

Generative Model for Zero-Shot Sketch-Based Image Retrieval

Vinay Kumar Verma*, Aakansha Mishra[‡], Ashish Mishra[†] and Piyush Rai*

*IIT-Kanpur, [‡]IIT-Guwahati, [†]IIT-Madras

vkverma@cse.iitk.ac.in, ak.kkb@iitg.ac.in, mishra@iitm.ac.in, piyush@cse.iitk.ac.in

Abstract

We present a probabilistic model for Sketch-Based Image Retrieval (SBIR) where, at retrieval time, we are given sketches from novel classes, that were not present at training time. Existing SBIR methods, most of which rely on learning class-wise correspondences between sketches and images, typically work well only for previously seen sketch classes, and result in poor retrieval performance on novel classes. To address this, we propose a generative model that learns to generate images, conditioned on a given novel class sketch. This enables us to reduce the SBIR problem to a standard image-to-image search problem. Our model is based on an inverse auto-regressive flow based variational autoencoder, with a feedback mechanism to ensure robust image generation. We evaluate our model on two very challenging datasets, Sketchy, and TU Berlin, with novel train-test split. The proposed approach significantly outperforms various baselines on both the datasets.

1. Introduction

The commonly used approaches to search for an image from a database of images are: (1) Text-based image retrieval, in which we search for an image using a text-based query and (2) Content-based image retrieval (CBIR), in which a related image is used as a query image. Image as the query has a much richer content as compared to text-based query. CBIR gives excellent search results but requires giving a *real* image as the query, which may not always be possible. Often it is more convenient to draw an outline sketch of the image and use that as a query to search for the desired image(s). The retrieval of images by giving the sketch as a query is termed as sketch-based image retrieval (SBIR) [3, 4, 51, 34]. The topic has drawn considerable attention recently. However, existing SBIR systems assume that the class represented by the input sketch at query time was also present in the image-sketch pairs used to train the SBIR model, and consequently, these systems suffer when the input sketch is from a previously *unseen/novel* class.

In this work, we present a method to handle the SBIR task for the unseen/novel class at test time. These novel

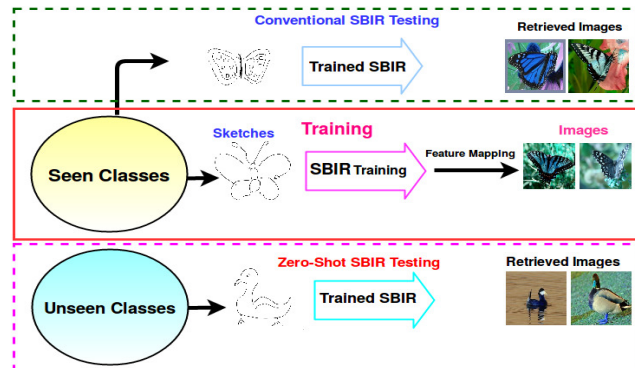


Figure 1. Illustration of Zero-Shot Sketch-Based Image Retrieval (ZS-SBIR)

classes are either absent at the training time or not used in training. This type of setup, to handle the previously unseen classes at test time is called Zero-Shot Learning (ZSL), and has been extensively investigated recently for problems, such as image classification [20, 31, 42, 1], action classification [22, 30], image tagging [54], and visual question answering [35] etc. To the best of our knowledge, the only works that have investigated SBIR in the zero-shot setting include [40, 50]. Among these, [40] used a hashing approach for the ZS-SBIR. This approach is motivated by other ZSL approaches where some side information about unseen classes is present, e.g., their textual description, word2vec or attribute based vectors are used for the knowledge transfer. Recently [50] proposed a vanilla conditional variational autoencoder (CVAE) architecture and adversarial autoencoder for the ZS-SBIR task.

In this paper, we address the drawbacks of existing approaches for SBIR to handle the retrieval of *novel/unseen* class examples. We propose a *conditional* generative model that can generate image features conditioned on the attributes (raw sketch or word2vec[28]) of a given class. Like [50] we also have a generative model, but our approach is significantly different from their model which uses a standard conditional VAE. In contrast, our proposed approach is built upon the *Inverse autoregressive flow* (IAF) based variational autoencoder [18], with a *feedback* based mechanism [17]. The IAF helps to learn the complex latent-space distri-

bution of the images while the feedback mechanism further helps in making the generated image distribution follow the original distribution more closely. The other recently proposed ZS-SBIR approach [40] requires side information in the form of description of the sketch, which may not always be available. In contrast, our proposed approach requires no side information and still performs significantly better than [40]. We also use a residual decoder that helps to learn a complex model with a deeper network. Notably, since we are able to generate images from any specified class, we are able to transform the zero-shot problem into a typical supervised learning problem. The main contributions of this paper can be summarized as follows:

- We propose a sketch-conditioned image generation scheme to solve the ZS-SBIR problem, using a generative model consisting of an inverse autoregressive flow based encoder.
- We leverage a feedback mechanism [17, 20] to encourage the synthesized distribution to be not too far from the original distribution of the observed unlabeled images.
- Unlike the other recently proposed approaches for ZS-SBIR [40], even without any side information (e.g., word2vec based attributes of the classes), our method yields significantly better results as compared to [40].

2. ZS-SBIR Setting

In the zero-shot setting, we partition image dataset into two parts based on sketch classes. One part is the training set which has paired seen-class (S) sketches and image. The second part is test set which has unseen-class (U) sketches only (and no images). Note that the training set is essentially labelled. The training and testing set are mutually exclusive in terms of the sketch classes. In zero-shot setting, we train our model in such a way that it can generalize to unseen class sketches. The mathematical formulation of the zero-shot problem for SBIR is given below:

Let $A = \{(\mathbf{x}_i^{skt}, \mathbf{x}_i^{img}, y_i) | y_i \in \mathcal{Y}\}$ be the triplet consisting of sketch, image and the class label, where \mathcal{Y} is the set of all class labels. We partition the class labels into two disjoint set \mathcal{Y}_{tr} and \mathcal{Y}_{te} for train and test set respectively. Let $A_{tr} = \{\mathbf{x}_i^{skt}, \mathbf{x}_i^{img}, y_i | y_i \in \mathcal{Y}_{tr}\}$ and $A_{te} = \{\mathbf{x}_i^{skt}, \mathbf{x}_i^{img}, y_i | y_i \in \mathcal{Y}_{te}\}$ be the partition of A into train and test set, respectively. Another assumption for the ZS-SBIR is that $A_{tr} \cap A_{te} = \emptyset$ i.e. train and test classes are disjoint. For simplicity, we will represent \mathbf{x}^{skt} as "a" and \mathbf{x}^{img} as "x" throughout this exposition.

