Edge Augmentation for Large-Scale Sketch Recognition without Sketches

Nikos Efthymiadis
VRG, Faculty of Electrical Engineering
Czech Technical University in Prague
Email: efthynik@fel.cvut.cz

Giorgos Tolias
VRG, Faculty of Electrical Engineering
Czech Technical University in Prague
Email: giorgos.tolias@cmp.felk.cvut.cz

Ondřej Chum
VRG, Faculty of Electrical Engineering
Czech Technical University in Prague
Email: chum@cmp.felk.cvut.cz

Abstract—This work addresses scaling up the sketch classification task into a large number of categories. Collecting sketches for training is a slow and tedious process that has so far precluded any attempts to large-scale sketch recognition. We overcome the lack of training sketch data by exploiting labeled collections of natural images that are easier to obtain. To bridge the domain gap we present a novel augmentation technique that is tailored to the task of learning sketch recognition from a training set of natural images. Randomization is introduced in the parameters of edge detection and edge selection. Natural images are translated to a pseudo-novel domain called "randomized Binary Thin Edges" (rBTE), which is used as a training domain instead of natural images. The ability to scale up is demonstrated by training CNN-based sketch recognition of more than 2.5 times larger number of categories than used previously. For this purpose, a dataset of natural images from 874 categories is constructed by combining a number of popular computer vision datasets. The categories are selected to be suitable for sketch recognition. To estimate the performance, a subset of 393 categories with sketches is also collected.

I. INTRODUCTION

Free-hand drawings, or *sketches*, have been a long-lasting means of human communication and expression. Nowadays, the prevalence of digital devices equipped with touch-screens has given free-hand sketches additional roles in a number of educational, business or leisure activities. As a result, computer vision research related to sketches has flourished in a variety of tasks including synthesis [1], perceptual grouping [2], sketch-based image retrieval [3], [4], [5], and sketch recognition [6], [7], [8]. In this paper, we focus on the task of sketch recognition, *i.e.* how to classify sketches into specific categories. In particular, we target a realistic application scenario, where the number of classes is as large as possible. Prior work commonly keeps the number of classes relatively low. The reason for that is simple - lack of training data. Since annotation effort for sketch recognition includes sketch drawing, this activity becomes prohibitively expensive. For example, the *Sketchy* dataset [9] required 3,921 hours of sketching for 125 categories. In order to obtain a seven times larger dataset, one would need **over 13 human-years** of sketch drawing (40 hours a week, 52 weeks a year). Therefore, the task of large-scale sketch recognition requires methods much less demanding on the training data.

To allow scalability in the number of classes, we propose a method that trains a deep network classifier without requiring a
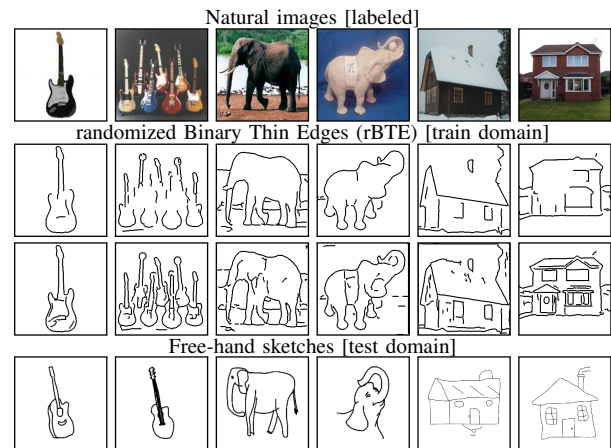

Fig. 1. The goal of this work is to recognize sketches at test time (bottom row) without using any sketches at training time. Labeled natural images (top row) are transformed to rBTEs with different level of details (the two middle rows are two instances of rBTE per natural image, thickened for visualization) to bridge the domain gap. Combined with geometric augmentations, the transformed dataset is used to train a deep network for sketch recognition.

single sketch during the training. Instead, only natural images with their labels are used to train the classifier. The method exploits the fact that human-drawn sketches often represent the 2D shape of depicted objects or of their parts. The sketch domain is approximated by detected 2D shapes of objects in natural images. In the following, terms *natural images*, *RGB images*, or simply *images* are used interchangeably.

A novel edge augmentation technique is used to map natural images to a pseudo-novel domain called randomized Binary Thin Edges (rBTE). This augmentation procedure randomly selects an edge detector and an edge selection strategy in order to generate a sketch-like output with different level of details (see Figure 1) and is combined with random geometric augmentations.

Sketch recognition is a standard and well defined task and, at the same time, collections of annotated images are available. It is possible to cast sketch recognition as a single-source domain generalization with natural images as the source domain. However, we show that a specific approach exploiting the specifics of natural images and sketches brings a relative recognition-rate improvement of more than 20% over unnecessarily generic single-source domain generalization approaches.

The proposed approach is a general augmentation scheme that can include any modern image-to-edge or image-to-sketch

method. In this work we demonstrate its potential by using edge detectors [10], [11], [12] trained on an extremely limited amount of non-sketch data, *i.e.* 200 natural images. Methods such as [13], [14] are trained on sketches, and methods such as [15] require more data. rBTEs form a rich training set allowing to train, without a single sketch, a CNN-based sketch classifier, which is the main contribution of this work.

Sketch synthesis is a popular task [16], [17], [18], [19], [20] in which sketches are generated from images. However, it is not applicable to the setup that this work explores since these approaches cannot work without training sketches.

