

Data-Free Sketch-Based Image Retrieval

Abhra Chaudhuri
 University of Exeter, UK
 ac1151@exeter.ac.uk

Ayan Kumar Bhunia, Yi-Zhe Song, Anjan Dutta
 Institute for People-Centred AI, University of Surrey, UK
 {a.bhunias, y.song, anjan.dutta}@surrey.ac.uk

Abstract

Rising concerns about privacy and anonymity preservation of deep learning models have facilitated research in data-free learning (DFL). For the first time, we identify that for data-scarce tasks like Sketch-Based Image Retrieval (SBIR), where the difficulty in acquiring paired photos and hand-drawn sketches limits data-dependent cross-modal learning algorithms, DFL can prove to be a much more practical paradigm. We thus propose Data-Free (DF)-SBIR, where, unlike existing DFL problems, pre-trained, single-modality classification models have to be leveraged to learn a cross-modal metric-space for retrieval without access to any training data. The widespread availability of pre-trained classification models, along with the difficulty in acquiring paired photo-sketch datasets for SBIR justify the practicality of this setting. We present a methodology for DF-SBIR, which can leverage knowledge from models independently trained to perform classification on photos and sketches. We evaluate our model on the Sketchy, TU-Berlin, and QuickDraw benchmarks, designing a variety of baselines based on state-of-the-art DFL literature, and observe that our method surpasses all of them by significant margins. Our method also achieves mAPs competitive with data-dependent approaches, all the while requiring no training data. Implementation is available at <https://github.com/abhrac/data-free-sbir>.

1. Introduction

Motivated by the high degree of expressiveness and flexibility provided by sketches, sketch-based image retrieval (SBIR) has emerged as a popular area of computer vision research [3,6,7,13,15]. SBIR is generally achieved by training photo and sketch encoders to respectively map photo and sketch inputs to a class or instance aligned common space. Training deep neural photo-sketch encoders for this task, however, requires datasets with matching photo-sketch pairs [45,58]. Unlike photos, sketches are fundamentally difficult to acquire as drawing them involve long time periods of laborious human participation. Driven by this practi-

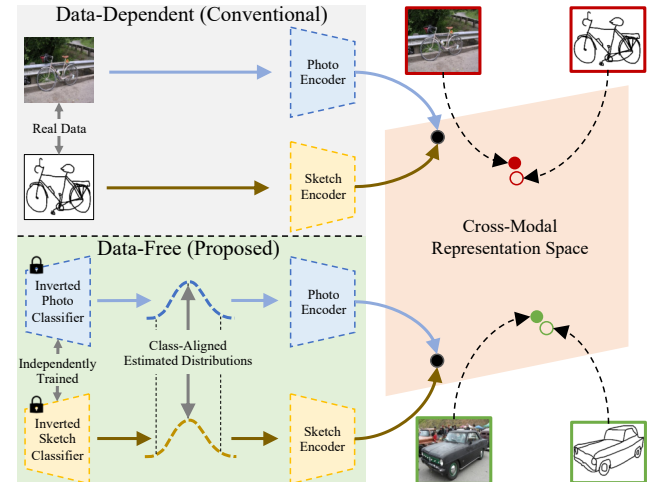


Figure 1. Our proposed Data-Free setting for SBIR does not need a real-world dataset of paired sketches and photos. Using only independently trained, modality-specific classifiers, it can estimate their train set distributions, as well as pair them at class-level for training the sketch and photo encoders.

cal constraint, the problem has been studied under a variety of data-scarce settings like semi-supervised [3], few-shot class incremental [5], zero-shot [14], and any-shot [15]. However, all such settings assume the availability of some amount of instance/class-aligned data for training the encoders. With the tremendous effort involved in acquiring such labelled photo-sketch pairs [3,13,14], as well as the rising concerns about privacy, security and anonymity preservation abilities of deep learning models [8,31,46,51], such assumptions may no longer be practical.

With this view, we propose Data-Free Sketch-Based Image Retrieval (DF-SBIR), a novel setting that requires training photo and sketch encoders for retrieval, but with no training data. Specifically, we only assume access to pre-trained photo and sketch classification models. In contrast to unsupervised cross-domain image retrieval [22] which only requires access to training data, but with no in-domain or cross-domain labels, our setting goes a step further and assumes access to no training data at all. Since classification does not require cross-modal pairings as in SBIR, and owing to the recent advances in domain generalization [29,52],

such pre-trained classifiers are widely available [40, 52], making our setting quite practical. Under this scenario, the problem of training the encoders can be posed in the light of data-free knowledge distillation (DFKD) [8, 30]. Classical knowledge distillation [21] aims to transfer knowledge from a pre-trained model (teacher) to a different model (student), by aligning the predictions of the latter with the former on train set inputs. Differently, DFKD aims to achieve this knowledge transfer without access to any form of training data. The conventional approach to DFKD involves the following two steps – (1) Reconstructing the train set distribution of the teacher; (2) Training the student network to match its predictions with that of the teacher, on samples from the reconstructed distribution. However, existing DFKD approaches so far have only been able to operate in a single modality, performing the same kind of task as that of the teacher. SBIR, being a cross-modal retrieval problem, cannot be tackled in the data-free setting by directly adapting the machineries developed for DFKD, the reasons for which we detail below.

First, the teachers (being classifiers) and the students (being encoders) operate in metric spaces of different nature, *i.e.*, probabilistic and Euclidean respectively. This restrains us from measuring the agreement between teachers and students in a straightforward way, thus preventing the direct application of state-of-the-art approaches from the DFKD literature like data-free adversarial distillation [10, 35]. We address this by designing a unified, class-proxy based interface via which the teachers and students can interact. *Second*, the sketch and the image classifiers that act as teachers are independently trained on modality-specific data. Their intermediate representations are thus modality sensitive. However, the representations learned by the encoders to be used for DF-SBIR need to be modality invariant. To this end, we introduce the concept of a modality guidance network, which constrains the reconstructions to belong to specific (photo/sketch) modalities. Training the encoders with such samples will ensure that they learn to eliminate unnecessary, modality-specific information. *Third*, the independent training of the classifiers also mean that their train set distributions may not have direct class-level correspondence. To address this, we design our distribution estimation process to reconstruct class-aligned samples, *i.e.*, ones that have class-level correspondence across the two modalities. This will guarantee the availability of matching photo-sketch pairs for the metric learning of the encoders. Our approach is hence able to perform Data-Free Learning Across Modalities and Metric-Spaces, which motivates us to abbreviate it as CrossX-DFL.

