Structured (De)composable Representations Trained with Neural Networks

Graham Spinks^(✉) and Marie-Francine Moens

Department of Computer Science, KU Leuven, Leuven, Belgium
{[graham.spinks](mailto:graham.spinks@cs.kuleuven.be), [sien.moens](mailto:sien.moens@cs.kuleuven.be)}@cs.kuleuven.be

Abstract. The paper proposes a novel technique for representing templates and instances of concept classes. A template representation refers to the generic representation that captures the characteristics of an entire class. The proposed technique uses end-to-end deep learning to learn structured and composable representations from input images and discrete labels. The obtained representations are based on distance estimates between the distributions given by the class label and those given by contextual information, which are modeled as environments. We prove that the representations have a clear structure allowing to decompose the representation into factors that represent classes and environments. We evaluate our novel technique on classification and retrieval tasks involving different modalities (visual and language data).

Keywords: Composable representations · Deep learning · Multimodal

1 Introduction

We propose a novel technique for representing templates and instances of concept classes that is agnostic with regard to the underlying deep learning model. Starting from raw input images, representations are learned in a classification task where the cross-entropy classification layer is replaced by a fully connected layer that is used to estimate a bounded approximation of the distance between each class distribution and a set of contextual distributions that we call ‘environments’. By defining randomized environments, the goal is to capture common sense knowledge about how classes relate to a range of differentiating contexts, and to increase the probability of encountering distinctive diagnostic features. This idea loosely resembles how human long-term memory might achieve retrieval [7] as well as how contextual knowledge is used for semantic encoding [6]. Our experiments confirm the value of such an approach.

In this paper, classes correspond to (visual) object labels, and environments correspond to combinations of contextual labels given by either object labels or image caption keywords. Representations for individual inputs, which we call ‘instance representations’, form a 2D matrix with rows corresponding to classes

Fig. 1. The last layer of a convolutional neural network is replaced with fully-connected layers that map to $n_c \times n_e$ outputs $f_{i,j}$ that are used to create instance representations that are interpretable along contextual dimensions, which we call ‘environments’. By computing the cosine similarity, rows are compared to corresponding class representations, which we refer to as ‘templates’.

and columns corresponding to environments, where each element is an indication of how much the instance resembles the corresponding class versus environment. The parameters for each environment are defined once at start by uniformly selecting a randomly chosen number of labels from the power set of all available contextual labels. The class representation, which we refer to as ‘template’, has the form of a template vector. It contains the average distance estimates between the distribution of a class and the distributions of the respective environments. By computing the cosine similarity between the instance representation and all templates, class membership can be determined efficiently (Fig. 1).

Template and instance representations are interpretable as they have a fixed structure comprised of distance estimates. This structure is reminiscent of traditional language processing matrix representations and enables operations that operate along matrix dimensions. We demonstrate this with a Singular Value Decomposition (SVD) which yields components that determine the values along the rows (classes) and columns (environments) respectively. Those components can then be altered to modify the information content, upon which a new representation can be reconstructed. The proposed representations are evaluated in four settings: (1) Multi-label image classification, *i.e.*, object recognition with multiple objects per image; (2) Image retrieval where we query images that look like existing images but contain altered class labels; (3) Single-label image classification on pre-trained instance representations for a previously unseen label; (4) Rank estimation with regard to compression of the representations.

Contributions. (1) We propose a new deep learning technique to create structured representations from images, entity classes and their contextual information (environments) based on distance estimates. (2) This leads to template representations that generalize well, as successfully evaluated in a classification task. (3) The obtained representations are interpretable as distances between a class and its environment. They are composable in the sense that they can be modified to reflect different class membership as shown in a retrieval task.

2 Background

We shortly discuss useful background related to different aspects of our research.

Representing Entities with Respect to Context. In language applications, structured matrices (*e.g.*, document-term matrices) have been used for a long time. Such matrices can be decomposed with SVD or non-negative matrix factorization. Low-rank approximations are found with methods like latent semantic indexing. Typical applications are clustering, classification, retrieval, etc. with the benefit that outcomes can usually be interpreted with respect to the contextual information. Contrary to our work, earlier methods build representations purely from labels and don't take deep neural network-based features into account. More recently [11] create an unsupervised sentence representation where each entity is a probability distribution based on co-occurrence of words.

Distances to Represent Features. The Earth Mover's Distance (EMD) also known as Wasserstein distance, is a useful metric based on the optimal transport problem to measure the distance between distributions. [3] use a similar idea to define the Word Mover's Distance (WMD) that measures the minimal amount of effort to move Word2Vec-based word embeddings from one document to another. The authors use a matrix representation that expresses the distance between words in respective documents. They note the structure is interpretable and performs well on text-based classification tasks.