To evaluate the proposed approach we introduce Im4Sketch, a dataset for large-scale sketch recognition without sketches for training. It consists of 1,007,878 natural images labeled into 874 classes used to train the sketch classifier. Testing is performed on 80,582 sketches coming from 393 classes that are a subset of the training classes. The dataset is a composition of existing popular image and sketch datasets, namely ImageNet [21], DomainNet (DN) [22], Sketchy [9], PACS [23], and TU-Berlin [24]. The classes are selected, so that classification by shape is meaningful. For example, ImageNet categories "Indian elephant" and "African elephant" are merged into category "Elephant". The dataset is described in detail in Section V.

To the best of our knowledge, this is the first work that delivers sketch recognition of the order of over 800 categories. The dataset with the second largest number of classes for sketch recognition is DomainNet [22] with 345 classes, *i.e.* more than 2.5 times smaller.

II. RELATED WORK

In this section, we review the prior work on three tasks that are relevant to the focus of this work, namely sketch recognition, sketch-based image retrieval, and domain generalization.

a) Sketch recognition: The crowd-sourced free-hand sketch dataset by Eitz *et al.* [24] is the first large-scale dataset on the domain of sketches. Early approaches [25], [24] focus on adapting hand-crafted features and encodings, such as SIFT [26], its variants, Fisher vectors [27], and SVM classifiers to solve the task. The application of deep network classifiers was stimulated by the enormous effort invested in annotating sketches. The *Sketch-a-Net* [28] approach demonstrates recognition performance surpassing human recognition ability. This is achieved with a tailored architecture and a labeled training set of sketches whose size is in the order of 10^6. In the recent work of Qi *et al.* [29], the combination of deep and hand-crafted features exhibits very good results. Some approaches [30], [31] exploit the additional information of stroke order, when available in the input, to further improve the recognition accuracy. We assume that this information is not available and deal with the more general problem.

The scarcity of training data in the domain of sketches is handled by some approaches by combining sketches and natural images during the training. Hua *et al.* [6] attempt to automatically learn the shared latent structures that exist between sketch images and natural images. Zhang *et al.* [32]

transfer the knowledge of a network learned on natural images to a sketch network. In both these methods the training set consists of both natural images and sketches. Even though these approaches are valuable in a few-shot setup, where only a few labeled sketches per category are available, the setup with no sketches has not been well studied before with a focus on the sketch domain. An exception is the recent work of Lamb *et al.* [33] where the *SketchTransfer* task is presented. Even though their work explores the setup of no available sketches too, promising results are achieved only when unlabeled sketches are available during the training. The authors conclude that the low resolution images of the benchmark is a limitation. Therefore, in our work we use benchmarks with higher resolution images that are more realistic.

b) Sketch-based image retrieval: Classical approaches use edge detection on natural images to bridge the domain gap and then handle both domains with hand-crafted descriptors or matching [34], [3], [35], [36], [37]. Deep learning methods mainly follow a different path. A two branch architecture is used [38], [39], [40], with a different branch per domain, where the expectation is to bridge the domain gap based on large amounts of training data with cross-domain labeling [9]. If learning is involved, the most realistic setup is the zero-shot sketch-based image retrieval [41]; which is a challenging task, that is related to, but different from, ours. Radenovic *et al.* [4] avoid cross-modal annotation by relying on training labels of natual images and using edge detection to bridge the domain gap. Their work focuses on learning shape similarity and does not attempt to directly generalize to category level recognition.

c) Domain generalization: The most common approach for domain generalization is invariant feature learning, based on the theoretical results of Ben-David *et al.* [42]. Representative approaches include kernel-based invariant feature learning by minimizing domain dissimilarity[43], multi-task autoencoders that transform the original image to other related domains, domain classifiers as adversaries to match the source domain distributions in the feature space [44], [45], and cross-domain non-contrastive learning as regularization [46].

Some methods specialize for single-source domain generalization. Examples include hard example generation in virtual target domains [47], style transfer using auxiliary datasets [48], and adversarial domain augmentation [49]. Narayanan *et al.* [50] argue that the shock graph of the contour map of an image is a complete representation of its shape content and use a Graph Neural Network as their model. Wang *et al.* [51] propose a style-complement module to create synthetic images from distributions that are complementary to the source domain.

Data augmentation techniques are commonly used for domain generalization. Zhou *et al.* [52] synthesize data from pseudo-novel domains under semantic consistency by using a data generator. Mancini *et al.* [53] use mixup [54] to combine different source domains. Carlucci *et al.* [55] train a model to solve jigsaw puzzles in a self-supervised manner in addition to the standard classification loss to improve the generalization ability of the model.

Fig. 2. Overview of the training pipeline. Natural images are transformed into rBTEs, which are used with class labels to train a network classifier with cross-entropy loss. The obtained network is used to classify free-hand sketches into the object categories.

III. Task formulation

In this section, we define the task and relate it to existing computer vision tasks. We follow the notation of transfer learning [56] and domain adaptation [57] literature.

A domain $\mathcal{D}$ is an ordered pair $\mathcal{D} = (\mathcal{X}, P(X))$ composed of a space of input examples $\mathcal{X}$ and a marginal probability distribution $P(X)$, where X is a random variable valued in $\mathcal{X}$. A task $\mathcal{T} = (\mathcal{Y}, P(Y|X))$ is defined by a label space $\mathcal{Y}$ and the conditional probability distribution $P(Y|X)$, where Y is a random variable with values in $\mathcal{Y}$.

