

Compositional Mixture Representations for Vision and Text

Stephan Alaniz^{1,2} Marco Federici³ Zeynep Akata^{1,2}

¹University of Tübingen ²Max Planck Institute for Informatics ³University of Amsterdam

{stephan.alaniz,zeynep.akata}@uni-tuebingen.de, m.federici@uva.nl

Abstract

Learning a common representation space between vision and language allows deep networks to relate objects in the image to the corresponding semantic meaning. We present a model that learns a shared Gaussian mixture representation imposing the compositionality of the text onto the visual domain without having explicit location supervision. By combining the spatial transformer with a representation learning approach we learn to split images into separately encoded patches to associate visual and textual representations in an interpretable manner. On variations of MNIST and CIFAR10, our model is able to perform weakly supervised object detection and demonstrates its ability to extrapolate to unseen combination of objects.

1. Introduction

For an artificial intelligence agent to gain an understanding of the world comparable to the one of humans', it should be able to connect the visual world with its semantic meaning. There has been a substantial effort in learning image representations [6, 20] as well as text representation [9, 26, 27], capturing the semantic meaning and disentangling the variation factors in a way similar to how humans learn their surroundings. However, learning unsupervised visual representations from the image data, (e.g., through image reconstruction), can be challenging because there is no guidance towards an informative content (e.g., presence of objects) and uninformative content (e.g., the exact pixel location of the horizon), when there is no supervision involved or the downstream task is unknown [36]. [2] hypothesises that language can be a good prior towards forming useful representations, i.e., things that are commonly described or talked about by people in visual data is information we would like to preserve in our representations.

Our goal is to build upon the idea of preserving the semantic information in the learned representations and learn an unified representation between the text and images, where we use the semantic structure of the text to disentangle the visual data. For instance, when a caption mentions a car, the

corresponding representation of the image should not only include the same information, but should be able to report as well, which part of the image caused the representation of "car". Naturally, such a representation is not limited to hold only a single concept, but can be composed of several individual components, both on the text as well as on the image side. As text is already highly structural, its representation can be viewed as the aggregation of its words' meanings. Images could be then decomposed into patches associated with the building blocks that encode the semantics within the text.

By directly modeling the compositionality of the representation, it becomes possible to obtain several desirable properties. The first property is generalizability. In other words, learning a representation of its constituents is usually simpler than inferring a meaningful joint representation of a complex example. Hence, in novel scenarios, being able to robustly embed partial entities, which are well known from the training set, allow representations to generalize more easily. The second property is combinatorial extrapolation. Certain objects in the images might have high co-occurrence probability. Previously unknown object combinations during test time can cause a deep model to fail. Compositionality can overcome this bias as it allows to arbitrarily combine partial representations. The third and the final property is interpretability. Associating the visual components of the image to the semantics encoded by the textual components allows humans to develop a better understanding of how deep models form their complex representations and might also assist in identifying the causes behind possible failures.

In this work, we implement the idea of using textual guidance to learn compositional image and text representations. One goal of the representation learning is to create representations of the raw input data, which are useful for an a priori unknown downstream tasks. We choose to evaluate our approach on image retrieval and weakly supervised object detection. The latter task is particularly fitting as an aim to decompose images into their respective objects as given by the textual labels without having any bounding box supervision. Therefore, it allows to give a good indication about the compositionality and disentanglement of the image

