

# How Stable are Transferability Metrics evaluations?

Andrea Agostinelli, Michal Pándy\*, Jasper Uijlings, Thomas Mensink, and Vittorio Ferrari

Google Research

{agostinelli, michalpandy, jrru, mensink, vittoferrari}@google.com

**Abstract.** Transferability metrics is a maturing field with increasing interest, which aims at providing heuristics for selecting the most suitable source models to transfer to a given target dataset, without fine-tuning them all. However, existing works rely on custom experimental setups which differ across papers, leading to inconsistent conclusions about which transferability metrics work best. In this paper we conduct a large-scale study by systematically constructing a broad range of 715k experimental setup variations. We discover that even small variations to an experimental setup lead to different conclusions about the superiority of a transferability metric over another. Then we propose better evaluations by aggregating across many experiments, enabling to reach more stable conclusions. As a result, we reveal the superiority of LogME at selecting good source datasets to transfer from in a semantic segmentation scenario,  $\mathcal{N}$ LEEP at selecting good source architectures in an image classification scenario, and GBC at determining which target task benefits most from a given source model. Yet, no single transferability metric works best in all scenarios.

## 1 Introduction

Transfer learning aims to re-use knowledge learned on a source task to help learning a target task, for which typically there is only little training data. The most prevalent method of transfer learning in computer vision is to pre-train a source model on a large source dataset (*e.g.*, ILSVRC’12 [66]), and then fine-tune it on the target dataset [5,17,29,31,40,69,94]. However, different target tasks benefit from using different source model architectures [13,32,55,64] or pre-training on different source datasets [51,56,88]. Hence, a key challenge is determining which source model is best suited for which target task, and doing so in a computationally efficient manner.

Transferability metrics [2,6,8,23,24,49,57,59,71,72,75,77,89] provide heuristics for selecting the most suitable source models for a given target dataset, without explicitly fine-tuning them all. These methods generally work by applying a source model to the target dataset to compute embeddings or predictions. Then

---

\* Currently at Waymo.

they efficiently assess how compatible these embeddings/predictions are with the target labels. This provides a proxy for how well the source model transfers to the target task.

Transferability metrics is a maturing field with numerous contributions and increasing interest [2,6,49,57,59,75,77,89]. However, existing works rely on custom experimental settings without standardized benchmarks, which leads to inconsistent conclusions across different papers. In particular, contradicting conclusions often appear when comparing findings across papers. For example, while NCE [77] consistently outperforms LEEP [57] in [89], LEEP outperforms NCE in both [75] and [57]. Furthermore LogME outperforms LEEP [89], but LEEP outperforms LogME in [59]. This raises questions about how stable conclusions are across different experimental setups and if there truly is a single best transferability metric.

The primary goal of our work is to evaluate the stability of experimental setup for transferability estimation. We observe that a single experiment typically consists of three components: (1) a choice of the pool of source models; (2) a choice of the target dataset; (3) a choice of the measure used to evaluate how well transferability metrics perform. We vary each of these components in a large-scale systematic study of two scenarios: selecting good source datasets for semantic segmentation (Sec. 5), and selecting good source model architectures for image classification (Sec. 6). In total we construct a large set of **715k experiments**, several orders of magnitude larger than previous works [57,49,89,59,6] (and in a computationally efficient way). The source code for our experiments and analysis is publicly available<sup>1</sup>. Based on these experiments: (A) We demonstrate that even small variations to an experimental setup leads to very different conclusions about the superiority of a transferability metric over another. (B) We provide a systematic analysis to investigate which of the three setup components contributes the most to the instability of experimental outcomes. (C) We propose better evaluations by aggregating outcomes from a broad set of diverse experiments, reducing the experimental uncertainty and enabling to reach more stable conclusions. Concretely, we reveal that LogME is the best transferability metric in our first scenario (selecting source datasets for semantic segmentation);  $\mathcal{N}$ LEEP is the best in our second scenario (selecting source model architectures for image classification). Moreover, we also consider a third, somewhat separate scenario about determining which target task benefits most from a given source model, among a pool of target tasks constructed by subsampling out of a large target dataset (in the supplementary material). In this scenario GBC is the best metric. So, no single transferability metric works best in all scenarios.

## 2 Related Work

**Transfer learning.** Training a deep neural network for a specific task often requires a large amount of data which can be difficult to obtain. The goal of

<sup>1</sup> [github.com/google-research/google-research/tree/master/stable-transfer](https://github.com/google-research/google-research/tree/master/stable-transfer)

transfer learning [9,62,76] is to leverage information from a source task with easily obtainable data, to improve performance on a target problem where data is scarce. Pre-training a neural network on large datasets and fine-tuning it on a target dataset is the most prevalent method of transfer learning in computer vision. For these reasons, there exists a wide array of large-scale source datasets, such as ILSVRC’12 [66], ImageNet21k [63], or Open Images [47]. In addition, recent works consider the use of unlabeled source datasets via self-supervised pre-training [14,15,30]. Other research studies settings in which transfer learning is effective. Taskonomy [91] develops connections between visual tasks (*e.g.* semantic segmentation, depth prediction, etc.), Mensink et al. [51] perform extensive experimental investigations in semantic segmentation settings, Mustafa et al. [53] focus on evaluating transfer learning for medical purposes, and Ding et al. [22] explore applications to human activity recognition. For a more general survey of transfer learning, we refer the reader to Weiss et al. [83].

**Transferability metrics.** Transferability metrics provide efficient heuristics for determining which pre-trained models are most suitable for a specific target task. To generate transferability metrics, *label comparison-based methods* leverage the labels of the source and target domains. These methods generally assume equivalence between source and target domain labels or obtain pseudo-labels by executing the source model on the target domain. Methods in this category include NCE [77] and LEEP [57]. *Source embedding-based methods* use the feature extractor of a pre-trained neural network to embed target domain images. Transferability metrics are then computed using the embeddings and their corresponding labels. These methods include GBC [59], LogME [89], and  $\mathcal{N}$ LEEP [49]. To conclude, *optimal transport-based methods* [75,4] develop cost functions usable within the optimal transport framework to determine transferability between two datasets. See Section 3 for a more detailed description of state-of-the-art transferability metrics.