We make the following contributions – (1) Propose Data-Free SBIR, a novel setting keeping in view the data-scarcity constraints arising from the collection of paired photos and sketches for SBIR, as well as concerns around pri-

vacy preservation; (2) A class-proxy based approach to perform data-free adversarial distillation from teachers with probabilistic outputs to students with outputs in the Euclidean space; (3) A novel technique to reconstruct class-aligned samples across independent modalities for cross-modal data-free knowledge distillation; (4) Introduce the concept of a modality guidance network to constrain the reconstructed sample distributions to specific modalities; (5) Extensive experiments on benchmark datasets and ablation studies that demonstrate the usefulness of our novel components in providing competitive performance relative to the data-dependent setting.

2. Related Work

Data-Free Learning Data-free learning (DFL) [8] aims to solve a learning problem without accessing any form of training data, possibly by leveraging the knowledge already learned by a different model, but for a related task. Such a setting provides practical benefits like data security [8], anonymity [32], privacy preservation [30], and model interpretability [33] to name a few. The standard approach for performing DFL is based on the principle of model inversion [33] – given a pre-trained model (*teacher*), the aim is to reconstruct its train set distribution by analyzing its activation patterns [8, 30, 35, 57]. Instead of real data, this reconstructed distribution is then used to train a different downstream model (*student*), thus making its learning process *data-free*. Such reconstructions can be ensured to have high fidelity and diversity by leveraging the teacher’s batch normalization statistics [56] and inducing contrastive teacher representations [17] respectively. Along these lines, Micaelli *et al.* [35] and Choi *et al.* [10] got a step further and introduced the idea of adversarial data-free distillation, which showed that training the downstream task networks (*students*) specifically on the worst-case inputs (for which its predictions diverged maximally from the teacher) helped ensure robustness to factors of variation [2]. However, by its very design, the teacher and the student have to operate in the exact same representation space (either probabilistic [8] or Euclidean [32]), and its extension to cross-space distillation has remained yet unexplored. In this work, for the first time, we design a methodology that allows students to adversarially learn from a teacher, even when they operate in representation spaces of different nature.

Applications of Data Free Learning: Originally, DFL was proposed with the view of performing knowledge distillation (KD) [21] in a data-free manner [8, 30]. This was motivated by the fact that traditional KD required access to the teacher’s train set [39, 41, 48, 54], while rising concerns about privacy and anonymity preservation in deep learning models brought the practicality of such a requirement under question [8, 32]. However, the potential of the underlying idea has helped DFL expand much beyond the horizons of

KD, with its current applications ranging all the way from incremental person re-ID [32], image super-resolution [61] and federated learning [60], to studying the security robustness of a model via black-box adversarial attacks [51, 59] or model stealing [46]. To address the non-IID nature of local data distributions at the individual clients in the federated learning setting [34], principles of DFL [8] allow local dataset reconstructions at the global level [60]. Going beyond such applications, Lu *et al.* [31] has showed that using online gradient update based reconstructions can be used for exposing privacy preservation gaps in federated learning systems. The latter idea of uncovering a model’s security deficits via DFL has further been studied in the context of black-box adversarial attacks [9] for designing a data-free generator that maximizes information leakage from the target model to the substitute model [59]. The same approach has also been applied in model stealing [46], or for learning a single dynamic substitute network for multiple target models with different downstream tasks [51]. However, none of the above works deal with cross-modal retrieval tasks (like SBIR), where generating paired instances across modalities is a challenging and time-intensive process. We address this issue through our novel construction of semantically-aligned, modality-specific estimators.

Sketch-Based Image Retrieval: SBIR methods relying on hand-crafted features focused on matching geometrically similar substructures via gradient field HOG [23], deformable parts model [27], histogram of edge local orientations [42], learned key shapes [43]. To address the sensitivity of such methods to the large sketch-photo modality gap, deep SBIR models employing a deep triplet network to learn a common embedding space were proposed in [45, 58], which was extended in [47] by the introduction of spatial attention, and in [4, 38] via self-supervised pretext tasks. Generative models have also shown promising results in SBIR, employing ideas like disentangling style and semantic contents [37, 44], cross-domain image synthesis [37], and reinforcement learning based sketch generation [3]. All the above methods rely on expensive continuous valued distance computation, hampering scalability, which [28] addressed by learning binary hash codes for sketches. Explainability [1] and robustness [6, 18] have also become recent areas of focus in SBIR research. The data scarcity issue in SBIR stemming from the difficulty in acquiring hand-drawn sketches has been addressed by studying the problem in semi-supervised [3] and zero-shot settings [13, 14, 26]. However, performing SBIR in a scenario where one has access to absolutely *no* training data, has never been explored before. With the widespread availability of pre-trained classification models on benchmark datasets [40, 52], and the costs and constraints involved in acquiring datasets with paired photo-sketch instances, data-free SBIR emerges as a promising research direction.

3. CrossX-DFL

Notations: Let t^p and t^s be photo and sketch classifiers (teachers) already trained over input distributions \mathcal{X}^p and \mathcal{X}^s respectively. Let g_θ^p and g_θ^s respectively be photo and sketch estimator networks that learn to map random Gaussian noise vectors $\xi \in [0, 1]^n$ to elements in the training domains of the photo and sketch classifiers. We denote by f_ψ^p and f_ψ^s , the photo and sketch encoders (students) that are trained to compute metric space embeddings for photos and sketches respectively, and can thus be used for the task of SBIR. For brevity, we will omit the parameters and the modality superscripts while referring to the networks in generic scenarios (*e.g.*, f_ψ^p will be denoted as f).

Overview: We approach the task of data-free SBIR by designing a bespoke process for performing data-free knowledge distillation from the classifiers to the encoders. The process consists of estimating the train set distribution of the classifiers followed by leveraging that estimation to train downstream photo and sketch encoders for retrieval. The steps are summarized below:

Step 1: Estimating the Input Distribution – We train the estimator networks g to produce approximations of the train set distributions of the classifiers t . Based on feedbacks from the classifiers t and the encoders f , the estimators generate distributions $\tilde{\mathcal{X}}$ such that they are as close to the true sample distributions \mathcal{X} as possible. Since the classifiers act as teacher networks, their weights are kept frozen throughout the process.