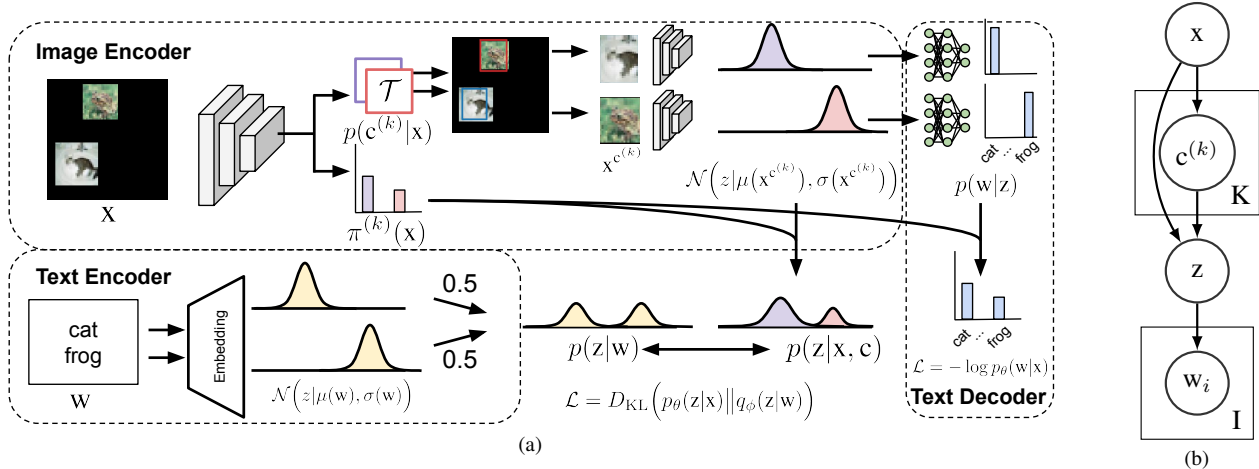


Figure 1. Compositional Mixture (CoMix) Model. The graphical model is shown in the (b) while (a) shows the architecture of our proposed CoMix model that learns two Gaussian mixtures, one is learned from the image and the other one is from the text data. The image encoder first uses a CNN to predict the Gaussian mixture weights $\pi^{(k)}(x)$ as well as the transformation parameters $p(c^{(k)}|x)$ used by the spatial transformer module to extract image patches $x^{c^{(k)}}$. Each patch is individually encoded by a second CNN into Gaussian distributions $\mathcal{N}(z|\mu(x^{c^{(k)}}), \sigma(x^{c^{(k)}}))$. A text decoder $p(w|z)$ is learned with a negative log likelihood loss and ensures that textual information is contained in the representation. A text encoder embeds individual words into Gaussian components and then mixes them into the textual Gaussian mixture representation $p(z|w)$. A KL-divergence loss allow to learn the correspondence between text tokens and image crops by matching the two representations. Without any bounding box supervision, our CoMix model learns to detect images.

representation that is achieved by our model.

Our contributions are as follows: we present a novel model that learns a compositional Gaussian mixture representation for both the image and the text and matches them with a KL divergence loss. The vision part of the model incorporates the spatial-transformer architecture to learn bounding boxes of the image parts corresponding to the different textual components. We evaluate our model on the altered MNIST and CIFAR10 datasets, where each image is composed of several images from the original dataset.

2. Related work

Representation Learning. We distinguish ourselves from most representation learning approaches on images in that we do not try to embed the complete image information in the latent code. Nonetheless, the goals of representation learning are shared across related work in this direction. MONet [5] also employs a compositional approach to scene understanding and reconstruction. It iteratively constructs a scene by its components such as objects and background elements and uses a VAE [20] architecture at each step. Similarly, IODINE [13] also employs a VAE iteratively to infer and reconstructs scene an object at a time. Related work has previously combined the spatial transformer architecture [16] with image representation learning both iteratively [10] and by processing the whole image at once [8].

Multi-modal Learning. When considering models working

with multi-modal data, MVAE [35] applies the VAE setting to multiple modalities such as image and text and can be trained semi-supervised, but it does not employ a compositionality approach. Learning shared representations between images and text is often done targeting a specific application such as retrieval [37] or grounding sentences in images [31]. Work on associating objects with parts of language through visual-semantic alignment [18, 19] usually uses greater level of supervision at least on object detection. One such example is VisualBERT [23] which combines the rich textual representations of BERT [9] with supervised object detectors to solve a variety of visual-language tasks. Similarly, [25] uses natural language explanations to improve the performance of a visual classifier. LXMERT [33] and UNITER [7] demonstrate that rich representations learned from multi-modal vision and text data can benefit diverse tasks such as Visual Question Answering (VQA). The learned representations of CLIP [28] achieve zero-shot object recognition capabilities by training on a large dataset of unstructured image-text pairs. Recently, several works [15, 17, 24, 32] build upon the contrastive learning objective similar to CLIP to learned multi-modal representations in an unsupervised manner.

Weakly Supervised Object Detection. Following the advances in supervised object detection in models such as YOLO [29], Faster-RCNN [30] and Mask-RCNN [14], more work is dedicated in solving the object detection task without bounding box labels, i.e. weakly supervised with only labels about which object are present in the image. WSDDN [3]

works by pooling spatial regions on the last convolutional feature layer of a CNN and has been challenged by similar approaches such as C-MIDN [12] and PredNet [1]. These approaches, however, do not offer the same level of introspection and interpretability as our CoMix.

3. Compositional Mixture Model

We present our proposed compositional mixture model (CoMix) in the following section. In our setting, we consider a joint data distribution $p(x, w)$ of images x and text w with vocabulary \mathcal{V} . Our goal is to use a deep learning model to encode the visual signal coming from the images into a compositional representation, which entails all the semantic information that the components of the textual counterpart contains. At the same time, the representation should be interpretable in terms of which image region contributes to its components' representations.