In the problem, two domains are considered: the *target* domain $\mathcal{D}^t = (\mathcal{X}^t, P(X^t))$ of sketches and the *source* domain $\mathcal{D}^s = (\mathcal{X}^s, P(X^s))$ of natural images, with tasks $\mathcal{T}^t = (\mathcal{Y}^t, P(Y^t|X^t))$ and $\mathcal{T}^s = (\mathcal{Y}^s, P(Y^s|X^s))$ respectively. The goal is to learn a predictor $f : \mathcal{X}^t \rightarrow \mathcal{Y}^t$ for the target domain without having access to any examples from that domain.

The input spaces of both the domains, target and source respectively, are images (RGB, fixed size), thus $\mathcal{X}^s = \mathcal{X}^t$. The same categories are to be recognized in the two domains, *i.e.* the label spaces are also identical $\mathcal{Y}^s = \mathcal{Y}^t$. However, the marginal distributions are significantly different, *i.e.* $P(X^s) \neq P(X^t)$. In this work, we advocate for bridging the domain gap by constructing a transformation $T : \mathcal{X} \rightarrow \mathcal{X}$ so that $P(T(X^s)) \approx P(X^t)$. With such a transformation, an approximation of $P(Y^t|X^t)$ in the form of $P(Y^s|T(X^s))$ can be learned with labeled examples from the source domain of natural images. In this work, we focus on designing the transformation based on prior knowledge about the two domains, see Section IV.

a) Relation to domain adaptation: In the domain adaptation task, similarly to our problem, $\mathcal{Y}^s = \mathcal{Y}^t$ and $P(X^s) \neq P(X^t)$. The main difference is that in domain adaptation, some examples from the target domain are available, either labeled or unlabeled for supervised or unsupervised domain adaptation respectively. For example, in unsupervised domain adaption from natural images to sketches, a labeled dataset of natural images is available, together with unlabeled sketch examples. The goal is to obtain a predictor for sketches.

b) Relation to domain generalization: The task of domain generalization is the closest one to our task. The essential difference is that most domain generalization methods are either unusable or they under-perform in the single-source task. Also domain generalization targets to perform well in every possible target domain. The domain label of each example is used by most approaches as additional supervision.
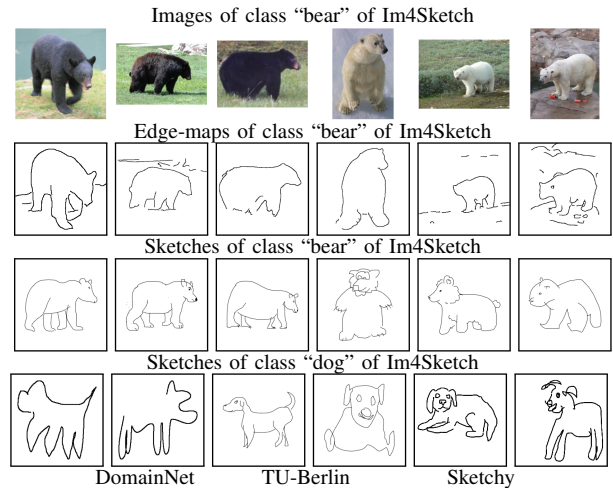

Fig. 3. Image and sketch examples from the Im4Sketch dataset. Class bear contains the original ImageNet classes "American Black Bear" and "Ice Bear" whose shape is indistinguishable. Sketches are collected from original datasets with different level of detail as in the case of class "dog"; original datasets are shown at the bottom.

It holds that $P(X_j^s) \neq P(X^t)$ and $\mathcal{Y}_j^s = \mathcal{Y}^t$ for $j = 1 \ldots d$, where d is the number of source domains. A sketch recognition example is the case where labeled datasets exist for the domain of natural images, artworks, and cartoons. The goal is to obtain a predictor for free-hand sketches. In contrast to our setup, the multiple domains allow for learning a domain invariant predictor, while in our task, exploiting prior knowledge is the only way to proceed.

c) Relation to attributed-based zero-shot learning: Zero-shot learning in visual applications is the following. The input spaces are the same $\mathcal{X}^s = \mathcal{X}^t$, but the marginal distributions are different $P(X^s) \neq P(X^t)$. The label spaces are disjoint $\mathcal{Y}^s \cap \mathcal{Y}^t = \emptyset$, and, therefore, the tasks are different too, *i.e.* $\mathcal{T}^s \neq \mathcal{T}^t$. The similarity to our task is that there are no input examples of the target domain during the learning. There is additional information, though, in the form of attributes. Each class, from both domains, is described by an attribute vector, whose dimensions correspond to high-level semantically meaningful properties [58]. This information is used to transfer knowledge from one domain to the other.

IV. Method

In this section, we describe the proposed method, the details of the construction of the rBTE domain, as well as the implementation details of the training.

A. Deep network classifier

Predictor $f : \mathcal{X}^t \rightarrow \mathcal{Y}^t$ is a deep convolutional neural network that takes as input a sketch and predicts the class. It is possible to obtain the predictor by minimizing empirical risk $\frac{1}{n}\sum_{i=1}^n \ell(f(x_i), y_i)$, with $(x_i, y_i) \in \mathcal{X}^t \times \mathcal{Y}^t$, when sketches are available during training. Instead, in this work, the predictor is obtained by minimizing $\frac{1}{n}\sum_{i=1}^n \ell(f(T(x_i)), y_i)$, with $(x_i, y_i) \in \mathcal{X}^s \times \mathcal{Y}^s$.