Step 2: Training the Encoders – The encoders f are then trained to produce metric space embeddings for the approximated training distributions $\tilde{\mathcal{X}}$ such that matching photo-sketch pairs $(\tilde{\mathbf{x}}^p, \tilde{\mathbf{x}}^s) \in (\tilde{\mathcal{X}}^p, \tilde{\mathcal{X}}^s)$ are mapped close to each other in the representation space, and non-matching pairs are mapped farther apart.

Step 3: Retrieval – Given the trained sketch encoder f_ψ^s and photo encoder f_ψ^p , a gallery of photos $\{\mathbf{x}_1^g, \mathbf{x}_2^g, \dots, \mathbf{x}_k^g\}$ and a query sketch \mathbf{x}^s , we return a ranking over the gallery such that \mathbf{x}_i^g is ranked above \mathbf{x}_j^g iff:

$$f^s(\mathbf{x}^s) \cdot f^p(\mathbf{x}_i^g) \geq f^s(\mathbf{x}^s) \cdot f^p(\mathbf{x}_j^g)$$

Step 3 is generic across all SBIR methods and is common to both data-free and data-driven settings. We further elaborate on Steps-1 and 2, which are specific to the data-free setting, below in Sections 3.1 and 3.2 respectively.

3.1. Estimating the Input Distribution

The primary constraint that the reconstructed samples of the estimated distributions need to satisfy is that their class and modality memberships must be unambiguous. Since the downstream encoders need positively paired photos and sketches for contrastive learning, there needs to be a constraint enforcing the semantic alignment of the two distribu-

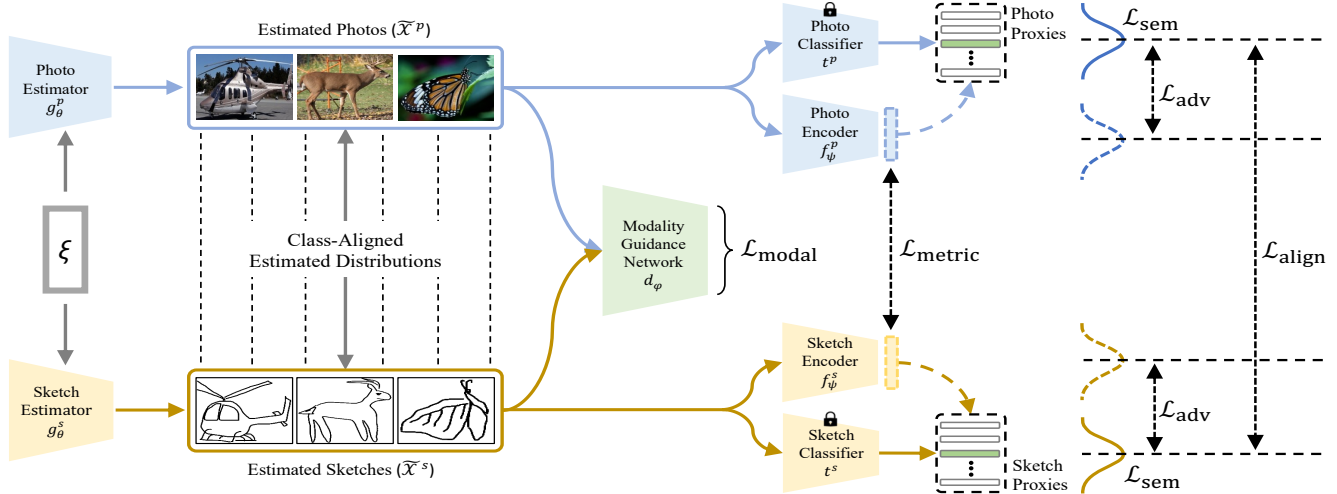


Figure 2. Our CrossX-DFL training pipeline for Data-Free SBIR – The photo and sketch estimators reconstruct the train set distributions of their corresponding classifiers. The estimation process is ensured to have semantic consistency and class-level, instance-wise correspondence via \mathcal{L}_{sem} and $\mathcal{L}_{\text{align}}$, while $\mathcal{L}_{\text{modal}}$ constrains the estimated distributions to belong to specific modalities. \mathcal{L}_{adv} incentivises the estimators to produce boundary samples that are the hardest for the encoders to correctly encode. Optimizing the performance of the encoders on such samples ensures its robustness to semantic variations. The reconstructed samples from the estimated photo and sketch distributions, *i.e.*, $\tilde{\mathcal{X}}^p$ and $\tilde{\mathcal{X}}^s$ respectively, are used for training the encoders in a data-free manner by minimizing $\mathcal{L}_{\text{metric}}$.

tions. One could also ensure the robustness of the encoders by presenting it only with the samples that are the hardest for it to correctly represent. Based on these requirements, we formulate the input distribution estimation of the two classifiers as follows.

Semantic Consistency: To keep the reconstructed elements of the estimated distribution semantically consistent, we constrain the estimation process to minimize the entropy of the classifier’s output distribution, while maximizing the activation values of the representation layer. For a reconstructed input $\tilde{x} = g(\xi)$ and its corresponding predicted probability distribution $\hat{Y} = t(\tilde{x}) = [\hat{y}_1, \hat{y}_2, \dots, \hat{y}_C]$ obtained from the classifier t over a set of C classes, we minimize the following:

$$\mathcal{L}_{\text{sem}} = \sum_{i=1}^C y_i \log \hat{y}_i - |t_{n-1}(\tilde{x})|,$$

where $y_i = 1$ if $\max \hat{Y} = \hat{y}_i$, or 0 otherwise, and $t_{n-1}(\tilde{x})$ is the output of the teacher’s representation layer. Minimizing the classification entropy helps to ensure that the reconstructions \tilde{x} come from distinct classes¹ that the classifier can recognize, while maximizing the magnitude of the classifiers’ representation layer activations incentivizes the estimators to introduce meaningful contents in the reconstructions, as observed in [8].

Class Alignment: Even though \mathcal{L}_{sem} incentivizes meaningful class memberships, it does not guarantee that the cor-

¹We ensure that set of estimated samples are class-balanced via the information entropy criterion as in [8].

responding samples between the two distributions would belong to the same class. Also, pairing samples solely on the basis of discrete class labels does not guarantee exact correspondence. This is because the degree/distribution of class information between the two samples may significantly differ, even if the class label happens to be the same. Hence, we direct the estimation process in a manner such that the corresponding samples across the reconstructed distributions are semantically-aligned. This ensures more fine-grained correspondence between the reconstructed sketch-photo pairs. For a given noise vector ξ , we enforce the photo and the sketch estimators to map it to their corresponding modality specific reconstructions, but in a way such that they represent elements of the same class. We achieve this by minimizing the symmetrized Kullback-Leibler (KL)-divergence between the output distributions of the photo and the sketch classifiers as follows:

$$\mathcal{L}_{\text{align}} = \frac{1}{2} \sum_{i=1}^C y_i^p \log \frac{y_i^p}{y_i^s} + y_i^s \log \frac{y_i^s}{y_i^p},$$

where y_i^p and y_i^s respectively stand for the probabilistic scores predicted by the photo and sketch classifiers for the i -th class.