3.1. CoMix Overview

In our CoMix model, we choose to model the latent representation as a Gaussian mixture. By being a multimodal distribution, it is able to model complex data while having the desired property of being compositional as it consists of several simple Gaussian distributions. Hence, each textual component as well as inferred image patches are embedded into individual Gaussians before being mixed to form the final representation. Our CoMix model consists of three parts as depicted by Figure 1a. A text encoder $p(z|w)$ takes as input the text tokens w and encodes each one of them into a latent mixture component.

Analogously, an image encoder maps the input image to another Gaussian mixture $p(z|x)$, where each Gaussian component corresponds to a different spatial image region. By aligning the two mixture distributions with a KL-divergence loss term, CoMix learns to associate components of the text to the concrete image regions. Attending to separate image patches is learned end-to-end by a spatial transformer module without requiring any extra supervision. Finally, a text decoder $p(w|z)$ ensures that the learned Gaussian mixture representation retains all the textual information, effectively preventing degenerate solutions. Thus, CoMix is trained by supplying image-text pairs (x, w) without any additional supervision as to how these two modalities are related, e.g. there is no bounding box supervision on where the text is grounded in the image.

3.2. Text Encoder

A string of text w , such as a sentence, is composed of entities of interest, such as objects, that are also present in the image. Each word is embedded into a Gaussian latent component $\mathcal{N}(z|\mu(w_i), \sigma(w_i))$, where $\mu(w_i)$ and $\sigma(w_i)$ are determined by an embedding layer. The latent variables of

all the words of the text are then combined into a Gaussian mixture model

$$\frac{1}{|w|} \sum_i \mathcal{N}(z|\mu(w_i), \sigma(w_i)), \quad (1)$$

where each Gaussian of the mixture is weighted equally. This allows us to express a complex multi-modal representation as the sum of simple uni-modal Gaussian components representing its constituent words. Given a one-hot vector for each word, the text encoder is a single matrix that stores a trainable set of Gaussian parameters for each word.

3.3. Image Encoder

A spatial transformer module [16] determines a total of k image regions, which are embedded into the parameters of a Gaussian distribution. Such an architecture allows us to crop, translate, and scale portions of the original input image by defining an affine transformation A_θ that maps image pixel locations of the source and the target image:

$$\begin{pmatrix} x_i^s \\ y_i^s \end{pmatrix} = A_\theta \begin{pmatrix} x_i^t \\ y_i^t \\ 1 \end{pmatrix} \quad (2)$$

where (x^s, y^s) and (x^t, y^t) are the source and target coordinates of pixel locations. The spatial transformer can be viewed as hard attention on the input image with differentiable transformation parameters A_θ .

We employ a convolutional neural network (CNN) to learn k different transformations $p(c^{(k)}|x)$ from the input image. By applying the spatial transformer on each of these, we get k different image crops $x^{c^{(k)}}$, which are then individually encoded with another CNN that, in its turn, provides the parameters μ and σ to a Gaussian distribution $\mathcal{N}(z|\mu(x^{c^{(k)}}), \sigma(x^{c^{(k)}}))$. Hence, each Gaussian component of the image is also associated to an image region by a hard-attention mechanism. The Gaussian components are combined through the categorical mixture distribution $\pi^{(k)}(x)$, which is anticipated from the original input image with the same network that predicts the k transformations $c^{(k)}$ to form the Gaussian mixture $\sum_k \pi^{(k)} \mathcal{N}(z|\mu(x^{c^{(k)}}), \sigma(x^{c^{(k)}}))$.

3.4. Representation Learning

Our requirements for learning a meaningful representation are twofold: 1) The data likelihood under our graphical model needs to be maximized. Since the prediction of w depends only on z , this forces z to capture all the required textual information from the corresponding images x ; 2) The representation should discard image details that are irrelevant for text prediction to obtain a correspondence of text and image components.

To address the former, we train a text decoder $p(w|z)$ on top of the image representation. This results in the full

graphical model depicted in Figure 1b. According to the graphical model, we can now define the joint distribution of our data $p(x, w) = p(w|x)p(x)$, where we are particularly interested in the conditional

$$p(w|x) = \prod_i p(w_i|x) \quad (3)$$

$$= \prod_i \iint p(w_i|z)p(z|x, c)p(c|x)dc dz \quad (4)$$

$$= \prod_i \iint p(w_i|z) \sum_k \pi^{(k)}(x) \mathcal{N}\left(z|\mu(x^{c^{(k)}}), \sigma(x^{c^{(k)}})\right) \quad (5)$$

$$p(c^{(k)}|x)dc^{(k)} dz \quad (6)$$

$$= \prod_i \sum_k \pi^{(k)}(x) \mathbb{E}_{z \sim \mathcal{N}(z|\mu(x^{c^{(k)}}), \sigma(x^{c^{(k)}}))} \left[p(w_i|z) \right] \quad (7)$$

where the crop parameters $c^{(k)}$ are deterministically determined as a parametric function of x , which is represented by the transformer architecture.