**Overview of experimental stability.** Several works study experimental settings and analyse methods’ performance with the aim of solidifying an area of research. In the field of graph neural networks (GNNs), Errica et al. [26] evaluate and compare GNNs across a large suite of experiments, while Dwivedi et al. [25] and Shchur et al. [68] point out issues with GNN evaluation and propose evaluation improvements. In reinforcement learning (RL), Whiteson et al. [84] discuss RL evaluation principles, Jordan et al. [42] discover flaws with RL evaluation metrics, and Colas et al. [19] study the necessary number of seeds for stable RL testing. In computer vision, Hoiem et al. [34] and Hosang et al. [35] analyse the stability of object detectors, Zendel et al. [92] discuss the quality of testing data, Xian et al. [86] evaluate state-of-the-art in zero-shot learning, and Abnar et al. [1] assess the trade-offs of upstream and downstream model performance. Our work provides clarity about good experimental settings for research in transferability metrics.

### 3 Background

We first discuss the basics of transferability metrics and how they are evaluated. At a high level, the overall pipeline in most papers is [2,6,49,57,59,75,77,89]: (1) a transferability metric considers a source model  $S$  and a target dataset  $T$  and predicts a transferability score  $M$ .  $M$  predicts how well  $S$  will transfer to  $T$  (Sec. 3.1). (2) compute the true accuracy of transferring from  $S$  on  $T$  by fine-tuning it on the target training set, then applying the fine-tuned model on the target test set, and finally evaluate accuracy  $A$  based on the ground-truth (Sec. 3.2). (3) evaluate the quality of the transferability metric by checking how well  $M$  predicted true accuracy  $A$  (Sec. 3.3).

#### 3.1 Transferability metrics

In our paper we compare the following transferability metrics  $M$ .

**H-score.** H-score [6] is based on the intuition that a model transfers well to a target dataset if the target embeddings have low inter-class variance and low feature redundancy. These quantities are computed by constructing the inter-class and data covariance matrices.

**LEEP &  $\mathcal{N}$ LEEP.** LEEP [57] first predicts pseudo-labels by applying the source model on the target images. These predictions are then used for computing a log-likelihood between the target labels and the source model predictions. The core idea is that if the predictions are ‘clustered’ around individual target labels, adaptation to the target dataset should be easier.  $\mathcal{N}$ LEEP [49] extends this idea by first embedding the target images using the source feature extractor and then fitting a Gaussian mixture model (GMM) on them. The GMM is said to provide a better density estimator for the pseudo-labels than the classification head in the original LEEP.

**LogME.** After embedding the target images using the source feature extractor, LogME [89] computes the probability of the target labels conditioned on these embeddings (i.e. the evidence of target labels). By setting up a graphical model and using independence assumptions between samples, the authors propose an efficient algorithm for computing such evidence.

**GBC.** GBC [59] measures the statistical overlap between classes of the target dataset, after representing the target images in the embedding space determined by the source feature extractor. The intuition is that the more classes overlap in that space, the more difficult it will be to achieve high accuracy after fine-tuning the source model. The overlap is estimated using the Bhattacharyya coefficient between multivariate Gaussians fitted to each class.

#### 3.2 True accuracy of transfer learning

We compute the true accuracy  $A$  of the source model  $S$  on the target test set after fine-tuning on the target training set.  $A$  represents how well  $S$  transfers to the target  $T$ . We note that we fine-tune the full source model on the target training dataset, rather than just the classification head.



### 3.3 Evaluating the quality of a transferability metric

We introduce common measures for evaluating transferability metrics. Suppose we wish to evaluate a transferability metric  $M$  given a source pool  $\mathcal{S}$  containing  $n$  source models, and a target dataset  $T$ . Hence, we have access to  $n$  transferability metric values  $M_i$ , each with its corresponding test accuracy  $A_i$  associated with source model  $S_i$ . An evaluation measure captures how well  $M_i$  relates to  $A_i$  (i.e. higher values of transferability  $M_i$  predict higher values of true accuracy  $A_i$ ).

**Pearson correlation coefficient** [7] ( $\rho$ ). It measures linear correlation between  $M_i$  and  $A_i$  across  $i \in [1 \dots n]$ . For transferability metrics that tend to be linear with respect to test scores (e.g, LEEP [57]), it provides a straightforward way for measuring performance. The disadvantage is that low  $\rho$  correlation does not imply a bad performing transferability metric. The transferability rankings could be correct even though  $M_i$  and  $A_i$  are non-linearly related.

**Kendall rank correlation coefficient** [43] ( $\tau$ ). The core idea is that we should have  $M_i > M_j$  if  $A_i > A_j$ . Based on this,  $\tau$  is computed as:

$$\tau = \frac{1}{\binom{n}{2}} \sum_{i < j} \text{sgn}(M_i - M_j) * \text{sgn}(A_i - A_j) \quad (1)$$

Eq. (1) can be interpreted as averaging over the agreements (+1) or disagreements (−1) in rankings between transferability metrics and test scores. Hence,  $\tau \in [-1, 1]$  and high  $\tau$  implies a strong correlation between the rankings of source models according to  $M_i$  and according to  $A_i$ .  $\tau$  does not rely on linearity assumptions and gives a holistic idea for the overall ranking performance. The downside is that  $\tau$  gives equal importance to all source models, regardless of their performance. However, practitioners ultimately care about a transferability metric’s ability to correctly rank only the few best performing models.

**Weighted Kendall rank correlation coefficient** ( $\tau_w$ ). To prioritize the top-performing models, we may assign higher weights to them. Models with higher  $A_i$  within the sum in Eq. (1) should be weighted more. Hence, two transferability metrics that would perform the same in terms of  $\tau$  can now be discriminated based on which metric ranks top models better in  $\tau_w$ . We still have  $\tau_w \in [-1, 1]$  and now the evaluation measure reflects practitioners’ priorities. We assume a hyperbolic drop-off in model importance with decreasing ranks in terms of  $A_i$ , which is implemented in the Sklearn library [61].