Random Features. The Word Mover's Embedding [14] is an unsupervised feature representation for documents, created by concatenating WMD estimates that are computed with respect to arbitrarily chosen feature maps. The authors calculate an approximation of the distance between a pair of documents with the use of a kernel over the feature map. The building blocks of the feature map are documents built from an arbitrary combination of words. This idea is based on Random Features approximation [9] that suggests that a randomized feature map is useful for approximating a shift-invariant kernel.

Our work can be viewed as a combination of the above ideas: we use distance estimates to create interpretable, structured representations of entities with respect to their contexts. The contextual dimension consists of features that are built from an arbitrary combination of discrete labels. Our work most importantly differs in the following manners: (1) We use end-to-end deep neural network training to include rich image features when building representations; (2) Information from different modalities (visual and language) can be combined.

3 CoDiR: Method

We first define some notions that are useful to understand the method, which we name Composable Distance-based Representation learning (CoDiR).

Setup and Notations. Given is a dataset with data samples $x \sim p_{\text{data}}$, with non-exclusive class labels $c_i, i \in \{1, \dots, n_c\}$ which in this work are visual object labels (*e.g.*, dog, ball, ...). Image instances s are fed through a (convolutional) neural network N . The outputs of N will serve to build templates $T_{i,:} \in \mathbb{R}^{n_e}$ and instance representations $D \in \mathbb{R}^{n_c \times n_e}$ with n_e a hyperparameter denoting the amount of environments. Each environment will be defined with the use of discrete environment labels $l_k, k \in \{1, \dots, n_l\}$, for which we experiment with two types: (1) the same visual object labels as used for the class labels (such that $n_l = n_c$) and (2) image caption keywords from the set of the n_l most common nouns, adjectives or verbs in the sentence descriptions in the dataset. We will refer to the first as ‘**CoDiR (class)**’ and the latter as ‘**CoDiR (capt)**’.

$\mathbf{1}_{c_i}$ is shorthand for the indicator function $\mathbf{1}_{c_i}(x) = 1$ if $x \in C_i$, 0 otherwise, with C_i the set of images with label c_i . Similarly we denote $\mathbf{1}_{l_k}$. Each element $D_{i,j}$ is a distance estimate between distributions p_{c_i} and p_{e_j} , where p_{c_i} is shorthand for $p(x = x, x \in C_i)$. Informally, p_{c_i} is the joint distribution modeling the data distribution and class membership c_i . To obtain p_{e_j} , several steps are performed before training. First, hyperparameter R is set, giving the maximum amount of labels per environment. For the j -th environment, we then (1) sample the actual amount of labels $r_j \sim U[1, R] \in \mathbb{N}$; (2) sample the labels $l_m^{(j)}$, with $m \in \{1, \dots, r_j\}$, uniformly without replacement from the set of all discrete environment labels $l_k, k \in \{1, \dots, n_l\}$. Now E_j , the set of images for environment e_j , is given by $E_j = \cup_{m=1}^{r_j} L_m^{(j)}$ with $L_m^{(j)}$ the set of images with label $l_m^{(j)}$. Thus, similarly to p_{c_i} , we have $p_{e_j} = p(x = x, x \in E_j)$. Note that by sampling a random amount of labels per environment, as inspired by [14], we ensure diversity in the type of composition of environments, with some holding many labels and some few.

Contextual Distance. We propose to represent each image as a 2D feature map that relates distributions of classes to environments. A suitable metric should be able to deal with neural network training as well as potentially overlapping distributions. A natural candidate is a Wasserstein-based distance function [1]. A key advantage is that the critic can be encouraged to maximize the distance between two distributions, whereas metrics based on Kullback-Leibler (KL) divergence are not well defined if the distributions have a negligible intersection [1]. In comparison to other neural network-based distance metrics, the Fisher IPM provides particularly stable estimates and has the advantage that any neural network can be used as f as long as the last layer is a linear, dense layer [5]. The Fisher GAN formulation bounds $\mathcal{F}$, the set of measurable, symmetric and bounded real valued functions by defining a data dependent constraint on its second order moments. The IPM is given by:

$$d\mathcal{F}_{\mathcal{F}}(p_{e_j}, p_{c_i}) = \sup_{f_{i,j} \in \mathcal{F}} \frac{\mathbb{E}_{x \sim p_{e_j}} [f_{i,j}(x)] - \mathbb{E}_{x \sim p_{c_i}} [f_{i,j}(x)]}{\sqrt{1/2 \mathbb{E}_{x \sim p_{e_j}} f_{i,j}^2(x) + 1/2 \mathbb{E}_{x \sim p_{c_i}} f_{i,j}^2(x)}} \quad (1)$$