By parameterizing the conditional distribution with our model parameters, we can define our objective to minimize the negative log-likelihood of the joint occurrence of pictures x and text w :

$$\min_{\theta} \mathbb{E}_{x, w \sim \bar{p}} \left[-\log p_{\theta}(x, w) \right] \quad (8)$$

$$= H(x) + \min_{\theta} \mathbb{E}_{x, w \sim \bar{p}} \left[-\log p_{\theta}(w|x) \right] \quad (9)$$

where $H(x)$ is the entropy of the image data and as a constant is irrelevant for the optimization procedure. Specifically, the second term on the right-hand side is the negative log-likelihood of the text data given the image under our model and will be denoted as NLL. By minimizing NLL, we are effectively capturing the information required to predict w since the predictive text distribution $p(w|z)$ depends only on the representation z .

Since we want the representation z to focus only on the image details which have a textual counterpart, the second condition aims to discard any image-specific detail from z . This can be done by minimizing the distance between the conditional distributions $p(z|x)$ and $p(z|w)$ to enforce consistency between encoded images and corresponding text. While we model $p(z|x)$ directly, $p(z|w)$ is intractable under our graphical model. Hence, we introduce a variational distribution $q(z|w)$ to approximate the representation induced by the observation of the text data. We model $q(z|w)$ as a Gaussian mixture distribution in which each component corresponds to a single word to capture the compositional nature of the text. By minimizing the Kullback-Leibler (KL) divergence between $p(z|x)$ and $q(z|w)$ we are effectively

minimizing the amount of image-specific information embedded into z , fulfilling our second requirement:

$$D_{\text{KL}}\left(p_{\theta}(z|x) \parallel q_{\phi}(z|w)\right) \geq D_{\text{KL}}\left(p(z|x) \parallel p(z|w)\right) \quad (10)$$

$$= D_{\text{KL}}\left(p(z|x, w) \parallel p(z|w)\right) \quad (11)$$

$$= I(z; x|w). \quad (12)$$

When $I(z; x|w)$ is minimal, the representation z must contain only the information that can be determined through the text (as z and x become conditionally independent), discarding picture specific nuisances not mentioned in the textual description w .

Combining the KL-divergence term with NLL, we arrive at the loss function

$$\min_{\theta, \phi} \mathcal{L} = \min_{\theta, \phi} E_{x, w \sim \bar{p}} \left[-\log p_{\theta}(w|x) \right] \quad (13)$$

$$+ D_{\text{KL}}\left(p_{\theta}(z|x) \parallel q_{\phi}(z|w)\right). \quad (14)$$

The KL-divergence between the two mixture models $p(z|x)$ and $q(z|w)$ is approximated by sampling from each Gaussian component using the reparameterization trick [20] resulting in a stochastic estimate of the loss.

3.5. Image Area Loss

In the previous sections, we provided a detailed description of how a compositional Gaussian mixture representation for both text and image may be learned. Learning the representation is the key part that makes CoMix implicitly split the image into elements that match the textual components. However, the crops that are learned by the spatial transformer are not guaranteed to be minimally enclosing the objects they represent. Since neural networks are able to learn arbitrary mappings from the image region to representation, the network can learn to attend to a larger image part than just the object, effectively including unnecessary information without harming the modeling performance.

We are specifically interested in the smallest bounding box per object to be encoded into a Gaussian component because a precise localization greatly helps model inspection and interpretability. Hence, we introduce another tunable loss term that penalizes the size of the image crops learned by the spatial transformer. Our final loss is given by

$$\min_{\theta, \phi} \mathcal{L} = \min_{\theta, \phi} E_{x, w \sim \bar{p}} \left[-\log p_{\theta}(w|x) \right] \quad (15)$$

$$+ D_{\text{KL}}\left(p_{\theta}(z|x) \parallel q_{\phi}(z|w)\right) \quad (16)$$

$$+ \lambda \left| \prod_i A_{\theta, ii} \right|. \quad (17)$$

where the product of the diagonal elements of the affine transformation matrix A_{θ} equal the area of the patches by

the transformer architecture. The coefficient λ regulates to which degree we would like to penalize the area of the image crops. In the experiments section below, we investigate the effect of this hyperparameter.

4. Experiments

In this section, we describe the datasets and experimental tasks and provide quantitative and qualitative results evaluating the representation learned by CoMix. Furthermore, we demonstrate the performance of our model on classification, image-retrieval, and weakly supervised object detection.

