

Compositional Learning of Image-Text Query for Image Retrieval

Muhammad Umer Anwaar (✉)^{1,2*}, Egor Labintcev^{1,2}*, Martin Kleinsteuber^{1,2}

¹Technische Universität München, Germany

²Mercateo AG, Germany

{umer.anwaar, egor.labintcev}@tum.de, martin.kleinsteuber@mercateo.com

Abstract

In this paper, we investigate the problem of retrieving images from a database based on a multi-modal (image-text) query. Specifically, the query text prompts some modification in the query image and the task is to retrieve images with the desired modifications. For instance, a user of an E-Commerce platform is interested in buying a dress, which should look similar to her friend's dress, but the dress should be of white color with a ribbon sash. In this case, we would like the algorithm to retrieve some dresses with desired modifications in the query dress. We propose an autoencoder based model, *ComposeAE*, to learn the composition of image and text query for retrieving images. We adopt a deep metric learning approach and learn a metric that pushes composition of source image and text query closer to the target images. We also propose a rotational symmetry constraint on the optimization problem. Our approach is able to outperform the state-of-the-art method TIRG [24] on three benchmark datasets, namely: MIT-States, Fashion200k and Fashion IQ. In order to ensure fair comparison, we introduce strong baselines by enhancing TIRG method. To ensure reproducibility of the results, we publish our code here: <https://github.com/ecom-research/ComposeAE>.

1. Introduction

One of the peculiar features of human perception is multi-modality. We unconsciously attach attributes to objects, which can sometimes uniquely identify them. For instance, when a person says *apple* it is quite natural that an image of an apple, which may be green or red in color, forms in their mind. In information retrieval, the user seeks information from a retrieval system by sending a query. Traditional information retrieval systems allow a unimodal

*indicates equal contribution. Umer contributed the rotation in complex space idea, complex projection module and rotationally symmetric loss. Egor and Umer contributed the idea of fusing different modalities before concatenation. Egor came up with the reconstruction loss idea and contributed by testing hypotheses and running numerous experiments.

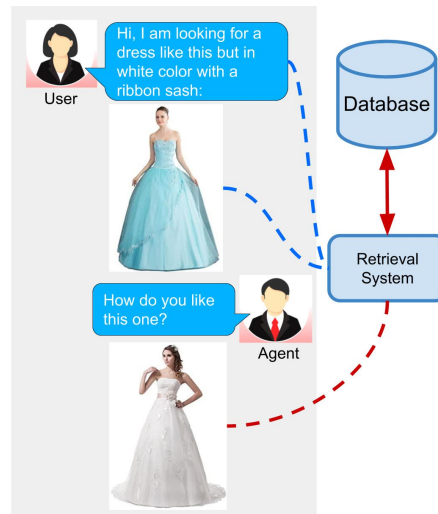


Figure 1. Potential application scenario of this task

query, i.e., either a text or an image. Advanced information retrieval systems should enable the users in expressing the concept in their mind by allowing a multi-modal query.

In this work, we consider such an advanced retrieval system, where users can retrieve images from a database based on a multi-modal query. Concretely, we have an image retrieval task where the input query is specified in the form of an image and natural language expressions describing the desired modifications in the query image. Such a retrieval system offers a natural and effective interface [19]. This task has applications in the domain of E-Commerce search, surveillance systems and internet search. Fig. 1 shows a potential application scenario of this task.

Recently, Vo *et al.* [24] have proposed the *Text Image Residual Gating* (TIRG) method for composing the query image and text for image retrieval. They have achieved state-of-the-art (SOTA) results on this task. However, their approach does not perform well for real-world application scenarios, i.e. with long and detailed texts (see Sec. 4.4). We think the reason is that their approach is too focused on changing the image space and does not give the query text

its due importance. The gating connection takes element-wise product of query image features with image-text representation after passing it through two fully connected layers. In short, TIRG assigns huge importance to query image features by putting it directly in the final composed representation. Similar to [20, 23], they employ LSTM for extracting features from the query text. This works fine for simple queries but fails for more realistic queries.

In this paper, we attempt to overcome these limitations by proposing ComposeAE, an autoencoder based approach for composing the modalities in the multi-modal query. We employ a pre-trained BERT model [1] for extracting text features, instead of LSTM. We hypothesize that by jointly conditioning on both left and right context, BERT is able to give better representation for the complex queries. Similar to TIRG [24], we use a pre-trained ResNet-17 model for extracting image features. The extracted image and text features have different statistical properties as they are extracted from independent uni-modal models. We argue that it will not be beneficial to fuse them by passing through a few fully connected layers, as typically done in image-text joint embeddings [25].

We adopt a novel approach and map these features to a complex space. We propose that the target image representation is an element-wise rotation of the representation of the source image in this complex space. The information about the degree of rotation is specified by the text features. We learn the composition of these complex vectors and their mapping to the target image space by adopting a deep metric learning (DML) approach. In this formulation, text features take a central role in defining the relationship between query image and target image. This also implies that the search space for learning the composition features is restricted. From a DML point of view, this restriction proves to be quite vital in learning a good similarity metric.