3. Background

As discussed earlier, our approach is based on turning the sketch-to-image search problem into an image-to-image

search problem. To this end, we need a model that can generate high-quality images, given a sketch of the class representing that image. This, essentially is a conditional image generation problem. To model the complex distribution of real-world images, we leverage the inverse auto-regressive flow (IAF) based variational autoencoder [18], and adapt it using a feedback mechanism to integrate the information provided by the sketch attribute. Before describing our architecture, we first provide a background of the components we build upon.

3.1. Variational Inference and Learning

Suppose $\mathbf{x} = \{\mathbf{x}^1, \dots, \mathbf{x}^N\}$ be a set of N i.i.d. observations (e.g., N images). Let us denote each sample by \mathbf{x} and assume \mathbf{z} be the latent variable associated with \mathbf{x} . For a given dataset \mathbf{X} , the marginal likelihood of observations is denoted as $\log p(\mathbf{x}) = \sum_{i=1}^N \log p(\mathbf{x}^i)$. The posterior over the latent variable is denoted by $q(\mathbf{z}|\mathbf{x})$. We can define a variational lower bound on the marginal log-likelihood

$$\log p(\mathbf{x}) \geq E_{q(\mathbf{z}|\mathbf{x})}[\log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z}|\mathbf{x})] = L(\mathbf{x}; \theta) \quad (1)$$

where p and q are distributions whose parameters are collectively denoted by θ , and L is the Evidence Lower Bound (ELBO), defined as

$$L(\mathbf{x}; \theta) = \log p(\mathbf{x}) - D_{kl}(q(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}|\mathbf{x})) \quad (2)$$

Maximizing the lower bound $L(\mathbf{x}; \theta)$ w.r.t. θ also maximizes $\log p(\mathbf{x})$ and minimizes $D_{kl}(q(\mathbf{z}|\mathbf{x}) || p(\mathbf{z}|\mathbf{x}))$, where $p(\mathbf{z}|\mathbf{x})$ is the true posterior over the latent variables and $q(\mathbf{z}|\mathbf{x})$ is the approximate posterior (often also called the inference network). In order to infer complex true posterior $p(\mathbf{z}|\mathbf{x})$, we need to have a sufficiently expressive approximation $q(\mathbf{z}|\mathbf{x})$. Normalizing Flows [10] is an idea that helps accomplish this by defining a series of transformations for a latent variable that enable learning sufficiently rich distribution for that variable.

3.2. Normalizing Flow

For the inference network $q(\mathbf{z}|\mathbf{x})$, we need a highly flexible method that captures the complex nature of the true posterior distribution. Normalizing flow is a popular approach used for the variational inference of posterior over latent space. Normalizing flow [10] depends on sequence of invertible mappings for transforming the initial probability density. Suppose z_0 be the initial random variable with a simple probability density function $q(z_0|\mathbf{x})$ and \mathbf{z}_t be the final output of a sequence of invertible transformations f_t on z_0 . \mathbf{z}_t can be computed as: $\mathbf{z}_t = f_t(\mathbf{z}_{t-1}, \mathbf{x}) \forall t = 1, \dots, T$. If Jacobian determinant of each f_t can be computed, then the final probability density function can be computed as:

$$\log q(\mathbf{z}_T|\mathbf{x}) = \log(q_0|\mathbf{x}) - \sum_{t=1}^T \log \det \left| \frac{\partial \mathbf{z}_t}{\partial \mathbf{z}_{t-1}} \right|$$

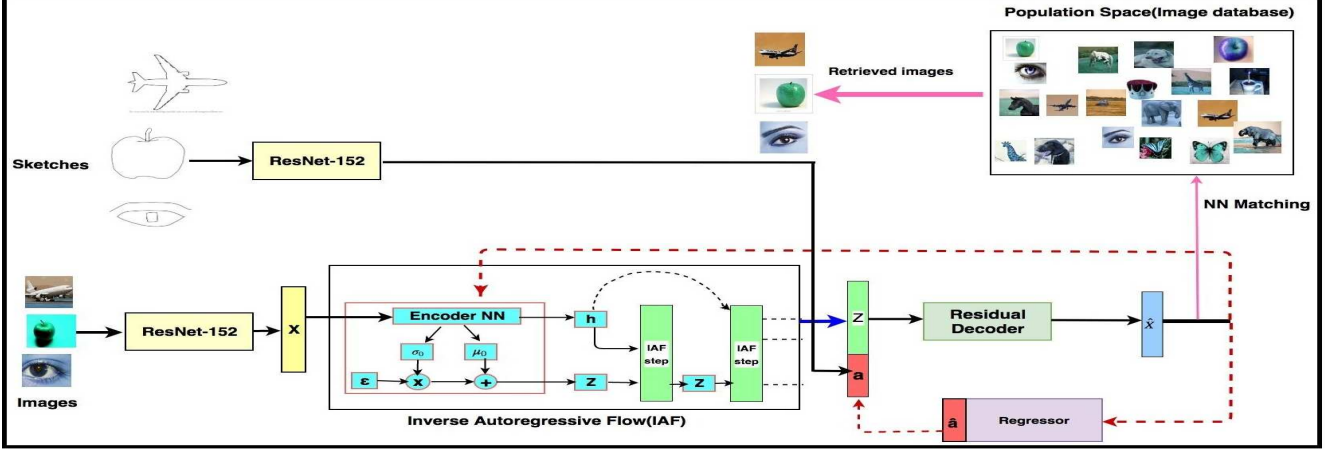


Figure 2. An illustration of our proposed model, based on the IAF architecture and feedback mechanism.