**Relative top-1 accuracy** [49] ( $Rel@1$ ). It measures how close model  $k$  with the highest predicted transferability ( $k = \text{argmax}_i M_i$ ) performs, in terms of accuracy  $A_k$ , compared to the highest performing model ( $\max_i A_i$ ). It is computed as  $Rel@1 = \frac{A_k}{\max_i A_i}$ . Since  $A_k \leq \max_i A_i$ , we have  $1 \geq Rel@1 \geq 0$ . The benefit of  $Rel@1$  is that there could be multiple models that perform within a close margin of the top model, which in terms of  $Rel@1$  would obtain similar performance scores. Since  $\tau$  and  $\tau_w$  are not sensitive to the actual value of the test scores of models, the evaluation of a particular transferability metric could vary drastically different even if similar-performing source models change order in the ranked lists (which can happen even for small fluctuations in the values of  $M_i$  or  $A_i$ ).

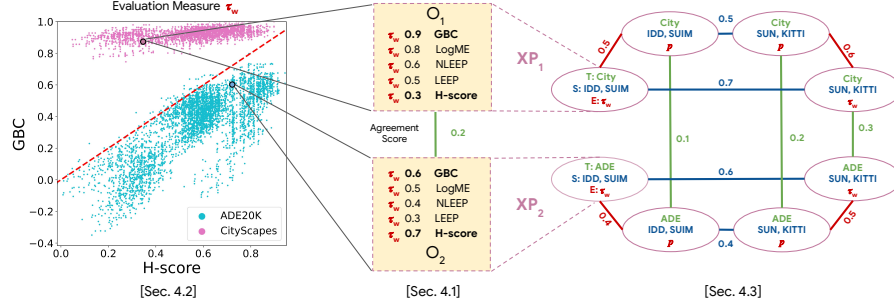


Fig. 1: An experiment  $XP$  is represented as a graph node, composed of source pool  $\mathcal{S}$ , target dataset  $T$  and evaluation measure  $E$ . We compare two experiments  $XP_1$  and  $XP_2$  by comparing their outcomes  $O_1$  and  $O_2$ . We do it qualitatively with a scatter plot (left) and quantitatively with a graph (right). For the latter we connect experiments that differ by a single variation in either  $\mathcal{S}$  ( $-$ ),  $T$  ( $\downarrow$ ) or  $E$  ( $\nearrow$ ). We compute the agreement score between the outcomes of each two connected experiments ( $XP_i, XP_j$ ) (edge value).

## 4 Methodology

Our goal is to investigate the experimental stability of protocols to evaluate transferability metrics. We first rigorously define what constitutes a single experiment and identify its three main components (Sec. 4.1). We then investigate the effects of small variations on each of these components both qualitatively (Sec. 4.2, Fig. 1 left) and quantitatively (Sec. 4.3, Fig. 1 right). Finally, in Sec. 4.4 we define how we create a huge set of experiments that allows us to explore the influence of varying each component on the evaluation of transferability metrics. At the same time, the sheer size of our experiments helps to reduce experimental uncertainty and improve their stability.

### 4.1 A single experiment with its three components

Given a set of transferability metrics, the outcome  $O$  of an experiment  $XP$  is a measure of quality for each transferability metric (Fig. 1, middle). A single experiment consists of three components whose choice influences its outcome: (1) A *pool of source models*  $\mathcal{S}$ . It can be created by training the same model architecture on different source datasets [2,59], or by training different model architectures on the same source dataset [57,89], or a combination of both [49]. (2) A *target dataset*  $T$ , which cannot be any dataset in the source pool. (3) An *evaluation measure*  $E$  to determine the quality of each transferability metric. Hence we can represent one experiment as  $XP(\mathcal{S}, T, E) \rightarrow O$  (Fig. 1).

We can compare two experiments  $XP_i$  and  $XP_j$  by looking at differences in their outcomes  $O_i$  and  $O_j$ . For example, the experiments  $XP_1$  and  $XP_2$  are composed of the same source pool  $\mathcal{S}$  (i.e. IDD, SUIM), the same evaluation measure

$E$  (i.e.  $\tau_w$ ), but different target datasets  $T$  (i.e. ADE20k vs CityScapes). These experiments produce divergent outcomes, as  $\text{GBC} > \text{H-score}$  on  $XP_1$ , whereas  $\text{H-score} > \text{GBC}$  on  $XP_2$ .

We use these definitions to investigate how varying the components of an experimental setup  $(\mathcal{S}, T, E)$  affects the outcome  $O$  across experiments. We note that an experimental setup is stable if small variations of  $\mathcal{S}$ ,  $T$ , or  $E$  produce similar outcomes (and unstable if small variations produces divergent outcomes).

## 4.2 Qualitative analysis of experimental setup stability

We qualitatively compare a large number of experiments by displaying their outcomes in a scatter plot (Fig. 1). Each point represents a single experiment and compares the quality of a transferability metric (*e.g.* H-score) to the quality of another (*e.g.* GBC) as assessed by a fixed evaluation measure  $E$  ( $\tau_w$  in this example). Each experiment is colored according to its target dataset  $T$ . The dotted red line represents  $x = y$ ; for points on this line both transferability metrics are equally good. If points are either all above or below the line, this means that one transferability metric consistently outperforms the other, suggesting the experiment is stable.

We can study how variations in the source pool  $\mathcal{S}$  influence the outcome by looking at the distribution of points of the same color (i.e. on the same target dataset  $T$ ). In Fig. 1, in all experiments where the target dataset is CityScapes (magenta), GBC consistently outperforms H-score. Hence this experiment is stable with respect to variations in the source pool. To study the influence of varying the target dataset  $T$ , we can compare whole point clouds (in different colors). In this case, for  $T = \text{ADE20k}$  H-score consistently outperforms GBC, which contradicts with what found for  $T = \text{CityScapes}$ . Hence this experiment setup is unstable when varying the target dataset. Finally, we can compare the influence of varying the evaluation measure  $E$  by comparing two scatter plots which differ only in the choice of  $E$  (*e.g.* Fig. 3b vs Fig. 5b).

## 4.3 Quantitative measures of experimental setup stability

**Setup Stability (SS).** We want to quantify the overall effect of varying exactly a single component  $\mathcal{S}$ ,  $T$  or  $E$  of an experimental setup. To do so, we consider pairs of experiments which differ in only one component. Each edge in Fig. 1 (right) connects two experiments and is colored based on which component is different between them. Experiments which differ in two or more components are not connected.

