

RESQUE: Quantifying Estimator to Task and Distribution Shift for Sustainable Model Reusability

Vishwesh Sangarya and Jung-Eun Kim*

Department of Computer Science, North Carolina State University

Abstract

As a strategy for sustainability of deep learning, reusing an existing model by retraining it rather than training a new model from scratch is critical. In this paper, we propose REpresentation Shift QUantifying Estimator (RESQUE), a predictive quantifier to estimate the retraining cost of a model to distributional shifts or change of tasks. It provides a single concise index for an estimate of resources required for retraining the model. Through extensive experiments, we show that RESQUE has a strong correlation with various retraining measures. Our results validate that RESQUE is an effective indicator in terms of epochs, gradient norms, changes of parameter magnitude, energy, and carbon emissions. The measures align well with RESQUE for new tasks, multiple noise types, and varying noise intensities. As a result, RESQUE enables users to make informed decisions for retraining to different tasks/distribution shifts and determine the most cost-effective and sustainable option, allowing for the reuse of a model with a much smaller footprint in the environment.

Introduction

With the rapid and ever increasing use of deep learning models in everyday applications, a crucial aspect that is often overlooked is the adaptation of a model to changes and new data. Models need to be adaptive and capable of learning from these new inputs. This dynamic adaptation is essential for maintaining model performance and relevance over time. Particularly, the continuous series of changes a model may need to undergo subject to distributional shifts, or need for primary task updates. For example, new data with better relevant features may be introduced, the environment of deployment may undergo a change leading to distribution shift of the input data, or the prediction classes and task may change. Retraining helps the model remain effective with satisfactory performance for the desired deployed application while shortening the adaptation time. While new models can be developed from scratch when the data undergoes a distribution shift or the task changes, leveraging existing learned features is a resource-efficient option. Reusing models not only saves time and resources but also aligns with the principles of *sustainable AI development*. This is

of primary importance with the overwhelmingly increasing deployments of AI models in the current era, which in turn has given rise to the socially important field of Green AI (Schwartz et al. 2019; Wu et al. 2022; Verdecchia, Salou, and Cruz 2023; Salehi and Schmeink 2024). It is not uncommon for models to have several billion parameters (Menghani 2023), and as these models grow in complexity, their computational demands also increase. With the increased computational requirements, there is an accompanying need for resources required to perform the computation and byproducts of the computation to deal with. These demands place significant strain on available resources and highlight the importance of addressing both the direct and indirect impacts of AI model use. Consequently, there exists a need for sustainable and efficient development of AI models so as to reduce the environmental impact, especially when looking at the *carbon footprint* (Ali et al. 2023; Strubell, Ganesh, and McCallum 2019; Bannour et al. 2021; Wu et al. 2024) and energy consumption (Dodge et al. 2022; Strubell, Ganesh, and McCallum 2019; Ali et al. 2023; Wu et al. 2024) involved in the research, evolution, and utilization of AI models.

With the goal of reusing models and reducing computation, energy consumption, and carbon emissions, we propose a novel estimating index called the REpresentation Shift QUantifying Estimator (RESQUE). This estimator provides a single index to predict the cost of retraining a model for new distributions or tasks before any computation is performed. RESQUE enables deep learning practitioners and researchers to quickly estimate the costs associated with adapting a model, offering valuable insights for informed decision-making. This quantifiable estimator helps them achieve sustainability goals when using and developing the models. For the context of distribution shifts, the estimator referred to as $RESQUE_{dist}$, measures the shift in the model’s representation outputs between the original and new distribution. To obtain $RESQUE_{dist}$, the input data is propagated through the model just *once* with *no backward* propagation or computation. For the case of a new target task, the estimator, specifically termed as $RESQUE_{task}$, quantifies the separation in class decision boundary of the new target task in the representation space, using the original task model. RESQUE obtained after the forward pass serves as an index to estimate the resources required to adapt the model to a

*Correspondence.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

new target task or new distribution. Through extensive experiments and by theoretical reasoning, we show that a lower value of RESQUE correlates with a lower cost of retraining to a new target task or new distribution.

Additionally, we show that RESQUE not only correlates with important sustainability costs such as energy consumption and carbon emissions, but also has strong correlation with computational related measures such as training epochs, gradients, and model parameter change magnitudes. Furthermore, RESQUE is model/architecture-agnostic, proving effective against either convolutional networks or transformers. Our experiments span a wide range of target tasks, various vision datasets, and different types and intensities of distribution shifts, demonstrating the practical value of RESQUE. Using RESQUE aids in meeting resource and sustainability targets while reducing the computational effort required for adapting AI models. The code for this work is available here:

<https://github.com/JEKimLab/AAAI2025ReSQuE>

Related Work

Sustainable computing and Green AI (Schwartz et al. 2019; Verdecchia, Sallou, and Cruz 2023; Wu et al. 2022; Salehi and Schmeink 2024) are gaining more and more attention due to their potential impacts on the modern AI era of large models. As the demand for AI grows (Liu, Liu, and Lee 2024; Wu et al. 2024), the need for efficient and sustainable computing practices has become important, leading to an important research domain. (Verdecchia, Sallou, and Cruz 2023; Wu et al. 2024) perform a detailed study regarding the potential quantification of undesirable effects of AI, they highlight the importance of reporting energy consumption and carbon emissions as key sustainability measures. Other studies (Wu et al. 2022; Schwartz et al. 2019; Salehi and Schmeink 2024) bring forth the need for a sustainable and greener development of AI, while highlighting the environmental consequences of neglecting the impact of inconsiderate use of resources.

