

When Deep Classifiers Agree: Analyzing Correlations between Learning Order and Image Statistics

Iuliia Pliushch¹, Martin Mundt², Nicolas Lupp¹, and Visvanathan Ramesh¹

¹ Goethe University Frankfurt, Germany
{plushch, vramesh}@em.uni-frankfurt.de
² TU Darmstadt and hessian.AI, Germany
martin.mundt@tu-darmstadt.de

Abstract. Although a plethora of architectural variants for deep classification has been introduced over time, recent works have found empirical evidence towards similarities in their training process. It has been hypothesized that neural networks converge not only to similar representations, but also exhibit a notion of empirical agreement on which data instances are learned first. Following in the latter works’ footsteps, we define a metric to quantify the relationship between such classification agreement over time, and posit that the agreement phenomenon can be mapped to core statistics of the investigated dataset. We empirically corroborate this hypothesis across the CIFAR10, Pascal, ImageNet and KTH-TIPS2 datasets. Our findings indicate that agreement seems to be independent of specific architectures, training hyper-parameters or labels, albeit follows an ordering according to image statistics.

Keywords: neural network learning dynamics, deep classifier agreement, data instance ordering, image dataset statistics

1 Introduction

Can we make the learning process of a neural network more transparent? Is there any order in which a network is learning dataset instances? Various prior works posit that certain data instances are easier to learn than others. Arpit *et al.*[3] argue that neural networks prioritize simple patterns first during the learning process by analyzing the nature of the decision boundary. Mangalam and Prabhu [35] extend the results, showing that neural networks first learn instances, classified correctly by shallow models like Random Forests and Support Vector Machines. Geirhos *et al.*[14] and Shah *et al.*[41] specify which patterns are simple for neural networks to learn, namely that they exhibit a *simplicity bias* during training by relying on simple-but-noisy features like color and texture, instead of shape. Recently, Hacoen *et al.*[15] have shown that there is agreement over learned examples throughout the entire learning process of a neural network, which is independent of initialization and batch-sampling and occurs also when changing different hyperparameter settings, like learning rate, optimizer,

weight-decay and architecture. If certain examples are easier for a neural network to learn than others, it can be hypothesized that the difficulty of an example depends on dataset and image statistics to a larger degree than the learner itself. Hence, different neural networks would learn the same (easy) examples at the same time, despite common epoch-wise random shuffling of the data during the training process and strong regularization through noise of mini-batch stochastic gradient descent. In this spirit, we extend the analysis to a stricter agreement metric and not only replicate the presence of agreement, but also correlate it to chosen image statistics, in an attempt to clarify the reason why certain image examples are more difficult to learn than others. Our contributions are:

- We design a strict instance-based agreement metric to quantify how learning progresses and data instances are classified correctly across neural networks.
- With our metric, we replicate insights on classifier agreement from Hacohen *et al.*[15] in terms of consistent neural network agreement across batch-sizes and architectures, as well as investigate the role of labels through a label randomization study.
- We select promising image statistics to correlate with neural network data instance agreement. We empirically corroborate correlations for metrics such as entropy, segment count, number of relevant frequency coefficients and summed edge strengths on popular image classification datasets: CIFAR10, Pascal, KTH-TIPS2b and ImageNet.

The analysis we conduct might help to better understand the tools - neural networks - we use for training. Further, as the next step, these insights could drive us towards the design of a *learning curriculum* [29,27,4,16], in which the dataset instances could be sampled according to a high correlation between observed agreement and the inspected metrics. We posit that this could lead to training speed ups and performance improvements. Our code is available at: https://github.com/ccf-frankfurt/intrinsic_ordering_nn_training.

2 Problem statement and motivation

Let us first recall the general procedure of neural network training and the rationale behind it. Assume that we have a training set $X = (x_n, y_n)_{n=1}^N$, where x_n are our (i.i.d.) dataset instances, y_n the corresponding labels and the number of dataset instances N . We also assume access to a similarly designed non-overlapping test set. We want to optimize a defined loss to measure and minimize the discrepancy between our network prediction and the ground truth. For that, we ideally want to integrate the loss L of our neural network - function f_θ with parameters θ , over the dataset distribution:

$$\int L(f_\theta(x), y) dP(x, y) \tag{1}$$

In practice, we only have a limited amount of samples from the dataset distribution. Hence we compute an *approximation*. Typically a noisy gradient estimate

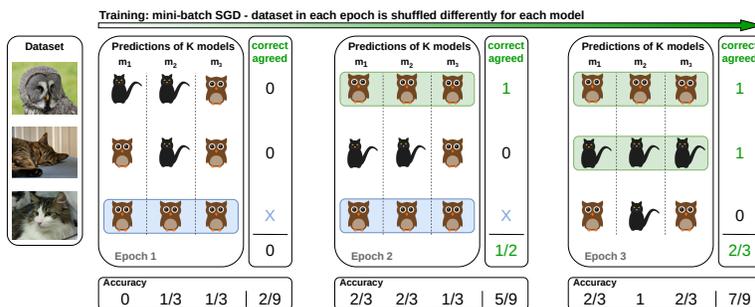


Fig. 1: **Agreement visualization:** For each image the classification results are compared across networks and, in addition to *average accuracy over networks*, *true positive agreement* calculated, which is the ratio of images that all networks classify correctly per epoch to those that at least one network classifies correctly.