3.3. Inverse Autoregressive Transformations (IAF)

Let v be a variable which is modeled by the Gaussian version of the autoregressive model. Suppose $[\mu(v), \sigma(v)]$ be the representation of function that maps v to the mean μ and variance σ . Due to the autoregressive structure, the Jacobian is lower triangular matrix with zeros on the diagonal. Mean and standard deviation of i^{th} element of v are computed from $v_{1:i-1}$ i.e., previous elements of v . To sample from such a model, we use a sequence of transformations from a noise vector $\epsilon \sim N(0, I)$ to the corresponding vector v as: $v_0 = \mu_0 + \sigma_0 \odot \epsilon_0$ and for $i > 0$ $v_i = \mu_i(v_{1:i-1}) + \sigma_i(v_{1:i-1})\epsilon_i$. Variational inference makes sampling from posterior, such models are not interesting to be directly used for the normalizing flow. Although, the inverse transformation is interesting for normalizing flows, as long as we have $\sigma_i > 0$ the transformation is one-to-one and it can be inverted as: $\epsilon_i = \frac{v_i - \mu_i(v_{1:i-1})}{\sigma_i(v_{1:i-1})}$. Two key observation for IAF as follows:

- As computation of every element ϵ_i does not depend on one another, inverse transformation can be parallelized $\epsilon = \frac{v - \mu(v)}{\sigma_v}$ (subtraction and division are element-wise).
- Inverse autoregressive operation has a simple Jacobian determinant. It is lower triangular matrix. As an outcome, the log-determinant of Jacobian of transformation is simple to compute: $\log \det \left| \frac{\partial \epsilon}{\partial v} \right| = \sum_{i=1}^D -\log \sigma_i(v)$

3.4. IAF step

As shown in Fig. 2, the output of initial encoder network is μ_0, σ_0 and one extra output h which is consider as one extra input to each subsequent step in the flow. In other word, the parameters of encoder are refined iteratively based on output of previous step μ_0, σ_0 and h . The

sampled vector from latent space of initial encoder is defined as: $z_0 = \mu_0 + \sigma_0 \odot \epsilon$. Where $\epsilon \sim N(0, I)$. After t steps the refinement of sample z_0 is recursively defined as: $z_t = \mu_t + \sigma_t \odot z_{t-1}$. In this sequential step the predicted posterior fits more closely to the true posterior.

Finding an appropriate latent space for sampling is a crucial part of generative models as in variational autoencoder (VAE). VAE based generative models compute latent space in one step which may not be sufficient to capture a complex distribution. So the distribution of the predicted posterior and true posterior could be different with adequate margin. Whereas in IAF based variational autoencoder, predicted posterior are transformed to the true posterior using some simple sequential transformation. This sequence of simple transformation can be reduced to any complex distribution. Therefore using an auto-regressive method we can reduce the difference between the distribution of the estimated posterior and true posterior as compare to standard VAE.

4. Zero-Shot Sketch-Based Image Retrieval

In this section, we describe the various components of our proposed model. Again, note that the goal is to learn to generate high-quality images, given the sketch and optionally other side information (e.g., word2vec description of the class).

4.1. Inverse Autoregressive Flow-Based Encoder

Learning the complex distribution of z in the high dimensional latent space is not feasible by a single step transformation. Therefore, in the plain VAE, the approximate posterior can be far away from the true posterior of z . IAF provides a way to learn the complex distribution by using the chain of simple transformation. The final latent variable z can be given as:

$$z_T = f_T(\dots f_2(f_1(f_0(z_0)))) \dots \quad (3)$$

Here each f_i is simple transformation function and are invertible in nature. Figure 2 shows the pipeline of IAF architecture.

4.2. VAE with feedback mechanism

In our model, the encoder consists of standard encoder coupled with an IAF module. The output of IAF based encoder is refinement of the latent code which is initialized by standard encoder, denoted as $p_E(z_t|x)$ with parameters θ_E . The regressor output distribution is denoted as $p_R(\mathbf{a}|x)$, and the VAE loss function is given by (assuming the regressor to be fixed):

$$\mathcal{L}_{VAE}(\theta_E, \theta_G) = -\mathbb{E}_{p_E(z_t|x), p(\mathbf{a}|x)}[\log p_G(\mathbf{x}|z_t, \mathbf{a})] + \text{KL}(p_E(z_t|x)||p(z_t)) \quad (4)$$

where the first term on the R.H.S. is generator’s reconstruction error and the second term promotes the estimated posterior to be close to the prior.

4.2.1 Regressor/Cyclic-consistency Loss

In our proposed model, the regressor, defined by a probabilistic model $p_R(\mathbf{a}|x)$ with parameters θ_R , is a feed-forward neural network that learns to project the example $x \in \mathbb{R}^D$ to its corresponding class-attribute vector $\mathbf{a} \in \mathbb{R}^L$. The objective of the regressor is to minimize the cyclic-consistency loss. The regressor is learned using two sources of data:

- Labeled examples $\{\mathbf{x}_n, \mathbf{a}_n\}_{n=1}^{N_S}$ from the seen classes, on which we can define a supervised loss, given by

$$\mathcal{L}_{Sup}(\theta_R) = -\mathbb{E}_{\mathbf{x}_n} [p_R(\mathbf{a}_n|\mathbf{x}_n)] \quad (5)$$

- Synthesized examples $\hat{\mathbf{x}}$ from the generator, for which we can define an unsupervised loss, given by

$$\mathcal{L}_{Unsup}(\theta_R) = -\mathbb{E}_{p_{\theta_G}(\hat{\mathbf{x}}|z_t)p(z_t)p(\mathbf{a})} [p_R(\mathbf{a}|\hat{\mathbf{x}})] \quad (6)$$

The weighted combination of supervised and unsupervised loss is defined as the overall objective to minimize the cyclic-consistency/regressor loss:

$$\min_{\theta_R} \mathcal{L}_R = \mathcal{L}_{Sup} + \lambda_R \cdot \mathcal{L}_{Unsup} \quad (7)$$

4.2.2 Regressor-Driven Learning

Regressor-Driven learning helps to minimize the cyclic-consistency loss and guide the generator to generate high-quality samples. The cyclic loss encourages the decoder/generator to generates example $\hat{\mathbf{x}}$ coherent with its sketch feature vector \mathbf{a} . This is done using a loss function described below.