We also propose an explicit rotational symmetry constraint on the optimization problem based on our novel formulation of composing the image and text features. Specifically, we require that multiplication of the target image features with the complex conjugate of the query text features should yield a representation similar to the query image features. We explore the effectiveness of this constraint in our experiments (see Sec. 4.5).

We validate the effectiveness of our approach on three datasets: MIT-States, Fashion200k and Fashion IQ. In Sec. 4, we show empirically that ComposeAE is able to learn a better composition of image and text queries and outperforms SOTA method. In DML, it has been recently shown that improvements in reported results are exaggerated and performance comparisons are done unfairly [11]. In our experiments, we took special care to ensure fair comparison. For instance, we introduce several variants of TIRG. Some of them show huge improvements over the

original TIRG. We also conduct several ablation studies to quantify the contribution of different modules in the improvement of the ComposeAE performance.

Our main contributions are summarized below:

- We propose a ComposeAE model to learn the composed representation of image and text query.
- We adopt a novel approach and argue that the source image and the target image lie in a common complex space. They are rotations of each other and the degree of rotation is encoded via query text features.
- We propose a rotational symmetry constraint on the optimization problem.
- ComposeAE outperforms the SOTA method TIRG by a huge margin, i.e., 30.12% on Fashion200k and 11.13% on MIT-States on the Recall@10 metric.
- We enhance SOTA method TIRG [24] to ensure fair comparison and identify its limitations.

2. Related Work

Deep metric learning (DML) has become a popular technique for solving retrieval problems. DML aims to learn a metric such that the distances between samples of the same class are smaller than the distances between the samples of different classes. The task where DML has been employed extensively is the cross-modal retrieval, i.e. retrieving images based on text query and getting captions from the database based on the image query [25, 9, 29, 2, 8, 27].

In the domain of Visual Question Answering (VQA), many methods have been proposed to fuse the text and image inputs [20, 18, 17]. We review below a few closely related methods. Relationship [20] is a method based on relational reasoning. Image features are extracted from CNN and text features from LSTM to create a set of relationship features. These features are then passed through a MLP and after averaging them the composed representation is obtained. FiLM [18] method tries to “influence” the source image by applying an affine transformation to the output of a hidden layer in the network. In order to perform complex operations, this linear transformation needs to be applied to several hidden layers. Another prominent method is parameter hashing [17] where one of the fully-connected layers in a CNN acts as the dynamic parameter layer.

In this work, we focus on the image retrieval problem based on the image and text query. This task has been studied recently by Vo *et al.* [24]. They propose a gated feature connection in order to keep the composed representation of query image and text in the same space as that of the target image. They also incorporate a residual connection which learns the similarity between concatenation of image-text features and the target image features. Another

simple but effective approach is Show and Tell[23]. They train a LSTM to predict the next word in the sequence after it has seen the image and previous words. The final state of this LSTM is considered the composed representation. Han *et al.* [6] presents an interesting approach to learn spatially-aware attributes from product description and then use them to retrieve products from the database. But their text query is limited to a predefined set of attributes. Nagarajan *et al.* [12] proposed an embedding approach, “Attribute as Operator”, where text query is embedded as a transformation matrix. The image features are then transformed with this matrix to get the composed representation.

This task is also closely related with interactive image retrieval task [4, 22] and attribute-based product retrieval task [28, 24]. These approaches have their limitations such as that the query texts are limited to a fixed set of relative attributes [28], require multiple rounds of natural language queries as input [4, 22] or that query texts can be only one word i.e. an attribute [6]. In contrast, the input query text in our approach is not limited to a fixed set of attributes and does not require multiple interactions with the user. Different from our work, the focus of these methods is on modeling the interaction between user and the agent.

3. Methodology

3.1. Problem Formulation

Let $\mathcal{X} = \{x_1, x_2, \dots, x_n\}$ denote the set of query images, $\mathcal{T} = \{t_1, t_2, \dots, t_n\}$ denote the set of query texts and $\mathcal{Y} = \{y_1, y_2, \dots, y_n\}$ denote the set of target images. Let $\psi(\cdot)$ denote the pre-trained image model, which takes an image as input and returns image features in a d -dimensional space. Let $\kappa(\cdot, \cdot)$ denote the similarity kernel, which we implement as a dot product between its inputs. The task is to learn a composed representation of the image-text query, denoted by $g(x, t; \Theta)$, by maximising

$$\max_{\Theta} \kappa(g(x, t; \Theta), \psi(y)), \quad (1)$$

where Θ denotes all the network parameters.

3.2. Motivation for Complex Projection

In deep learning, researchers aim to formulate the learning problem in such a way that the solution space is restricted in a meaningful way. This helps in learning better and robust representations. The objective function (Equation 1) maximizes the similarity between the output of the composition function of the image-text query and the target image features. Thus, it is intuitive to model the query image, query text and target image lying in some common space. One drawback of TIRG is that it does not emphasize the importance of text features in defining the relationship between the query image and the target image.