B. Edge Augmentation

A natural image x is transformed to $T(x)$, called rBTE, through a sequence of randomized transformations that are described in the following.

a) Edge detection: is performed to map the input image to an edge-map with values in $[0, 1]$. Various edge detectors are used to extract the edges: The *Structures Edges* (SE) [10], the *Holistically-Nested Edge Detection*[1] [11], and the *Bi-Directional Cascade Network* [12]. While SE uses random-forest classifier, the other two are DNN based. Using multiple detectors, the size of available training examples is expanded, in this case by a factor of 3.

All three edge detectors are originally trained with clean edge-maps obtained from ground truth segmentation masks, in particular the BSDS500 [59] segmentation dataset containing only 500 images, 200 of which are for training. The dataset consists of natural images, and the ground truth masks are obtained by manual labeling. This is an additional labeled dataset, not including free-hand sketches of objects, that is indirectly used in our work to improve sketch recognition. Even though not considered in this work, one could possible make better use of this dataset even during learning predictor f to obtain further improvements.

b) Geometric augmentations: Edge-maps are geometrically transformed with a set of commonly used geometric augmentations for CNNs: Zero-padding to obtain square edge-maps, re-sample to 256×256, rotation by angle uniformly sampled in $[-5, 5]$ degrees, crop with relative size uniformly sampled in $[0.8, 1.0]$ and aspect ratio in $[\frac{3}{4}, \frac{4}{3}]$ resized to 224×224, and finally horizontal flip with a probability of 0.5.

c) Thinning - NMS: Edge thinning is performed by non-maximum suppression that finds the locations with the sharpest change of the edge strength. All values that are not larger than the neighboring values along the positive and negative gradient directions are suppressed. This is a standard process which is also part of the well known Canny detector [60] and commonly used as post-processing for the SE detector.

d) Hysteresis thresholding: Hysteresis thresholding, which is another common step of the Canny [60] detector, that transforms the input to a binary images using two thresholds (low and high). Pixels whose value is above (below) the high (low) threshold are (are not) considered edges, while pixels with values between the two thresholds are considered as edges only if they belong to a connected-component containing an edge pixel. Setting a fixed threshold that operates well for

[1]A reimplementation is used https://github.com/sniklaus/pytorch-hed

the large range of images and objects considered in this work is not possible. Instead, a threshold t is estimated by standard approaches that perform binarization of grayscale images, and set the low and high thresholds to $0.5t$ and $1.5t$, respectively. The thresholding approach is randomly chosen from a candidate pool comprising methods of Otsu [61], Yen [62], Li [63], Isodata [64], and the mean approach [65].

e) Large connected-components: In the last stage, pixels belonging to small (less than 10 pixels) connected components, estimated with 8-neighbor connectivity, are discarded.

C. Testing

The relative size of sketches with respect to the image dimensions varies from dataset to dataset. In order to tackle this we perform inference in the following two ways. single-scale testing, which is the testing of each sketch in its original relative size and multi-scale testing. In multi-scale testing the sketch is cropped to its bounding box, padded to have aspect ratio 1:1 and then resized to be 90%, 65% and 45% of the 224x224 network input size. The prediction of the three resized sketches are then ensembled by averaging before the softmax function.

D. Overview

A visual overview of the pipeline is shown in Figure 2. The proposed approach uses the sequence of transformations on natural images during the learning stage, and trains a deep convolutional neural network with rBTEs. Randomness is involved in the selection of the edge-map, in the geometric augmentation, and in the selection of the threshold estimation method. Each time an image participates in a batch, only one of the edge-maps and one of the threshold estimation methods is used, chosen with equal probability. A set of thorough ablations is presented in the experimental section by discarding parts of the overall pipeline. During inference, a sketch is fed to the exact same network, after simply performing thinning.

E. Implementation details

a) Training details: ResNet-101 [66] is used as the backbone network for our ablation study on Sketchy and for the core Im4Sketch experiments. The parameters of the network are initialized with the standard weights from training on ImageNet, *i.e.* training with rBTEs starts with the network trained on ImageNet with RGB images. Adam optimizer is used with batch size equal to 64 for training on rBTEs. The *range test* [67] is used to identify the initial learning rate. This process starts with a high learning rate and keeps decreasing it every 50 updates. The initial learning rate is picked in the range of the steepest ascending accuracy on the training set. This method indicates a learning rate of 10^{-4} for all of our experiments. The learning rate is decayed by a factor of 10 every 10 epochs for a total of 30 epochs.

b) Comparison with domain generalization methods: For a fair comparison with the existing domain generalization methods we train a ResNet-18 [66] with SGD optimizer with momentum 0.9, batch size 64 and learning rate 0.004. We train for 30 epochs without a scheduler.

Dataset	Im4Sketch	ImageNet	DomainNet	Sketchy	PACS	TU-Berlin
ImageNet	690	-	162	99	4	141
DomainNet	345	162	-	98	7	197
Sketchy	125	99	98	-	5	99
PACS	7	4	7	5	-	7
TU-Berlin	223	141	197	99	7	-

TABLE I

EACH ROW REPRESENTS THE ORIGINAL DATASET AS A COMPONENT OF IM4SKETCH. THE IM4SKETCH COLUMN SHOWS THE NUMBER OF FINAL CLASSES THAT EXIST IN EACH DATASET. THE REST OF THE COLUMNS SHOW THE NUMBER OF COMMON FINAL CLASSES, AFTER OUR MERGING, BETWEEN THE DIFFERENT DATASETS. IMAGENET IS WITH OUR MERGED VERSION. TU-BERLIN IS WITHOUT THE DELETED CLASSES.