4.1. Datasets and Setup

Datasets. We conduct experiments on two datasets, MultiMNIST and MultiCIFAR10, which represent variations of MNIST [22] and CIFAR-10 [21], respectively. MNIST consists of 60K/10K training/test examples from 10 handwritten digits and CIFAR-10 contains 50K/10K training/test natural image examples from 10 classes.

In contrast, MultiMNIST and MultiCIFAR10 combine several original images randomly sampled onto the same canvas. For each example, we sample between one to four original images, scale them with a factor drawn from $\mathcal{N}(1.3, 0.1)$, and place them at random locations on the canvas. The canvas size is 56×56 for MultiMNIST and 80×80 for MultiCIFAR10. The presence of an object class in the image is indicated with the help of labels, regardless of how many times the same object appears. We sample training and test data once according to the original data sizes and then keep the datasets fixed across all the experiments. See Figure 3 for the examples of images from these datasets.

Experimental setting. All the compared models consist of the same neural-network architectures with the similar or same number of parameters. For MultiMNIST, the image encoder is a 3-layer convolutional network with BatchNorm and ReLU after each layer. For MultiCIFAR10, we use ResNet18 as the image encoder. For both datasets, we use an embedding layer for the text encoder, a 2-layer MLP with ReLUs as the text decoder, and a spatial transformer network consisting of a 3-layer convolutional network with BatchNorm and ReLU.

The number of the mixture components k learned by the spatial transformer network is set to 5 across all the datasets such that the model has to learn to actively choose to use or not to use components. The decision about using a particular component is regulated by the mixture weights $\pi^{(k)}$, which is a learned output of the same network. The area loss coefficient λ is set to 4 for both datasets (see Section 4.3). We randomly split 10% of the training data as a validation set to tune the remaining hyperparameters. The model’s code, experiments, and data generation will be made publicly available.

	MultiMNIST		MultiCIFAR10	
	Cls	Det	Cls	Det
CNN	95.87	n/a	73.99	n/a
WSDDN	99.63	88.27	75.42	49.47
CoMix	99.39	87.04	90.74	75.83

Table 1. Classification and object detection performance measured in mean average precision (mAP, in percent) on MultiMNIST and MultiCIFAR10 datasets.

4.2. Classification and Object Detection

Task. We evaluate our model on its ability to detect and classify objects in an image. For the classification, since there are multiple target classes per image, we measure the mean average precision (mAP) of the predictions of our model’s text decoder. While our model has a direct supervision on the classification task, there is no supervision signal or loss on object detection. Solely the learning of a compositional representation as well as the spatial-transformer architecture facilitates the fact that our model naturally learns to detect objects without the ground-truth bounding-box supervision. Thus, the task setup is identical to weakly supervised object detection (WSOD).

The predicted bounding boxes come from the transformer networks predictions that select a region of the image to be used as one component in our mixture model. We evaluate on the ground-truth bounding boxes that are known from the data generation process, i.e., location and boundary of each individual original image. Following [11], object detection is measured by mAP, where a detection is considered to be a true positive whenever the intersection over union (IOU) of the predicted and the ground-truth bounding box is greater than 0.5.

Baselines. Two baselines are introduced for the classification and object-detection tasks. Firstly, we isolate the image encoder of our model to do the classification directly from the input image to the label output without any representation learning denoted as CNN. This baseline is exclusively trained on a binary cross-entropy loss to predict the presence of the object classes in the image. Our model should match the CNNs classification performance to ensure that the representation learning does not hinder the prediction performance.

Secondly, we compare against Weakly Supervised Deep Detection Network (WSDDN) [3] serving as an object-detection baseline. WSDDN is an established backbone for the weakly supervised object-detection networks that relies on adaptive pooling over a convolutional feature map similar to modern supervised object detection networks such as Faster-RCNN [30]. Contrary to our approach, bounding-box proposals are not learned, but generated algorithmically




Figure 2. Mean average precision of our CoMix model on both classification as well as object detection with varying values of λ . The hyperparameter λ enables an area loss term that tightens bounding box predictions around objects.

using the selective windows search strategy [34].

Results. We report the classification and object detection results in Table 1. On classification, our CoMix model outperforms both baselines on MultiCIFAR10 by a large margin (90.7% vs 75.4%/74%) and is on par with WSDDN on MNIST. The close results on MultiMNIST can most likely be explained by getting close to the perfect classification results rather than choices in the model. We believe, the considerably higher performance of our model on MultiCIFAR10 can be attributed to the compositional modelling of the individual objects in the image. For this reason, mAP approaches ResNet18’s classification performance of around 93% on the individual CIFAR10 images. Since our model processes image parts separately, we can generalize from the single image CIFAR10 classification performance to our more difficult MultiCIFAR10 dataset. Both baseline models lack this property and, therefore, fall short in classification.