Figure 2. Conceptual Diagram of Rotation of the Images in Complex Space. Blue and Red Circle represent the query and the target image respectively. δ represents the rotation in the complex space, learned from the query text features. δ^* represents the complex conjugate of the rotation in the complex space.

Based on these insights, we restrict the compositional learning of query image and text features in such a way that: (i) query and target image features lie in the same space, (ii) text features encode the transition from query image to target image in this space and (iii) transition is symmetric, i.e. *some function of the text features* must encode the reverse transition from target image to query image.

In order to incorporate these characteristics in the composed representation, we propose that the query image and target image are rotations (transitions) of each other in a complex space. The rotation is determined by the text features. This enables incorporating the desired text information about the image in the common complex space. The reason for choosing the complex space is that *some function of text features* required for the transition to be symmetric can easily be defined as the complex conjugate of the text features in the complex space (see Fig. 2).

Choosing such projection also enables us to define a constraint on the optimization problem, referred to as *rotational symmetry constraint* (see Equations 12, 13 and 14). We will empirically verify the effectiveness of this constraint in learning better composed representations. We will also explore the effect on performance if we fuse image and text information in the real space. Refer to Sec. 4.5.

An advantage of modelling the reverse transition in this way is that we do not require captions of query image. This is quite useful in practice, since a user-friendly retrieval system will not ask the user to describe the query image for it. In the public datasets, query image captions are not always available, e.g. for Fashion IQ dataset. In addition to that, it also forces the model to learn a good “internal” representation of the text features in the complex space.

Interestingly, such restrictions on the learning problem serve as implicit regularization. *e.g.*, the text features only influence angles of the composed representation. This is in line with recent developments in deep learning theory [15, 14]. Neyshabur *et al.* [16] showed that imposing simple but global constraints on the parameter space of deep networks is an effective way of analyzing learning theoretic properties and aids in decreasing the generalization error.

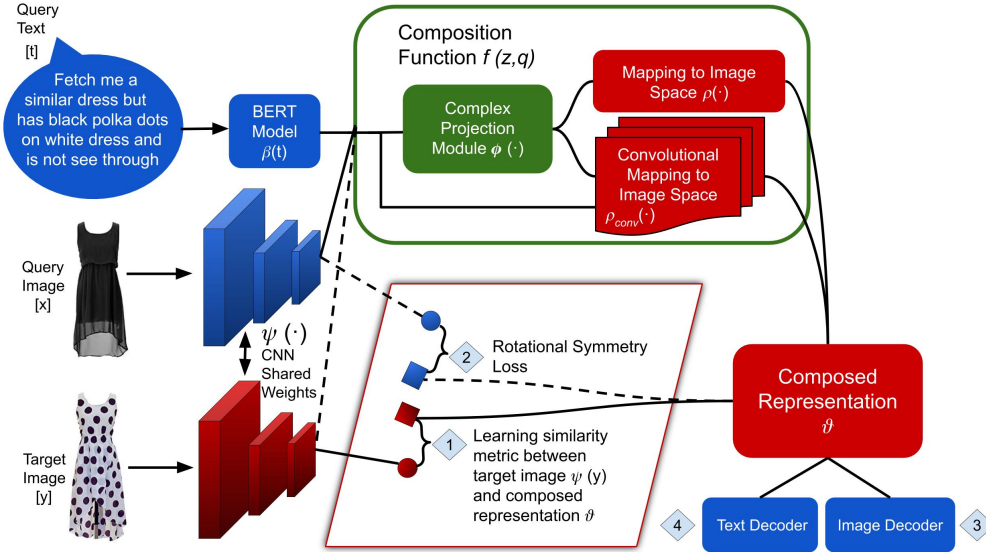


Figure 3. ComposeAE Architecture: Image retrieval using text and image query. Dotted lines indicate connections needed for calculating rotational symmetry loss (see Equations 12, 13 and 14). Here 1 refers to L_{BASE} , 2 refers to L_{SYM}^{BASE} , 3 refers to L_{RI} and 4 refers to L_{RT} .

3.3. Network Architecture

Now we describe ComposeAE, an autoencoder based approach for composing the modalities in the multi-modal query. Figure 3 presents the overview of the ComposeAE architecture.

For the image query, we extract the image feature vector living in a d -dimensional space, using the image model $\psi(\cdot)$ (e.g. ResNet-17), which we denote as:

$$\psi(x) = z \in \mathbb{R}^d. \quad (2)$$

Similarly, for the text query t , we extract the text feature vector living in an h -dimensional space, using the BERT model [1], $\beta(\cdot)$ as:

$$\beta(t) = q \in \mathbb{R}^h. \quad (3)$$

Since the image features z and text features q are extracted from independent uni-modal models; they have different statistical properties and follow complex distributions. Typically in image-text joint embeddings [24, 25], these features are combined using fully connected layers or gating mechanisms.