(Georgiou et al. 2022; Strubell, Ganesh, and McCallum 2019) conduct detailed studies regarding energy usage of neural networks. In particular (Strubell, Ganesh, and McCallum 2019; Bannour et al. 2021; Georgiou et al. 2022) explore the increasing energy demands of neural networks from the perspective of computing power, operational costs, and training time. Along with energy consumption, (Henderson et al. 2020; Patterson et al. 2021) highlight the absence of carbon emissions reporting associated with deep learning research. Besides these studies, other works such as (Xu et al. 2023; García-Martín et al. 2019; McDonald et al. 2022) discuss the environmental impact of training models from a carbon emissions perspective. To aid in tracking carbon emissions and energy usage associated with deep learning development, (Schmidt et al. 2021; Anthony, Kanding, and Selvan 2020; Lacoste et al. 2019) propose useful quantifying strategies, including frameworks and libraries which can accurately log all energy consumption and resulting carbon emissions from training. (Dodge et al. 2022) demonstrate that training Vision Transformers (Dosovitskiy et al. 2021)

results in significantly higher carbon emissions than convolutional networks due to their greater energy requirements from extensive training characteristics. On a similar stream, (Strubell, Ganesh, and McCallum 2019) assess the substantial energy consumption of transformer-based models.

Changes in data due to varying input distributions can arise due to various factors, as elaborated in (Hendrycks and Dietterich 2019; Arjovsky et al. 2020; Hendrycks et al. 2021). To adapt to the specific distributional changes, augmentation strategies (Hendrycks et al. 2020; Liu et al. 2022; Kim, Choo, and Song 2020; Lee et al. 2020) were effective to a limited extent. Although these techniques improve performance compared to a standard trained model against certain types of shifts, they are computationally costly and require training models from scratch. Moreover, the performance improvements are for a limited number of shifts in the distributions, while providing no benefits or a drop in performance for other types of distribution shifts. Studies such as (Geirhos et al. 2020; Yin et al. 2020; Ford et al. 2019) detail the non-uniform performance gains, where improvements in performance cannot be controlled and may result in lower than acceptable performance at the end of training. Adaptation during test time, such as (Lim et al. 2023; Niu et al. 2022; Goyal et al. 2022; Wang et al. 2022), addresses the computational cost concerns by having additional tuned components that can be added to the model. However, these approaches heavily rely on batch data and tend to provide uncertain and low-confidence outputs when there is a greater variation of distribution shifts in the input.

Gradients play a crucial role in evaluating the complexity and difficulty of learning when training models, as demonstrated in (Agarwal, D’souza, and Hooker 2022; Lee and AlRegib 2020; Huang, Geng, and Li 2021). These studies tell us that difficult samples produce greater gradients and thus require larger computation costs for convergence. (Sangarya, Bradford, and Kim 2024, 2023) provide measures for evaluating distribution shifts when a shifted input sample is only derived from a corresponding original sample, which is not a realistic scenario. In particular, (Sangarya, Bradford, and Kim 2023) only operates on individual shifts without the ability to evaluate multiple types of shifts simultaneously. Moreover, both studies only focused on shifts in data but did not take into account the practical scenario where the task itself undergoes a change. Additionally, studies such as (Stacke et al. 2020) measure changes in layer-wise representations to analyze different shifts in pathology data.

Adjusted Rand Index (Hubert and Arabie 1985), is a useful metric for comparing clustering algorithms. However, it can also be used to assess the performance of supervised classification and feature selection as demonstrated by (Santos and Embrechts 2009). With regard to the case of task change, we refer to research in various subfields of deep learning, such as (Ravi and Larochelle 2016; Oreshkin, Rodríguez López, and Lacoste 2018; Caruana 1997), which defines a new task as any change in the class data or the introduction of completely new classes of associated data. (Achille et al. 2019) is designed to measure task similarity and perform an evaluation of pretrained models. However, adapting to a new task from a current task model re-

quires dataset and task-specific analysis to estimate retraining costs. Likewise, although (Li et al. 2023) consider factors such as dataset and model characteristics, they do not specifically address the magnitude of retraining needed to adapt a model to a new task.

RESQUE

In this section, we provide details of $\text{RESQUE}_{\text{dist}}$ for distributional shifts and $\text{RESQUE}_{\text{task}}$ for change of task.

Distribution Shift

In this scenario, we consider the case when a model needs to be updated and adapted as it saw a distributional shift in data - for example, because some ground truths might be changed, some data samples might become stale, or some new data samples might need to come into play, etc. We explain how to obtain and use $\text{RESQUE}_{\text{dist}}$, which is a predictive quantifier for model retraining to distributional shifts. It quantifies the distance between representations of the original and the shifted distribution so as to determine how *familiar* or *foreign* the distribution is.

Normalized Embedding Vectors for Distributions For a realistic scenario where a distribution-shifted sample has no corresponding clean image, or the number of distributed shifted samples may not be equal to the number of original clean samples, we make use of all the samples to generate representation embedding vectors. For the case of distribution shifts, since there is no change in the original task and therefore no change in the number of classes, we obtain the class-wise representation embedding vectors for each class individually. To obtain the data representations, we use the output representation from the final convolutional layer in convolutional layers and the final dense layer in Vision Transformers. We perform a forward pass of each dataset and obtain a summed embedding representation vector of each class as follows,

$$V_{l,sum}^D = \sum_{i=0}^{n_l} V_{l,i}^D \quad (1)$$

where, $V_{l,i}^D$ represents the flattened representation of sample i with class label l within dataset D . For each class l , where n_l is the number of samples within the class, the representation vectors are summed to obtain the summed embedding representation vector $V_{l,sum}^D$. The summed embedding vector for each class label l is then normalized as follow,

$$V_{l,norm}^D = \frac{V_{l,sum}^D}{\|V_{l,sum}^D\|} \quad (2)$$

RESQUE_{dist} from Normalized Embeddings The final representation angle is obtained by averaging the inverse cosine angle between each class' normed representation vector for the original dataset and the distribution-shifted dataset. Let O represent the original dataset, S represent the distribution shifted dataset, and k represent the number of classes. $\text{RESQUE}_{\text{dist}}$ is then obtained as,

$$\text{RESQUE}_{\text{dist}} = \frac{\sum_{i=0}^k \arccos(V_{i,norm}^O, V_{i,norm}^S)}{k} \quad (3)$$