We now detail how we measure the Setup Stability for a single component (*e.g.*  $T$  in green). First, for each pair of experiments which differ only in  $T$ , we calculate an agreement score between their outcomes: the rank correlation  $\tau$  between the two lists of transferability metrics as ranked by their quality. For example, the agreement score between  $XP_1$  and  $XP_2$  in our example is 0.2 (middle of Fig. 1). Then we average agreement scores over all experiment pairs which differ only in  $T$  (i.e. green edges in Fig. 1, right side). We call this

|                                      |               | Ours | LEEP | $\mathcal{N}$ LEEP | LogME | GBC  | OTCE |
|--------------------------------------|---------------|------|------|--------------------|-------|------|------|
|                                      |               |      | [57] | [49]               | [89]  | [59] | [75] |
| <i>Selecting</i>                     | #sources      | 17   | 0    | 16                 | 0     | 17   | 4    |
| <i>source datasets</i> (Sec.5)       | #targets      | 17   | 0    | 4                  | 0     | 17   | 5    |
| <i>Selecting</i>                     | #sources      | 19   | 9    | 13                 | 24    | 9    | 0    |
| <i>source architectures</i> (Sec.6)  | #targets      | 9    | 1    | 4                  | 17    | 8    | 0    |
| <i>Evaluation measures</i> (Sec.3.3) | # $E$         | 4    | 1    | 6                  | 1     | 2    | 1    |
| <i>Constructing XPs</i> (Sec.4.4)    | # experiments | 715k | 1    | 48                 | 19    | 50   | 500  |

Table 1: Volume of experiments for evaluating transferability metrics in various papers. We systematically construct 715k experiments by varying the source pool  $\mathcal{S}$ , target dataset  $T$  and evaluation measure  $E$ . In contrast, previous works have at most 500 experiments. All stats are about the scenario of selecting source models for a given target. Some papers [57,6,75,59,49] also report the converse case (determining which target task benefits most from a given source model). Even when including it, no previous work reports more than 508 experiments [75]. The OTCE paper [75] constructs 500 experiments by randomly sampling 100 target tasks within each of the 5 target datasets used. We note that the H-score paper [6] only reports in this converse scenario (and only 1 experiment).

aggregated measure the *Setup Stability* (SS). It is 1 when all experiments produce the same outcome, and indicates a stable setup. In contrast, a score of 0 implies there is no agreement at all across the experiment outcomes (very unstable). In our example in Fig. 1, the SS for  $T$  equals to  $\frac{1}{4} (.2 + .1 + .2 + .3) = 0.20$ .

**Win rate.** We introduce the *win rate* to assess whether a particular transferability metric is the best across many experiments. The win rate of a transferability metric is the percentage of experiments where it is the best metric (i.e. outperforms all other metrics in that experiment). We report the win rate for each of the four evaluation measures  $E$  individually, which provides another view of its influence. We also consider the win rate across all experiments, as the final, most aggregated assessment of a transferability metric.

#### 4.4 Constructing 715k experiments

We need to construct a set of experiments which vary along each of the three components  $\mathcal{S}$ ,  $T$ , and  $E$ . This set needs to be extremely large to properly investigate the influence of each component and to be able to reduce experimental uncertainty when deciding which transferability metrics work best. In Sec. 5 and 6 we will study in-depth two scenarios about selecting good source models for the tasks of semantic segmentation and image classification. In total we construct *715k experiments*. This is several orders of magnitude more than other works on transferability metrics, which all have at most 500 experiments when selecting source models (Tab. 1).

## 5 Scenario 1: selecting a good source dataset in semantic segmentation

**Experimental setup.** We consider the scenario of selecting good source datasets to transfer from on the task of semantic segmentation. We use a total of 17 datasets from a wide range of image domains: Pascal Context [52], Pascal VOC [27], ADE20K [93], COCO [12,50,45], KITTI [3], CamVid [10], CityScapes [20], IDD [78], BDD [90], MVD [54], ISPRS [65], iSAID [81,85], SUN RGB-D [73], ScanNet [21], SUIM [41], vKITTI2 [11,28] and vGallery [82]. We use a fixed model architecture composed of a HRNetV2-W48 backbone [80] with a linear classifier on top, as this model achieves state-of-the-art performances on dense prediction tasks [48,51,80]. We create a total of 17 source models by training HRNetV2-W48 on the full training set of each source dataset.

Now we create many setup variations. First, each dataset plays the role of target, in turn. For each target the other 16 datasets are sources, leading to 16 source models. To construct many different source pools  $\mathcal{S}$ , we consider all combinations of 11 out of 16 source models, resulting in  $\binom{16}{11}$  source pools. Finally, we consider four evaluation measures  $E$ . Hence we have a total of  $17 \times \binom{16}{11} \times 4 = 297k$  experiments, spanning variations across source pool  $\mathcal{S}$ , target dataset  $T$  and evaluation measure  $E$ . Importantly, while we consider a very large space of setup variations, this only requires training 17 source models and fine-tuning each 16 times, which is computationally efficient.

When fine-tuning on the target training set, we follow the low-shot setup of [51]: we limit it to 150 images for each dataset (except COCO and ADE20k, on which we use 1000 images as they contain a large number of classes). We follow [2] and apply the transferability metrics to semantic segmentation by subsampling 1000 pixels per image. As in [2], we sample pixels inversely proportionally to the frequency of their class labels in the target dataset.

**Qualitative analysis.** We perform here the qualitative analysis described in Sec.4.2. We fix the evaluation measure  $E$  to be  $\tau_w$  and visualize experiments with variations of source pool  $\mathcal{S}$  and target dataset  $T$ . Fig. 2 shows comparisons for all possible pairs of transferability metrics. If we compare LogME to LEEP (Fig.2-left), the experiments evaluated on the target dataset  $T = \text{iSAID}$  (red) form a group below the red line, indicating LEEP consistently outperforms LogME. Conversely, for both  $T = \text{COCO}$  (brown) and  $T = \text{SUIM}$  (pink), LogME consistently outperforms LEEP. This shows the experimental instability w.r.t. variations of  $T$ . If we look at  $T = \text{ADE}$  (light blue), the experiments are centered on the red line and stray quite far from it on *both sides*. This means that contradicting outcomes are found also when varying the source pool  $\mathcal{S}$ . The same observation holds for  $T = \text{vGallery}$  (purple).