In practice, the Fisher IPM is estimated with neural network training where the numerator in Eq. 1 is maximized while the denominator is expressed as a con-

straint, enforced with a Lagrange multiplier. While the Fisher IPM is an estimate of the chi-squared distance, the numerator can be viewed as a bounded estimate of the inter-class distance, closely related to the Wasserstein distance [5]. From now on, we denote this approximation of the inter-class distance as the ‘distance’. During our training, critics $f_{i,j}$ are trained from input images to maximize the Fisher IPM for distributions p_{c_i} and p_{e_j} , $\forall i \in \{1, \dots, n_c\}, \forall j \in \{1, \dots, n_e\}$. The numerator then gives the distance between p_{c_i} and p_{e_j} . We denote $\mathbf{T} \in \mathbb{R}^{n_c \times n_e}$, with $T_{i,j} = \mathbb{E}_{x \sim p_{e_j, \text{train}}} [f_{i,j}(x)] - \mathbb{E}_{x \sim p_{c_i, \text{train}}} [f_{i,j}(x)]$, i.e., the evaluation of the estimated distances over the training set. Intuitively, one can see why a matrix $\mathbf{T}$ with co-occurrence data contains useful information. A subset of images containing ‘cats’, for example, will more closely resemble a subset containing ‘dogs’ and ‘fur’ than one containing ‘forks’ and ‘tables’.

Template and Instance Representations. As the template representation for class c_i , we simply use the corresponding row of the learned distance matrix: $\mathbf{T}_{i,:}$. Each element $T_{i,j}$ gives an average distance estimate for how a class c_i relates to environment e_j , where smaller values indicate that class and environment are similar or even (partially) overlap. For the instance representation for an input s we then propose to use $\mathbf{D} \in \mathbb{R}^{n_c \times n_e}$ with elements given by Eq. 2:

$$D_{i,j}^{(s)} = \mathbb{E}_{x \sim p_{e_j, \text{train}}} [f_{i,j}(x)] - f_{i,j}(s) \quad (2)$$

where $f_{i,j}(s)$ is simply the output of critic $f_{i,j}$ for the instance s . The result is that for an input s with class label c_i , $\mathbf{D}_{i,:}^{(s)}$ is correlated to $\mathbf{T}_{i,:}$, as its distance estimates with respect to all different environments should be similar. Therefore, the cosine similarity between vector $\mathbf{D}_{i,:}^{(s)}$ and the template $\mathbf{T}_{i,:}$ will be large for input samples from class i , and small otherwise.

Such templates can be evaluated, for example, in multi-label classification tasks (see Sect. 4). Finding the classes for an image is then simply calculated by computing whether $\forall c_i, \cos(\mathbf{D}_{i,:}^{(s)}, \mathbf{T}_{i,:}) > t_{c_i}$ with t_{c_i} a threshold (the level of which is determined during training). From here on we will use a shorthand notation $\mathbf{D}^{(s)} \subset c_i$ to denote $\cos(\mathbf{D}_{i,:}^{(s)}, \mathbf{T}_{i,:}) > t_{c_i}$, and $\mathbf{D}^{(s)} \not\subset c_i$ otherwise.

Implementation. Training $n_c \times n_e$ critics is not feasible in practice, so we pass input images through a common neural network for which the classification layer is replaced by $n_c \times n_e$ single layer neural networks, the outputs of which constitute $f_{i,j}$ (see Fig. 1). During training, any given mini-batch will contain inputs with many different c_i and e_j . To maximize Eq. 1 efficiently, instead of feeding a separate batch for the samples of $x \sim p_{c_i}$ and $x \sim p_{e_j}$, we use the same mini-batch. Additionally, instead of directly sampling $x \sim p_{c_i}$ we multiply each output $f_{i,j}$ with a mask $\mathbf{M}_{i,j}^c$ where $\mathbf{M}_{i,j}^c = \mathbf{1}_{c_i}$. Similarly, for $x \sim p_{e_j}$ we multiply each output $f_{i,j}$ with a mask $\mathbf{M}_{i,j}^e$ where $\mathbf{M}_{i,j}^e = \sum_{m=1}^{r_j} \mathbf{1}_{l_m^{(j)}}$. The result is that instances then are weighted according to their label prevalence as

Algorithm 1. Algorithm of the training process. For matrices and tensors, $\times$ refers to matrix multiplication and $*$ refers to element-wise multiplication.