Algorithm 1: Initialize cluster representation centroids

Input: Number of samples n_s ,
Number of labels k ,
Model M producing representation output (R, Y) ,
Output: Centroids matrix $\text{Cent}M$
 $\text{Cent}M \leftarrow \text{Empty Vector of size } k$
for $i \leftarrow 1$ to k **do**
 $\text{current_centroid} \leftarrow \vec{0}$
 $\text{label_count} \leftarrow 0$
 for $j \leftarrow 1$ to n_s **do**
 $(R, Y) \leftarrow M(X_j)$ $\triangleright X_j$ is the j th sample
 if $Y == i$ **then**
 $\text{current_centroid} \leftarrow \text{current_centroid} + R$
 $\text{label_count} \leftarrow \text{label_count} + 1$
 end if
 end for
 $\text{Cent}M_i \leftarrow \text{current_centroid}/\text{label_count}$
end for

Change of Task

For the first, we consider a case in which a model is required to be updated for application toward a new target task. We show how a model can be retrained and reused to meet such a requirement. To characterize the resources and cost of retraining to a new task, we make use of the separation of classes within the latent space of representations. The separation of classes in the representation latent space is an effective gauge of the model's ability to provide accurate outputs. We relate this to the distance between decision boundaries for classes within the representation space, where distinct boundaries result in correct and confident predictions, whereas class boundaries that overlap or are close to each other lead to low confidence and incorrect predictions. We use this trait of deep learning models to evaluate how well a model's current learned features and representation space translate into learning a new task.

Quantifying Class Separation from Known Task Data to New Task Data

We use the representation of the data produced by the model and a clustering algorithm to assign cluster labels for each data sample in the representation latent space. Using the assigned cluster labels, we derive our estimator, $\text{RESQUE}_{\text{task}}$, by applying the Adjusted Rand Index (Hubert and Arabie 1985). $\text{RESQUE}_{\text{task}}$ can quantify the effectiveness of class separation by using the cluster labels and the true data labels. To obtain the cluster labels, we first retrain the original model with the new task just for a *single* epoch. This creates a representation that is primarily of the original task but incorporates features of the new task and updates the output classes of the model to match the new task classes. Following this step, the new task data is used to perform a forward pass and assign a label, which is a representation label, obtained via clustering on all samples.

Clustering Mechanism for the New Unknown Task The data representation is obtained from the final convolution layer for a convolutional network' case while from the final dense layer in the last transformer encoder block for a

vision transformer’s case. To obtain the representation labels by clustering, we use KMeans with 3 different centroid initialization techniques as follows:

1. Using the original data to obtain centroids as detailed in Algorithm 1,
2. Initializing by Kmeans++ (Arthur, Vassilvitskii et al. 2007), and
3. Initializing by random cluster assignment, and selecting the cluster with least entropy among 20 random initialization seeds.

The above three initialization schemes result in similar clustering labels. We use Algorithm 1 due to its computational efficiency as it does not require random re-initialization or iterative assignment of centroids.

In Algorithm 1, we begin by creating an empty vector that has a size equal to the number of class labels, as depicted in line 1. $CentM$ is a vector of size k when initialized, but as the classes’ centroids are obtained, each entry in $CentM$ is a vector itself. In the end, $CentM$ is a matrix of size $k \times$ size of the flattened representation.

Adjusted Rand Index to Quantify Class Separation

Adjusted Rand Index takes the value 0 for purely random clustering, and 1 for identical clustering. For our estimator, it is required to have a low value for decision boundaries which are well separated and high value for boundaries which overlap and result in incorrect representation cluster labels. Hence, for our estimator, $RESQUE_{task}$, we take the complement of Adjusted Rand Index and define it in Eq. (4). In Eq. (4), tl represents the total count of true labels for each label in the contingency table of true labels vs. representation labels via clustering. rc represents the summed values of representation cluster labels in the contingency table. A contingency table in this scenario is a matrix that summarizes the number of samples belonging to the same cluster or having the same label in both clustering scenarios. Here, by ‘both clustering scenarios,’ we refer to the representation-based clustering labels and the true labels. n_c is the number of cluster labels, which is equal to the number of class labels. $n_{i,j}$ represents the value in each entry of the contingency table, which is common to both cluster labels for a given label i and j . n_s is the total number of samples.

The overall formulation $RESQUE_{task}$ is as follows,

$$RESQUE_{task} = \frac{\frac{1}{2} \left[\sum_{i=0}^{n_c} \binom{rc_i}{2} + \sum_{j=0}^{n_c} \binom{tl_j}{2} \right] - \sum_{i,j=0}^{n_c} \binom{n_{i,j}}{2}}{\frac{1}{2} \left[\sum_{i=0}^{n_c} \binom{rc_i}{2} + \sum_{j=0}^{n_c} \binom{tl_j}{2} \right] - \left[\sum_{i=0}^{n_c} \binom{rc_i}{2} \sum_{j=0}^{n_c} \binom{tl_j}{2} \right]} / \binom{n}{2} \quad (4)$$

Retraining Measures

Retraining measures refer to the resources and costs expended when a model is adapted to a new task or distribution. We demonstrate that $RESQUE$ is strongly correlated with these measures and serves as an effective estimator. Using $RESQUE$, users can estimate the resource expenditure required for retraining the models. We evaluate several measures including carbon emissions, energy consump-

| Model | Correlation | Epochs | GradNorm | Param. change |
|---------------------|---------------------|------------------------------|------------------------------|------------------------------|
| ResNet18 (CIFAR100) | Pearson Spearman | 0.77 (6e-07) 0.97 (1e-19) | 0.74 (2e-06) 0.97 (1e-20) | 0.70 (1e-05) 0.95 (6e-17) |
| ViT (CIFAR10) | Pearson Spearman | 0.87 (2e-10) 0.88 (9e-11) | 0.88 (1e-10) 0.88 (8e-11) | 0.84 (3e-09) 0.88 (1e-10) |
| VGG16 (SVHN) | Pearson Spearman | 0.93 (2e-08) 0.92 (5e-13) | 0.95 (1e-16) 0.95 (9e-16) | 0.94 (1e-15) 0.93 (3e-14) |

Table 1: Correlation coefficients (and associated p-value) for multiple models

tion, epochs, total gradient norm, and normalized parameter change.