Modality Guidance: We ensure that the estimators produce samples that are specific to their corresponding modalities by incorporating feedback from a modality guidance network. This guarantees that the reconstructions have sufficient modality-specific information, which the encoder needs to be able to eliminate in order to learn robust representations for retrieval. The modality guidance network

d is a classifier parameterized by ϕ that discriminates between photos and sketches. The modality guidance loss is formulated as follows:

$$\mathcal{L}_{\text{modal}} = -m \log d(\tilde{x}) - (1 - m) \log(1 - d(\tilde{x})),$$

where $m = 1$ if \tilde{x} comes from the photo estimator, and 0 if it comes from the sketch estimator. We independently optimize the weights of the estimators and the modality guidance network to minimize $\mathcal{L}_{\text{modal}}$. The convergence of this criterion will imply that the estimators are able to incorporate modality-specific information into the reconstructions \tilde{x} , since it is the only premise on which d_ϕ can distinguish the two modalities (as they are already semantically-aligned as ensured by minimizing \mathcal{L}_{sem} and $\mathcal{L}_{\text{align}}$).

Metric-Agnostic Adversarial Estimation: Inspired by robust optimization [2], existing literature on data-free knowledge distillation [10,35] has shown that the efficiency of the distillation process can be significantly improved by identifying the samples which are the hardest for the student to classify. Concretely, this objective translates to tasking the estimators to produce boundary samples, *i.e.*, the ones that maximize the *disagreement* between the teachers (classifiers) and the students (encoders). However, in data free SBIR, the teacher being a classifier outputs probability distributions in $[0, 1]^C$, while the student being an encoder produces representation vectors in \mathbb{R}^n , which restrains us from directly measuring the *disagreement* between the teachers and the students in a straightforward way. To address this issue, we treat the final layer neurons of the teacher (classifier) networks as class-proxies [36] against which we perform metric learning with the students’ representations. We thus adapt the proxies to provide an interface between the metric spaces of the teachers and the students, allowing outputs from both networks to be compared with each other. Let $K^p = [\mathbf{k}_1^p, \mathbf{k}_2^p, \dots, \mathbf{k}_C^p]$ and $K^s = [\mathbf{k}_1^s, \mathbf{k}_2^s, \dots, \mathbf{k}_C^s]$ be the sets of class-proxies for the photo and the sketch modalities respectively (generically denoted as K). We align the metric spaces of the classifiers and the encoders by mapping representations obtained from the encoders $f(\tilde{x})$ to the corresponding class proxy vectors $\mathbf{k} \in K$, minimizing the following Proxy Anchor loss [25]:

$$\mathcal{L}_{\text{metric}} = \frac{1}{|K^+|} \sum_{\mathbf{k} \in K^+} \log(1 + \sum_{\tilde{x} \in \tilde{X}_{\mathbf{k}}^+} e^{-\alpha(s(f(\tilde{x}), \mathbf{k}) - \delta)}) + \frac{1}{|K^-|} \sum_{\mathbf{k} \in K^-} \log(1 + \sum_{\tilde{x} \in \tilde{X}_{\mathbf{k}}^-} e^{\alpha(s(f(\tilde{x}), \mathbf{k}) + \delta)})$$

where α and δ respectively denote the margin and scaling factors, K^+ denote the set of positive proxies corresponding to the samples in a minibatch, $\tilde{X}_{\mathbf{k}}^+$ and $\tilde{X}_{\mathbf{k}}^-$ respectively denote the matching and non-matching samples for the proxy \mathbf{k} , $f^p(\tilde{x}) = \bar{\mathbf{0}}$ if \tilde{x} is obtained from the sketch

estimator, and $f^s(\tilde{x}) = \bar{\mathbf{0}}$ if \tilde{x} is obtained from the photo estimator.

The estimators g and encoders f then adversarially optimize the following objective:

$$\mathcal{L}_{\text{adv}} = \min_f \max_g \{D_{\text{KL}}(\text{softmax}(f(g(\xi)) \cdot K), \hat{Y})\},$$

where D_{KL} stands for KL divergence. The maxima of the D_{KL} between the output distributions of the classifier and encoder corresponds to the samples that are the hardest for the encoders to correctly encode. By aiming for this maximum, the estimators learn to generate those hard samples, while the encoders aim to optimize their performance on such samples by getting closer to the classifiers’ predictions.

3.2. Training the Encoders

As ensured by minimizing $\mathcal{L}_{\text{align}}$, the reconstructed distributions \tilde{X}^p and \tilde{X}^s are class-aligned, *i.e.*, for each index i , $\tilde{x}^p = \tilde{X}^p(i)$ and $\tilde{x}^s = \tilde{X}^s(i)$ belong to the same class. With such positive pairs available out of the box, we train the photo and the sketch encoders to minimize the following queue-based version of the InfoNCE loss [49]:

$$\mathcal{L}_{\text{enc}} = -\log \frac{\exp(f^s(\tilde{x}^s) \cdot f^p(\tilde{x}^p)/\tau)}{\sum_{\mathbf{z}^p \in \mathcal{Q}} \exp(f^s(\tilde{x}^s) \cdot \mathbf{z}^p/\tau)},$$

where \mathcal{Q} is a gradient-free queue of photo sample representations from previous mini-batches and τ is a hyperparameter controlling the spread of the embedding distribution [19,53].

Learning objective: We optimize the estimators g , encoders f , and the modality discriminator d parameterized by θ, ψ , and ϕ respectively via the following gradient updates:

$$\begin{aligned} \theta &\leftarrow \theta - \eta \nabla_{\theta} (\mathcal{L}_{\text{sem}} + \mathcal{L}_{\text{align}} + \mathcal{L}_{\text{modal}} + \mathcal{L}_{\text{adv}}), \\ \psi &\leftarrow \psi - \eta \nabla_{\psi} (\mathcal{L}_{\text{metric}} + \mathcal{L}_{\text{adv}} + \mathcal{L}_{\text{enc}}), \\ \phi &\leftarrow \phi - \eta \nabla_{\phi} \mathcal{L}_{\text{modal}}, \end{aligned}$$

where η is the learning rate. Upon convergence, we only retain the encoders f_ψ for performing retrieval.