In the first case, suppose the generator generates low-quality samples. Then the regressor will incur a high cyclic

loss for these samples. In this case, the regressor assumes that it has optimal parameters and will not regress to the correct value. This loss occurs because of the bad quality samples generated by the generator. Minimizing this loss w.r.t θ_G helps generator to improve the samples quality. The objective function is given by

$$\mathcal{L}_c(\theta_G) = -\mathbb{E}_{p_G(\hat{\mathbf{x}}|z_t, \mathbf{a})p(z_t)p(\mathbf{a})} [\log p_R(\mathbf{a}|\hat{\mathbf{x}})] \quad (8)$$

The other loss which acts as a regularizer that encourages the generator to generate a good class-specific sample even from a random z_t drawn from the prior distribution $p(z_t)$ and combined with the sketch from $p(\mathbf{a})$ is

$$\mathcal{L}_{Reg}(\theta_G) = -\mathbb{E}_{p(z_t)p(\mathbf{a})} [\log p_G(\hat{\mathbf{x}}|z_t, \mathbf{a})] \quad (9)$$

The above two loss functions help us increase the coherence of $\hat{\mathbf{x}} \sim p_G(\hat{\mathbf{x}}|z, \mathbf{a})$ with class-attribute \mathbf{a} . A third loss function is used to ensure that the sampling distribution $p(z_t)$ and the distribution obtained from the generated examples $p_E(z_t|\hat{\mathbf{x}})$ follow the same distribution.

$$\mathcal{L}_E(\theta_G) = -\mathbb{E}_{\hat{\mathbf{x}} \sim p_G(\hat{\mathbf{x}}|z_t, \mathbf{a})} \text{KL}[(p_E(z_t|\hat{\mathbf{x}})||q(z_t))] \quad (10)$$

Hence the complete learning objective for the generator and encoder is given by,

$$\min_{\theta_G, \theta_E} \mathcal{L}_{VAE} + \lambda_c \cdot \mathcal{L}_c + \lambda_{reg} \cdot \mathcal{L}_{Reg} + \lambda_E \cdot \mathcal{L}_E \quad (11)$$

4.3. Residual Decoder

The proposed decoder is a combination of the deep and shallow network. The deep network is responsible for the better reconstruction of visual space while shallow network reduces over-fitting. This architecture is motivated by ResNet [13] where the network has skip connections. These skip connections provide more paths to the network for information propagation. While some paths are deeper, others are shallow [14]. If in the deeper path the gradient vanishing or explosion problem occurs, the shallow paths still work, and proper gradient flows in the backward direction. In the residual network, the output of a neural network layer is given by $f_o(x) = f_{in}(x) + x$, (here $f_{in}(x)$, is the direct output), i.e., the output does not only depend on the current layer neural network, but it depends on input as well.

5. Related Work

Images have rich and vibrant content, while a sketch only provides rough information like shape and size. It is easy for a human to match the sketch from the image, but for machines, this is a very complex task. Since for an algorithm, it is very difficult to learn the features that are invariant to color, shape, size, pose, etc. The common pipeline for SBIR is to project the images and sketches in common subspace such that the same class images and sketches are close to each other on some metric space. Then any similarity metric

can be used for the retrieval task. Most of the traditional approaches for SBIR have used hand-crafted features such as gradient field HOG descriptor [6], SIFT [25] and SURF [2] etc. [32] proposed a dynamic programming based method for SBIR which is effective in translation, rotation, and scale (similarity). Recent advancement of deep learning provides an automatic feature extraction technique which learns the pose and color invariant feature. Recently [52, 38, 40, 50] have used deep feature for SBIR task. Instead of finding the common subspace other approach projects the sketch space to image space or vice versa such that the information gap between the sketches and the real images are minimum [16, 32].

Recently zero-shot learning drew more attention due to its capability of classifying a novel class object during the test phase. In the ZSL each class is associated with side information like description of the class, attribute or unsupervised word embedding (Word2vec [28], Glove [33], etc.). This side information of the class is called the semantic features/attributes. In ZSL, the core concept is to learn projection between class feature and side information, using labeled seen class data only. We can categories all proposed models for ZSL in three types based on projection. The most popular work learns the projection between visual space to semantic space and vice-versa [49, 1, 31, 45, 42, 29]. Another popular approach projects the visual and semantic features in a shared subspace such that same class visual features and semantic attributes map closer, whereas different class visual features and semantic attributes are well-separated [44].

Recently generative models are emerging as the most popular approach for zero-shot image classification. This type of approach gaining popularity because of its ability to synthesize the unseen class sample and can reduce the ZSL problem to a supervised learning problem. These approach learns the data distribution based on the given conditions [11, 42, 20, 46]. Most of the previous methods for zero-shot learning are focused on image classification. However, a few models are used for zero-shot action classification, zero-shot image tagging and zero-shot multi-label learning as well [21, 48, 30, 49, 54, 9].

Recently [40, 50] have proposed a model for the ZS-SBIR. [40] proposed a hashing based approach for the ZS-SBIR. The hashing architecture is based on the multi-model deep network. [50] proposed a generative model for the ZS-SBIR based on the CVAE architecture. The proposed approach is also a generative in nature based on IAF to get the improved variational inference [18]. Here our encoder is based on the IAF architecture that learns the complex latent encoding of the input into the latent space. It can learn the complex distribution with the simple sequential transformation. Also, we are using the β -VAE [15] architecture for the disentangled representation. The residual decoder is

used that gives the better generation of the sample because it can flow the gradient with the deeper layers. In the proposed approach the external feedback mechanism provides the feedback to the encoder about the generation quality. Hence the generator has better guidance for generating the robust sample.