In contrast to this we propose that the source image and target image are rotations of each other in some complex space, say, \mathbb{C}^k . Specifically, the target image representation is an element-wise rotation of the representation of the source image in this complex space. The information of how much rotation is needed to get from source to target image is encoded via the query text features. During training, we learn the appropriate mapping functions which give us the composition of z and q in \mathbb{C}^k .

More precisely, to model the text features q as specifying element-wise rotation of source image features, we learn a mapping $\gamma: \mathbb{R}^k \rightarrow \{D \in \mathbb{R}^{k \times k} \mid D \text{ is diagonal}\}$ and obtain the coordinate-wise complex rotations via

$$\delta = \mathcal{E}\{j\gamma(q)\},$$

where \mathcal{E} denotes the matrix exponential function and j is square root of -1 . The mapping γ is implemented as a multilayer perceptron (MLP) i.e. two fully-connected layers with non-linear activation.

Next, we learn a mapping function, $\eta: \mathbb{R}^d \rightarrow \mathbb{C}^k$, which maps image features z to the complex space. η is also implemented as a MLP. The composed representation denoted by $\phi \in \mathbb{C}^k$ can be written as:

$$\phi = \delta \eta(z) \quad (4)$$

The optimization problem defined in Eq. 1 aims to maximize the similarity between the composed features and the target image features extracted from the image model. Thus, we need to learn a mapping function, $\rho: \mathbb{C}^k \mapsto \mathbb{R}^d$, from the complex space \mathbb{C}^k back to the d -dimensional real space where extracted target image features exist. ρ is implemented as MLP.

In order to better capture the underlying cross-modal similarity structure in the data, we learn another mapping, denoted as ρ_{conv} . The convolutional mapping is implemented as two fully connected layers followed by a single convolutional layer. It learns 64 convolutional filters and adaptive max pooling is applied on them to get the representation from this convolutional mapping. This enables learning effective local interactions among different features. In

addition to ϕ , ρ_{conv} also takes raw features z and q as input. ρ_{conv} plays a really important role for queries where the query text asks for a modification that is spatially localized. *e.g.*, a user wants a t-shirt with a different logo on the front (see second row in Fig. 4).

Let $f(z, q)$ denote the overall composition function which learns how to effectively compose *extracted image and text features* for target image retrieval. The final representation, $\vartheta \in \mathbb{R}^d$, of the composed image-text features can be written as follows:

$$\vartheta = f(z, q) = a \rho(\phi) + b \rho_{conv}(\phi, z, q), \quad (5)$$

where a and b are learnable parameters.

In autoencoder terminology, the encoder has learnt the composed representation of image and text query, ϑ . Next, we learn to reconstruct the extracted image z and text features q from ϑ . Separate decoders are learned for each modality, *i.e.*, image decoder and text decoder denoted by d_{img} and d_{txt} respectively. The reason for using the decoders and reconstruction losses is two-fold: first, it acts as regularizer on the learnt composed representation and secondly, it forces the composition function to retain relevant text and image information in the final representation. Empirically, we have seen that these losses reduce the variation in the performance and aid in preventing overfitting.

3.4. Training Objective

We adopt a deep metric learning (DML) approach to train ComposeAE. Our training objective is to learn a similarity metric, $\kappa(\cdot, \cdot) : \mathbb{R}^d \times \mathbb{R}^d \mapsto \mathbb{R}$, between composed image-text query features ϑ and extracted target image features $\psi(y)$. The composition function $f(z, q)$ should learn to map semantically similar points from the data manifold in $\mathbb{R}^d \times \mathbb{R}^h$ onto metrically close points in \mathbb{R}^d . Analogously, $f(\cdot, \cdot)$ should push the composed representation away from non-similar images in \mathbb{R}^d .

For sample i from the training mini-batch of size N , let ϑ_i denote the composition feature, $\psi(y_i)$ denote the target image features and $\psi(\tilde{y}_i)$ denote the randomly selected negative image from the mini-batch. We follow TIRG [24] in choosing the base loss for the datasets.

So, for MIT-States dataset, we employ triplet loss with soft margin as a base loss. It is given by:

$$L_{ST} = \frac{1}{MN} \sum_{i=1}^N \sum_{m=1}^M \log \left\{ 1 + \exp\{\kappa(\vartheta_i, \psi(\tilde{y}_{i,m})) - \kappa(\vartheta_i, \psi(y_i))\} \right\}, \quad (6)$$

where M denotes the number of triplets for each training sample i . In our experiments, we choose the same value as mentioned in the TIRG code, *i.e.* 3.