Similarly on the object detection, CoMix obtains a comparable performance as WSDDN on MultiMNIST, but outperforms it on MultiCIFAR10 (75.8% vs. 49.5%). Being able to flexibly learn the location of the objects instead of relying on the statically generated bounding boxes helps our model to more accurately pinpoint the image region responsible for a class label. Both results show the benefit of modelling image parts in isolation instead of the whole image as once (CNN) and of learning bounding boxes for our components with the spatial transformer module.

4.3. Inspecting the Area Loss

We argue that our area loss is essential for our model to learn bounding boxes that enclose the object tightly, which is important for both the detection performance and model

inspection. In order to validate this claim, we train our model on both datasets with varying values of λ . A λ of 0 indicates that no area-loss term was applied while the bigger λ gets, the more bounding boxes will be penalized for their size. In Figure 2, we report mAP for both the classification and object detection for different values of λ . Interestingly, the classification performance is largely unaffected with increasing λ and it is even higher for MultiCIFAR10 for any $\lambda > 0$ compared to no area loss.

The mAP of the weakly supervised object-detection task steadily increases for both datasets until λ reaches a value of 4. A value greater than 4 enforces bounding boxes that are too small and, thus, detection accuracy diminishes. Based on this analysis, we set $\lambda = 4$ for our experiments. Most importantly, introducing this loss term only benefits both the detection and classification in the range that we tested, so that its use can be easily justified.

4.4. Image Retrieval

Task. To evaluate our model’s representation learning capabilities, we introduce an image-retrieval task. Given a text sample, the task is to find the best matching image from a set of images. If our model learns a good shared representation of image-text pairs, it should be possible to retrieve the image matching a query text by finding the image representation closest to the inferred text representation. With the use of our model, we first encode both the text and all the images into Gaussian mixtures independently. Then, we score the text representation to all the image representations. Since these representations are multi-modal distributions, we resort to calculating the piece-wise distance between Gaussian components of all the possible text-image pairs. The distance is defined by the average euclidean distance of each text component’s mean to the closest image component’s mean.

Based on the distance score, we can rank the images from the high to low similarity. To evaluate retrieval performance, we report the average rank of the ground-truth matching image as well as the recall at ranks 1, 5, and 10.

Ablation Study. A key advantage of our model is that its representations are compositional. To study the impact of having a compositional representation, we create an ablation model (Gauss), where the latent variable z is a Gaussian instead of a mixture of Gaussians. In this ablation model, the image encoder directly infers the latent variable z from the input image without the compositional spatial transformer. Moreover, the text encoder’s embedding layer is replaced by a 2-layer MLP such that the representation of the full sentence can be learned as it can no longer be composed of simple embeddings. Since the ablation model only uses a single Gaussian, the retrieval distance is calculated by taking the euclidean distance between the means of the text and image latents.

	MultiMNIST				MultiCIFAR10			
	avg. Rank	R@1	R@5	R@10	avg. Rank	R@1	R@5	R@10
<i>Regular</i>								
Gauss	6.85	21.09	64.84	86.72	16.47	14.06	32.03	55.47
CoMix	2.33	64.06	91.41	96.88	2.5	46.09	93.75	97.66
<i>Skewed</i>								
Gauss	15.96	31.25	40.63	53.13	22.35	7.81	35.16	49.22
CoMix	2.45	57.03	89.84	98.44	5.45	47.66	77.34	85.94

Table 2. Image retrieval results of our CoMix model compared to the Gauss baseline where we replace our composition Gaussian mixture latent representation with a single Gaussian distribution. *Regular* refers to the normal version of our datasets MultiMNIST and MultiCIFAR10, whereas *skewed* indicates a more difficult version where the combination of seen objects is highly correlated during training. Scores refer to the average retrieval rank and recall percentage at rank 1, 5 and 10.

Skewed Datasets. To better contrast two different representation-learning approaches, we conduct experiments on a skewed version of MultiMNIST and MultiCIFAR10. During the training, there are three groups of object classes, where each object only co-occurs with objects from the same group. For instance, for MultiMNIST, only the digits (0, 1, 2) would be seen together on the same canvas and similarly for digit groups (3, 4, 5) as well as (6, 7, 8, 9).

During the test phase, all the digits can co-occur with any other digits as in our initial MultiMNIST and MultiCIFAR10 definition. The purpose of this skewed dataset versions is to show how compositionality can help overcome dataset bias, where objects are highly correlated in the training data, but might occur in the novel combinations during the test time.