6. Experiments and Results

To show the effectiveness of our proposed model we use two challenging datasets: Sketchy [38] and TU-Berlin [7]. Originally, Sketchy dataset [38] contains 75471 hand-drawn sketches and 12500 corresponding images from 125 classes. [23] have provided 60502 more real images from all 125 classes, which extends the original dataset. TU-Berlin extended [7] is a large scale dataset having 20000 sketches and 204489 images from 250 different categories provided by [23, 53].

The visual features for images and sketches are extracted using ResNet-152 [13], pretrained on ImageNet [37] dataset. The sketches and image features are extracted from the last fully connected layer. It gives 2048-dimensional feature vectors. We believe that further finetuning on this dataset on ResNet-152 architecture will give better performance. The visual features of the sketches are used as a class attributes in our proposed generative model.

6.1. Sketchy Dataset (Extended)

For fair comparison with the recent work [40, 50], we have two splits of the dataset. [40] randomly selected 25 classes of sketches as the test set (A_{te}) and the remaining labeled 100 classes are used as the training set (A_{tr}). Here $A_{tr} \cap A_{te} = \phi$, i.e. train and test class are disjoint. We have another split of the dataset similar to [50], this contains 104 classes in training, and used 21 classes images and sketches as a test set.

Random split proposed by [40] is not the realistic Zero-Shot setting [47]. Since in random split test set may have some classes that are present in the ImageNet class. Since we are using the ImageNet pre-trained model for the feature extraction and this training is done in a supervised manner. This violates the assumption that A_{te} are the unseen classes. Therefore it is not the exact Zero-Shot setting. The split proposed by [50] is the realistic setup for ZS-SBIR, where the split is done in such a way that any of the A_{te} classes are not present in the ImageNet dataset. In our setup for training, we need a paired image and sketch set. To make the paired set, we selected a random image and sketch from the same training class and paired them. This process repeated 1000 time, i.e., each class in the training set has 1000 pair of data point. We are comparing our model with the previous approach in their original setup; therefore, we are using both the split.

6.2. TU Berlin Dataset (Extended)

Similar to [40] for the fair comparison randomly 30 classes are selected for the A_{te} and remaining 220 classes are used for the A_{tr} . This dataset is highly biased; few classes have large examples while few have only limited samples. In the Zero-shot setup, learning with biased data is a very hard problem. Therefore from training we removed the biases. For doing so, we are sampling the equal number of image and sketch sample pairs from each class. In the testing, we selected the class that has more than 400 samples. For making image and sketch pair, we follow the same pattern mentioned in the previous section. Here again, each class has a 1500 pair of image and sketch.

6.3. Implementation details

In our model, we have three components, Encoder (E), Generator (G) and Regressor (R). The encoder is based on the IAF architecture, refer to figure-2. The encoder contains the two fully connected layers of size 4096 followed by one layer that gives the μ and σ this passed to 3 layers IAF architecture. The generator has five layers of the fully connected neural network (NN) with the residual connection. It is a combination of the deep and shallow network. Here sigmoid activation is used. All layers are of the same size of 6144. Regressor takes the reconstructed samples \hat{x} and regresses the sketch. It uses the two-layer fully connected NN of size 4096. The learning rate (η) is set as a stepwise decreasing rate. Initially for the 5 epoch $\eta = 0.001$ then after each 10 epoch it changed to [0.0005,0.0001,0.00001]. Here instead of $\mathcal{N}(0, I)$ prior, we found from the validation data $p \sim \mathcal{N}(0, 0.005)$ gives the better performance. Also, for the ablation, we experiment with the plain autoencoder. The autoencoder used contains the same architecture as the IAF encoder with feedback connection; only the difference is that the dimension of z is zero. Therefore the generated sample is deterministic and depends only on the given sketch feature α .

6.4. Training and Testing

There are two modules in the model, IAF-VAE and regressor. We are alternately optimizing the IAF-VAE and regressor. These two module helps each other to learn the robust generator. In the VAE training, we are minimizing the loss w.r.t. E and G 's parameters, and for regressor training, we are minimizing the regressor's loss w.r.t R 's parameter only. The alternate optimization is done Until convergence. The complete setup is for the zero-shot learning; therefore the testing is performed from the unseen class sketch to unseen class image. In the testing phase each x^{skt} is concatenated with the $z \sim \mathcal{N}(0, 0.005)$ and generate the c samples using generator G . Now from these c , samples find the image that gives maximum cosine similarity. The similarity of




Figure 3. Top-6 retrieved images for randomly chosen five sketches for ZS-SBIR.

the sample x_i from the query sketch x^{skt} can be given as:

$$s(x^{skt}, x_i) = \max_{t=1:c} \text{cosine} \left(G(\theta(x_t^{skt})), \theta(x_i) \right) \quad (12)$$

Here, x_i is the image in the query database. Equation-12 is repeated for each image and find the K samples with maximum similarity scores for the $top@K$ retrieval. θ is the ResNet-152 model and gives the feature vector for each image, and G is the generator.

Result Analysis with existing methods

Since best of our knowledge, only two very recent works ZSIH [40] and CVAE [50] have been proposed for ZS-SBIR. Therefore for supporting the performance of the proposed model, we compare the performance of our model with several other state-of-the-art. We have analyzed two types of baselines methods, 1- Sketch-Based Image Retrieval (SBIR) Baselines and 2- Zero-Shot Learning (ZSL) baselines.

Sketch base image retrieval baselines (SBIR)

Several approaches have been proposed for SBIR. We compare our model with Siamese-1 [12], Siamese-2 [34], Coarse-grained triplet [39], Fine-grained triplet [38], DSH [24], SaN [52], GN Triplet [38], 3D Shape [43], Siamese CNN [?]. Since these baselines are not originally proposed for zero-shot setting, [40] provides these baseline for the zero-shot setting. We have taken these baseline result directly from the paper [40].