(see appendix) is leveraged for such empirical optimization, by presenting our data in mini-batches over several epochs t and shuffling the dataset after every epoch, i.e. when the network has seen all the data (at least) once. Of course, when the networks have fully converged and learnt a sufficiently large amount of data, the dataset instances they have learnt trivially overlap. It is however not self-evident that different approximators would learn the data in a similar way or in other words that despite shuffling and mini-batch updating, neural networks would learn data in the same order. We first provide a definition for such agreement and then consider when networks necessarily start to agree on the dataset instances during learning. Hacohen *et al.*[15] define one form of *agreement* as “the largest fraction of classifiers that predict the same label” for the same data instance, as well as *true positive agreement* as an “average accuracy of a single example over multiple models”. In order to take the next step and try to quantify the difficulty of the images for training, we step back from agreement as an average and define (*true positive*) agreement of an instance per epoch as an *exact match*, such that all K networks classify the same instance correctly in epoch t , where K is the number of networks we have trained. *True positive agreement per epoch* can then be computed as the sum of instances classified correctly by *all* classifiers in that epoch, normalized by the sum of instances classified correctly by *any* classifier in that epoch. Formally, true positive agreement TPa per epoch t is defined as:

$$TPa^{(t)}(x, y) = \frac{\sum_{n \in N} \prod_{k \in K} \mathbb{1}_{f_k^{(t)}(x_n) = y_n}}{\sum_{n \in N} \max_{k \in K} \mathbb{1}_{f_k^{(t)}(x_n) = y_n}} \quad (2)$$

During training, we now monitor the true positive agreement in every epoch for each training instance. Suppose that we train K networks, as in fig. 1. In the first epoch, some models classify some dataset instances correctly (indicator function, $\mathbb{1}_{f_k^{(t)}(x_n) = y_n}$ being the condition for a prediction match), but for no instance it is the case that all models classify it correctly. In the second epoch -

one dataset instance is classified correctly by all models. As during training more models classify instances correctly, if they learn the same instances first, then as soon as all models agree, it will be reflected in the agreement scores. Perfect agreement of 100% can be reached if all models learn the same instances. So, agreement - fraction of correctly classified instances over all models compared to learned at all by at least one model - can be higher than average model accuracy per epoch, which is the fraction of correctly classified instances from the whole train set, averaged over the trained models. False positive agreement (all models misclassify an instance in the same way), as in the case of the partly white cat in the first two epochs (blue shaded box in fig. 1), is left out of the true positive agreement. However, it could be analyzed in future work.

To assess, how trivial it is that neural networks agree *before* convergence we consider the minimum possible agreement - the *lower bound*. When do the models necessarily start to agree to learn the same instances? Let us assume that we train 2 networks and every one is 50% correct at some epoch t . Then every network can learn the fraction of data that the other one did not. The same applies if we train 3 networks and the accuracy is $2/3$ per network, because if we split the data into 3 portions, every network can learn 2 portions of the data such that there is not one portion common to all 3 networks. The same occurs for K networks with accuracy being $\frac{K-1}{K}$. In all these cases the sum of errors err_k^e all networks make is $K * \frac{1}{K} = 1$ or 100%. As the accuracy rises higher than that and the error gets lower, the networks will necessarily start to learn the fractions of data others are learning. Hence, the lower bound is 0% when the sum of errors the networks make is bigger or equal to 100%. The lower bound fraction, reported in the remainder of the paper in percent, can be defined as:

$$LBA^t(x, y) = 1 - \min\left(\sum_{k \in K} \underbrace{1 - acc_k^e}_{err_k^e}, 1\right). \quad (3)$$

In particular, the difference between agreement and lower bound shows us the portion of agreement which could not have been predicted on the basis of the lower bound alone.

Lastly, since one of the datasets we test our hypotheses on is *multilabel*, in this scenario we extend the definition of agreement and lower bound such that agreement and accuracy are calculated on the basis of *exact match*, meaning that for all present and absent labels in an image, the prediction should exactly match: the presence and absence of a label should be predicted correctly for all labels. This is a criterion which is non-forgiving: if one of the two labels has been predicted correctly, the exact match still classifies this image as wrongly predicted. On the other hand, it is a strong criterion to test agreement on, since what we are after is a *full* and not *partial* agreement. It is also a much stronger criterion than investigated in Hacohen *et al.*[15]. Notably, our TP-agreement is thus also different from an *observed agreement*, i.e. the sum over instances given estimators classify correctly (true positives) and incorrectly (true negatives) divided by the total number of instances [18]. Our choice to separate out true positive agreement is intended to avoid potential confusion, as the number of

true positives and negatives is not equal during training and can complicate the assessment of reliability across estimators. The alternative Cohen’s kappa, for example, suffers from the prevalence problem: it is non-representatively low in case of uneven frequencies of events. An alternative PABAK measure counteracting the prevalence bias [6] is a linear function of observed agreement, which takes both true positives and negatives into account. We provide a deeper discussion on alternatives in the appendix, in context of the preliminary results shown in the upcoming section.

3 Agreement for different batch-sizes and architectures, as well as for random labels

To reiterate, the general procedure is to train several (in our case 5) networks on the same dataset and track examples, classified correctly per epoch per network. Upon training several networks for the same amount of epochs, *agreement per epoch* is defined as the sum of images *every* network classified correctly, normalized by the sum of images *any* network classified correctly. We have seen on the example of the *lower bound* that a quite high accuracy is necessary for the networks to unavoidably learn the same instances and that this relation depends on the number of networks trained: the more networks, the higher accuracy each of those has to obtain to start learning the same data portion. Hence, it is not self-evident that agreement does occur during neural network training.