In fact, we can find many cases where the outcome of an experiment is wildly different even when varying only a single setup component (Fig. 3). In Fig. 3a, we see that for  $T = \text{ADE20K}$  H-score outperforms LogME, whereas for  $T = \text{CityScapes}$  the opposite is true. Next, Fig. 3b show two examples where changing only the evaluation measure  $E$  leads to conflicting outcomes.







Fig. 2: Each point reports the quality of two transferability metrics in terms of  $\tau_w$  (within a single experiment). In each plot we vary the source pool  $\mathcal{S}$  and target dataset  $T$ . Points on the red line mean that the two transferability metrics have equal quality. For some comparisons the setup is not stable w.r.t  $T$  and  $\mathcal{S}$ , as different points fall above and below the line, hence they produce divergent outcomes.



(a) Effects of varying the target dataset  $T$ .



(b) Effects of varying the evaluation measure  $E$ .



(c) Effects of varying the source pool  $\mathcal{S}$ .

Fig. 3: Examples of how changing experimental setup components produces different winning metrics across experiments.

|                   | SS score |                | %W $\tau_w$ | %W $\tau$ | %W $p$    | %W $Rel@1$ | Avg       |
|-------------------|----------|----------------|-------------|-----------|-----------|------------|-----------|
| $\neq \mathbf{T}$ | 0.23     | <b>GBC</b>     | 20          | 16        | 10        | 5          | 13        |
| $\neq \mathbf{E}$ | 0.53     | <b>LEEP</b>    | 11          | 13        | 11        | 2          | 9         |
| $\neq \mathbf{S}$ | 0.60     | <b>NLEEP</b>   | 7           | 3         | 5         | 2          | 4         |
|                   |          | <b>LogME</b>   | <b>52</b>   | <b>55</b> | <b>68</b> | <b>49</b>  | <b>56</b> |
|                   |          | <b>H-score</b> | 10          | 13        | 6         | 42         | 18        |

(a)
(b)

Table 2: Aggregate quantitative analysis for Sec. 5. (a) *Setup Stability* (SS) score for variations of source pool  $\mathcal{S}$ , target dataset  $T$  or evaluation measure  $E$ . (b) *Win rate* (%W) for each transferability metric and evaluation measure. LogME is on average the best metric.

On  $T = \text{COCO}$ , for  $E = \tau_w$  NLEEP is better than H-score (most points are above the line), whereas for  $E = \rho$  the opposite is true (large dense cluster below the line). A similar contradicting result is observed on  $T = \text{Pascal VOC}$  for LEEP and GBC. Finally, we look at varying the source pool  $\mathcal{S}$  in Fig. 3c. We see that in each of these example plots the points are scattered on both sides of the  $x = y$  line, indicating contradicting outcomes within each plot.

To conclude, we observed that varying even just a single component of an experimental setup can produce conflicting results. This means it is important to consider variations of all components when designing experiments to evaluate transferability metrics.

**Quantitative analysis.** We now apply the *Setup Stability* (SS) score defined in Sec. 4.3. Tab. 2a shows results when varying each component of the experimental setup, aggregating results over our experiments. Variations of the target dataset  $T$  are the main factor of instability: with a low Setup Stability (SS) score of 0.23, most experiments have different outcomes as to which transferability metric performs better. Results are more stable to variations in the evaluation measure  $E$  (0.53 SS), and most stable to variations in the source pool  $\mathcal{S}$  (0.60 SS).

Finally, we compute the *win rate* of each transferability metric (Sec. 4.3). Results split by evaluation measure are shown in Tab. 2b. LogME outperforms the other metrics on average 56% of the times. In contrast, all other metrics have much lower winning rates (4-18%). Hence, by aggregating over all variations across setup components, we can reduce the experimental uncertainty and reveal that, on average, LogME is the best transferability metric for selecting source datasets. In contrast, when evaluating on a limited set of experiments as in previous works (Tab. 1), results vary wildly and lead to contradictory conclusions (Fig. 2 and Fig. 3).



## 6 Scenario 2: selecting a good source model architecture in image classification

**Experimental setup.** We consider the scenario of selecting good source model architectures to transfer from on the task of image classification, as in [57,49,89,59]. We fix the source dataset to ImageNet. The goal is to select which architecture would lead to the best transfer learning results on a given target dataset. We use 9 target datasets: CIFAR10 & 100 [46], Imagenette [38], Oxford IIIT Pets [60], Caltech-USCD Birds 2011 [79], Stanford Dogs [44], Oxford Flowers 102 [58], SUN-397 [87], and DTD [18]. We consider a total of 19 source architectures: ResNet-50 & ResNet-101 [32], ResNetV2-50, ResNetV2-101 & ResNetV2-152 [33], DenseNet-121, DenseNet-169 & DenseNet-201 [39], MobileNet [37], MobileNetV2 [67], MobileNetV3 [36], EfficientNetB0, EfficientNetB1, EfficientNetB2 & EfficientNetB3 [74], NASNet Mobile [95], VGG16 & VGG19 [70], Xception [16]. We train 19 source models, each with a *different* model architecture but all trained on the *same* source dataset ImageNet [66].

Now we create many setup variations. To construct many different source pools, we consider all combinations of 14 out of 19 source models, resulting in  $\binom{19}{14}$  source pools. Finally, we consider four evaluation measures  $E$ , and the 9 target datasets above. Hence we have a total of  $9 \times \binom{19}{14} \times 4 = 418k$  experiments.

**Qualitative analysis.** We perform the qualitative analysis described in Sec. 4.2. We fix the evaluation measure  $E$  to be  $\rho$  and we visualize experiments with variations of source pool  $\mathcal{S}$  and target dataset  $T$ . Fig. 4 shows comparisons for all possible pairs of transferability metrics. Similar to Sec. 5 we observe that many experiments within the same plot fall on either sides of the red line with a large spread. Hence outcomes are not stable. In particular, on the left-side plot, for  $T = \text{Oxford Pets}$  (pink) GBC consistently outperforms LogME, whereas for  $T = \text{Imagenette}$  (purple) LogME consistently outperforms GBC.