Zero-Shot baselines (ZSL)

The most of the existing approaches for zero-shot learning are proposed for the zero-shot image classification. We select a set of zero-shot learning approaches as baseline to

Type	Method	Sketchy Dataset		TU Berlin Dataset	
		Precision@100	mAP@all	Precision@100	mAP@all
SBIR	Softmax Baseline	0.176	0.099	0.139	0.083
	Siamese CNN [34]	0.183	0.143	0.153	0.122
	SaN [52]	0.129	0.104	0.112	0.096
	GN Triplet [38]	0.310	0.211	0.241	0.189
	3D Shape [43]	0.070	0.062	0.063	0.057
	DSH (64 bits) [24]	0.227	0.164	0.198	0.122
Zero-Shot	CMT [41]	0.096	0.084	0.082	0.065
	DeViSE [8]	0.078	0.071	0.075	0.067
	SSE [55]	0.154	0.108	0.133	0.096
	JLSE [56]	0.178	0.126	0.165	0.107
	SAE [55]	0.302	0.210	0.210	0.161
	DSH [24]	0.217	0.165	0.174	0.139
	ZSIH [40]	0.340	0.254	0.291	0.220
Feedback-Auto	GZS-SBIR(Our)	0.305	0.253	0.281	0.187
Feedback-VAE	GZS-SBIR(Our)	0.358	0.289	0.334	0.238

Table 1. Precision@100 and mAP@all results on the traditional SBIR and ZSL method on the ZS-SBIR setup. Feedback-Auto is the IAF autoencoder with the feedback mechanism and Feedback-VAE is the IAF-VAE with the feedback mechanism.

compare with our proposed model. These ZSL baseline approaches are CMT [41], DeVISE [8], SAE [19], SSE [55], ESZSL [36], CAAE [27], JLSE [56], DAP [21]. The baseline results are borrowed from the [50, 40]. We again reproduce the the baseline results reported in the table-2.

The recent work on the ZS-SBIR is ZSIH [40] and CVAE [50]. ZSIH [40] shows the experiment on the TU-Berlin and Sketchy dataset. They reported the result of precision@100 and mAP@all for all datasets. [40] using the word2vec[28] as side information in their model. As mentioned earlier we are not using any side information but have significantly better result compare to ZSIH. The comparison result with the baseline and ZSIH are shown in the table-1. We can see without any side information our approach performs significantly better than all the previous approach that used the side information. Also, we experimented the proposed approach with autoencoder only and found that the proposed VAE model significantly outperforms the autoencoder model. Please refer to table-1 for the more details.

Another approach CVAE [50] has proposed a generative model for ZS-SBIR; they showed the result on the Sketchy dataset. CVAE suggested the realistic train-test split similar to [47] for the ZS-SBIR. [50] evaluated the performance of the model over precision@200 and mAP@200 metric. We are following the same setup to compare our result with CVAE. Our result on the Sketchy-dataset shows that the proposed approach is significantly better compare to CVAE. CVAE not using any side information and without using any side information (e.g., word2vec), our method shows the 3.0% and 5.8% relative improvement over precision@200 and mAP@200 metric.

In the Figure-3 we have illustrated the top-6 retrieved result using the unseen class sketches from the image database. Retrieved images are closely matched to the out-

Sketchy Dataset			
Type	Method	Precision@200	mAP@200
SBIR	Baseline	0.106	0.054
	Siamese-1 [5]	0.243	0.134
	Siamese-2 [34]	0.251	0.149
	Coarse-grained triplet [39]	0.169	0.083
	Fine-grained triplet [38]	0.155	0.081
	DSH ¹ [24]	0.153	0.059
ZS-SBIR	DAP [21]	0.066	0.022
	ESZSL [36]	0.187	0.117
	SAE [19]	0.238	0.136
	CAAE CAAE [27]	0.260	0.156
	CVAE [50]	0.333	0.225
Feedback-Auto	GZS-SBIR(our)	0.288	0.191
Feedback-VAE	GZS-SBIR(our)	0.343	0.238

Table 2. Precision@200 and mAP@200 results on the traditional SBIR and ZSL method on the ZS-SBIR setup. This table follow the realistic train-test split.

line of the sketches. Since our model learns the mapping between sketches and images based on components. So it may retrieve some other class images which are significantly similarity with the sketch. In Figure-3 we can see that for helicopter sketch our model retrieve the fish because the outline of sketches of the helicopter and fish are very similar. Also, we have shown the t-SNE [26] plot of the original data and the reconstructed data for the Sketchy [38] dataset. In the t-SNE plot, we can observe that the generated samples for the novel classes are not as good as the original one. But the generated sample nearly follow the same distribution as the original one. The few class samples are as good as the original samples. Please refer to figure-4 for the t-SNE plot.

Ablation Study

We now show the significance of the different components of the model as compared to the basic VAE model.




Figure 4. t-SNE plot for the original and the reconstructed sample for the sketchy dataset without using any side information.

We have found that the proposed approach outperforms by a significant margin across all the dataset. Even though without using any side information we are performing better than the previous approach [40] that used the side information. In the section-[6.5] we are showing the ablation with and without VAE. Also in section-[6.6] we are showing the significance of the IAF component.

6.5. With/Without VAE

We also perform the ablation analysis over the different component of the proposed approach. In the first experiment, we compare the performance of autoencoder with the proposed VAE architecture. We found that the proposed model with the VAE component always outperforms compare to autoencoder architecture. Using the feedback-VAE architecture on the Sketchy-dataset the model shows the 20% and 16% relative improvement on the precision@100 and mAP@all metric compare to plain autoencoder architecture. The similar pattern we observe for the Tu-Berlin dataset also. The feedback-VAE shows the 17.8% and 26.3% relative improvement over the plain autoencoder architecture on the precision@100 and mAP@all metric.