Carbon Emissions As highlighted in earlier sections, carbon emissions reporting (Henderson et al. 2020) is vital for sustained development as it represents the environmental cost linked to model training. We utilize a carbon tracking library (Schmidt et al. 2021), to track the carbon emissions associated with training. The library estimates the carbon footprint by taking into account the method of energy generation in the region where the GPUs are used for model training, and, along with the energy consumed, calculates the estimated carbon compound byproducts.

Energy Consumption Energy consumption is another vital sustainability cost that needs to be measured and reported. We utilize the same library (Schmidt et al. 2021) used for carbon tracking to measure the energy consumed for the entire training and retraining process. To measure the energy consumed, the energy of processes that operate on the GPU, CPU, and RAM are taken into account to provide the total energy consumption.

Epochs Along with carbon emissions and energy consumption, the training cost can be quantified by the number of epochs required to reach the desired performance when adapting to new distributions or tasks. We provide detailed empirical results showing that $RESQUE$ consistently aligns with the number of epochs, confirming its effectiveness as an estimator. The epochs reported are the number of epochs a model requires to reach a predefined test accuracy. Due to common hyperparameters, as additional termination conditions, training is halted after 25 or 50 epochs if the accuracy falls within 0.5% or 1% of the selected cutoff, respectively.

Total Gradient Norm Gradient norm is a common measure of the difficulty of learning algorithms. Research studies such as (Agarwal, D’souza, and Hooker 2022; Huang, Geng, and Li 2021; Lee and AlRegib 2020) use gradient norm as a proxy for sample difficulty. In this paper, we report the gradient norm by aggregating the gradient norm at each time interval during retraining. The final gradient norm value represents the total magnitude of gradients the model experienced from the beginning to the end of the training.

Normalized Parameter Change The change in the model parameters denotes the extent of update a model is supposed to undergo when estimating the retraining measure. A larger normalized parameter change indicates significant model deviation and higher resource use, conversely, a lower change suggests lower costs and faster convergence.




Figure 1: Retraining vs. Training from scratch on the SVHN dataset with Gaussian noise using VGG16

Similar to (Zhang, Bengio, and Singer 2022), we calculate parameter change but instead of comparing initial and final values, we aggregate changes between consecutive time steps of training and retraining. For parameters of layer l at time t , the normalized change $N_{l,t}$, between current parameters and the previous time step’s parameters is given by

$$N_{l,t} = \|W_{l,t} - W_{l,t-1}\|_2 / \sqrt{\|W_{l,t}\|} \quad (5)$$

Here, t (≥ 1) represents the current time interval, while $t-1$ represents the previous time interval. The distance between current parameters $W_{l,t}$ and the parameters from the previous time interval $W_{l,t-1}$ is calculated for each layer in the model. The final value of the normalized parameter change is obtained by summing up the total change per time interval, per layer, and averaged by the number of layers.

Experiments

All experiments were carried out on NVIDIA RTX GPUs, specifically utilizing the RTX 2060, 2070, 2080, 3060Ti, 4060Ti, and A100 models. Carbon emissions and energy were measured on RTX 2060. All experiment results are an average of three runs.

Distribution Shift

In this section, we assess $\text{RESQUE}_{\text{dist}}$ in the context of distributional shifts using three datasets : CIFAR10 (Krizhevsky, Hinton et al. 2009), CIFAR100 (Krizhevsky, Hinton et al. 2009), and SVHN (Netzer et al. 2011). We utilize 3 types of noises - Gaussian noise, Image Blur, and Salt-Pepper noise, which represent various real world noises that can occur due to hardware changes, environmental factors, or artifacts in the data, respectively. For each noise type, we generate 10 levels of noise intensity, with level 1 corresponding to minimal noise and level 10 aligning with severity 4 as described in (Hendrycks and Dietterich 2019). To represent the reusability of pre-existing models, we initially train a randomly initialized model on the original data distribution until it achieves the minimum required accuracy for each dataset for fair comparisons. While tuning hyper-parameters for higher noise levels could expedite convergence, it may hinder comparisons across different noise types and levels. Therefore, we maintain consistent hyper-parameters across all experiments for consistency.

It is important to note that, for the initial model training on original data without distributional shifts, we utilize 70% of the entire dataset. For retraining to data with distributional

| Model | Correlation | Epochs | Grad Norm | Param. change |
|------------|-------------|--------------|--------------|---------------|
| ResNet18 | Pearson | 0.86 (0.012) | 0.82 (0.022) | 0.65 (0.100) |
| (CIFAR10) | Spearman | 0.96 (4e-04) | 0.92 (0.002) | 0.75 (0.052) |
| ResNet18 | Pearson | 0.75 (0.050) | 0.74 (0.053) | 0.71 (0.071) |
| (CIFAR100) | Spearman | 0.89 (0.006) | 0.85 (0.014) | 0.78 (0.036) |
| ResNet18 | Pearson | 0.93 (0.002) | 0.83 (0.020) | 0.93 (0.002) |
| (GTSRB) | Spearman | 0.75 (0.052) | 0.75 (0.052) | 0.64 (0.119) |
| ViT B/16 | Pearson | 0.93 (7e-04) | 0.93 (7e-04) | 0.92 (0.001) |
| (CIFAR10) | Spearman | 0.83 (0.009) | 0.85 (0.006) | 0.71 (0.046) |

Table 2: Correlation coefficients (and p-values) of ResNet and ViT for different new target tasks.

shifts, we utilize 50% of the dataset with added noise. The 20% excess is the overlap data that was common in the original distribution and the new distribution shifted data.

Fig. 1 illustrates the retraining measures for VGG16 trained from scratch versus retrained on SVHN under varying intensities of Gaussian noise. The results clearly show that retraining a model requires significantly fewer resources than training from scratch. All retraining measures for retraining are substantially lower compared to those for training a new model, highlighting the efficiency of retraining to distribution shifts.

For evaluating $\text{RESQUE}_{\text{dist}}$ as an estimator for the retraining measures, we evaluate convolutional networks and vision transformers, retrained to various noise types and levels on different datasets. We train and retrain ResNet18 (He et al. 2016) on CIFAR100, VGG16 (Simonyan 2014) on SVHN, and ViT (Dosovitskiy et al. 2021) on CIFAR10. For the Vision Transformer, we utilize a ViT model with a patch size of 4, comprising 8 transformer blocks, a latent vector size of 512, 8 attention heads, and an MLP with a hidden layer size of 1024.