As before, there exists multiple examples where changing a single experimental setup component is sufficient to produce divergent outcomes across experiments (Fig. 5). For example, in Fig. 5a for  $T = \text{CIFAR10}$ , LogME outperforms GBC, while for  $T = \text{Caltech Birds}$  the opposite is true. Moreover, in Fig. 5b we can clearly see the mass of the points move when changing  $E$ . For example, on the left LogME performs about equally to GBC when using  $E = \tau_w$ , whereas GBC is the best when using  $E = \rho$ . Finally, when varying only the source pool  $\mathcal{S}$  in Fig. 5c we see points scattered on both sides of the line within each individual plot. Hence the choice of each experimental setup component can have a large effect on the outcome.

**Quantitative analysis.** We now quantify the influence of varying experimental setup components by computing the *Setup Stability* (SS) score defined in Sec. 4.3. Results are shown in Tab. 3a. Varying the target dataset has the biggest impact (0.62 SS) on the instability of the experimental outcomes, followed by the evaluation measure (0.69 SS) and the source pool (0.80 SS). This is in line with the observations in Sec. 5, except that the overall scores are higher, indicating that the outcomes across experiments are broadly more consistent.









Fig. 4: Each point reports the quality of two transferability metrics in terms of  $\rho$  (within a single experiment). In each plot we vary the source pool  $\mathcal{S}$  and target dataset  $T$ . Points on the red line mean that the two transferability metrics have equal quality. For some comparisons the setup is not stable w.r.t  $T$  and  $\mathcal{S}$ , as different points fall above and below the line, hence they produce divergent outcomes (i.e LEEP vs H-score and GBC vs LogME).



(a) Effects of varying the target dataset  $T$ .



(b) Effects of varying the evaluation measure  $E$ .



(c) Effects of varying the source pool  $\mathcal{S}$ .

Fig. 5: Examples of how changing a single experimental setup component produces different winning metrics across experiments.

|                   | SS score |                | %W $\tau_w$ | %W $\tau$ | %W $p$    | %W $Rel@1$ | Avg       |
|-------------------|----------|----------------|-------------|-----------|-----------|------------|-----------|
| $\neq \mathbf{T}$ | 0.62     | <b>GBC</b>     | 12          | 4         | 5         | 4          | 6         |
| $\neq \mathbf{E}$ | 0.69     | <b>LEEP</b>    | 0           | 0         | 0         | 2          | 1         |
| $\neq \mathbf{S}$ | 0.80     | <b>NLEEP</b>   | <b>80</b>   | <b>91</b> | <b>93</b> | 30         | <b>73</b> |
|                   |          | <b>LogME</b>   | 8           | 5         | 2         | <b>53</b>  | 17        |
|                   |          | <b>H-score</b> | 0           | 0         | 0         | 11         | 3         |

(a)
(b)

Table 3: Aggregate quantitative analysis for Sec. 6. (a) *Setup Stability* (SS) score for variations of  $\mathcal{S}$ ,  $T$  &  $E$ . (b) *Win rate* (%W) for each transferability metric and evaluation measure. NLEEP is on average the best metric.

Tab.3b shows the *win rate* of each transferability metric against all others, for each evaluation measure. NLEEP outperforms on average the other metrics in 73% of the experiments, while LEEP and H-score performs poorly (1 and 3% winning rates, respectively). Hence, by aggregating over many experiments we reveal a clear winner transferability metric in this scenario: NLEEP. Moreover, LogME can be considered a good second place since it is the best for  $E = Rel@1$ .

Importantly, the conclusion in this scenario differs substantially from the previous scenario: In Sec. 5, LogME was the best while NLEEP was the worst. Now, NLEEP is the best while LogME is a good second. Hence even though we aggregate over many experiments, we still find different conclusions for even larger changes in our setup (*scenarios*): changing both the nature of the source pool (i.e. different source datasets vs different source model architectures) and the nature of the task (semantic segmentation vs image classification).

## 7 Conclusion

Our work investigates for the first time the stability of experiments to evaluate transferability metrics. We base our investigation on a systematic analysis of over 715k experiments. We show that even small variations across experiment setups leads to divergent conclusions about the superiority of a transferability metric over another. Then we improve stability by aggregating outcomes across many experiments. As a result, we reveal the superiority of LogME at selecting good source datasets to transfer from, and NLEEP at selecting good source architectures. Finally, in the third scenario, we show that GBC is best for determining which target task benefits most from a given source model (see supp. mat.). Hence, no single transferability metric works best in all scenarios.

In this work all scenarios feature source and target tasks of identical types (i.e., both segmentation, or both classification). However, the proposed experimental protocols and analysis methods are not restricted to this setting. Hence, we leave to future work investigations of the experimental stability of transferability metrics when source and target task types differ.