For Fashion200k and Fashion IQ datasets, the base loss is the softmax loss with similarity kernels, denoted as L_{SMAX} . For each training sample i , we normalize the similarity between the composed query-image features (ϑ_i) and target image features by dividing it with the sum of similarities between ϑ_i and all the target images in the batch. This is equivalent to the classification based loss in [24, 3, 21, 10].

$$L_{SMAX} = \frac{1}{N} \sum_{i=1}^N -\log \left\{ \frac{\exp\{\kappa(\vartheta_i, \psi(y_i))\}}{\sum_{j=1}^N \exp\{\kappa(\vartheta_i, \psi(y_j))\}} \right\}, \quad (7)$$

In addition to the base loss, we also incorporate two reconstruction losses in our training objective. They act as regularizers on the learning of the composed representation. The image reconstruction loss is given by:

$$L_{RI} = \frac{1}{N} \sum_{i=1}^N \left\| z_i - \hat{z}_i \right\|_2^2, \quad (8)$$

where $\hat{z}_i = d_{img}(\vartheta_i)$.

Similarly, the text reconstruction loss is given by:

$$L_{RT} = \frac{1}{N} \sum_{i=1}^N \left\| q_i - \hat{q}_i \right\|_2^2, \quad (9)$$

where $\hat{q}_i = d_{txt}(\vartheta_i)$.

3.4.1 Rotational Symmetry Loss

As discussed in [subsection 3.2](#), based on our novel formulation of learning the composition function, we can include a *rotational symmetry loss* in our training objective. Specifically, we require that the composition of the target image features with the complex conjugate of the text features should be similar to the query image features. In concrete terms, first we obtain the complex conjugate of the text features projected in the complex space. It is given by:

$$\delta^* = \mathcal{E}\{-j\gamma(q)\}. \quad (10)$$

Let $\tilde{\phi}$ denote the composition of δ^* with the target image features $\psi(y)$ in the complex space. Concretely:

$$\tilde{\phi} = \delta^* \eta(\psi(y)) \quad (11)$$

Finally, we compute the composed representation, denoted by ϑ^* , in the following way:

$$\vartheta^* = f(\psi(y), q) = a \rho(\tilde{\phi}) + b \rho_{conv}(\tilde{\phi}, \psi(y), q) \quad (12)$$

The *rotational symmetry constraint* translates to maximizing this similarity kernel: $\kappa(\vartheta^*, z)$. We incorporate this

constraint in our training objective by employing softmax loss or soft-triplet loss depending on the dataset.

Since for Fashion datasets, the base loss is L_{SMAX} , we calculate the rotational symmetry loss, L_{SYM}^{SMAX} , as follows:

$$L_{SYM}^{SMAX} = \frac{1}{N} \sum_{i=1}^N -\log \left\{ \frac{\exp\{\kappa(\vartheta_i^*, z_i)\}}{\sum_{j=1}^N \exp\{\kappa(\vartheta_i^*, z_j)\}} \right\}, \quad (13)$$

Analogously, the resulting loss function, L_{SYM}^{ST} , for MIT-States is given by:

$$L_{SYM}^{ST} = \frac{1}{MN} \sum_{i=1}^N \sum_{m=1}^M \log \left\{ 1 + \exp\{\kappa(\vartheta_i^*, \tilde{z}_{i,m}) - \kappa(\vartheta_i^*, z_i)\} \right\}, \quad (14)$$

The total loss is computed by the weighted sum of above mentioned losses. It is given by:

$$L_T = L_{BASE} + \lambda_{SYM} L_{SYM}^{BASE} + \lambda_{RI} L_{RI} + \lambda_{RT} L_{RT}, \quad (15)$$

where $BASE \in \{SMAX, ST\}$ depending on the dataset.

4. Experiments

4.1. Experimental Setup

We evaluate our approach on three real-world datasets, namely: MIT-States[7], Fashion200k [6] and Fashion IQ [5]. For evaluation, we follow the same protocols as other recent works [24, 6, 18]. We use recall at rank k , denoted as $R@k$, as our evaluation metric. We repeat each experiment 5 times in order to estimate the mean and the standard deviation in the performance of the models.

To ensure fair comparison, we keep the same hyperparameters as TIRG [24] and use the same optimizer (SGD with momentum). Similar to TIRG, we use ResNet-17 for image feature extraction to get 512-dimensional feature vector. In contrast to TIRG, we use pretrained BERT [1] for encoding text query. Concretely, we employ BERT-as-service [26] and use Uncased BERT-Base which outputs a 768-dimensional feature vector for a text query. Further implementation details can be found in the code: <https://github.com/ecom-research/ComposeAE>.

4.2. Baselines

We compare the results of ComposeAE with several methods, namely: Show and Tell, Parameter Hashing, Attribute as Operator, Relationship, FiLM and TIRG. We explained them briefly in Sec. 2.

In order to identify the limitations of TIRG and to ensure fair comparison with our method, we introduce three

	MIT-States	Fashion200k	Fashion IQ
Total images	53753	201838	62145
# train queries	43207	172049	46609
# test queries	82732	33480	15536
Average length of complete text query	2	4.81	13.5
Average # of target images per query	26.7	3	1

Table 1. Dataset statistics

variants of TIRG. First, we employ the BERT model as a text model instead of LSTM, which will be referred to as *TIRG with BERT*. Secondly, we keep the LSTM but text query contains full target captions. We refer to it as *TIRG with Complete Text Query*. Thirdly, we combine these two variants and get *TIRG with BERT and Complete Text Query*. The reason for complete text query baselines is that the original TIRG approach *generates* text query by finding one word difference in the source and target image captions. It disregards all other words in the target captions.

