x}}$ is smaller than $\theta_{\mathbf{w}_2, \mathbf{x}}$, the network will classify the sample \mathbf{x} as class 1. However, in order to further minimize the cross-entropy loss after making $\theta_{\mathbf{w}_1, \mathbf{x}}$ smaller than $\theta_{\mathbf{w}_2, \mathbf{x}}$, the network has a trivial solution: increasing the feature norm $\|\mathbf{x}\|$ instead of further minimizing the $\theta_{\mathbf{w}_1, \mathbf{x}}$. It is obviously a much more difficult task to minimize $\theta_{\mathbf{w}_1, \mathbf{x}}$ rather than increasing $\|\mathbf{x}\|$. Therefore, the network will tend to increase the feature norm $\|\mathbf{x}\|$ to minimize the cross-entropy loss, which is equivalent to maximizing the Model Confidence in class 1. In fact, this also matches our empirical observation that the feature norm keeps increasing during training. Most importantly, one can notice that AVH will stay unchanged no matter how large the feature norm $\|\mathbf{x}\|$ is. Moreover, this also matches our empirical result that AVH easily gets saturated while Model Confidence can keep improving. Therefore, AVH is able to better characterize the visual hardness, since it is trivial for the network to increase feature norm. This is the fundamental difference between Model Confidence and AVH.

To get a more intuitive sense of how feature norm can affect the Model Confidence, we plot the value of the Model Confidence for two scenarios: $\theta_{\mathbf{w}_1, \mathbf{x}} < \theta_{\mathbf{w}_2, \mathbf{x}}$ and $\theta_{\mathbf{w}_1, \mathbf{x}} > \theta_{\mathbf{w}_2, \mathbf{x}}$. Under the case that the sample \mathbf{x} belongs to class 1, once we have $\theta_{\mathbf{w}_1, \mathbf{x}} < \theta_{\mathbf{w}_2, \mathbf{x}}$, then we only need to increase the feature norm and can easily get nearly perfect confidence on this sample. In contrast, AVH will stay unchanged during the entire process and therefore is a more robust indicator for visual hardness than Model Confidence.




Figure 20. The comparison between AVH and Model Confidence when the feature norm keeps increasing. The figure is plotted according to the binary classification example discussed above. We assume $\|\mathbf{w}_1\| = \|\mathbf{w}_2\|$. When $\theta_{\mathbf{w}_1, \mathbf{x}} < \theta_{\mathbf{w}_2, \mathbf{x}}$, we use $\theta_1 = \pi/4 - 0.05$ and $\theta_2 = \pi/4 + 0.05$. When $\theta_{\mathbf{w}_1, \mathbf{x}} > \theta_{\mathbf{w}_2, \mathbf{x}}$, we use $\theta_1 = \pi/4 + 0.05$ and $\theta_2 = \pi/4 - 0.05$. Note that, unlike Model Confidence, the smaller AVH is, the more confident the network is (i.e., the easier the sample is).

F. Experimental Details

Self-training and domain adaptation: As mentioned in section 5, a major challenge of self-training is the amplification of error due to misclassified pseudo-labels. Therefore, traditional self-training methods such as CBST often use Model Confidence as the measure to select confidently labeled examples. The hope is that higher confidence potentially implies lower error rate. While this generally proves useful, the model tends to focus on the “less informative” samples, whereas ignoring the “more informative”, harder ones near classier boundaries that could be essential for learning a better classifier. Figure 21 shows examples of what AVH selects and labels correctly but the softmax score does not select in CBST. We can see they are all visually confusing examples which can better help with the iterative self-training process when pseudo labeled correctly. The left one has the true label “Truck” but is easy to be confused with the “Car”. The right one has the true label “Person” but is easy to be confused with the “Motor”.



Figure 21. Two example images which AVH selects but softmax score do not. The left one has the true label “Truck” but is easy to be confused with the “Car”. The right one has the true label “Person” but is easy to be confused with the “Motor”.