## References

1. Abnar, S., Dehghani, M., Neyshabur, B., Sedghi, H.: Exploring the limits of large scale pre-training. In: ICLR (2022)
2. Agostinelli, A., Uijlings, J., Mensink, T., Ferrari, V.: Transferability metrics for selecting source model ensembles. In: CVPR (2021)
3. Alhaija, H., Mustikovela, S., Mescheder, L., Geiger, A., Rother, C.: Augmented reality meets computer vision : Efficient data generation for urban driving scenes. IJCV (2018)
4. Alvarez-Melis, D., Fusi, N.: Geometric dataset distances via optimal transport. In: NeurIPS (2020)
5. Azizpour, H., Razavian, A.S., Sullivan, J., Maki, A., Carlsson, S.: Factors of transferability for a generic convnet representation. IEEE Trans. on PAMI (2015)
6. Bao, Y., Li, Y., Huang, S.L., Zhang, L., Zheng, L., Zamir, A., Guibas, L.: An information-theoretic approach to transferability in task transfer learning. In: ICIP (2019)
7. Benesty, J., Chen, J., Huang, Y., Cohen, I.: Pearson correlation coefficient. In: Noise reduction in speech processing. Springer (2009)
8. Bolya, D., Mittapalli, R., Hoffman, J.: Scalable diverse model selection for accessible transfer learning. In: NeurIPS (2021)
9. Bozinovski, S., Fulgosi, A.: The influence of pattern similarity and transfer of learning upon training of a base perceptron b2 (original in croatian). In: Proceedings of the Symposium Informatica. pp. 3–121 (1976)
10. Brostow, G.J., Fauqueur, J., Cipolla, R.: Semantic object classes in video: A high-definition ground truth database. Patt. Rec. Letters **30**(2), 88–97 (2009)
11. Cabon, Y., Murray, N., Humenberger, M.: Virtual kitti 2. arXiv (2020)
12. Caesar, H., Uijlings, J., Ferrari, V.: COCO-Stuff: Thing and stuff classes in context. In: CVPR (2018)
13. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE Trans. on PAMI (2018)
14. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: ICML (2020)
15. Chen, X., He, K.: Exploring simple siamese representation learning. In: CVPR (2021)
16. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: CVPR (2017)
17. Chu, B., Madhavan, V., Beijbom, O., Hoffman, J., Darrell, T.: Best practices for fine-tuning visual classifiers to new domains. In: ECCV (2016)
18. Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., Vedaldi, A.: Describing textures in the wild. In: CVPR (2014)
19. Colas, C., Sigaud, O., Oudeyer, P.Y.: How many random seeds? statistical power analysis in deep reinforcement learning experiments. arXiv preprint arXiv:1806.08295 (2018)
20. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: CVPR (2016)
21. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In: CVPR (2017)

22. Ding, R., Li, X., Nie, L., Li, J., Si, X., Chu, D., Liu, G., Zhan, D.: Empirical study and improvement on deep transfer learning for human activity recognition. *Sensors* **19**(1), 57 (2019)
23. Dwivedi, K., Huang, J., Cichy, R.M., Roig, G.: Duality diagram similarity: A generic framework for initialization selection in task transfer learning. In: *ECCV* (2020)
24. Dwivedi, K., Roig, G.: Representation similarity analysis for efficient task taxonomy and transfer learning. In: *CVPR* (2019)
25. Dwivedi, V.P., Joshi, C.K., Laurent, T., Bengio, Y., Bresson, X.: Benchmarking graph neural networks. *arXiv preprint arXiv:2003.00982* (2020)
26. Errica, F., Podda, M., Bacciu, D., Micheli, A.: A fair comparison of graph neural networks for graph classification. *arXiv preprint arXiv:1912.09893* (2019)
27. Everingham, M., Van Gool, L., Williams, C.K.I., Winn, J., Zisserman, A.: The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. <http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html> (2012)
28. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: *CVPR* (2016)
29. Girshick, R.: Fast R-CNN. In: *ICCV* (2015)
30. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *CVPR* (2020)
31. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: *ICCV* (2017)
32. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR* (2016)
33. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: *ECCV* (2016)
34. Hoiem, D., Chodpathumwan, Y., Dai, Q.: Diagnosing error in object detectors. In: *ECCV* (2012)
35. Hosang, J., Benenson, R., Dollár, P., Schiele, B.: What makes for effective detection proposals? *IEEE Trans. on PAMI* **38**(4), 814–830 (2015)
36. Howard, A., Sandler, M., Chu, G., Chen, L.C., Chen, B., Tan, M., Wang, W., Zhu, Y., Pang, R., Vasudevan, V., et al.: Searching for mobilenetv3. In: *ICCV* (2019)
37. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *Tech. rep.*, *arXiv* (2017)
38. Howard, J.: Imagenette. [github.com/fastai/imagenette/](https://github.com/fastai/imagenette/)
39. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *CVPR* (2017)
40. Huh, M., Agrawal, P., Efros, A.: What makes imagenet good for transfer learning? In: *NeurIPS LSCVS workshop* (2016)
41. Islam, M.J., Edge, C., Xiao, Y., Luo, P., Mehtaz, M., Morse, C., Enan, S.S., Sattar, J.: Semantic Segmentation of Underwater Imagery: Dataset and Benchmark. In: *IROS* (2020)
42. Jordan, S., Chandak, Y., Cohen, D., Zhang, M., Thomas, P.: Evaluating the performance of reinforcement learning algorithms. In: *ICML* (2020)
43. Kendall, M.G.: A new measure of rank correlation. *Biometrika* **30**(1/2), 81–93 (1938)
44. Khosla, A., Jayadevaprakash, N., Yao, B., Fei-Fei, L.: Novel dataset for fine-grained image categorization. In: *CVPR Workshops* (2011)
45. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: *CVPR* (2019)