4. Experiments

Datasets: We evaluate our model on the three most common large scale SBIR datasets, namely Sketchy [28,45], TU-Berlin [16,28] and QuickDraw-Extended [13]. We provide details on their statistics, as well as train-test splits in the supplementary. For Sketchy, the sketches and photos have instance-level correspondences. However, TU-Berlin [16] and QuickDraw [24] were originally sketch-only datasets, which were extended for SBIR in [28] and

Method	Sketchy		TU-Berlin		QuickDraw	
	mAP@all	Prec@200	mAP@all	Prec@200	mAP@all	Prec@200
Classifier Only [20]	0.530	0.542	0.330	0.338	0.160	0.180
Uni-Modal Distillation [8]	0.529	0.537	0.291	0.295	0.130	0.140
Gaussian Prior [8]	0.365	0.391	0.110	0.126	0.080	0.110
Averaging Weights [52]	0.625	0.630	0.450	0.473	0.300	0.320
Meta-Data [30]	0.573	0.576	0.380	0.395	0.200	0.221
Alternative Data [12,50]	0.656	0.680	0.510	0.530	0.290	0.330
Ours (CrossX-DFL)	0.827	0.831	0.680	0.693	0.400	0.410

Table 1. Quantitative comparison of CrossX-DFL with baselines. [13] by adding photos from ImageNet [12] and Flickr respectively, which only provide class-level, but *no* instance-level correspondence. Therefore, experiments on these two classes of datasets pose very different challenges for our proposed model to handle. We use mean average precision (mAP@all) and precision at 200 (Prec@200) as evaluation metrics in accordance with the convention in existing works [18, 28]. In the supplementary, we additionally provide details on the pre-trained classifiers (teachers), hyperparameter choices, and our experimentation platform.

4.1. Comparison with Baseline Approaches

In Table 1, we compare our method with various baseline approaches derived from existing state-of-the-art DFL literature. Below, we discuss our observations, with additional implementation details and analyses in the supplementary.

Classifier Representations: The pre-trained ResNet50 [20] classifiers are made to act as encoders. We remove the final classification layer and use the output from the representation layer of the classifiers to perform SBIR. Since their representations spaces have not been aligned by a contrastive objective like \mathcal{L}_{enc} , directly using the classifiers as encoders result in sub-par retrieval performance.

Uni-Modal Knowledge Distillation: Since all state-of-the-art data-free knowledge distillation approaches operate on uni-modal data [8, 10, 17, 30], we adapt them to train the photo and sketch students, but without any cross-modal interaction. Thus, during the distillation phase, the photo and the sketch students learn to match their outputs with those of their modality-specific teachers for classification, but do not interact with each other via class-alignment or contrastive losses for optimizing retrieval performance. We then directly use the representations obtained from such uni-modal networks for SBIR. Although promising, the performance of this approach has to be upper-bounded by what is achieved using the teachers directly. This phenomenon is also supported by our experimental results. Related literature suggests that it is possible to train students to surpass the performance of their teacher by making the former larger in size than the latter [54]. However, we used same-sized backbones for both to ensure fair comparison.

Sampling from a Gaussian Prior: Instead of estimating the train set distributions by inverting the classifiers, inspired by related approaches in data-free learning [8], we assume a Gaussian prior for their distributions. Considering the pixels of the photos and sketches as independently and

Objective	Data-Dependent	Data-Free	Δ
Siamese [11]	0.715	0.679	0.036
Triplet [45]	0.772	0.750	0.022
MIB [18]	0.871	0.815	0.056
Ours (CrossX-DFL) [53]	0.862	0.827	0.035

Table 2. Data-Dependent vs. Data-Free settings on Sketchy. identically distributed Gaussians turns out to be too general an assumption. Identifying the semantic structures in such a manifold would require dense sampling of these distributions, which is computationally intractable. Hence, with a similar training time-scale as CrossX-DFL, this approach performs significantly worse.

Averaging Weights: Averaging the weights of all the networks in an ensemble has recently been shown to outperform even the best model in the ensemble, providing better robustness and domain generalization [52]. In similar vein, we can pose SBIR as a domain generalization problem for the aggregated model. We average the weights of the photo and sketch classifiers, remove the classification layer, and obtain a single network for encoding both photos and sketches. This achieves around a 9% improvement over the individual classifier. However, the performance gap with our CrossX-DFL still remains significant.

Meta-Data Based Reconstruction: We follow [30] by using statistical metadata for obtaining photo/sketch reconstructions. This provides better performance than the classifier-only setting. However, the reconstructions are not very efficient due to the lack of semantic consistency, modality specificity, sample hardness, etc.

Training with Alternative Datasets: Due to the limited availability of bespoke datasets for training SBIR models, we leverage more easily available datasets for related tasks like classification, domain generalization, robustness, etc. We use the photos from ImageNet [12] and their corresponding sketches from ImageNet-Sketch [50] to train the photo and sketch encoders respectively. We use the class-level correspondence information from the two datasets to generate positive photo-sketch pairs for contrastive learning. This setting gives the highest accuracy among all baselines. However, it still remains much lower than what is achieved by our CrossX-DFL. We speculate that although training on related datasets may help in learning task-specific (here retrieval) representations, the large error rate stems from the distribution shifts across different datasets. Our CrossX-DFL, by the virtue of the semantic consistency, class-alignment, modality guidance and metric-agnostic adversarial estimation criteria, is able to reconstruct a much more faithful approximation of the train set distribution, thereby providing significantly superior performance.

4.2. Comparison with the Data-Dependent Setting

The upper-bound for the performance of our model is given by the mAP@all obtained using encoders directly




Figure 3. Data-free photo and sketch reconstructions of the Sketchy and TU-Berlin datasets produced by our estimator networks.

trained on the original datasets. We empirically measure these upper-bounds for various commonly used, state-of-the-art contrastive learning objective functions, and analyze how close our algorithm gets to each of them. To this end, we train our encoders using both the original train set photos and sketches from Sketchy (data-dependent), as well the estimated distributions obtained by our method (data-free). We summarize our findings in Table 2. For all objectives, we are able to get quite close to the data-dependent setting, while using *no* training data at all. For instance, with the Triplet loss, our method achieves a difference (Δ) of only 0.022 with the data-dependent setting. However, with Queued-InfoNCE, our method achieves the highest mAP@all of 0.827, with a Δ of 0.035 with the data-dependent setting, which is also quite competitive.