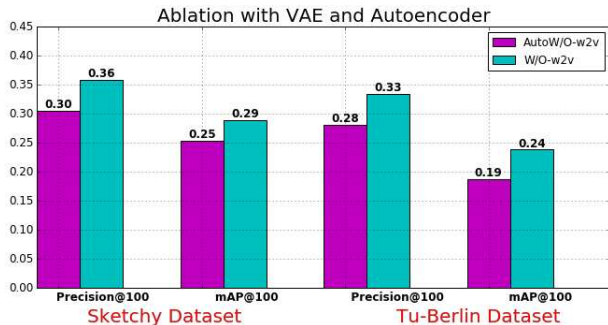


Figure 5. Ablation of our proposed VAE with the plain autoencoder without any side information. *AutoW/O-w2v*: Autoencoder without word2vec, *W/O-w2v*: Proposed approach without word2vec.

Please refer to figure-5 for comparison details.

6.6. With/Without IAF

We present the ablation study with IAF and without IAF component, without using any side information. We have found that if we remove the IAF component, the performance drop is significant as compared to with-IAF. For the Sketchy dataset, we reported in Table-3 that with-IAF component, the precision@100 and mAP@all are 0.358 and 0.289, respectively. If we remove the IAF component, the performance drop is significant, and precision@100 and mAP@all are 0.313 and 0.261, respectively. Therefore we have 12.6% and 9.7% relative drop in the performance without-IAF. We also observed a similar pattern on the TU-Berlin dataset. Here in Table-3 with-IAF we have 0.334 and 0.238, precision@100 and mAP@all, respectively. But if we drop the IAF component, our precision@100 and mAP@all are 0.294 and 0.198, respectively. Therefore we have 12.0% and 16.8% performance drop, respectively. Please refer to Table-3 for the more details.

Type	Sketchy Dataset		TU Berlin Dataset	
	Precision@100	mAP@all	Precision@100	mAP@all
W/O-w2v	0.313	0.261	0.294	0.198
W/O-w2v+IAF	0.358	0.289	0.334	0.238
Improvement (%)	12.6%	9.7%	12.0%	16.8%

Table 3. **Ablation study:** Precision@100 and mAP@all results on the Sketchy and TU-Berlin dataset without-IAF and with-IAF.

7. Conclusion

In this paper, we addressed the Zero-Shot Sketch-Based Image Retrieval problem, which is a challenging and more realistic setting as compared to the conventional SBIR. The proposed generative approach can solve the SBIR problem when the classes are growing with time, and does not require all classes to be present at the training time. We have found that the proposed approach, based on the IAF architecture with the feedback mechanism, generates high-quality samples of the novel classes. Moreover, without using any side information, our proposed generative model can retrieve novel class examples and gives state-of-art results on benchmark datasets. In this work, we assume that the test query comes from the unseen classes only. In future, it will be interesting to explore the *Generalized ZS-SBIR* problem where test query can come from the seen as well as unseen classes. Also, the domain shift is a critical problem in the ZSL. It will be an interesting direction of future work to handle the domain shift for zero-shot SBIR. The recent model shows a significant improvement in ZS-SBIR using the side information. In future, it will also be exciting to explore the model with the help of side information.