Dataset	Original Datasets				Im4Sketch		
	Train	Val	Test	Classes	Train	Val	Test
IN	1,281,167	50,000	100,000	1,000	885,946	34,500	0
DN Real	120,906	0	52,041	345	61,039	26,303	0
DN Qdr	120,750	0	51,750	345	0	0	51,750
Sketchy Im	11,250	0	1,250	125	72	18	0
Sketchy Sk	68,418	0	7,063	125	0	0	7,063
PACS Sk	——— 3,929 ———			7	0	0	3,929
TUB	——— 20,000 ———			250	0	0	17,840
Total					947,057	60,821	80,582

TABLE II

NUMBER OF NATURAL IMAGES AND SKETCHES FOR THE ORIGINAL DATASETS AND FOR THE IM4SKETCH DATASET PARTITIONED INTO TRAIN, VALIDATION AND TEST SUBSETS.

V. THE IM4SKETCH DATASET

We present a large-scale dataset with shape-oriented set of classes for image-to-sketch generalization called "Im4Sketch"[2]. It consists of a collection of natural images from 874 categories for training and validation, and sketches from 393 categories (a subset of natural image categories) for testing.

The images and sketches are collected from existing popular computer vision datasets. The categories are selected having shape similarity in mind, so that object with same shape belong to the same category.

The natural-image part of the dataset is based on the ILSVRC2012 version of ImageNet (IN) [21]. The original ImageNet categories are first merged according to the shape criteria. Object categories for objects whose shape, *e.g.* how a human would draw the object, is the same are merged. For this step, semantic similarity of categories, obtained through WordNet [68] and category names, is used to obtain candidate categories for merging. Based on visual inspection of these candidates, the decision to merge the original ImageNet classes is made by a human. For instance, "Indian Elephant" and "African Elephant", or "Laptop" and "Notebook" are merged. An extreme case of merging is the new class "dog" that is a union of 121 original ImageNet classes of dog breeds.

In the second step, classes from datasets containing sketches are used. In particular, DomainNet (DN) [22], Sketchy [9], PACS [23], and TU-Berlin [24]. Note that merging is not necessary for classes in these datasets, because the shape criteria are guaranteed since they are designed for sketches. In this step, a correspondence between the merged ImageNet categories and categories of the other datasets is found. As in the merging step, semantic similarity is used to guide the correspondence search. A summary of the common classes per dataset pairs is shown in Table I. Sketch categories that are not present in the merged ImageNet are added to the overall category set, while training natural images of those categories are collected from either DomainNet or Sketchy. In the end, ImageNet is used for 690 classes, DomainNet for 183 classes, and Sketchy for 1 class, respectively. An example of merging and mapping is shown in Figure 3.

Training images of Im4Sketch come from training sets of ImageNet and DomainNet(rel), and from 80% of training set of Sketchy. The validation set is obtained from the validation set of ImageNet, the test set of DomainNet(rel) and 20% of the training set of Sketchy. To avoid large imbalance when collecting images from ImageNet we keep at most 1350 images per class, chosen uniformly from all corresponding original ImageNet classes.

Almost all sketch categories from the four datasets are covered in Im4Sketch. We exclude 27 classes of the TU-Berlin dataset in order to either avoid class conflicts, *e.g.*"flying bird" and "standing bird", or because we are unable to map them to any existing category with natural images in another dataset, *e.g.*"sponge bob". All sketches assigned to any of the final set of categories are used to form the Im4Sketch test set, with an exception for Sketchy and DomainNet where we keep only the sketches from the test set; see more details in Table II.

VI. EXPERIMENTS

In this section, experimental settings, datasets, and results are presented[3]. Firstly, we perform an extensive set of ablations by training and testing on Sketchy dataset. Then, we train and test the proposed approach on PACS dataset to compare with prior domain generalization results. Lastly, we perform large-scale training and testing on the newly introduced Im4Sketch dataset. Recognition performance on Sketchy and Im4Sketch is evaluated by mean prediction accuracy, first estimated per class and then averaged for the whole dataset. For the comparison with the existing domain generalization methods, mean accuracy is used in order to be consistent with the literature.

A. Ablation study on Sketchy

Sketchy dataset [9] is originally created for the purpose of fine-grained sketch-based image retrieval, while we define a new evaluation protocol for our task. It consists of 125 object categories with natural images. Sketches of objects are drawn with these images as reference. The training part consists of 11,250 images and the corresponding 68,418 sketches, while the testing part consists of 1,250 images and the corresponding 7,063 sketches. Since there is no designated validation set, we randomly split the training part keeping 80% for training and 20% for validation. This policy is followed both for images and sketches.