Figure 2 displays the correlation between $\text{RESQUE}_{\text{dist}}$ and the retraining measures, for all three models on the three datasets. Across various datasets and architectures, $\text{RESQUE}_{\text{dist}}$ consistently aligns with the retraining measures for different types of distribution shifts. Furthermore, Table 1 provides the Pearson and Spearman correlation coefficients, along with associated p-values, between $\text{RESQUE}_{\text{dist}}$ and all the retraining measures, demonstrating a strong correlation between $\text{RESQUE}_{\text{dist}}$ and the retraining measures.

Task Change

We evaluate the effectiveness of RESQUE to estimate resources for learning a new task from an original task. First, we compare training a model to a target task from scratch vs.

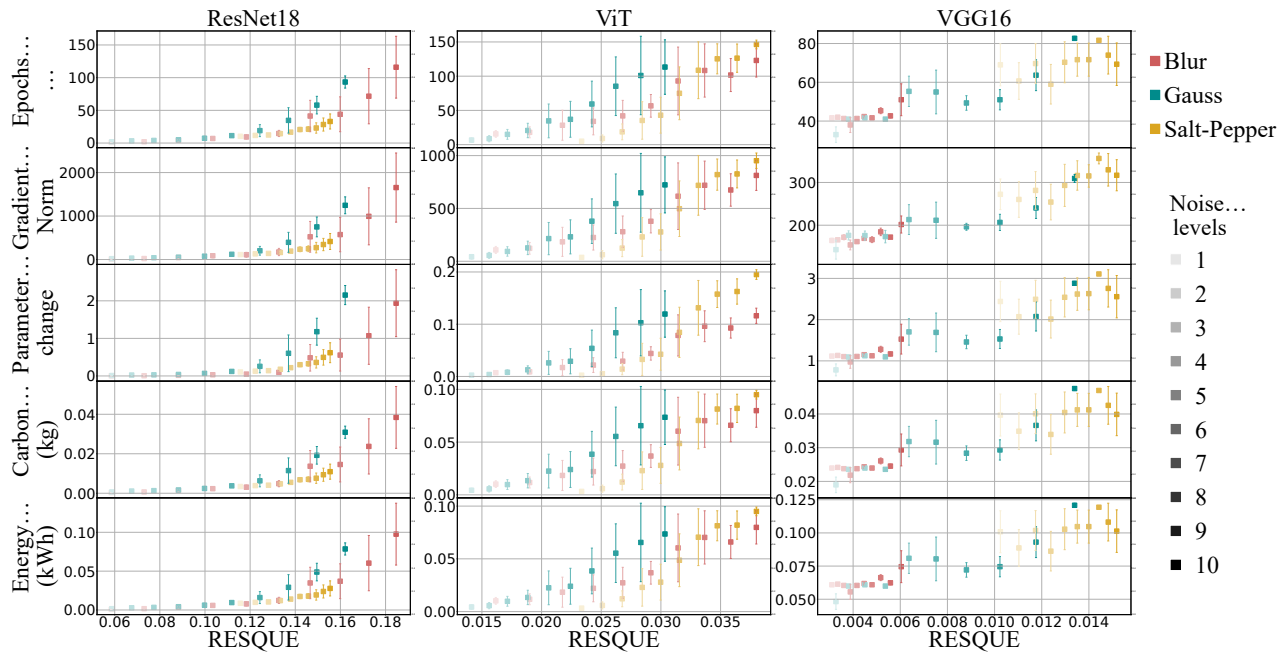


Figure 2: RESQUE and retraining measures for different models and datasets




Figure 3: As for the new target task, CIFAR10, comparisons of retraining from Food101 vs. training from scratch. Retraining consumes significantly less resources, epochs, energy, and carbon, than training from scratch.

retraining a model from the original task. For training from scratch, we randomly initialize a new model and train it on the target task. We set a common cutoff accuracy for a fair comparison. Fig. 3 displays the cost indicators for ResNet18, which was trained on Food101 for the cases of retraining vs. training just from scratch. The target task is CIFAR10. It is clear that for all retraining measures, retraining a model requires significantly fewer resources than training a model from scratch. When retraining an original task model to a new task, since the features of the new task are not known, the model experiences higher initial gradient norm and parameter change. However, this is not the case when retraining to distribution shifts, since the model has already learned a good amount of features of the data, resulting in lower gradient norm and parameter changes, as shown in Fig. 1.

Experiments Across Different Original Tasks We show how well $\text{RESQUE}_{\text{task}}$ and the retraining measures are aligned for different original tasks. The original tasks evaluated in Fig. 4 are trained on CIFAR10 (Krizhevsky, Hinton et al. 2009), CIFAR100 (Krizhevsky, Hinton et al. 2009),

EMNIST (Cohen et al. 2017), Fashion MNIST (Xiao, Rasul, and Vollgraf 2017), Food101 (Bossard, Guillaumin, and Van Gool 2014), GTSRB (Stallkamp et al. 2012), MNIST (LeCun et al. 1998), and SVHN (Netzer et al. 2011). We use ResNet18 and ViT-B/16 (Dosovitskiy et al. 2021), and perform clustering to obtain $\text{RESQUE}_{\text{task}}$ using the last convolutional layer of ResNet18, and the final dense layer in the last transformer encoder block of the ViT. During training and retraining, images are resized to 224x224 for ViT-B/16 and to 32x32 for ResNet18. To achieve better performance and to ensure consistent training conditions across all new target tasks, we retrain all layers of the model, which provides better accuracy performance than retraining only the final few layers as outlined in (Deng et al. 2023).

Fig. 4 depicts the relation between $\text{RESQUE}_{\text{task}}$ and all the retraining measures for ResNet and ViT on different target tasks. For the target tasks, CIFAR10 and GTSRB, we use multiple cutoff accuracies to evaluate how learning differs from the early stage to the later stage. From the figure, it is evident that $\text{RESQUE}_{\text{task}}$ aligns with the retraining measures and has a strong correlation.