Results. Image-retrieval results are reported in Table 2. We observe, learning a Gaussian mixture representation achieves a lower retrieval rank on average and a higher recall at all the measured levels as compared to a single Gaussian representation. There are two reasons for these results. Firstly, the Gaussian mixture of CoMix is multimodal and, thus, can model more complex latent distributions. Secondly, the compositionality of CoMix allows it to extrapolate to the unseen combination of objects during the test time. We observe this property when comparing the performance loss between the regular and skewed datasets. CoMix can maintain a low mean rank (2.3 \rightarrow 2.5; 2.5 \rightarrow 5.5), while average rank for the Gauss ablation gets considerably worse (6.9 \rightarrow 16.0; 16.5 \rightarrow 22.4) on both datasets. In conclusion, the explicit modelling of a compositional representation greatly benefits CoMix’s generalization and extrapolation performance of the downstream tasks.

4.5. Qualitative Results

To illustrate the datasets and inspect our model, we present qualitative example of CoMix applied to MultiMNIST and MultiCIFAR10 in Figure 3. For each dataset, we show an example for 1 to 4 objects in the image. In each

data image, we visualize the ground-truth bounding boxes of the objects (dashed yellow line) as well as the object detections by our CoMix model (blue, orange, green, red, purple boxes).

Each image is accompanied by two additional columns. The first column displays the textual prediction of the text decoder applied onto the representation of the patch with the same color. The opaqueness of the bounding boxes and the bar plots is proportional to the categorical mixture distribution π that is predicted by the spatial transformer network. In other words, CoMix predicts the number of objects in the image through the π values of each component and sets a components’ $\pi^{(k)}$ (close) to zero when less than the maximum number of components are needed. The second additional column is the joint categorical class prediction of the whole image that is obtained by mixing the individual component prediction by their respective π weights.

CoMix transparently shows which image patch causes which label prediction as visualized by the color in the last column. As an interesting side effect, the joint classification prediction reflects the count of objects in the image. For example, in the third row of the MultiMNIST data, the joint prediction probability of number 3 is twice the probability of number 9. Naturally, this is caused by the number 3 occurring twice in the image. Hence, our model learns to count without ever having received the supervision in this regard because the textual labels only indicate the presence of objects and not the number of them.

Being able to backtrack exactly how predictions are composed and caused by the concrete image regions is another strength of our model that makes it more interpretable. Due to the hard-attention mechanism of the spatial transformer module, CoMix only uses pixel information inside the bounding box to form its representations and the predictions. In contrast, while WSDN also predicts bounding boxes, there is no exclusive relationship to the part of the picture the bounding box captures. The relationship is on a feature level




Figure 3. Qualitative results on MultiMNIST and MultiCIFAR10. For each dataset there are three columns: 1) data example visualizing the ground truth bounding boxes (yellow, dotted line) and the predicted bounding boxes, one for each component (various colors, solid line); 2) Categorical prediction of object for each individual component corresponding to bounding box with sample color; 3) Categorical mixture prediction of the whole image with colors indicating the contribution of each component.

that often has a receptive field covering the whole image.

5. Conclusion

We introduced our CoMix model that learns a compositional Gaussian mixture representation for both images and text. It utilizes the spatial-transformer architecture to decompose the raw image and expose transparently, which patches are responsible for the Gaussian components of the representation. The information content of the representation is guided by the textual data by employing a likelihood and KL-divergence loss. Without any additional supervision, CoMix learns to detect objects solely facilitated by an additional area loss.

We demonstrate the advantages of learning a compositional representation on the tasks of weakly supervised object detection and image retrieval, where we can validate that

our model can generalize and extrapolate to an unseen combination of objects while, at the same time, being easier to inspect and interpret. Although our current results focus on synthetically generated datasets, a plausible next step would be to scale to the natural images with the natural-language captions, where the employment of recent language models such as BERT [9], GPT-3 [4] could enable richer textual representations. Taking further advantage from advances in contrastive learning such as in CLIP [28], our model could extend these approaches to be more interpretable and inherently exhibit compositionality in their representations.

Acknowledgments

This work has been partially funded by the ERC (853489 - DEXIM) and by the DFG (2064/1 – Project number 390727645).