---

[2]Im4Sketch is publicly available on http://cmp.felk.cvut.cz/im4sketch/

[3]The code and the proposed evaluation protocols are publicly available on https://github.com/NikosEfth/im2rbte

ID	0	1	2	3	4 rBTE	5 rBTE	6	7
Geometric	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Edge detector		HED	HED	HED	All	All	HED	HED
NMS	N/A	Yes	Yes	Yes	Yes	Yes	No	Yes
Thresholder		Fixed	Fixed	All	All	All	No	No
Pretrained	RGB/IN	RGB/IN	RGB/IN	RGB/IN	RGB/IN	BTE/I4S	RGB/IN	RGB/IN
Single scale	11.4 ±3.5	39.1 ±0.8	47.0 ±0.3	47.9 ±0.6	49.7 ±0.5	57.2 ±0.5	5.9 ±2.9	40.2 ±0.6
Multi scale	11.6 ±3.9	42.4 ±1.5	49.8 ±0.6	50.9 ±0.4	52.3 ±0.5	59.6 ±0.8	8.0 ±2.2	41.1 ±0.8

TABLE III
SKETCHY: ABLATION STUDY FOR TRAINING ON NATURAL IMAGES AND TESTING ON SKETCHES.

Method	Photo $\longrightarrow$ Sketch
RGB	32.6 ± 2.4
SelfReg [46]	33.7 ± 2.6
L2D [51]	58.5 ± 3.9
SagNet [69]	40.7
rBTE (Ours)	**70.6 ± 2.2**

TABLE IV
PACS: IMAGE-TO-SKETCH GENERALIZATION.

We choose the Sketchy benchmark to perform an ablation study for computational reasons. Each ablation includes a 30 epoch training procedure followed by the evaluation; the reported numbers are averaged over five randomized runs. The results of the study are summarized in Table III. Training using single edge-maps, and fixed threshold, without geometric augmentation results in poor performance (ID=1), while simply adding geometric augmentations provides a noticeable boost (ID=2). In these two ablations the threshold is fixed and set equal to the average threshold estimated by the Otsu method on the whole dataset. Then, using all the threshold estimation methods per image increases performance (ID=3). Using all edge-maps (ID=4) provide further boost. This variant constitutes the main approach of this work. However, instead of initializing with the result of training on ImageNet with RGB images, we also perform an experiment where the initialization is the result of training with rBTEs on Im4Sketch (described at the end of this section). This kind of large-scale pre-training is shown very essential (ID=5) and increases performance by 7.5% with single scale and 7.3 in multi scale testing. Two additional ablations show that edge-map binarization is needed (ID=7 versus ID=3), and that without the NMS-based thinning the performance is very low (ID=6). The baseline approach of training on RGB images of the Sketchy dataset and then testing on sketches (ID=0) performs very poorly. This comparison demonstrates the large benefit of rBTEs for the Sketchy dataset which includes detailed and well drawn sketches.

We additionally use the Canny edge detector employed with geometric augmentations, Otsu's method to adaptively set Canny's thresholds [70], and Gaussian smoothing with σ randomly chosen in $[1, 5]$. It achieves 44.9% and 47.9% accuracy in single and multi-scale testing, which is 2.1% and 1.9% lower than HED ablation with no adaptive threshold (ID=2).

B. Single-source domain generalization comparison on PACS

Table IV summarises the performance of our approach compared to the current state-of-the-art in domain generalization

arg max over	RGB (single)		rBTE (single)		rBTE (multi)	
	all	subset	all	subset	all	subset
Im4Sketch	5.3	5.3	11.3	12.7	12.4	14.0
DomainNet	3.9	3.9	5.2	6.9	6.6	8.5
Sketchy	3.9	12.4	26.0	42.6	26.8	43.0
TUB	12.0	13.9	25.8	31.5	27.8	33.9
PACS	11.6	42.2	24.7	62.7	23.8	64.8

TABLE V
RESULTS FOR TRAINING ON IM4SKETCH. TESTING IS PERFORMED ON THE TEST SET OF IM4SKETCH OR ITS SUBSETS THAT COME FROM THE DIFFERENT ORIGINAL DATASETS.

and to the baseline trained directly on RGB natural images. SelfReg [46] performs poorly as expected; it is intended for multi-source domain generalization. For L2D [51] which is designed specifically for the single-source task, we run the provided code and ensure that optimal learning rate according to validation performance on the source task is used; the reported score is reproduced. The reported numbers are averaged over twenty randomized runs. Our approach outperforms all other generic domain generalization methods by a large margin.

C. Training on Im4Sketch and testing on all

The proposed approach for learning without sketches is to train on the corresponding image training part of Im4Sketch by transforming them into rBTEs. In this way, both the backbone network and the soft-max classifier are trained only with rBTEs. This is performed either to obtain a sketch classifier for 874 categories, or as pre-training to obtain a better backbone network, tailored for shape-based representation (see ID=5 on Table III).

There are 874 training classes in Im4Sketch, while only 393 classes have sketches for testing, since drawing all the classes is impractical at best. Besides results at single and mutli-scale, two evaluating scenarios are reported. First, "argmax over-all", is the testing over all 874 possible classes. This should be seen as an estimate of the overall performance, as the test-classes are unknown during training. The other scenario, "argmax over-subset" is the testing over the classes that appears in the sketch test set. This corresponds to an unrealistic situation, when posterior probability of classes not in the test set is known to be zero. We only report these results to provide some intuition. The results of our Im4Sketch trained model are summarized in Table V to allow future comparisons. Training on rBTEs performs significantly better than training on RGB images on all setups.

VII. CONCLUSIONS

In this work, we are the first to train a large-scale sketch classifier that is able to recognize up to 874 categories. Due to the absence of such a large training dataset, the learning is performed without any sketches. Instead, we proposed a novel edge augmentation technique to translate natural images to a pseudo-novel domain and use it to train a network classifier. This tailored image-to-sketch method is noticeably better than generic single-source domain generalization approaches.

REFERENCES

[1] D. Ha and D. Eck, "A neural representation of sketch drawings," in *ICLR*, 2018.