A lower value of $\text{RESQUE}_{\text{task}}$ indicates that the model will require fewer resources to converge. For the target tasks, CIFAR10 and GTSRB, during the early learning stage, the trends exhibited are linear for all the cases of the original tasks. However, as training progresses, the number of epochs required by models with higher $\text{RESQUE}_{\text{task}}$ increases substantially. We relate this to the difficulty of learning and the usefulness of task and feature similarity.

In an original task model with aligned and well-learned initial features, the increase in epochs from low accuracy to high cutoff accuracy is smaller. In contrast, for a model with poorly learned initial features, there is a substantial increase

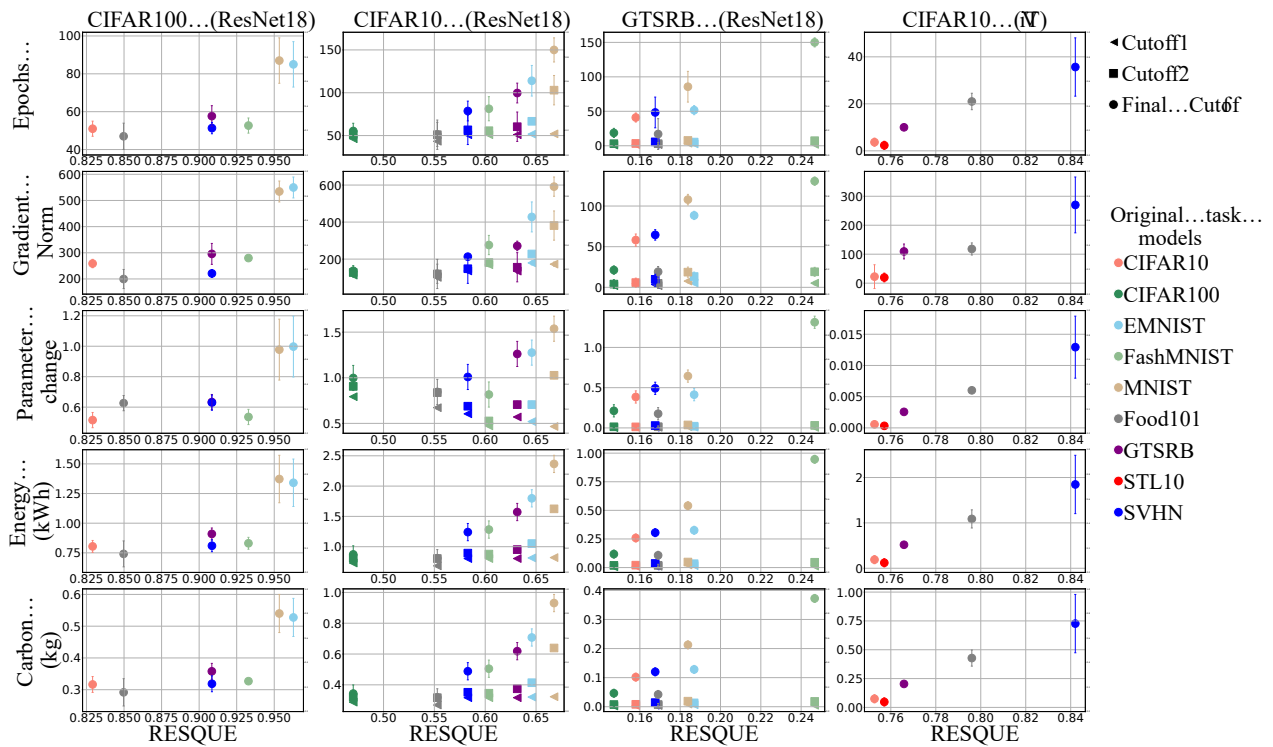


Figure 4: RESQUE vs. resource measures for ResNet and ViT retrained to different new target tasks. A positive relation between RESQUE and resource measures is exhibited.




Figure 5: CIFAR10 model trained to different target tasks. Error bars represent deviations. RESQUE has a strong linear relation with the peak performance a model can achieve on new tasks.

in epochs from low accuracy to high cutoff accuracy.

This is due to the fact that learning becomes progressively more difficult for models with less aligned or poorly learned features. This implies two crucial learning challenges - not only does the model lack relevant prior knowledge, but additionally, it also struggles to extract effective features from its initial training, leading to steeper learning curves and thus resulting in a larger number of epochs and resources expended. For ViT B/16 in Fig. 4, some of the prior task models fail to reach a high cutoff accuracy. CNNs overfit and achieved low accuracy on STL10 due to the limited sample size and the need for resizing to 32x32, hence, we did not use STL10-trained CNNs for new target task retraining.

Table 2 provides numerical correlation values as evidence

of a strong correlation between $RESQUE_{task}$ and the retraining measures. $RESQUE_{task}$ has a strong positive Pearson correlation and Spearman correlation coefficient, with a low p-value for both experiments. This indicates that there is a strong and statistically significant relationship between $RESQUE_{task}$ and the retraining measures.

Experiments Across Different Target Tasks $RESQUE_{task}$ can also accurately estimate the performance an original task model can achieve on different new target tasks. This can help profile and categorize task similarity and provide useful information for retraining a current model in future instances. For retraining to different target tasks, we retrain the original task model for a fixed number of epochs across all new target tasks. Fig. 5 illustrates the relation between $RESQUE_{task}$ and the highest test accuracy reached on each of the new task datasets (original test task is CIFAR10). Using $RESQUE_{task}$, we obtain an accurate estimation regarding the peak performance a model can attain on the new task.

Conclusion

We introduced a novel metric, RESQUE, to estimate the various resources that would be expended when reusing a model by adapting to distributional shifts or retraining to new target tasks. We validate the effectiveness of RESQUE on CNNs and ViTs. All the results consistently validate that RESQUE is an effective estimator of various retraining measures, including energy consumption and carbon emission, enabling sustainable decisions with regard to model reusability.