While such formulation of queries may be effective on some datasets, but the restriction on the specific form (or length) of text query largely constrain the information that a user can convey to benefit the retrieval process. Thus, such an approach of generating text query has limited applications in real life scenarios, where a user usually describes the modification text with multiple words. This argument is also supported by several recent studies [5, 4, 22]. In our experiments, Fashion IQ dataset contains queries *asked by humans* in natural language, with an average length of 13.5 words. (see Table 1). Due to this reason, we can not get results of original TIRG on this dataset.

4.3. Datasets

Table 1 summarizes the statistics of the datasets. The train-test split of the datasets is the same for all the methods. **MIT-States** [7] dataset consists of $\sim 60k$ diverse real-world images where each image is described by an adjective (state) and a noun (categories), e.g. “ripe tomato”. There are 245 nouns in the dataset and 49 of them are reserved for testing. This split ensures that the algorithm is able to learn the composition on the unseen nouns (categories). The input image (say “unripe tomato”) is sampled and the text query asks to change the state to ripe. The algorithm is considered successful if it retrieves the correct target image (“ripe tomato”) from the pool of *all test images*.

Fashion200k [6] consists of $\sim 200k$ images of 5 different fashion categories, namely: pants, skirts, dresses, tops and jackets. Each image has a human annotated caption, e.g. “blue knee length skirt”.

Fashion IQ[5] is a challenging dataset consisting of 77684

Method	R@1	R@5	R@10
Show and Tell	11.9 \pm 0.1	31.0 \pm 0.5	42.0 \pm 0.8
Att. as Operator	8.8 \pm 0.1	27.3 \pm 0.3	39.1 \pm 0.3
Relationship	12.3 \pm 0.5	31.9 \pm 0.7	42.9 \pm 0.9
FiLM	10.1 \pm 0.3	27.7 \pm 0.7	38.3 \pm 0.7
TIRG	12.2 \pm 0.4	31.9 \pm 0.3	43.1 \pm 0.3
TIRG with BERT	12.3 \pm 0.6	31.8 \pm 0.3	42.6 \pm 0.8
TIRG with Complete Text Query	7.9 \pm 1.9	28.7 \pm 2.5	34.1 \pm 2.9
TIRG with BERT and Complete Text Query	<u>13.3</u> \pm 0.6	<u>34.5</u> \pm 1.0	<u>46.8</u> \pm 1.1
ComposeAE	13.9 \pm 0.5	35.3 \pm 0.8	47.9 \pm 0.7

Table 2. Model performance comparison on MIT-States. The best number is in bold and the second best is underlined.

images belonging to three categories: dresses, top-tees and shirts. Fashion IQ has two human written annotations for each target image. We report the performance on the validation set as the test set labels are not available.

4.4. Discussion of Results

Tables 2, 3 and 4 summarize the results of the performance comparison. In the following, we discuss these results to gain some important insights into the problem.

First, we note that our proposed method ComposeAE outperforms other methods by a significant margin. On Fashion200k, the performance improvement of ComposeAE over the original TIRG and its enhanced variants is most significant. Specifically, in terms of R@10 metric, the performance improvement over the second best method is 6.96% and 30.12% over the original TIRG method. Similarly on R@10, for MIT-States, ComposeAE outperforms the second best method by 2.35% and by 11.13% over the original TIRG method. For the Fashion IQ dataset, ComposeAE has 2.61% and 3.82% better performance than the second best method in terms of R@10 and R@100 respectively.

Second, we observe that the performance of the methods on MIT-States and Fashion200k datasets is in a similar range as compared to the range on the Fashion IQ. For instance, in terms of R@10, the performance of *TIRG with BERT and Complete Text Query* is 46.8 and 51.8 on MIT-States and Fashion200k datasets while it is 11.5 for Fashion IQ. The reasons which make Fashion IQ the most challenging among the three datasets are: (i) the text query is quite complex and detailed and (ii) there is only one target image per query (See Table 1). That is even though the algorithm retrieves semantically similar images but they will not be considered correct by the recall metric. For instance, for the first query in Fig.4, we can see that the second, third and fourth image are semantically similar and modify the image as described by the query text. But if the third image