[2] L. Yang, J. Zhuang, H. Fu, K. Zhou, and Y. Zheng, "Sketchgcn: Semantic sketch segmentation with graph convolutional networks," in *arXiv*, 2020.

[3] T. Bui and J. Collomosse, "Scalable sketch-based image retrieval using color gradient features," in *ICCV*, 2015.

[4] F. Radenovic, G. Tolias, and O. Chum, "Deep shape matching," in *ECCV*, 2018.

[5] F. Wang, L. Kang, and Y. Li, "Sketch-based 3d shape retrieval using convolutional neural networks," in *CVPR*, 2015.

[6] H. Zhang, S. Liu, C. Zhang, W. Ren, R. Wang, and X. Cao, "Sketchnet: Sketch classification with web images," in *CVPR*, 2016.

[7] Q. Jia, X. Fan, M. Yu, Y. Liu, D. Wang, and L. J. Latecki, "Coupling deep textural and shape features for sketch recognition," in *ACM Multimedia*, 2020.

[8] F. Liu, X. Deng, Y.-K. Lai, Y.-J. Liu, C. Ma, and H. Wang, "Sketchgan: Joint sketch completion and recognition with generative adversarial network," in *CVPR*, 2019.

[9] P. Sangkloy, N. Burnell, C. Ham, and J. Hays, "The sketchy database: learning to retrieve badly drawn bunnies," *ACM Transactions on Graphics*, vol. 35, no. 4, p. 119, 2016.

[10] P. Dollár and C. L. Zitnick, "Structured forests for fast edge detection," in *ICCV*, 2013.

[11] S. Xie and Z. Tu, "Holistically-nested edge detection," *IJCV*, vol. 125, no. 1–3, p. 3–18, Dec. 2017.

[12] J. He, S. Zhang, M. Yang, Y. Shan, and T. Huang, "Bi-directional cascade network for perceptual edge detection," in *CVPR*, 2019.

[13] A. K. Bhunia, P. N. Chowdhury, A. Sain, Y. Yang, T. Xiang, and Y.-Z. Song, "More photos are all you need: Semi-supervised learning for fine-grained sketch based image retrieval," *CVPR*, 2021.

[14] J. Song, K. Pang, Y.-Z. Song, T. Xiang, and T. Hospedales, "Learning to sketch with shortcut cycle consistency," *CVPR*, 2018.

[15] J. Lim, L. Zitnick, and P. Dollár, "Sketch tokens: A learned mid-level representation for contour and object detection," in *CVPR*, 2013.

[16] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," *CVPR*, 2017.

[17] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *European Conference on Computer Vision*, 2016.

[18] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, 2014.

[19] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017.

[20] Y. Zhang, G. Su, Y. Qi, and J. Yang, "Unpaired image-to-sketch translation network for sketch synthesis," in *2019 IEEE Visual Communications and Image Processing (VCIP)*, 2019, pp. 1–4.

[21] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, "Imagenet large scale visual recognition challenge," *IJCV*, 2015.

[22] X. Peng, Q. Bai, X. Xia, Z. Huang, K. Saenko, and B. Wang, "Moment matching for multi-source domain adaptation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[23] D. Li, Y. Yang, Y.-Z. Song, and T. Hospedales, "Deeper, broader and artier domain generalization," in *ICCV*, 2017.

[24] M. Eitz, J. Hays, and M. Alexa, "How do humans sketch objects?" *ACM Trans. on Graphics*, 2012.

[25] R. G. Schneider and T. Tuytelaars, "Sketch classification and classification-driven analysis using fisher vectors," *ACM TOG*, vol. 33, no. 6, pp. 1–9, 2014.

[26] D. G. Lowe, "Distinctive image features from scale-invariant keypoints," *IJCV*, vol. 60, no. 2, pp. 91–110, Nov. 2004.

[27] F. Perronnin and C. R. Dance, "Fisher kernels on visual vocabularies for image categorization," in *CVPR*, Jun. 2007.

[28] Q. Yu, Y. Yang, Y.-Z. Song, T. Xiang, and T. M. Hospedales, "Sketch-a-Net that beats humans," in *BMVC*, 2015.

[29] Q. Jia, X. Fan, M. Yu, Y. Liu, D. Wang, and L. J. Latecki, "Coupling deep textural and shape features for sketch recognition," in *ACM Multimedia*, 2020.

[30] P. Xu, C. K. Joshi, and X. Bresson, "Multi-graph transformer for free-hand sketch recognition," in *arXiv*, 2019.

[31] L. Li, C. Zou, Y. Zheng, Q. Su, H. Fu, and C.-L. Tai, "Sketch-r2cnn: An attentive network for vector sketch recognition," in *arXiv*, 2018.

[32] K. Zhang, W. Luo, L. Ma, and H. Li, "Cousin network guided sketch recognition via latent attribute warehouse," in *AAAI*, 2019.

[33] A. Lamb, S. Ozair, V. Verma, and D. Ha, "Sketchtransfer: A challenging new task for exploring detail-invariance and the abstractions learned by deep networks," in *WACV*, 2020.

[34] R. Hu and J. Collomosse, "A performance evaluation of gradient field hog descriptor for sketch based image retrieval," *CVIU*, 2013.

[35] J. M. Saavedra, J. M. Barrios, and S. Orand, "Sketch based image retrieval using learned keyshapes (LKS)," in *BMVC*, 2015.