If networks agree on what to learn, which factors does this agreement depend on? To provide an intuition for above paragraphs, as well as to replicate prior insights obtained by Hacohen *et al.*[15], we first test for the overall presence of agreement. Then, we extend our experiments to different batch-sizes and architectures. Presence of agreement in all these diverse conditions supports the hypothesis that agreement is not dependent on the *model*, but on the (sampled) dataset population itself in terms of the nature of the true unknown underlying distribution and the intrinsic characteristic of the difficulty of classification [7]. To further investigate whether the *joint* dataset and label distribution leads to agreement, or whether it is independent of labels, we test the presence of agreement in case of random labels.

As hypothesized, mirroring prior hypothesis of Hacohen *et al.*[15], but using our strict agreement criterion, we found visually clear agreement during training of CIFAR10 [28] on DenseNet121 [22], presented in fig. 2a. Particularly in the first epochs, as the accuracy grows, the area between agreement and the lower bound, shaded in blue, is the most prominent. It shows that throughout training we observe growing agreement on the learned instances, i.e. certain data instances are labelled correctly in earlier stages than others, which is most remarkable in the training epochs before networks converge and the accuracy plateaus.

Agreement persists also when training different architectures. We have chosen 4 diverse architectures, ranging from simple ones like LeNet5 [30] and VGG16 [43] to the more complex ones like ResNet50 [19] and DenseNet121 [22]. Training CIFAR10 on them for the same amount of epochs, we could observe visually

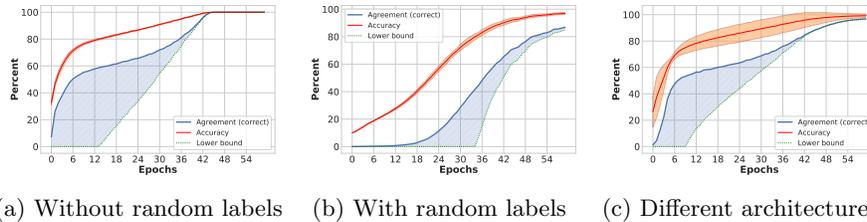


Fig. 2: Ablation study on **CIFAR10**: training DenseNet121 with and without randomization, as well as with different architectures: LeNet5, VGG16, ResNet50, DenseNet121. The **blue area** demonstrates the difference between agreement and lower bound. The **red area** is the epoch-wise standard deviation from average accuracy across trained networks.

prominent agreement as well, see fig. 2c (training details and additional plots in the appendix). The agreement is similar to that for the same architecture. The differences in learning speed are reflected in the standard deviation across different model accuracies, shaded in red. Note that accuracy deviation has been almost negligible when the same architectures have been trained.

We have also replicated the presence of agreement for different batch-sizes in the appendix. A comparison shows that agreement curves are similar for simpler architectures and smaller batch-sizes, as well as more complex architectures and larger batch-sizes, suggesting that model capacity and enhanced randomness when training with smaller batch sizes slow down the learning and agreement.

Above experiments suggest that batch-size or architecture type are not the underlying causes for agreement during training. Next, we test whether it is the structure of the data itself, or the relationship between the training data and its human assigned labels that account for it. To test this hypothesis, we assess agreement under label randomization. State-of-the-art convolutional networks can fit random labels with ease [49]. Maennel *et al.*[34] argue further that during training with random labels an alignment between the principal components of data and network parameters takes place. Hence, if the dataset structure is responsible for agreement, it should be visible also in case of random labels. fig. 2b supports this hypothesis. We observe that the accuracy grows at a slower pace and that agreement grows more slowly than during training with ground truth labels. This may reflect the fact that it takes longer for the network to start learning the dataset structure without label guidance. Nonetheless, there is still sufficient agreement once accuracy starts to rise. Interestingly, here we disagree with [15], who did not find agreement with randomized labels.

To conclude, we have observed a clear gap between theoretical lower-bound and observed agreement, independence on semantic labels and architecture, even coherence in between them. This supports the existence of a fundamental core mechanism linked to more elemental dataset properties. To further support and strengthen our early results, we show in the appendix that observed trends of

true positive agreement clearly surpassing the lower-bound persist even when comparing it to the *expected random agreement*. The latter is computed as the product of accuracies for a given epoch. It is based on the assumption that networks classify instances independently of each other, which does not seem to be the case, partly because, as we show, the dataset structure plays an important role. In addition, we also compute the standard deviation on agreement for Pascal in the appendix, showing that it is negligibly small.

4 Dataset metrics

In the face of the insight that dataset properties are the probable cause of agreement, we proceed by choosing several diverse datasets, as well as dataset metrics to establish a correlation between them and training agreement. We have chosen the following datasets to validate our agreement hypothesis: a tiny-sized CIFAR10 [28] for ablation experiments, a diverse dataset with objects differing in illumination, size and scale - Pascal Visual Object Classes (VOC) 2007 and 2012 [11,10], a large-scale ILSVRC-2012 (ImageNet) [9,40], as well as a texture dataset KTH-TIPS2b [8]. CIFAR10, Pascal and ImageNet have been gathered by means of search engines and then manual clean-up. In case of ImageNet, classes to search for were obtained from the hierarchical structure of WordNet and the manual clean-up proceeded by means of the crowdsourcing platform Amazon Mechanical Turk. On the example of the person category, Yang *et al.*[48] elaborate that a dataset gathered in such a manner is as strong as the semantic label assumptions and distinctions on which it is based, the quality of the images obtained using search engines (e.g. lack of image diversity), as well as the quality of the clean-up (annotation) procedure. In distinction to the above three datasets, though KTH-TIPS2b is a quite small dataset, it has been carefully designed, controlling for several illumination and rotation conditions, and varying scales.