Method	R@1	R@10	R@50
Han <i>et al.</i> [6]	6.3	19.9	38.3
Show and Tell	12.3 \pm 1.1	40.2 \pm 1.7	61.8 \pm 0.9
Param Hashing	12.2 \pm 1.1	40.0 \pm 1.1	61.7 \pm 0.8
Relationship	13.0 \pm 0.6	40.5 \pm 0.7	62.4 \pm 0.6
FiLM	12.9 \pm 0.7	39.5 \pm 2.1	61.9 \pm 1.9
TIRG	14.1 \pm 0.6	42.5 \pm 0.7	63.8 \pm 0.8
TIRG with BERT	14.2 \pm 1.0	41.9 \pm 1.0	63.3 \pm 0.9
TIRG with Complete Text Query	18.1 \pm 1.9	<u>52.4</u> \pm 2.7	<u>73.1</u> \pm 2.1
TIRG with BERT and Complete Text Query	19.9 \pm 1.0	51.7 \pm 1.5	71.8 \pm 1.3
ComposeAE	22.8 \pm 0.8	55.3 \pm 0.6	73.4 \pm 0.4

Table 3. Model performance comparison on Fashion200k. The best number is in bold and the second best is underlined.

Method	R@10	R@50	R@100
TIRG with Complete Text Query	3.34 \pm 0.6	9.18 \pm 0.9	9.45 \pm 0.8
TIRG with BERT and Complete Text Query	<u>11.5</u> \pm 0.8	<u>28.8</u> \pm 1.5	<u>28.8</u> \pm 1.6
ComposeAE	11.8 \pm 0.9	29.4 \pm 1.1	29.9 \pm 1.3

Table 4. Model performance comparison on Fashion IQ. The best number is in bold and the second best is underlined.

which is the labelled target image did not appear in top-5, then R@5 would have been zero for this query. This issue has also been discussed in depth by Nawaz et al.[13].

Third, for MIT-States and Fashion200k datasets, we observe that the TIRG variant which replaces LSTM with BERT as a text model results in slight degradation of the performance. On the other hand, the performance of the TIRG variant which uses complete text (caption) query is quite better than the original TIRG. However, for the Fashion IQ dataset which represents a real-world application scenario, the performance of TIRG with complete text query is significantly worse. Concretely, TIRG with complete text query performs 253% worse than ComposeAE on R@10. The reason for this huge variation is that the average length of complete text query for MIT-States and Fashion200k datasets is 2 and 3.5 respectively. Whereas average length of complete text query for Fashion IQ is 12.4. It is because TIRG uses the LSTM model and the composition is done in a way which underestimates the importance of the text query. This shows that TIRG approach does not perform well when the query text description is more realistic and complex.

Fourth, one of the baselines (*TIRG with BERT and Complete Text Query*) that we introduced shows significant improvement over the original TIRG. Specifically, in terms of R@10, the performance gain over original TIRG is 8.58%



Figure 4. Qualitative Results: Retrieval examples from FashionIQ Dataset

and 21.65% on MIT-States and Fashion200k respectively. This method is also the second best performing method on all datasets. We think that with more detailed text query, BERT is able to give better representation of the query and this in turn helps in the improvement of the performance.

Qualitative Results: Fig.4 presents some qualitative retrieval results for Fashion IQ. For the first query, we see that all images are in “blue print” as requested by text query. The second request in the text query was that the dress should be “short sleeves”, four out of top-5 images fulfill this requirement. For the second query, we can observe that all retrieved images share the same semantics and are visually similar to the target images. Qualitative results for other two datasets are given in the supplementary material.

Method	Fashion200k	MIT-States	Fashion IQ
ComposeAE	55.3	47.9	11.8
- without L_{SYM}	51.6	47.6	10.5
- Concat in real space	48.4	46.2	09.8
- without ρ_{conv}	52.8	47.1	10.7
- without ρ	52.2	45.2	11.1

Table 5. Retrieval performance (R@10) of ablation studies.

4.5. Ablation Studies

We have conducted various ablation studies, in order to gain insight into which parts of our approach helps in the high performance of ComposeAE. Table 5 presents the quantitative results of these studies.

Impact of L_{SYM} : on the performance can be seen on Row 2. For Fashion200k and Fashion IQ datasets, the decrease in performance is quite significant: 7.17% and 12.38% respectively. While for MIT-States, the impact of incorporating L_{SYM} is not that significant. It may be because the text query is quite simple in the MIT-states case, i.e. 2 words. This needs further investigation.

Efficacy of Mapping to Complex Space: ComposeAE has a complex projection module, see Fig. 3. We removed this module to quantify its effect on the performance. Row 3 shows that there is a drop in performance for all three datasets. This strengthens our hypothesis that it is better to map the extracted image and text features into a common complex space than simple concatenation in real space.

Convolutional versus Fully-Connected Mapping: ComposeAE has two modules for mapping the features from complex space to target image space, i.e., $\rho(\cdot)$ and the second with an additional convolutional layer $\rho_{conv}(\cdot)$. Rows 4 and 5 show that the performance is quite similar for fashion datasets. While for MIT-States, ComposeAE without $\rho_{conv}(\cdot)$ performs much better. Overall, it can be observed that for all three datasets both modules contribute in improving the performance of ComposeAE.