[36] S. Parui and A. Mittal, "Similarity-invariant sketch-based image retrieval in large databases," in *ECCV*, 2014.

[37] G. Tolias and O. Chum, "Asymmetric feature maps with application to sketch based retrieval," in *CVPR*, 2017.

[38] T. Bui, L. Ribeiro, M. Ponti, and J. Collomosse, "Generalisation and sharing in triplet convnets for sketch based visual search," in *arXiv*, 2016.

[39] O. Seddati, S. Dupont, and S. Mahmoudi, "Quadruplet networks for sketch-based image retrieval," in *ICMR*, 2017.

[40] Q. Yu, F. Lie, Y.-Z. Song, T. Xian, T. Hospedales, and C. C. Loy, "Sketch me that shoe," in *CVPR*, 2016.

[41] S. Dey, P. Riba, A. Dutta, J. Llados, and Y.-Z. Song, "Doodle to search: Practical zero-shot sketch-based image retrieval," in *CVPR*, 2019.

[42] S. Ben-David, J. Blitzer, K. Crammer, and F. Pereira, "Analysis of representations for domain adaptation," in *Advances in Neural Information Processing Systems*. MIT Press, 2007.

[43] K. Muandet, D. Balduzzi, and B. Schölkopf, "Domain generalization via invariant feature representation," in *ICML*, 2013.

[44] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-adversarial training of neural networks," *JMLR*, vol. 17, no. 1, pp. 2096–2030, 2016.

[45] Y. Li, X. Tian, M. Gong, Y. Liu, T. Liu, K. Zhang, and D. Tao, "Deep domain generalization via conditional invariant adversarial networks," in *Proceedings of the European Conference on Computer Vision (ECCV)*, September 2018.

[46] D. Kim, S. Park, J. Kim, and J. Lee, "Selfreg: Self-supervised contrastive regularization for domain generalization," 2021.

[47] R. Volpi, H. Namkoong, O. Sener, J. Duchi, V. Murino, and S. Savarese, "Generalizing to unseen domains via adversarial data augmentation," in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, ser. NIPS'18, 2018, p. 5339–5349.

[48] X. Yue, Y. Zhang, S. Zhao, A. Sangiovanni-Vincentelli, K. Keutzer, and B. Gong, "Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[49] F. Qiao, L. Zhao, and X. Peng, "Learning to learn single domain generalization," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 12 556–12 565.

[50] M. Narayanan, V. Rajendran, and B. Kimia, "Shape-biased domain generalization via shock graph embeddings," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021.

[51] Z. Wang, Y. Luo, R. Qiu, Z. Huang, and M. Baktashmotlagh, "Learning to diversify for single domain generalization," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2021, pp. 834–843.

[52] K. Zhou, Y. Yang, T. Hospedales, and T. Xiang, "Learning to generate novel domains for domain generalization," in *16th European Conference on Computer Vision, ECCV 2020*, 2020.

[53] M. Mancini, Z. Akata, E. Ricci, and B. Caputo, "Towards recognizing unseen categories in unseen domains," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.

[54] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *ICML*, 2018.

[55] F. M. Carlucci, A. D'Innocente, S. Bucci, B. Caputo, and T. Tommasi, "Domain generalization by solving jigsaw puzzles," in *CVPR*, 2019.

[56] K. Weiss, T. M. Khoshgoftaar, and D. Wang, "A survey of transfer learning," *Journal of Big data*, vol. 3, no. 1, p. 9, 2016.

[57] G. Csurka, "Domain adaptation for visual applications: A comprehensive survey," in *arXiv*, 2017.

[58] C. H. Lampert, H. Nickisch, and S. Harmeling, "Attribute-based classification for zero-shot visual object categorization," *PAMI*, vol. 36, no. 3, pp. 453–465, 2013.

[59] P. Arbelaez, M. Maire, C. Fowlkes, and J. Malik, "Contour detection and hierarchical image segmentation," *PAMI*, vol. 33, no. 5, pp. 898–916, May 2011.

[60] J. Canny, "A computational approach to edge detection," *PAMI*, 1986.

[61] N. Otsu, "A threshold selection method from gray-level histograms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, no. 1, pp. 62–66, 1979.

[62] Y. Jui-Cheng, C. Fu-Juay, and C. Shyang, "A new criterion for automatic multilevel thresholding," IEEE *Transactions on Image Processing*, vol. 4, no. 3, pp. 370–378, 1995.

[63] C. H. Li and C. Lee, "Minimum cross entropy thresholding," *Pattern Recognition*, vol. 26, no. 4, pp. 617 – 625, 1993.

[64] R. TW and C. S, "Picture thresholding using an iterative selection method," *Transactions on Systems, Man, and Cybernetics*, vol. 8, no. 8, pp. 630–632, 1978.

[65] C. A. Glasbey, "An analysis of histogram-based thresholding algorithms," *CVGIP*, vol. 55, no. 6, pp. 532 – 537, 1993.

[66] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, 2016.

[67] L. N. Smith, "Cyclical learning rates for training neural networks," in *WACV*, 2017.

[68] C. Fellbaum, *WordNet: An Electronic Lexical Database*. Bradford Books, 1998.

[69] H. Nam, H. Lee, J. Park, W. Yoon, and D. Yoo, "Reducing domain gap by reducing style bias," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[70] M. Fang, G. Yue, and Q. Yu, "The study on an application of otsu method in canny operator," *Proceedings of the 2009 International Symposium on Information Processing*, 2009.