References

- Achille, A.; Lam, M.; Tewari, R.; Ravichandran, A.; Maji, S.; Fowlkes, C. C.; Soatto, S.; and Perona, P. 2019. Task2vec: Task embedding for meta-learning. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6430–6439.
- Agarwal, C.; D’souza, D.; and Hooker, S. 2022. Estimating Example Difficulty Using Variance of Gradients. arXiv:2008.11600.
- Ali, A. N. F.; Sulaima, M. F.; Razak, I. A. W. A.; Kadir, A. F. A.; and Mokhlis, H. 2023. Artificial Intelligence Application in Demand Response: Advantages, Issues, Status, and Challenges. *IEEE Access*, 11: 16907–16922.
- Anthony, L. F. W.; Kanding, B.; and Selvan, R. 2020. Carbontracker: Tracking and Predicting the Carbon Footprint of Training Deep Learning Models. arXiv:2007.03051.
- Arjovsky, M.; Bottou, L.; Gulrajani, I.; and Lopez-Paz, D. 2020. Invariant Risk Minimization. arXiv:1907.02893.
- Arthur, D.; Vassilvitskii, S.; et al. 2007. k-means++: The advantages of careful seeding. In *Soda*, volume 7, 1027–1035.
- Bannour, N.; Ghannay, S.; Névéol, A.; and Ligozat, A.-L. 2021. Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools. In Moosavi, N. S.; Gurevych, I.; Fan, A.; Wolf, T.; Hou, Y.; Marasović, A.; and Ravi, S., eds., *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*, 11–21. Virtual: Association for Computational Linguistics.
- Bossard, L.; Guillaumin, M.; and Van Gool, L. 2014. Food-101—mining discriminative components with random forests. In *Computer vision—ECCV 2014: 13th European conference, zurich, Switzerland, September 6–12, 2014, proceedings, part VI 13*, 446–461. Springer.
- Caruana, R. 1997. Multitask learning. *Machine learning*, 28: 41–75.
- Cohen, G.; Afshar, S.; Tapson, J.; and Van Schaik, A. 2017. EMNIST: Extending MNIST to handwritten letters. In *2017 international joint conference on neural networks (IJCNN)*, 2921–2926. IEEE.
- Deng, A.; Li, X.; Hu, D.; Wang, T.; Xiong, H.; and Xu, C.-Z. 2023. Towards Inadequately Pre-trained Models in Transfer Learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 19397–19408.
- Dodge, J.; Prewitt, T.; Tachet des Combes, R.; Odmark, E.; Schwartz, R.; Strubell, E.; Luccioni, A. S.; Smith, N. A.; DeCario, N.; and Buchanan, W. 2022. Measuring the Carbon Intensity of AI in Cloud Instances. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency, FAccT ’22*, 1877–1894. New York, NY, USA: Association for Computing Machinery. ISBN 9781450393522.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Hounsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Ford, N.; Gilmer, J.; Carlini, N.; and Cubuk, D. 2019. Adversarial Examples Are a Natural Consequence of Test Error in Noise. arXiv:1901.10513.
- García-Martín, E.; Rodrigues, C. F.; Riley, G.; and Grahn, H. 2019. Estimation of energy consumption in machine learning. *Journal of Parallel and Distributed Computing*, 134: 75–88.
- Geirhos, R.; Temme, C. R. M.; Rauber, J.; Schütt, H. H.; Bethge, M.; and Wichmann, F. A. 2020. Generalisation in humans and deep neural networks. arXiv:1808.08750.
- Georgiou, S.; Kechagia, M.; Sharma, T.; Sarro, F.; and Zou, Y. 2022. Green AI: do deep learning frameworks have different costs? In *Proceedings of the 44th International Conference on Software Engineering, ICSE ’22*, 1082–1094.
- Goyal, S.; Sun, M.; Raghunathan, A.; and Kolter, Z. 2022. Test-Time Adaptation via Conjugate Pseudo-labels. arXiv:2207.09640.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Henderson, P.; Hu, J.; Romoff, J.; Brunskill, E.; Jurafsky, D.; and Pineau, J. 2020. Towards the Systematic Reporting of the Energy and Carbon Footprints of Machine Learning. *Journal of Machine Learning Research*, 21(248): 1–43.
- Hendrycks, D.; Basart, S.; Mu, N.; Kadavath, S.; Wang, F.; Dorundo, E.; Desai, R.; Zhu, T.; Parajuli, S.; Guo, M.; Song, D.; Steinhardt, J.; and Gilmer, J. 2021. The Many Faces of Robustness: A Critical Analysis of Out-of-Distribution Generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 8340–8349.
- Hendrycks, D.; and Dietterich, T. 2019. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations. arXiv:1903.12261.
- Hendrycks, D.; Mu, N.; Cubuk, E. D.; Zoph, B.; Gilmer, J.; and Lakshminarayanan, B. 2020. AugMix: A Simple Data Processing Method to Improve Robustness and Uncertainty. arXiv:1912.02781.
- Huang, R.; Geng, A.; and Li, Y. 2021. On the Importance of Gradients for Detecting Distributional Shifts in the Wild. In Beygelzimer, A.; Dauphin, Y.; Liang, P.; and Vaughan, J. W., eds., *Advances in Neural Information Processing Systems*.
- Hubert, L.; and Arabie, P. 1985. Comparing partitions. *Journal of classification*, 2: 193–218.
- Kim, J.-H.; Choo, W.; and Song, H. O. 2020. Puzzle Mix: Exploiting Saliency and Local Statistics for Optimal Mixup. In *International Conference on Machine Learning (ICML)*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Lacoste, A.; Luccioni, A.; Schmidt, V.; and Dandres, T. 2019. Quantifying the Carbon Emissions of Machine Learning. arXiv:1910.09700.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11): 2278–2324.