Prior work: What surrogate image statistics may be used to study the agreement between classifiers on the order of learning? Let us look at some prior works in this direction. First, some dataset properties make learning difficult in general, namely the diverse nature of the data itself: the difficulty of assigning image categories due to possible image variation, e.g. rotation, lighting, occlusion, deformation [38]. Second, learning algorithms also show preferences for particular kinds of data. Russakovsky *et al.*[39] analyze the impact onto ImageNet classification and localization performance of several dataset properties on the image and instance level, like object size, whether the object is human- or man-made, whether it is textured or deformable. Their insights are consistent with the intuition that object classification algorithms rely more on texture and color as cues than shape. They further argue that object classification accuracy is higher for natural than man-made objects. Hoiem *et al.*[21] analyze the impact of object characteristics on errors in several non-neural network object detectors, coming to the conclusion that the latter are sensitive to object size and confusion with semantically similar objects.

In an attempt to mimic the decisions of a human learner, several image features have been related to image memorability and object importance. Isola *et al.* [25] investigate memorability of an image as a stable property across viewers and its relation to basic image and object statistics, like mean hue and number of objects in the image. Spain and Perona [45], on the other hand, establish the connection between image object features and the importance of this object in the image, which is the probability that a human observer will name it upon seeing the image. Berg *et al.* [5] extends this analysis to encompass semantic features like object categories and scene context.

Knowing which image cues the learning process is sensitive to gives room for improving it. Alexe *et al.* [2] design an objectness measure generic over classes from image features (multi-scale saliency, color contrast, edge density and superpixels straddling) which can be used as a location prior for object detection. Extending the latter work, Lee and Grauman [31] use this objectness measure, as well as an additional familiarity metric (whether it belongs to a familiar category) to design a learning procedure which first considers easy objects. Liu *et al.* [33] fit a linear regression model with several image features, like color, gradient and texture, to estimate the difficulty of segmenting an image. In [46] an active learner is designed, which partly on the basis of edge density and color histogram metrics, proposes which instances to annotate and estimates the annotation cost for the multi-label learning task. Notably the question of whether the networks learn the same examples first is different from the question whether they learn the same representations [47,32], but if the former is correct, then the latter is more probable.

Our choice: Several works have correlated basic image statistics to various human-related concepts, like memorability [25], importance [45] or image difficulty [24], or directly attempted to find out the influence of such metrics onto object classification [39,2]. Inspired by these approaches, we have chosen 4 image statistics to correlate agreement to: segment count [12], (sum of) edge strengths [26], (mean) image intensity entropy [13,44] and percentage of coefficients needed to reconstruct the image based on the DCT coefficient matrix [1]. First 3 metrics are shown in fig. 3 (DCT coefficients matrix in the appendix).

The choice of the first two is inspired by [24], who correlate several image properties, including segment count and (sum of) edge strengths, to the *image difficulty score*, defined as a normalized response time needed for human annotations to detect the objects in the image. Their hypothesis is that segments divide the image into homogeneous textural regions, such that the more regions - the more cluttered an image might be (and the more difficult an object to find). The same line of reasoning goes for the sum of edge strengths: the more edges, the more time might be needed to get a grasp of the image. If the way humans search for objects in the image corresponds to some degree to the way a neural networks learns to predict their presence, segment count and the (sum of) edge strengths might be predictive for the agreement during training. Mean image entropy and DCT coefficients have been chosen for a similar line of reasoning. Both entropy and DCT coefficients, similar to edge strengths and segment count, provide a

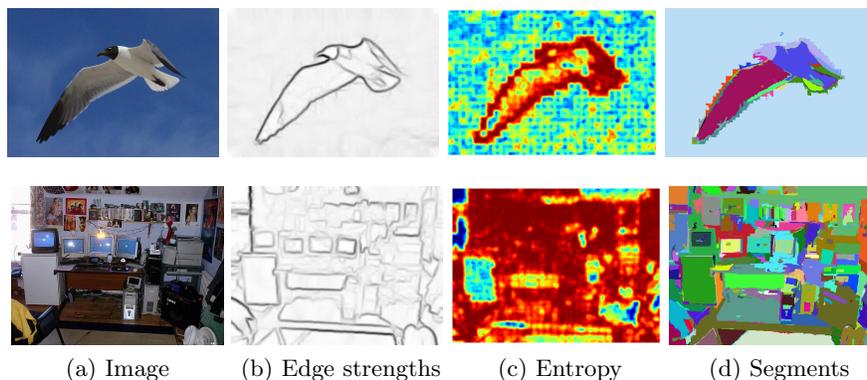


Fig. 3: Visualization of selected metrics on Pascal examples. For intuition, we have chosen an "easy" and a "difficult" image according to Ionescu *et al.*[24]. Implementation details can be found in the appendix.

measure of variability: how uniform the variation in image pixels is in the case of entropy and how many (vertical and horizontal) frequencies are needed to describe an image in the case of DCT coefficients. Recently, [36] have analyzed how certain dataset features relate to classification, stating, for instance, that for MNIST and ImageNet the decision boundary is small for low frequency and high for high frequency components. In other words, the classifier develops a strong invariance along high frequencies. Hence, frequency-related information might influence agreement.