46. Krizhevsky, A.: Learning multiple layers of features from tiny images. Tech. rep., University of Toronto (2009)
47. Kuznetsova, A., Rom, H., Alldrin, N., Uijlings, J., Krasin, I., Pont-Tuset, J., Kamali, S., Popov, S., Mallocci, M., Kolesnikov, A., Duerig, T., Ferrari, V.: The open images dataset v4. *IJCV* **128**(7), 1956–1981 (2020)
48. Lambert, J., Liu, Z., Sener, O., Hays, J., Koltun, V.: MSeg: A composite dataset for multi-domain semantic segmentation. In: *CVPR* (2020)
49. Li, Y., Jia, X., Sang, R., Zhu, Y., Green, B., Wang, L., Gong, B.: Ranking neural checkpoints. In: *CVPR* (2021)
50. Lin, T.Y., Maire, M., Belongie, S., Bourdev, L., Girshick, R., Hays, J., Perona, P., Ramanan, D., Zitnick, C.L., Dollár, P.: Microsoft COCO: Common objects in context. In: *ECCV* (2014)
51. Mensink, T., Uijlings, J., Kuznetsova, A., Gygli, M., Ferrari, V.: Factors of influence for transfer learning across diverse appearance domains and task types. *IEEE Trans. on PAMI* (2021)
52. Mottaghi, R., Chen, X., Liu, X., Cho, N.G., Lee, S.W., Fidler, S., Urtasun, R., Yuille, A.: The role of context for object detection and semantic segmentation in the wild. In: *CVPR* (2014)
53. Mustafa, B., Loh, A., Freyberg, J., MacWilliams, P., Wilson, M., McKinney, S.M., Sieniek, M., Winkens, J., Liu, Y., Bui, P., et al.: Supervised transfer learning at scale for medical imaging. *arXiv preprint arXiv:2101.05913* (2021)
54. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: *ICCV* (2017)
55. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: *ECCV*. Springer (2016)
56. Ngiam, J., Peng, D., Vasudevan, V., Kornblith, S., Le, Q.V., Pang, R.: Domain adaptive transfer learning with specialist models. In: *arXiv* (2018)
57. Nguyen, C., Hassner, T., Seeger, M., Archambeau, C.: LEEP: A new measure to evaluate transferability of learned representations. In: *ICML* (2020)
58. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: *Indian Conf. on CVGIP* (2008)
59. Pándy, M., Agostinelli, A., Uijlings, J., Ferrari, V., Mensink, T.: Transferability estimation using bhattacharyya class separability. In: *CVPR* (2022)
60. Parkhi, O.M., Vedaldi, A., Zisserman, A., Jawahar, C.: Cats and dogs. In: *CVPR* (2012)
61. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825–2830 (2011)
62. Pratt, L.Y.: Discriminability-based transfer between neural networks. In: *NeurIPS* (1992)
63. Ridnik, T., Ben-Baruch, E., Noy, A., Zelnik-Manor, L.: Imagenet-21k pretraining for the masses. *arXiv preprint arXiv:2104.10972* (2021)
64. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: *MICCAI* (2015)
65. Rottensteiner, F., Sohn, G., Gerke, M., Wegner, J.D., Breikopf, U., Jung, J.: Results of the isprs benchmark on urban object detection and 3d building reconstruction. *ISPRS journal of photogrammetry and remote sensing* (2014)
66. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., Berg, A., Fei-Fei, L.: ImageNet large scale visual recognition challenge. *IJCV* (2015)

67. Sandler, M., Howard, A.G., Zhu, M., Zhmoginov, A., Chen, L.: Mobilenetv2: Inverted residuals and linear bottleneck. In: CVPR (2018)
68. Shchur, O., Mumme, M., Bojchevski, A., Günnemann, S.: Pitfalls of graph neural network evaluation. arXiv preprint arXiv:1811.05868 (2018)
69. Shelhamer, E., Long, J., Darrell, T.: Fully convolutional networks for semantic segmentation. IEEE Trans. on PAMI (2016)
70. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: ICLR (2015)
71. Song, J., Chen, Y., Wang, X., Shen, C., Song, M.: Deep model transferability from attribution maps. In: NeurIPS (2019)
72. Song, J., Chen, Y., Ye, J., Wang, X., Shen, C., Mao, F., Song, M.: Depara: Deep attribution graph for deep knowledge transferability. In: CVPR (2020)
73. Song, S., Lichtenberg, S., Xiao, J.: SUN RGB-D: A RGB-D scene understanding benchmark suite. In: CVPR (2015)
74. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: ICML (2019)
75. Tan, Y., Li, Y., Huang, S.L.: OTCE: A transferability metric for cross-domain cross-task representations. In: CVPR (2021)
76. Thrun, S., Pratt, L.: Learning to learn. Springer Science & Business Media (2012)
77. Tran, A.T., Nguyen, C.V., Hassner, T.: Transferability and hardness of supervised classification tasks. In: CVPR (2019)
78. Varma, G., Subramanian, A., Namboodiri, A., Chandraker, M., Jawahar, C.: Idd: A dataset for exploring problems of autonomous navigation in unconstrained environments. In: Proc. WACV (2019)
79. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The Caltech-UCSD Birds-200-2011 Dataset. Tech. Rep. CNS-TR-2011-001, California Institute of Technology (2011)
80. Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., Xiao, B.: Deep high-resolution representation learning for visual recognition. IEEE Trans. on PAMI (2020)
81. Waqas Zamir, S., Arora, A., Gupta, A., Khan, S., Sun, G., Shahbaz Khan, F., Zhu, F., Shao, L., Xia, G.S., Bai, X.: isaid: A large-scale dataset for instance segmentation in aerial images. In: CVPR Workshops (2019)
82. Weinzaepfel, P., Csurka, G., Cabon, Y., Humenberger, M.: Visual localization by learning objects-of-interest dense match regression. In: CVPR (2019)
83. Weiss, K., Khoshgoftaar, T.M., Wang, D.: A survey of transfer learning. Journal of Big data **3**(1), 1–40 (2016)
84. Whiteson, S., Littman, M.L.: Introduction to the special issue on empirical evaluations in reinforcement learning. Machine learning **84**(1-2), 1 (2011)
85. Xia, G.S., Bai, X., Ding, J., Zhu, Z., Belongie, S., Luo, J., Datcu, M., Pelillo, M., Zhang, L.: Dots: A large-scale dataset for object detection in aerial images. In: CVPR (2018)
86. Xian, Y., Lampert, C.H., Schiele, B., Akata, Z.: Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. IEEE Trans. on PAMI **41**(9), 2251–2265 (2019)
87. Xiao, J., Hays, J., Ehinger, K., Oliva, A., Torralba, A.: SUN database: Large-scale scene recognition from Abbey to Zoo. In: CVPR (2010)
88. Yan, X., Acuna, D., Fidler, S.: Neural data server: A large-scale search engine for transfer learning data. In: CVPR (2020)
89. You, K., Liu, Y., Wang, J., Long, M.: LogME: Practical assessment of pre-trained models for transfer learning. In: ICML (2021)

90. Yu, F., Chen, H., Wang, X., Xian, W., Chen, Y., Liu, F., Madhavan, V., Darrell, T.: Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In: CVPR (2020)
91. Zamir, A.R., Sax, A., Shen, W., Guibas, L.J., Malik, J., Savarese, S.: Taskonomy: Disentangling task transfer learning. In: CVPR (2018)
92. Zendel, O., Murschitz, M., Humenberger, M., Herzner, W.: How good is my test data? introducing safety analysis for computer vision. IJCV **125**(1), 95–109 (2017)
93. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ADE20K dataset. In: CVPR (2017)
94. Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. In: arXiv (2019)
95. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: CVPR (2018)