- Lee, J.; and AlRegib, G. 2020. Gradients as a Measure of Uncertainty in Neural Networks. In *IEEE International Conference on Image Processing (ICIP)*.
- Lee, J.-H.; Zaheer, M. Z.; Astrid, M.; and Lee, S.-I. 2020. Smoothmix: a simple yet effective data augmentation to train robust classifiers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 756–757.
- Li, H.; Fowlkes, C.; Yang, H.; Dabeer, O.; Tu, Z.; and Soatto, S. 2023. Guided recommendation for model fine-tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3633–3642.
- Lim, H.; Kim, B.; Choo, J.; and Choi, S. 2023. TTN: A Domain-Shift Aware Batch Normalization in Test-Time Adaptation. arXiv:2302.05155.
- Liu, M.; Liu, H.-F.; and Lee, C.-C. 2024. An empirical study on the response of the energy market to the shock from the artificial intelligence industry. *Energy*, 288: 129655.
- Liu, X.; Shen, F.; Zhao, J.; and Nie, C. 2022. RandomMix: A mixed sample data augmentation method with multiple mixed modes. arXiv:2205.08728.
- McDonald, J.; Li, B.; Frey, N.; Tiwari, D.; Gadepally, V.; and Samsi, S. 2022. Great Power, Great Responsibility: Recommendations for Reducing Energy for Training Language Models. In Carpuat, M.; de Marneffe, M.-C.; and Meza Ruiz, I. V., eds., *Findings of the Association for Computational Linguistics: NAACL 2022*, 1962–1970. Seattle, United States: Association for Computational Linguistics.
- Menghani, G. 2023. Efficient deep learning: A survey on making deep learning models smaller, faster, and better. *ACM Computing Surveys*, 55(12): 1–37.
- Netzer, Y.; Wang, T.; Coates, A.; Bissacco, A.; Wu, B.; Ng, A. Y.; et al. 2011. Reading digits in natural images with unsupervised feature learning. In *NIPS workshop on deep learning and unsupervised feature learning*, volume 2011, 4. Granada.
- Niu, S.; Wu, J.; Zhang, Y.; Chen, Y.; Zheng, S.; Zhao, P.; and Tan, M. 2022. Efficient Test-Time Model Adaptation without Forgetting. arXiv:2204.02610.
- Oreshkin, B.; Rodríguez López, P.; and Lacoste, A. 2018. Tadam: Task dependent adaptive metric for improved few-shot learning. *Advances in neural information processing systems*, 31.
- Patterson, D.; Gonzalez, J.; Le, Q.; Liang, C.; Munguia, L.-M.; Rothchild, D.; So, D.; Texier, M.; and Dean, J. 2021. Carbon Emissions and Large Neural Network Training. arXiv:2104.10350.
- Ravi, S.; and Larochelle, H. 2016. Optimization as a model for few-shot learning. In *International conference on learning representations*.
- Salehi, S.; and Schmeink, A. 2024. Data-Centric Green Artificial Intelligence: A Survey. *IEEE Transactions on Artificial Intelligence*, 5(5): 1973–1989.
- Sangarya, V.; Bradford, R.; and Kim, J.-E. 2023. Aggregate Representation Measure for Predictive Model Reusability. In *Computational sustainability at NeurIPS*.
- Sangarya, V.; Bradford, R.; and Kim, J.-E. 2024. Estimating Environmental Cost Throughout Model’s Adaptive Life Cycle. In *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*.
- Santos, J. M.; and Embrechts, M. 2009. On the use of the adjusted rand index as a metric for evaluating supervised classification. In *International conference on artificial neural networks*, 175–184. Springer.
- Schmidt, V.; Goyal, K.; Joshi, A.; Feld, B.; Conell, L.; Laskaris, N.; Blank, D.; Wilson, J.; Friedler, S.; and Lucioni, S. 2021. CodeCarbon: estimate and track carbon emissions from machine learning computing (2021). DOI: <https://doi.org/10.5281/zenodo.4658424>.
- Schwartz, R.; Dodge, J.; Smith, N. A.; and Etzioni, O. 2019. Green AI. arXiv:1907.10597.
- Simonyan, K. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Stacke, K.; Eilertsen, G.; Unger, J.; and Lundström, C. 2020. Measuring domain shift for deep learning in histopathology. *IEEE journal of biomedical and health informatics*, 25(2): 325–336.
- Stallkamp, J.; Schlipsing, M.; Salmen, J.; and Igel, C. 2012. Man vs. computer: Benchmarking machine learning algorithms for traffic sign recognition. *Neural networks*, 32: 323–332.
- Strubell, E.; Ganesh, A.; and McCallum, A. 2019. Energy and Policy Considerations for Deep Learning in NLP. arXiv:1906.02243.
- Verdecchia, R.; Sallou, J.; and Cruz, L. 2023. A Systematic Review of Green AI. arXiv:2301.11047.
- Wang, Q.; Fink, O.; Gool, L. V.; and Dai, D. 2022. Continual Test-Time Domain Adaptation. arXiv:2203.13591.
- Wu, C.-J.; Acun, B.; Raghavendra, R.; and Hazelwood, K. 2024. Beyond Efficiency: Scaling AI Sustainably. *IEEE Micro*, 1–8.
- Wu, C.-J.; Raghavendra, R.; Gupta, U.; Acun, B.; Ardalani, N.; Maeng, K.; Chang, G.; Aga, F.; Huang, J.; Bai, C.; Gschwind, M.; Gupta, A.; Ott, M.; Melnikov, A.; Candido, S.; Brooks, D.; Chauhan, G.; Lee, B.; Lee, H.-H.; Akyildiz, B.; Balandat, M.; Spisak, J.; Jain, R.; Rabbat, M.; and Hazelwood, K. 2022. Sustainable AI: Environmental Implications, Challenges and Opportunities. In Marculescu, D.; Chi, Y.; and Wu, C., eds., *Proceedings of Machine Learning and Systems*, volume 4, 795–813.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Xu, Y.; Martínez-Fernández, S.; Martínez, M.; and Franch, X. 2023. Energy Efficiency of Training Neural Network Architectures: An Empirical Study. arXiv:2302.00967.
- Yin, D.; Lopes, R. G.; Shlens, J.; Cubuk, E. D.; and Gilmer, J. 2020. A Fourier Perspective on Model Robustness in Computer Vision. arXiv:1906.08988.
- Zhang, C.; Bengio, S.; and Singer, Y. 2022. Are All Layers Created Equal? arXiv:1902.01996.