5. Conclusion

In this work, we propose ComposeAE to compose the representation of source image rotated with the modification text in a complex space. This composed representation is mapped to the target image space and a similarity metric is learned. Based on our novel formulation of the problem, we introduce a *rotational symmetry loss* in our training objective. Our experiments on three datasets show that ComposeAE consistently outperforms SOTA method on this task. We enhance SOTA method TIRG [24] to ensure fair comparison and identify its limitations.

Acknowledgments

This work has been supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the *WoWNet* project IUK-1902-003// IUK625/002.

References

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
- [2] Fartash Faghri, David J Fleet, Jamie Ryan Kiros, and Sanja Fidler. Vse++: Improving visual-semantic embeddings with hard negatives. 2018.
- [3] Jacob Goldberger, Geoffrey E Hinton, Sam T Roweis, and Russ R Salakhutdinov. Neighbourhood components analysis. In *Advances in neural information processing systems*, pages 513–520, 2005.
- [4] Xiaoxiao Guo, Hui Wu, Yu Cheng, Steven Rennie, Gerald Tesauro, and Rogerio Feris. Dialog-based interactive image retrieval. In *Advances in Neural Information Processing Systems*, pages 678–688, 2018.
- [5] Xiaoxiao Guo, Hui Wu, Yupeng Gao, Steven Rennie, and Rogerio Feris. The fashion iq dataset: Retrieving images by combining side information and relative natural language feedback. *arXiv preprint arXiv:1905.12794*, 2019.
- [6] Xintong Han, Zuxuan Wu, Phoenix X Huang, Xiao Zhang, Menglong Zhu, Yuan Li, Yang Zhao, and Larry S Davis. Automatic spatially-aware fashion concept discovery. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1463–1471, 2017.
- [7] Phillip Isola, Joseph J. Lim, and Edward H. Adelson. Discovering states and transformations in image collections. In *CVPR*, 2015.
- [8] Kuang-Huei Lee, Xi Chen, Gang Hua, Houdong Hu, and Xiaodong He. Stacked cross attention for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 201–216, 2018.
- [9] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of IEEE CVPR*, 2016.
- [10] Yair Movshovitz-Attias, Alexander Toshev, Thomas K Leung, Sergey Ioffe, and Saurabh Singh. No fuss distance metric learning using proxies. In *Proceedings of the IEEE ICCV*, pages 360–368, 2017.
- [11] Kevin Musgrave, Serge Belongie, and Ser-Nam Lim. A metric learning reality check, 2020.
- [12] Tushar Nagarajan and Kristen Grauman. Attributes as operators: factorizing unseen attribute-object compositions. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 169–185, 2018.
- [13] Shah Nawaz, Kamran Janjua, Ignazio Gallo, Arif Mahmood, Alessandro Calefati, and Faisal Shafait. Do cross modal systems leverage semantic relationships?, 2019.
- [14] Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nati Srebro. Exploring generalization in deep learning. In *Advances in Neural Information Processing Systems*, pages 5947–5956, 2017.
- [15] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *CoRR*, abs/1412.6614, 2015.
- [16] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. Norm-based capacity control in neural networks. In Peter Grünwald, Elad Hazan, and Satyen Kale, editors, *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 1376–1401, Paris, France, 03–06 Jul 2015. PMLR.
- [17] Hyeonwoo Noh, Paul Hongsuck Seo, and Bohyung Han. Image question answering using convolutional neural network with dynamic parameter prediction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 30–38, 2016.
- [18] Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [19] Amrita Saha, Mitesh M Khapra, and Karthik Sankaranarayanan. Towards building large scale multimodal domain-aware conversation systems. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [20] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In *Advances in neural information processing systems*, pages 4967–4976, 2017.
- [21] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
- [22] Fuwen Tan, Paola Cascante-Bonilla, Xiaoxiao Guo, Hui Wu, Song Feng, and Vicente Ordonez. Drill-down: Interactive retrieval of complex scenes using natural language queries. In *Advances in Neural Information Processing Systems*, pages 2647–2657, 2019.
- [23] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015.
- [24] N. Vo, L. Jiang, C. Sun, K. Murphy, L. Li, L. Fei-Fei, and J. Hays. Composing text and image for image retrieval - an empirical odyssey. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6432–6441, 2019.
- [25] Liwei Wang, Yin Li, and Svetlana Lazebnik. Learning deep structure-preserving image-text embeddings. In *Proceedings of the IEEE CVPR*, pages 5005–5013, 2016.
- [26] Han Xiao. bert-as-service. <https://github.com/hanxiao/bert-as-service>, 2018.
- [27] Ying Zhang and Huchuan Lu. Deep cross-modal projection learning for image-text matching. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 686–701, 2018.
- [28] Bo Zhao, Jiashi Feng, Xiao Wu, and Shuicheng Yan. Memory-augmented attribute manipulation networks for interactive fashion search. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1520–1528, 2017.

- [29] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, and Yi-Dong Shen. Dual-path convolutional image-text embeddings with instance loss. *ACM TOMM*, 2017.