Domain generalization: For domain generalization, we use the PACS benchmark dataset (Li et al., 2017) which contains consists of art painting, cartoon, photo and sketch domains. Each domain has the same 7 classes. Our experimental settings basically follow (Li et al., 2017). Specifically, we pick one domain as the unseen testing domain and train our model on the remaining three domains. The testing accuracy is evaluated on the unseen testing domain. Therefore, we will have 4 testing accuracies in total and we can use the average accuracy as the final evaluation metric. We use a convolutional neural network similar to (Liu et al., 2017c) with the detailed structure of $[7 \times 7, 64] \Rightarrow 2 \times 2 \text{ Max Pooling} \Rightarrow [3 \times 3, 64] \times 3 \Rightarrow 2 \times 2 \text{ Max Pooling} \Rightarrow [3 \times 3, 128] \times 3 \Rightarrow 2 \times 2 \text{ Max Pooling} \Rightarrow [3 \times 3, 256] \times 3 \Rightarrow 2 \times 2 \text{ Max Pooling} \Rightarrow 512\text{-dim Fully Connected}$. For example, $[3 \times 3, 64] \times 3$ denotes 3 cascaded convolution layers with 64 filters of size 3×3 . We use momentum SGD with momentum as 0.9 and batch size 40. Batch normalization and ReLU activation are used by default. Following the existing methods (Li et al., 2017; 2018; Balaji et al., 2018; Li et al., 2019), we will first pretrain our network on ImageNet with standard learning rate and decay, and then finetune on the PACS dataset with batch size 40 and smaller learning rate ($1e - 3$).

G. Extensions and Applications

Adversarial Example: A Counter Example?

Our claim about the stronger correlation between AVH score and human visual hardness does not apply on non-natural images such as adversarial examples. For such examples, the human can not tell the difference visually, but the adversarial example has a worse AVH than the original image, which runs counter to our claim that AVH has strong correlation with human visual hardness. So this claim is limited to distribution of natural images. However, on a positive note, we do find that AVH is slower to change compared to the embedding norm during the dynamics of adversarial training.

We show a special case in Figure 22 to illustrate how the norm and the angle change when one sample switches from one class to another. Specifically, we change the sample from one class to another using adversarial perturbation. It is essentially performing gradient ascent to the ground truth class. In Figure 22, the purple line denotes the trajectory of an adversarial sample switching from one class to another. We can see that the sample will first shrink its norm towards origin and then push its angle away from the ground truth class. Such a trajectory indicates that the adversarial sample will first approach to the origin in order to become a hard sample for this class. Then the sample will change the angle in order to switch its label. This special example fully justifies the importance of both norm and angle in terms of the hardness of samples.

Measuring Human Visual Hardness is Hard: Measuring Human Visual Hardness is non-trivial and dependent on many factors such as (i) How much are the annotators penalized for wrong answers and how much time are they given? (ii) What are the cultural and language differences that can cause annotators to be confused about the label categories. Figure 23 shows an example of groom from ImageNet dataset. Since a large contingent of Mturk users are from India, they have high confidence for this image, but the answer would be very different if asked a different population. The proxies we used in this paper, Human Selection Frequency and Image Degradation Level are best efforts.

Connection to deep metric learning: Measuring the hardness of samples is also of great importance in the field of deep metric learning (Oh Song et al., 2016; Sohn, 2016; Wu et al., 2017). For instance, objective functions in deep metric learning consist of *e.g.*, triplet loss (Schroff et al., 2015) or contrastive loss (Hadsell et al., 2006), which requires data pair/triplet mining in order to perform well in practice. One of the most widely used data sampling strategies is semi-hard negative sample mining (Schroff et al., 2015) and hard negative sample mining. These negative sample mining techniques highly depend on how one defines the hardness of samples. AVH can be potentially useful in this setting.

Connections to fairness in machine learning: Easy and hard samples can implicitly reflect imbalances in latent attributes in the dataset. For example, the CASIA-WebFace dataset (Yi et al., 2014) mostly contains white celebrities, so the neural network trained on CASIA-WebFace is highly biased against the other races. (Buolamwini & Gebru, 2018) demonstrates a performance drop of faces of darker people due to the biases in the training dataset. In order to ensure fairness and remove dataset biases, the ability to identify hard samples automatically can be very useful. We would like to test if AVH is effective in these settings.

Connections to knowledge transfer and curriculum learning: The efficiency of knowledge transfer (Hinton et al., 2015) is partially determined by the sequence of input training data. (Liu et al., 2017a) theoretically shows feeding easy samples first and hard samples later (known as curriculum learning) can improve the convergence of model. (Bengio et al., 2009) also show that the curriculum of feeding training samples matters in terms of both accuracy and convergence. We plan to investigate the use of AVH metric in such settings.




Figure 22. Trajectory of an adversarial example switching from one class to another. The purple line denotes the trajectory of the adversarial example.



Figure 23. An image of an Indian Groom from ImageNet.