4.3. Qualitative Results

Figure 3 shows sample photo and sketch reconstructions produced by our estimator networks. We would like to emphasise that our end-goal is *not* to generate realistic train set reconstructions of the classifiers, but rather to produce sample estimations that faithfully represent their manifold. However, we observe that for the sketch modality, all the estimations are quite realistic and akin to what a human amateur sketcher might draw. This is also true for a significant number of reconstructed photos, however for a few (like the Frog), some structural details are missing. This shows that our photo estimators can sometimes be more biased towards replicating textural details over structures, if the former plays a dominant role in the train set of the photo classifier. Also, our estimated photos can be seen to generate distinct objects, but not much background context. We conjecture that this happens because to produce robust predictions, the classifiers themselves learn to be invariant to background information, thereby being unable to provide much feedback about the same.

4.4. Ablation Studies

We present the results of ablating the key components of our model in Table 3. We group them into two categories. Rows 1 to 4 show the effect of incremental addition of the three novel, domain-specific components of our methodol-

ID	Semantic Consistency	Class Alignment	Modality Guidance	Adversarial Estimation	Encoding Loss	Sketchy (mAP@all)	TU-Berlin (mAP@all)
1.	✓					0.600	0.416
2.	✓	✓				0.660	0.535
3.	✓	✓	✓			0.705	0.577
4.	✓	✓	✓	✓		0.740	0.630
5.		✓	✓	✓	✓	0.551	0.360
6.	✓	✓	✓	✓	✓	0.690	0.526
7.	✓	✓	✓	✓	✓	0.773	0.665
8.	✓	✓	✓	✓	✓	0.758	0.641
9.	✓	✓	✓	✓	✓	0.827	0.680

Table 3. Ablation Studies.

ogy (from Section 3.1), namely Class-Alignment, Modality Guidance, and Metric-Agnostic Adversarial Estimation.

Rows 5 to 8 demonstrate how individually removing each of these key components affects the overall performance of the complete model. Row 9 shows the performance of our model with all the components included. We present more details on the individual observations below.

Semantic Consistency: The semantic consistency loss (\mathcal{L}_{sem}) is the most fundamental component of a data-free SBIR pipeline [8, 10, 17], which ensures that the reconstructed train set distributions represent meaningful classes. Row-1 represents a baseline model trained only using the semantic consistency loss. The relative contributions of each of our novel components can be evaluated on top of this baseline. Row-5 shows the performance of our model if the semantic consistency criterion is removed. The results show a significant drop in accuracy, as the reconstructed inputs no longer have any meaningful class information and is equivalent to random noise with some modality specific information (coming from the modality guidance network).

Class-Alignment: The class-alignment loss ($\mathcal{L}_{\text{align}}$) ensures that the corresponding photos and sketches in the estimated distributions belong to the same class. Without this, the degree/distribution of class information between a photo-sketch pair may diverge arbitrarily, even if their hard-labels (as predicted by the teacher) happen to be the same. $\mathcal{L}_{\text{align}}$ ensures more fine-grained correspondence between the pairs by minimizing the symmetrized KL-divergence between the classifiers’ outputs. Rows 2 and 6 show that this leads to a significant improvement in retrieval accuracy.

Modality Guidance: The Modality Guidance Network restricts the estimated distributions to belong to specific modalities, *i.e.*, photos and sketches ($\mathcal{L}_{\text{modal}}$). This is important for training the encoders, as they need to learn to

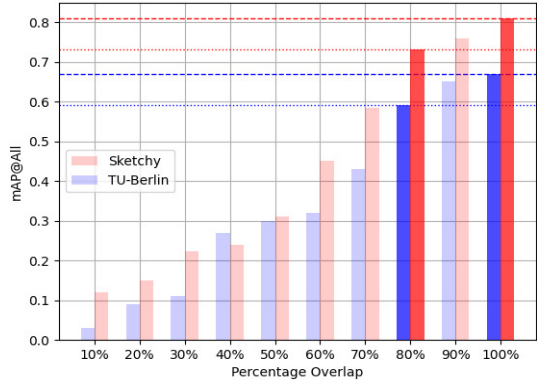


Figure 4. DF-SBIR performance of our model when the classifiers (teachers) are trained on only partially overlapping sets of classes. eliminate irrelevant, modality-specific information, and embed their inputs based on semantic content. The results in Rows 3 and 7 demonstrate the contribution of the modality guidance network to the process of learning cross-modal invariants in the encoding process.

Adversarial Estimation: Our proposed technique for Metric-Agnostic Adversarial Estimation of train set distributions directs the estimators towards producing samples that are the hardest for the encoder to correctly encode by communicating feedback across representation spaces through \mathcal{L}_{adv} . As previously observed in the data-free knowledge distillation literature [10, 35], training the encoder on such samples ensures its robustness to semantic variations. Rows 4 and 8 illustrate the significant improvement in the encoders’ ability to deal with semantic variations when trained with such samples.

Encoding Loss: Based on the observations in the Row groups 1-4 and 5-8, as well as the difference between Rows 4 and 9, we found that the queue-based version of the InfoNCE loss (\mathcal{L}_{enc}) works best for training the encoders with the estimated distributions. The Triplet loss as an ablative alternative to \mathcal{L}_{enc} resulted in degraded performance.

4.5. Further Constraints on Classifiers

Teachers with Partial Class Overlap: We propose a highly challenging novel setting for the evaluation of cross-modal DFL models. Here, the classifiers (teachers) from the two modalities (photo and sketch) are being trained on *partially overlapping* sets of classes. In other words, not all classes between the photo and the sketch classifiers are common. However, the retrieval has to be performed on the *union* of the training classes of both the classifiers. Formally, if the photo classifier has been trained on the set of classes \mathcal{C}_p and the sketch classifier on \mathcal{C}_s , where \mathcal{C}_p and \mathcal{C}_s are partially overlapping, the retrieval test set will constitute of instances from classes $\mathcal{C}_p \cup \mathcal{C}_s$.

We make minor modifications to the training process of our model for evaluation under this scenario, the details of which are provided in the supplementary. Figure 4 shows

its performance as we vary the degree of class overlap between the two classifiers. Even when the encoders of a particular modality do not get direct information about 20% of the classes (80% overlap setting), by virtue of our novel class-aligned estimation procedure, our model is able to leverage complementary information across modalities to achieve strong performance in this challenging scenario.