References

- [1] Aditya Arun, C. V. Jawahar, and M. Pawan Kumar. Dissimilarity coefficient based weakly supervised object detection. In *CVPR*, 2019. 3
- [2] Yoshua Bengio. The consciousness prior. *CoRR*, 2017. 1
- [3] Hakan Bilen and Andrea Vedaldi. Weakly supervised deep detection networks. In *CVPR*, 2016. 2, 5
- [4] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *NeurIPS*, 2020. 8
- [5] Christopher P. Burgess, Loïc Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matthew Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *CoRR*, abs/1901.11390, 2019. 2
- [6] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NeurIPS*, 2016. 1
- [7] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. UNITER: universal image-text representation learning. In *ECCV*, 2020. 2
- [8] Eric Crawford and Joelle Pineau. Spatially invariant unsupervised object detection with convolutional neural networks. In *AAAI*, 2019. 2
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, 2019. 1, 2, 8
- [10] S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, Koray Kavukcuoglu, and Geoffrey E. Hinton. Attend, infer, repeat: Fast scene understanding with generative models. In *NeurIPS*, 2016. 2
- [11] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. 5
- [12] Yan Gao, Boxiao Liu, Nan Guo, Xiaochun Ye, Fang Wan, Haihang You, and Dongrui Fan. C-midn: Coupled multiple instance detection network with segmentation guidance for weakly supervised object detection. In *ICCV*, 2019. 3
- [13] Klaus Greff, Raphaël Lopez Kaufman, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *ICML*, 2019. 2
- [14] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *ICCV*, 2017. 2
- [15] Yuqi Huo, Manli Zhang, Guangzhen Liu, Haoyu Lu, Yizhao Gao, Guoxing Yang, Jingyuan Wen, Heng Zhang, Baogui Xu, Weihao Zheng, Zongzheng Xi, Yueqian Yang, Anwen Hu, Jinming Zhao, Ruichen Li, Yida Zhao, Liang Zhang, Yuqing Song, Xin Hong, Wanqing Cui, Dan Yang Hou, Yingyan Li, Junyi Li, Peiyu Liu, Zheng Gong, Chuhao Jin, Yuchong Sun, Shizhe Chen, Zhiwu Lu, Zhicheng Dou, Qin Jin, Yanyan Lan, Wayne Xin Zhao, Ruihua Song, and Ji-Rong Wen. Wenlan: Bridging vision and language by large-scale multi-modal pre-training. *CoRR*, abs/2103.06561, 2021. 2
- [16] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *NeurIPS*, 2015. 2, 3
- [17] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 2
- [18] Andrej Karpathy, Armand Joulin, and Fei-Fei Li. Deep fragment embeddings for bidirectional image sentence mapping. In *NeurIPS*, 2014. 2
- [19] Andrej Karpathy and Fei-Fei Li. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 2
- [20] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014. 1, 2, 4
- [21] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009. 5
- [22] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11), 1998. 5
- [23] Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. Visualbert: A simple and performant baseline for vision and language. *CoRR*, abs/1908.03557, 2019. 2
- [24] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *CoRR*, abs/2110.05208, 2021. 2
- [25] Weixin Liang, James Zou, and Zhou Yu. ALICE: active learning with contrastive natural language explanations. In *EMNLP*, 2020. 2
- [26] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality. In *NeurIPS*, 2013. 1
- [27] Jeffrey Pennington, Richard Socher, and Christopher D. Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014. 1
- [28] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 2, 8
- [29] Joseph Redmon, Santosh Kumar Divvala, Ross B. Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, 2016. 2
- [30] Shaoqing Ren, Kaiming He, Ross B. Girshick, and Jian Sun. Faster R-CNN: towards real-time object detection with region proposal networks. In *NeurIPS*, 2015. 2, 5

- [31] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, and Bernt Schiele. Grounding of textual phrases in images by reconstruction. In *ECCV*, 2016. [2](#)
- [32] Sheng Shen, Liunian Harold Li, Hao Tan, Mohit Bansal, Anna Rohrbach, Kai-Wei Chang, Zhewei Yao, and Kurt Keutzer. How much can CLIP benefit vision-and-language tasks? *CoRR*, abs/2107.06383, 2021. [2](#)
- [33] Hao Tan and Mohit Bansal. LXMERT: learning cross-modality encoder representations from transformers. In *EMNLP-IJCNLP*, 2019. [2](#)
- [34] Koen E. A. van de Sande, Jasper R. R. Uijlings, Theo Gevers, and Arnold W. M. Smeulders. Segmentation as selective search for object recognition. In *ICCV*, 2011. [6](#)
- [35] Mike Wu and Noah Goodman. Multimodal generative models for scalable weakly-supervised learning. In *NeurIPS*, 2018. [2](#)
- [36] Shengjia Zhao, Jiaming Song, and Stefano Ermon. Towards deeper understanding of variational autoencoding models. *CoRR*, 2017. [1](#)
- [37] Zhedong Zheng, Liang Zheng, Michael Garrett, Yi Yang, Mingliang Xu, and Yi-Dong Shen. Dual-path convolutional image-text embeddings with instance loss. *ACM Trans. Multim. Comput. Commun. Appl.*, 2020. [2](#)