In addition, for Pascal, which contains further annotations, we considered the *number of object instances* (since it is a *multilabel* dataset such that several label instances may be present in the same image) and the *bounding box area* (ratio of the area taken by objects divided by the image size), as well as the *image difficulty scores* described above, computed by [24]. KTH-TIPS2b, is constructed in such a way that for each texture type, there are 4 texture samples, each of which varies in *illumination*, *scale* and *rotation*, which we also considered. Finally, for the CIFAR10 test set, Peterson *et al.*[37] have computed soft labels, reflecting human uncertainty that the given target class is in the image, to test whether networks trained on soft labels are more robust to adversarial attacks and generalize better than those trained on hard one-hot labels. *Soft label entropy* shows weak negative correlation with test agreement (see Fig. 5 of the appendix).

5 Do basic image statistics correlate with training learning dynamics?

Our initial investigation of section 3 suggests that neither the labels, nor the precisely chosen neural architecture or optimization hyper-parameters seem to be the primary source for agreement, eliminating all factors of eq. (1) other

than the data distribution itself. As such, we investigate the question “*do image statistics provide a sufficient description in correlation for agreement?*” on four datasets: Pascal [11], CIFAR10 [28], KTH-TIPS2b [8], and ImageNet [9].

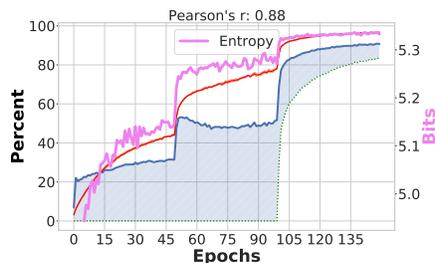


Fig. 4: **Pascal DenseNet**: Agreement visualization on *train set*. In addition to accuracy (**red**), agreement (**blue curve**) and its difference to lower-bound (shaded **blue area**) (left y-axis), dataset metric values are shown in **purple** (right y-axis). *Pearson’s correlation coefficient* between agreement (**blue curve**) and metric values (in **purple**) quantitatively supports the visible correlation.

To make the full upcoming results easier to follow, we first start by visualizing and discussing one example metric, namely entropy on Pascal in fig. 4. In a given epoch, the agreement (blue) and accuracy (red) curves can be compared to the average entropy of the agreed upon instances (in purple). Again, the shaded blue area accentuates the difference between lower bound and agreement. Since we have chosen a step-wise learning rate scheduler for training, in order to reach roughly 50% exact match accuracy on the test set, the step-wise learning is reflected in the accuracy and agreement curve. As a new addition to our previously shown figures, we now notably also observe a strong positive correlation with the dataset entropy metric. This correlation between network agreement and the entropy of the correspondingly agreed upon instances is further quantified through a high Pearson correlation coefficient of 0.88.

We continue our analysis in this form for all other mentioned metrics and dataset combinations in fig. 6 and fig. 7. To better evaluate the shown correlations, we have visualized the distribution (in the form of a histogram) of the values for each dataset metric on the train sets in the appendix. Since the metrics fluctuate a lot in the first epochs due to predictions being primarily random, we omit plotting them for the first 5 epochs. For reference, we provide the agreement, accuracy and lower-bound curves in the style of our previous figures in fig. 5. To clarify potential correlations, these curves are then followed by visualizations and quantitative Pearson’s r values of only agreement to the set of chosen image metrics in fig. 6 and fig. 7.

For the **Pascal** dataset, shown in fig. 6a, the correlations between agreement and average dataset metrics are apparent (apart from sum of edge strengths).

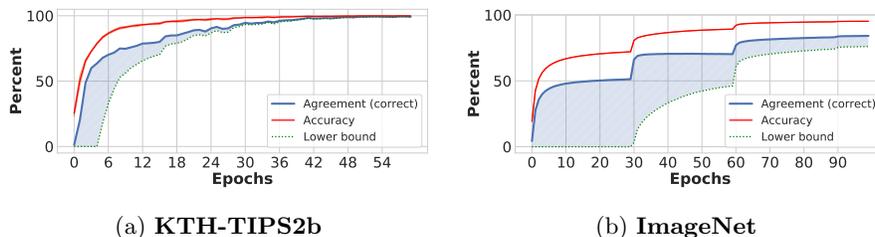


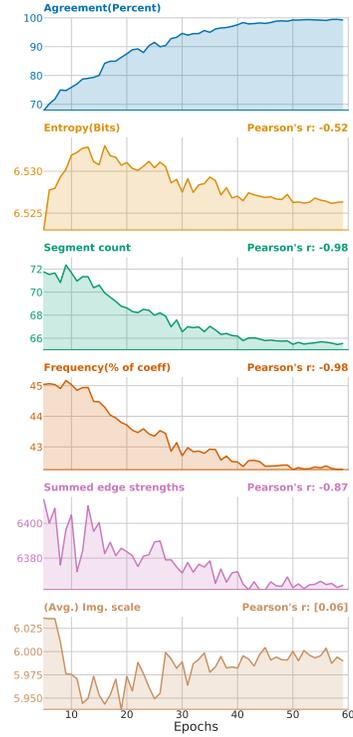
Fig. 5: **On DenseNet**: Agreement, accuracy and lower bound as in fig. 4

The correlations suggest that as the models learn, they in progression first learn dataset instances with lower entropy, segment count, number of significant DCT coefficients and number of object instances, meaning that less labels are present in the same image. The distribution for the number of instances (in the appendix) shows that the number of labels in most images is only 1 (equivalent to the single label classification), less for 2 and 3. This is reflected in the correlation, which goes up from 1.1 to 1.4. There is also a correlation between image difficulty (how much time humans need to find the objects in the image) and agreement, which is consistent with the other metrics, namely that easy examples are learnt first. This result is further supported by the bounding box size, where we see an inverse correlation, such that large objects are learnt first, in agreement with insights presented in the dataset metrics section. However, note that particularly for segment count, frequency coefficients and "image difficulty" the dataset metric curve first goes down, before it reverses its direction for the remainder of epochs. What we observe is that when accuracy is low and the model is trying to find an optimal trajectory to learn, there are more random fluctuations, due to the stochasticity of the learning process. In the appendix we show the correlations for Pascal ResNet with the same training setup. There, the trend of metric values first going down in first epochs (until the agreement approximatively reaches 20%) and then up is even more pronounced for some metrics.