Cross-Dataset Teachers: To test the ability of our model to deal with the *independence* in the classifiers’ training, we evaluate it using photo and sketch classifiers that are each trained on different datasets, but with overlapping classes. With the photo classifier from TU-Berlin and sketch classifier from Sketchy, our model provides an mAP@all of 0.607 on TU-Berlin. With the photo-classifier from Sketchy and sketch classifier from TU-Berlin, we get an mAP@all of 0.759 on Sketchy. This shows that our model is truly agnostic to the independence in the classifiers’ training while learning a cross-modal metric space in a data-free manner.

5. Limitations

Figure 3 shows that our estimators are unable to incorporate much background information in the reconstructions due to lack of feedback from the classifiers, as the classifiers themselves had learned to become invariant to background information for robust prediction performance. However, this is necessary as the downstream encoders also need to develop this invariance in order to tackle real-world scenarios. We address this by instantiating the encoders with the weights of the classifiers. However, as learning progresses, this knowledge may not be retained given the distribution shift in the estimations. We believe that this limitation can be overcome by leveraging the recent advances in generalized source-free domain adaptation [55].

6. Conclusion

Motivated by the practical constraints of privacy and anonymity preservation, as well as the difficulty involved in obtaining pairs of photos and hand-drawn sketches, for the first time, we studied the problem of SBIR in the data-free setting. We designed an approach for performing data-free learning across modalities and metric spaces (CrossX-DFL), that allows distilling knowledge from independently trained photo and sketch classifiers, into encoders that learn a unified photo-sketch metric space in a completely data-free manner. We illustrate the efficacy of our model by comparison against an extensive set of baselines, as well as exhaustive ablation studies and imposing constraints that challenge its ability to learn complementary cross-modal information. We hope that our practically motivated problem formulation, along with the methodological design choices that come with it, will propel further research on both DF-SBIR as well as general cross-modal data-free learning.

References

- [1] Stephan Alaniz, Massimiliano Mancini, Anjan Dutta, Diego Marcos, and Zeynep Akata. Abstracting sketches through simple primitives. In *ECCV*, 2022. 3
- [2] Aharon Ben-Tal, Laurent Ghaoui, and Arkadi Nemirovski. *Robust Optimization*. Princeton University Press, 2009. 2, 5
- [3] Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Aneeshan Sain, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. More photos are all you need: Semi-supervised learning for fine-grained sketch based image retrieval. In *CVPR*, 2021. 1, 3
- [4] Ayan Kumar Bhunia, Pinaki Nath Chowdhury, Yongxin Yang, Timothy Hospedales, Tao Xiang, and Yi-Zhe Song. Vectorization and rasterization: Self-supervised learning for sketch and handwriting. In *CVPR*, 2021. 3
- [5] Ayan Kumar Bhunia, Viswanatha Reddy Gajjala, Subhadeep Koley, Rohit Kundu, Aneeshan Sain, Tao Xiang, and Yi-Zhe Song. Doodle it yourself: Class incremental learning by drawing a few sketches. In *CVPR*, 2022. 1
- [6] Ayan Kumar Bhunia, Subhadeep Koley, Abdullah Faiz Ur Rahman Khilji, Aneeshan Sain, Pinaki Nath Chowdhury, Tao Xiang, and Yi-Zhe Song. Sketching without worrying: Noise-tolerant sketch-based image retrieval. In *CVPR*, 2022. 1, 3
- [7] Abhra Chaudhuri, Massimiliano Mancini, Yanbei Chen, Zeynep Akata, and Anjan Dutta. Cross-modal fusion distillation for fine-grained sketch-based image retrieval. In *BMVC*, 2022. 1
- [8] Hanting Chen, Yunhe Wang, Chang Xu, Zhaohui Yang, Chuanjian Liu, Boxin Shi, Chunjing Xu, Chao Xu, and Qi Tian. Daff: Data-free learning of student networks. In *ICCV*, 2019. 1, 2, 3, 4, 6, 7
- [9] Shuyu Cheng, Yinpeng Dong, Tianyu Pang, Hang Su, and Jun Zhu. Improving black-box adversarial attacks with a transfer-based prior. In *NeurIPS*, 2019. 3
- [10] Yoojin Choi, Jihwan Choi, Mostafa El-Khamy, and Jungwon Lee. Data-free network quantization with adversarial knowledge distillation. In *CVPRW*, 2020. 2, 5, 6, 7, 8
- [11] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005. 6
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *CVPR*, 2009. 6
- [13] Sounak Dey, Pau Riba, Anjan Dutta, Josep Lladós, and Yi-Zhe Song. Doodle to search: Practical zero-shot sketch-based image retrieval. In *CVPR*, 2019. 1, 3, 5, 6
- [14] Anjan Dutta and Zeynep Akata. Semantically tied paired cycle consistency for zero-shot sketch-based image retrieval. In *CVPR*, 2019. 1, 3
- [15] Anjan Dutta and Zeynep Akata. Semantically tied paired cycle consistency for any-shot sketch-based image retrieval. *IJCV*, 2020. 1
- [16] Mathias Eitz, James Hays, and Marc Alexa. How do humans sketch objects? *SIGGRAPH*, 2012. 5
- [17] Gongfan Fang, Jie Song, Xinchao Wang, Chengchao Shen, Xingen Wang, and Mingli Song. Contrastive model inversion for data-free knowledge distillation. In *IJCAI*, 2021. 2, 6, 7
- [18] Marco Federici, Anjan Dutta, Patrick Forré, Nate Kushman, and Zeynep Akata. Learning robust representations via multi-view information bottleneck. In *ICLR*, 2020. 3, 6
- [19] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross B. Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 5
- [20] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CVPR*, 2016. 6
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. In *NIPS*, 2015. 2
- [22] Conghui Hu and Gim Hee Lee. Feature representation learning for unsupervised cross-domain image retrieval. In *ECCV*, 2022. 1
- [23] Rui Hu and John Collomosse. A performance evaluation of gradient field hog descriptor for sketch based image retrieval. *CVIU*, 2013. 3
- [24] Jonas Jongejan, Henry Rowley, Takashi Kawashima, Jongmin Kim, and Nick Fox-Gieg. The quick, draw! - a.i. experiment. <https://quickdraw.withgoogle.com>, 2016. 5
- [25] Sungyeon Kim, Dongwon Kim, Minsu Cho, and Suha Kwak. Proxy anchor loss for deep metric learning. In *CVPR*, 2020. 5
- [26] Sasi Kiran Yelamarthi, Shiva Krishna Reddy, Ashish Mishra, and Anurag Mittal. A zero-shot framework for sketch based image retrieval. In *ECCV*, 2018. 3
- [27] Yi Li, Timothy Hospedales, Yi-Zhe Song, and Shaogang Gong. Fine-grained sketch-based image retrieval by matching deformable part models. In *BMVC*, 2014. 3
- [28] L. Liu, F. Shen, Y. Shen, X. Liu, and L. Shao. Deep sketch hashing: Fast free-hand sketch-based image retrieval. In *CVPR*, 2017. 3, 5, 6
- [29] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 1
- [30] Raphael Gontijo Lopes, Stefano Fenu, and Thad Starner. Data-free knowledge distillation for deep neural networks. *arXiv*, 2017. 2, 6
- [31] Jiahao Lu, Xi Sheryl Zhang, Tianli Zhao, Xiangyu He, and Jian Cheng. April: Finding the achilles' heel on privacy for vision transformers. In *CVPR*, 2022. 1, 3
- [32] Yichen Lu, Mei Wang, and Weihong Deng. Augmented geometric distillation for data-free incremental person reid. In *CVPR*, 2022. 2, 3
- [33] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *CVPR*, 2015. 2
- [34] H. B. McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *AIS-TATS*, 2017. 3
- [35] Paul Micaelli and Amos J Storkey. Zero-shot knowledge transfer via adversarial belief matching. In *NeurIPS*, 2019. 2, 5, 8