References

- [1] Z. Akata, S. Reed, D. Walter, H. Lee, and B. Schiele. Evaluation of output embeddings for fine-grained image classification. In *CVPR*, pages 2927–2936, 2015. 1, 5
- [2] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006. 5
- [3] X. Cao, H. Zhang, S. Liu, X. Guo, and L. Lin. Sym-fish: A symmetry-aware flip invariant sketch histogram shape descriptor. In *ICCV*, pages 313–320, 2013. 1
- [4] Y. Cao, C. Wang, L. Zhang, and L. Zhang. Edgel index for large-scale sketch-based image search. In *CVPR*, 2011. 1
- [5] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, pages 539–546, 2005. 7
- [6] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *CVPR*, pages 886–893, 2005. 5
- [7] M. Eitz, J. Hays, and M. Alexa. How do humans sketch objects? *ACM Trans. Graph.*, 31:44–1, 2012. 5
- [8] A. Frome, G. S. Corrado, J. Shlens, S. Bengio, J. Dean, T. Mikolov, et al. Devise: A deep visual-semantic embedding model. In *NIPS*, pages 2121–2129, 2013. 7
- [9] A. Gaure, A. Gupta, V. K. Verma, and P. Rai. A probabilistic framework for zero-shot multi-label learning. *The Conference on Uncertainty in Artificial Intelligence (UAI)*, 2017. 5
- [10] M. Germain, K. Gregor, I. Murray, and H. Larochelle. Made: Masked autoencoder for distribution estimation. In *ICML*, pages 881–889, 2015. 2
- [11] Y. Guo, G. Ding, J. Han, and Y. Gao. Synthesizing samples for zero-shot learning. *IJCAI*, 2017. 5
- [12] R. Hadsell, S. Chopra, and Y. LeCun. Dimensionality reduction by learning an invariant mapping. In *null*, pages 1735–1742, 2006. 6
- [13] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 4, 5
- [14] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. 4
- [15] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017. 5
- [16] R. Hu and J. Collomosse. A performance evaluation of gradient field hog descriptor for sketch based image retrieval. *CVIU*, 117(7):790–806, 2013. 5
- [17] Z. Hu, Z. Yang, X. Liang, R. Salakhutdinov, and E. P. Xing. Toward controlled generation of text. 2017. 1, 2
- [18] D. P. Kingma, R. Salimans, T.J. X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *NIPS*, pages 4743–4751, 2016. 1, 2, 5
- [19] E. Kodirov, T. Xiang, and S. Gong. Semantic autoencoder for zero-shot learning. *arXiv preprint arXiv:1704.08345*, 2017. 7
- [20] V. Kumar Verma, G. Arora, A. Mishra, and P. Rai. Generalized zero-shot learning via synthesized examples. In *CVPR*, 2018. 1, 2, 5
- [21] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *PAMI*, 36(3):453–465, 2014. 5, 7
- [22] J. Liu, B. Kuipers, and S. Savarese. Recognizing human actions by attributes. In *CVPR*, pages 3337–3344, 2011. 1
- [23] L. Liu, F. Shen, Y. Shen, X. Liu, and L. Shao. Deep sketch hashing: Fast free-hand sketch-based image retrieval. In *CVPR*, pages 2862–2871, 2017. 5
- [24] L. Liu, F. Shen, Y. Shen, X. Liu, and L. Shao. Deep sketch hashing: Fast free-hand sketch-based image retrieval. *CoRR*, abs/1703.05605, 2017. 6, 7
- [25] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999. 5
- [26] L. v. d. Maaten and G. Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9:2579–2605, 2008. 7
- [27] A. Makhzani, J. Shlens, N. Jaitly, and I. Goodfellow. Adversarial autoencoders. In *ICLR*, 2016. 7
- [28] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, pages 3111–3119, 2013. 1, 5, 7
- [29] A. Mishra, M. Reddy, A. Mittal, and H. A. Murthy. A generative model for zero shot learning using conditional variational autoencoders. *CVPRW*, 2018. 5
- [30] A. Mishra, V. K. Verma, M. S. K. Reddy, A. Subramaniam, P. Rai, and A. Mittal. A generative approach to zero-shot and few-shot action recognition. *WACV*, pages 372–380, 2018. 1, 5
- [31] M. Norouzi, T. Mikolov, S. Bengio, Y. Singer, J. Shlens, A. Frome, G. S. Corrado, and J. Dean. Zero-shot learning by convex combination of semantic embeddings. *arXiv preprint arXiv:1312.5650*, 2013. 1, 5
- [32] S. Parui and A. Mittal. Similarity-invariant sketch-based image retrieval in large databases. In *ECCV*, pages 398–414, 2014. 5
- [33] J. Pennington, R. Socher, and C. Manning. Glove: Global vectors for word representation. In *EMNLP*, pages 1532–1543, 2014. 5
- [34] Y. Qi, Y.-Z. Song, H. Zhang, and J. Liu. Sketch-based image retrieval via siamese convolutional neural network. In *ICIP*, 2016. 1, 6, 7
- [35] S. K. Ramakrishnan, A. Pal, G. Sharma, and A. Mittal. An empirical evaluation of visual question answering for novel objects. *CoRR*, abs/1704.02516, 2017. 1
- [36] B. Romera-Paredes and P. Torr. An embarrassingly simple approach to zero-shot learning. In *ICML*, 2015. 7
- [37] 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*, 115(3):211–252, 2015. 5
- [38] P. Sangkloy, N. Burnell, C. Ham, and J. Hays. The sketchy database: learning to retrieve badly drawn bunnies. *ACM Transactions on Graphics (TOG)*, 35:119, 2016. 5, 6, 7
- [39] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Pro-*

- ceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. 6, 7
- [40] Y. Shen, L. Liu, F. Shen, and L. Shao. Zero-shot sketch-image hashing. In *CVPR*, June 2018. 1, 2, 5, 6, 7, 8
 - [41] R. Socher, M. Ganjoo, C. D. Manning, and A. Ng. Zero-shot learning through cross-modal transfer. In *NIPS*, pages 935–943, 2013. 7
 - [42] V. K. Verma and P. Rai. A simple exponential family framework for zero-shot learning. In *ECML-PKDD*, 2017. 1, 5
 - [43] F. Wang, L. Kang, and Y. Li. Sketch-based 3d shape retrieval using convolutional neural networks. *CoRR*, abs/1504.03504, 2015. 6, 7
 - [44] Q. Wang and K. Chen. Zero-shot visual recognition via bidirectional latent embedding. *IJCV*, 124(3):356–383, 2017. 5
 - [45] W. Wang, Y. Pu, V. K. Verma, K. Fan, Y. Zhang, C. Chen, P. Rai, and L. Carin. Zero-shot learning via class-conditioned deep generative models. *AAAI*, 2018. 5
 - [46] Y. Xian, T. Lorenz, B. Schiele, and Z. Akata. Feature generating networks for zero-shot learning. 2018. 5
 - [47] Y. Xian, B. Schiele, and Z. Akata. Zero-shot learning - the good, the bad and the ugly. In *CVPR*, 2017. 5, 7
 - [48] X. Xu, T. Hospedales, and S. Gong. Semantic embedding space for zero-shot action recognition. In *ICIP*, 2015. 5
 - [49] X. Xu, T. Hospedales, and S. Gong. Transductive zero-shot action recognition by word-vector embedding. *IJCV*, 123(3):309–333, 2017. 5
 - [50] S. K. Yelamathi, S. K. Reddy, A. Mishra, and A. Mittal. A zero-shot framework for sketch-based image retrieval. *ECCV*, 2018. 1, 5, 6, 7
 - [51] Q. Yu, F. Liu, Y.-Z. Song, T. Xiang, T. M. Hospedales, and C. C. Loy. Sketch me that shoe. In *CVPR*, pages 799–807. IEEE Computer Society, 2016. 1
 - [52] Q. Yu, Y. Yang, F. Liu, Y.-Z. Song, T. Xiang, and T. M. Hospedales. Sketch-a-net: A deep neural network that beats humans. *IJCV*, pages 411–425, 2017. 5, 6, 7
 - [53] H. Zhang, S. Liu, C. Zhang, W. Ren, R. Wang, and X. Cao. Sketchnet: Sketch classification with web images. In *CVPR*, pages 1105–1113, 2016. 5
 - [54] Y. Zhang, B. Gong, and M. Shah. Fast zero-shot image tagging. *CVPR*, 2016. 1, 5
 - [55] Z. Zhang and V. Saligrama. Zero-shot learning via semantic similarity embedding. In *CVPR*, pages 4166–4174, 2015. 7
 - [56] Z. Zhang and V. Saligrama. Zero-shot learning via joint latent similarity embedding. In *CVPR*, pages 6034–6042, 2016. 7