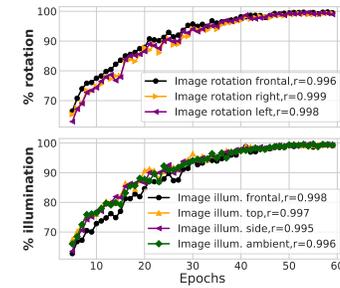
Correlations are also present on **KTH-TIPS2b**, visualized in fig. 6b. However, for this texture dataset, the tendency observed in Pascal is reversed, such that entropy, summed edge strengths, segment count and frequency percentage are inversely correlated with the dataset metric. In addition to these metrics, the way in which the dataset has been designed allows to extract additional ones, namely several illuminations, rotations and scales. The corresponding correlations are visualized in fig. 6c. For illumination and rotation, instead of building an average over the metric values per epoch, we calculate for each illumination kind and rotation direction the fraction of values agreed upon, normalized by all metric values of that type. For frontal illumination, for example, we count the instances that models agree on and divide by the number of instances of that type in the train set. We observe that texture patterns *illuminated* from the front are agreed on slower than other illumination types, while texture patterns



(a) **Pascal DenseNet:** Correlation between agreement (in blue) to per epoch averaged dataset metric values for correctly agreed upon training instances, on a shared x-axis (in analogy to the purple curve in fig. 4). Pearson correlation coefficient between agreement and the metric is provided. It is in rectangular brackets if the 2-tailed p-value is ≥ 0.001 , see appendix for details.



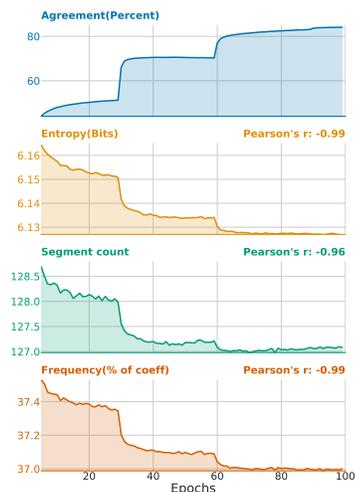
(b) **KTH-TIPS2b DenseNet:** Agreement and metric values in analogy to fig. 6a



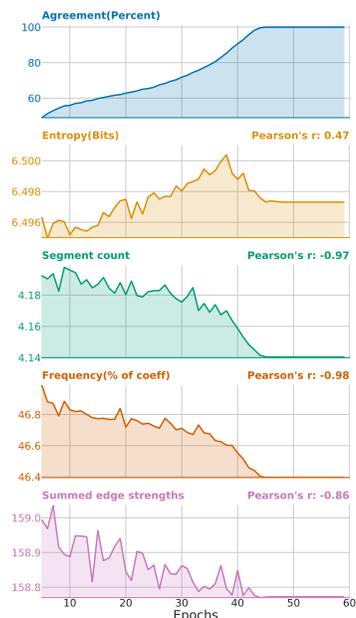
(c) **KTH-TIPS2b DenseNet:** For illumination and rotation types, the percentage learned per category per epoch has been visualized separately for immediate comparison.

Fig. 6: Visualization of correlations between agreement and dataset metrics on *train sets* for **Pascal** and **KTH-TIPS2b**.

captured from the front are agreed on quicker than other rotation directions. Correlation for texture scale seems absent.



(a) **ImageNet DenseNet:** In analogy to fig. 6a.



(b) **CIFAR10 DenseNet:** In analogy to fig. 6a.

Fig. 7: Visualization of correlations between agreement and dataset metrics on *train sets* for **ImageNet** and **CIFAR10**.

We have seen that correlations on Pascal and KTH-TIPS2b diverge. In fig. 7a we also see that on **ImageNet** correlations, albeit on a small scale, are present and they seem to be congruent with those on KTH-TIPS2b rather than Pascal. Namely, first images with higher segment count, entropy and number of relevant frequency coefficients are learned. Pascal and ImageNet are datasets of objects, while KTH-TIPS2b of texture. Why the difference? If we follow recent hypotheses on neural networks exhibiting simplicity bias and primarily using texture to discriminate [41,14], then correlation directions on KTH-TIPS2b and ImageNet would be the same (which our results indicate). On both Pascal and ImageNet random-crop is used to get train images of the same size, which has been argued by [20] to enhance texture bias. The way objects are presented in Pascal and ImageNet differs however, for ImageNet objects are centered and not much background is present, which is not the case for Pascal, making latter classification more challenging, but also effects of random-crop eventually different.

On **CIFAR10**, the direction of correlations is more consistent with KTH-TIPS2b and ImageNet (see fig. 7b). The correlations are marginal, however, given that metric value range is very small. The presence of agreement with little correlation can also be an indication that there are some other metrics which explain it, though it may be some humanly non-interpretable noise patterns [23].