- [36] Yair Movshovitz-Attias, Alexander Toshev, Thomas K. Leung, Sergey Ioffe, and Saurabh Singh. No fuss distance metric learning using proxies. In *ICCV*, 2017. 5
- [37] Kaiyue Pang, Yi-zhe Song, Tony Xiang, and Timothy Hospedales. Cross-domain generative learning for fine-grained sketch-based image retrieval. In *BMVC*, 2017. 3
- [38] Kaiyue Pang, Yongxin Yang, Timothy M Hospedales, Tao Xiang, and Yi-Zhe Song. Solving mixed-modal jigsaw puzzle for fine-grained sketch-based image retrieval. In *CVPR*, 2020. 3
- [39] Wonpyo Park, Dongju Kim, Yan Lu, and Minsu Cho. Relational knowledge distillation. In *CVPR*, 2019. 2
- [40] Pytorch. Pre-trained models in pytorch. <https://pytorch.org/vision/stable/models.html>. 2, 3
- [41] Karsten Roth, Timo Milbich, Bjorn Ommer, Joseph Paul Cohen, and Marzyeh Ghassemi. Simultaneous similarity-based self-distillation for deep metric learning. In *ICML*, 2021. 2
- [42] Jose M Saavedra. Sketch based image retrieval using a soft computation of the histogram of edge local orientations (shelo). In *ICIP*, 2014. 3
- [43] Jose M Saavedra, Juan Manuel Barrios, and S Orand. Sketch based Image Retrieval using Learned KeyShapes (LKS). In *BMVC*, 2015. 3
- [44] Aneeshan Sain, Ayan Kumar Bhunia, Yongxin Yang, Tao Xiang, and Yi-Zhe Song. Stylemeup: Towards style-agnostic sketch-based image retrieval. In *CVPR*, 2021. 3
- [45] Patsorn Sangkloy, Nathan Burnell, Cusuh Ham, and James Hays. The sketchy database: Learning to retrieve badly drawn bunnies. *ACM SIGGRAPH*, 2016. 1, 3, 5, 6
- [46] Sunandini Sanyal, Sravanti Addepalli, and R. Venkatesh Babu. Towards data-free model stealing in a hard label setting. In *CVPR*, 2022. 1, 3
- [47] Jifei Song, Qian Yu, Yi-Zhe Song, Tao Xiang, and Timothy M. Hospedales. Deep spatial-semantic attention for fine-grained sketch-based image retrieval. In *ICCV*, 2017. 3
- [48] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive representation distillation. In *ICLR*, 2020. 2
- [49] Aäron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv*, 2018. 5
- [50] Haohan Wang, Songwei Ge, Zachary Lipton, and Eric P Xing. Learning robust global representations by penalizing local predictive power. In *NeurIPS*, 2019. 6
- [51] Wenxuan Wang, Xuelin Qian, Yanwei Fu, and Xiangyang Xue. Dst: Dynamic substitute training for data-free black-box attack. In *CVPR*, 2022. 1, 3
- [52] Mitchell Wortsman, Gabriel Ilharco, Samir Ya Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. In *ICML*, 2022. 1, 2, 3, 6
- [53] Zhirong Wu, Yuanjun Xiong, Stella X. Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 5, 6
- [54] Qizhe Xie, Minh-Thang Luong, Eduard Hovy, and Quoc V. Le. Self-training with noisy student improves imagenet classification. In *CVPR*, 2020. 2, 6
- [55] Shiqi Yang, Yaxing Wang, Joost van de Weijer, Luis Herranz, and Shangling Jui. Generalized source-free domain adaptation. In *ICCV*, 2021. 8
- [56] Hongxu Yin, Pavlo Molchanov, Jose M. Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K. Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *CVPR*, 2020. 2
- [57] Jaemin Yoo, Minyong Cho, Taebum Kim, and U Kang. Knowledge extraction with no observable data. In *NeurIPS*, 2019. 2
- [58] Qian Yu, Feng Liu, Yi-Zhe Song, Tao Xiang, Timothy M. Hospedales, and Chen Change Loy. Sketch me that shoe. In *CVPR*, 2016. 1, 3
- [59] Jie Zhang, Bo Li, Jianghe Xu, Shuang Wu, Shouhong Ding, Lei Zhang, and Chao Wu. Towards efficient data free black-box adversarial attack. In *CVPR*, 2022. 3
- [60] Lin Zhang, Li Shen, Liang Ding, Dacheng Tao, and Ling-Yu Duan. Fine-tuning global model via data-free knowledge distillation for non-iid federated learning. In *CVPR*, 2022. 3
- [61] Yiman Zhang, Hanting Chen, Xinghao Chen, Yiping Deng, Chunjing Xu, and Yunhe Wang. Data-free knowledge distillation for image super-resolution. In *CVPR*, 2021. 3