6 Discussion on limitations and prospects

Although agreement is present for *every* dataset, the previous sections have exposed a dataset dependency with respect to how precisely correlations between agreement and dataset metrics manifest. This means that the order, in which dataset instances are learned, is to a large degree independent of network parameters, but instead dependent on the general dataset statistics. This seems congruent with the ideas that a classifier compactly encodes the data in the first layers [42], as well as that agreement results from the principal component bias [17] in that learning order is influenced by the principal components of the data at earlier stages. Hence, in order to better understand the learning process of neural networks (and make a leap from correlation to causation, see appendix), agreement should be computed for carefully designed datasets and suitable metrics, for which the generative process and dataset statistics are well known.

Another research direction is to assess generalization and analyze the agreement on unknown *test sets*. Though these are often assumed to follow the same distribution as their training counterparts, it is practically not always the case. Supposedly, if the subset statistics are sufficiently similar, then agreement and metrics correlations should similarly manifest, as we exemplify in the appendix.

Last but not least, as mentioned in the introduction, the insights on agreement could help us design a more efficient *learning curriculum*, for which an appropriate pacing function [16] should be chosen with care.

7 Conclusions and outlook

In this paper we have defined a new notion of agreement, characterising the learning process of neural networks in a more detailed way. We have demonstrated agreement on the train (and test set) for CIFAR10, Pascal, ImageNet and KTH-TIPS dataset. We have further correlated agreement on these datasets to several image statistics, in an attempt to explain why neural networks prefer to learn dataset instances in the way they do. Our results have shown several positive and negative correlations to dataset metrics, though different for each dataset. For future research is left the opportunity to test the results on further datasets, to test the correlation for further metrics, as well as to design curricula for training neural networks based on these insights.

Acknowledgements. This work was supported by the German Federal Ministry of Education and Research (BMBF) funded project 01IS19062 "AISEL" and the European Union's Horizon 2020 project No. 769066 "RESIST".



References

1. Ahmed, N., Natarajan, T., Rao, K.R.: Discrete Cosine Transform. *IEEE Transactions on Computers* **C-24**(1), 90–93 (1974)
2. Alexe, B., Deselaers, T., Ferrari, V.: Measuring the objectness of image windows. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **34**(11), 2189–2202 (2012)
3. Arpit, D., Jastrzebski, S., Bailas, N., Krueger, D., Bengio, E., Kanwal, M.S., Maharaaj, T., Fischer, A., Courville, A., Bengio, Y., Lacoste-Julien, S.: A closer look at memorization in deep networks. *International Conference on Machine Learning (ICML)* (2017)
4. Bengio, Y., Louradour, J., Collobert, R., Weston, J.: Curriculum learning. *International Conference on Machine Learning (ICML)* (2009)
5. Berg, A.C., Berg, T.L., Daume, H., Dodge, J., Goyal, A., Han, X., Mensch, A., Mitchell, M., Sood, A., Stratos, K., Yamaguchi, K.: Understanding and predicting importance in images. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2012)
6. Byrt, T., Bishop, J., Carlin, J.B.: Bias, prevalence and kappa. *Journal of Clinical Epidemiology* **46**(5), 423–429 (1993)
7. Caelen, O.: A Bayesian interpretation of the confusion matrix. *Annals of Mathematics and Artificial Intelligence* **81**(3-4), 429–450 (2017)
8. Caputo, B., Hayman, E., Mallikarjuna, P.: Class-specific material categorisation. In: *International Conference on Computer Vision (ICCV)* (2005)
9. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009)
10. Everingham, M., Eslami, S.M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The Pascal Visual Object Classes Challenge: A Retrospective. *International Journal of Computer Vision* **111**(1), 98–136 (2015)
11. Everingham, M., Van-Gool, L., Williams, C.K.L., Winn, J., Zisserman, A.: The Pascal Visual Object Classes (VOC) Challenge. *International Journal of Computer Vision* **88**(2), 303–338 (2010)
12. Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient Graph-Based Image Segmentation. *International Journal of Computer Vision* **59**, 167–181 (2004)
13. Frieden, B.R.: Restoring with maximum likelihood and maximum entropy. *Journal of the Optical Society of America (JOSA)* **62**(4), 511 – 518 (1972)
14. Geirhos, R., Michaelis, C., Wichmann, F.A., Rubisch, P., Bethge, M., Brendel, W.: Imagenet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness. *International Conference on Learning Representations (ICLR)* (2019)
15. Hacoheh, G., Choshen, L., Weinshall, D.: Let’s Agree to Agree: Neural Networks Share Classification Order on Real Datasets. *International Conference on Learning Representations (ICLR)* (2020)
16. Hacoheh, G., Weinshall, D.: On the power of curriculum learning in training deep networks. *International Conference on Machine Learning (ICML)* (2019)
17. Hacoheh, G., Weinshall, D.: Principal Components Bias in Over-parameterized Linear Models, and its Manifestation in Deep Neural Networks. *Journal of Machine Learning Research* **23**, 1–46 (2022)
18. Hallgren, K.A.: Computing Inter-Rater Reliability for Observational Data: An Overview and Tutorial. *Tutor Quant Methods Psychol.* **8**(1), 23–34 (2012)

19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 770–778 (2016)
20. Hermann, K.L., Chen, T., Kornblith, S.: The Origins and Prevalence of Texture Bias in Convolutional Neural Networks. *Neural Information Processing Systems (NeurIPS)* **34** (2020)
21. Hoiem, D., Chodpathumwan, Y., Dai, Q.: Diagnosing error in object detectors. *European Conference on Computer Vision (ECCV)* (2012)
22. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 2261–2269 (2017)
23. Ilyas, A., Santurkar, S., Tsipras, D., Engstrom, L., Tran, B., Madry, A.: Adversarial Examples Are Not Bugs, They Are Features gradient science. *Neural Information Processing Systems (NeurIPS)* (2019)
24. Ionescu, R.T., Alexe, B., Leordeanu, M., Popescu, M., Papadopoulos, D.P., Ferrari, V.: How hard can it be? Estimating the difficulty of visual search in an image. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 2157–2166 (2016)
25. Isola, P., Xiao, J., Parikh, D., Torralba, A., Oliva, A.: What makes a photograph memorable? *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)* **36**(7), 1469–1482 (2014)
26. Isola, P., Zoran, D., Krishnan, D., Adelson, E.H.: Crisp boundary detection using pointwise mutual information. *Lecture Notes in Computer Science (including sub-series Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **8691 LNCS(PART 3)**, 799–814 (2014)
27. Jiang, L., Meng, D., Zhao, Q., Shan, S., Hauptmann, A.G.: Self-paced curriculum learning. *Proceedings of the National Conference on Artificial Intelligence* **4**, 2694–2700 (2015)
28. Krizhevsky, A.: Learning Multiple Layers of Features from Tiny Images. Tech. rep., Toronto (2009)
29. Kumar, M., Packer, B., Koller, D., Kumar, P., Packer, B., Koller, D.: Self-paced learning for latent variable models. *Neural Information Processing Systems (NeurIPS)* (2010)
30. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2323 (1998)
31. Lee, Y.J., Grauman, K.: Learning the easy things first: Self-paced visual category discovery. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 1721–1728 (2011)
32. Li, Y., Yosinski, J., Clune, J., Lipson, H., Hopcroft, J.: Convergent Learning: Do different neural networks learn the same representations? *International Conference on Learning Representations (ICLR)* (2016)
33. Liu, D., Xiong, Y., Pulli, K., Shapiro, L.: Estimating image segmentation difficulty. *Machine Learning and Data Mining in Pattern Recognition* **6871 LNAI**, 484–495 (2011)
34. Maennel, H., Baldock, R.J., Alabdulmohsin, I., Bousquet, O., Keysers, D., Tolstikhin, I., Gelly, S.: What Do Neural Networks Learn When Trained With Random Labels? *Neural Information Processing Systems (NeurIPS)* (2020)
35. Mangalam, K., Prabhu, V.: Do deep neural networks learn shallow learnable examples first? *International Conference on Machine Learning (ICML), Deep Phenomena Workshop* (2019)

36. Ortiz-Jiménez, G., Modas, A., Moosavi-Dezfooli, S.M., Frossard, P.: Hold me tight! Influence of discriminative features on deep network boundaries. *Neural Information Processing Systems (NeurIPS)* (2020)
37. Peterson, J.C., Battleday, R.M., Griffiths, T.L., Russakovsky, O.: Human uncertainty makes classification more robust. *International Conference on Computer Vision (ICCV)* (2019)
38. Pinto, N., Cox, D.D., DiCarlo, J.J.: Why is real-world visual object recognition hard? *PLoS Computational Biology* **4**(1), 0151–0156 (2008)
39. Russakovsky, O., Deng, J., Huang, Z., Berg, A.C., Fei-Fei, L.: Detecting avocados to zucchinis: What have we done, and where are we going? *International Conference on Computer Vision (ICCV)* (2013)
40. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A.C., Fei-Fei, L.: ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* **115**(3), 211–252 (2015)
41. Shah, H., Tamuly, K., Raghunathan, A., Jain, P., Netrapalli, P.: The Pitfalls of Simplicity Bias in Neural Networks. *Neural Information Processing Systems (NeurIPS)* (2020)
42. Shwartz-Ziv, R., Tishby, N.: Opening the Black Box of Deep Neural Networks via Information. *Why & when Deep Learning works: looking inside Deep Learning ICRI-CI* (2017)
43. Simonyan, K., Zisserman, A.: Very Deep Convolutional Networks for Large-Scale Image Recognition. In: *International Conference on Learning Representations (ICLR)* (2015)
44. Skilling, J., Bryan, R.: Maximum entropy image reconstruction: general algorithm. *Monthly notices of the royal astronomical society* **211**, 111–124 (1984)
45. Spain, M., Perona, P.: Some objects are more equal than others: measuring and predicting importance. *European Conference on Computer Vision (ECCV)* (2008)
46. Vijayanarasimhan, S., Grauman, K.: What’s it going to cost you?: Predicting effort vs. informativeness for multi-label image annotations. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2009)
47. Wang, L., Hu, L., Gu, J., Wu, Y., Hu, Z., He, K., Hopcroft, J.: Towards understanding learning representations: To what extent do different neural networks learn the same representation. *Neural Information Processing Systems (NeurIPS)* (2018)
48. Yang, K., Qinami, K., Fei-Fei, L., Deng, J., Russakovsky, O.: Towards fairer datasets: Filtering and balancing the distribution of the people subtree in the ImageNet hierarchy. *Conference on Fairness, Accountability, and Transparency (FAT)* pp. 547–558 (2020)
49. Zhang, C., Bengio, S., Hardt, M., Recht, B., Vinyals, O.: Understanding deep learning requires rethinking generalization. *International Conference on Learning Representations (ICLR)* (2017)