Inputs: images s , class labels c , environment labels l
 $\forall j \in \{1, \dots, n_e\}$: $r_j \sim U[1, R] \in \mathbb{N} \wedge \forall m \in \{1, \dots, r_j\}$, $l_m^{(j)} \sim U[1, n_l]$
 Create $V \in \mathbb{N}^{n_l \times n_e}$ which has value 1 for each uniformly selected label, 0 otherwise.
 Init $\lambda = 0 \in \mathbb{R}^{n_c \times n_e} \wedge$ Init weights in neural network N
while Training **do**
 Sample a mini-batch b , with batch size n_b , containing images s and binary class labels $C_b \in \mathbb{N}^{n_b \times n_c}$ and binary environment labels $L_b \in \mathbb{N}^{n_b \times n_l}$.
 Create masks
 Expand C_b into $M^c \in \mathbb{N}^{n_b \times n_c \times n_e}$, s.t. $M^c_{k,i,:} = \mathbf{1}_{c_i}(s_k)$ for the k -th sample s_k .
 Multiply L_b and V , then expand the result into $M^e \in \mathbb{N}^{n_b \times n_c \times n_e}$, s.t. $M^e_{k,:,j} = \sum_{m=1}^{r_j} \mathbf{1}_{l_m^{(j)}}(s_k)$ for the k -th sample s_k .
 Calculate the FISHER GAN loss
 Propagate b through N to obtain $O_f \in \mathbb{R}^{n_c \times n_e}$ containing all outputs $f_{i,j}$.
 Apply masks to N 's outputs: $O_E = O_f * M^e$ and $O_C = O_f * M^c$.
 $E_{fE} = \text{mean}(O_E, \text{dim} = 0)$
 $E_{fEs} = \text{mean}(O_E * O_E, \text{dim} = 0)$
 $E_{fC} = \text{mean}(O_C, \text{dim} = 0)$
 $E_{fCs} = \text{mean}(O_C * O_C, \text{dim} = 0)$
 $\text{constraint} = 1 - (0.5 * E_{fEs} + 0.5 * E_{fCs})$
 Minimize $\text{loss} = -\text{sum}(E_{fE} - E_{fC} + \lambda * \text{constraint} - \rho/2 * \text{constraint}^2)$
end while

required. From these quantities, the Fisher IPM can be calculated and optimized. Algorithm 1 explains all the above in detail.¹ When comparing to similar neural network-based methods, the last layer imposes a slightly larger memory footprint ($O(n^2)$ vs $O(n)$) but training time is comparable as they have the same amount of layers. After training completes we perform one additional pass through the training set where we use 2/3rd of the samples to calculate the templates and the remaining 1/3rd to set the thresholds for classification.²

(De-)composing Representations. As the CoDiR representations have a clear structure, a Singular Value Decomposition of D : $D = USV$ can be performed, such that the rows of U and the columns of V can be interpreted as the corresponding factors as contributed by the c_i and e_j respectively. This leads to two applications: (1) Composition: by modifying the elements of U , one can easily obtain $\tilde{U}$ with modified information content. By building a new representation $\tilde{D}$ from $\tilde{U}$, S and V , one thus obtains a similar representation to the original but with modified class membership. This will be further explained in this section. (2) Compression: The spectral norm for instance representations is large with a non-flat spectrum. One can thus compress the representations substantially by retaining only the first k eigenvectors of U and V , thus

¹ Implementation code can be found at <https://github.com/GR4HAM/CoDiR>.

² All models are trained on a single 12Gb gpu.

creating representations in a lower k dimensional space (rank k) without significant loss of classification accuracy. If $k = 1$, the new representations are $(91 + 300)/(91 * 300) = 1.4\%$ the size of the original representations. We call this method C-CoDiR(k).

Let us consider in detail how to achieve composition. To keep things simple, we only discuss the case for ‘CoDiR (capt)’. Given an image s for which $\mathbf{D}^{(s)} \subset c_+$ and $\mathbf{D}^{(s)} \not\subset c_-$. The goal is now to modify $\mathbf{D}^{(s)}$ such that it represents an image $\tilde{s}$ for which $\mathbf{D}^{(\tilde{s})} \not\subset c_+$ and $\mathbf{D}^{(\tilde{s})} \subset c_-$ while preserving the contextual information in the environments of $\mathbf{D}^{(s)}$. As an example, for a $\mathbf{D}^{(s)}$ of an image where $\mathbf{D}^{(s)} \subset c_{dog}$ and the discrete labels from which the environments are built indicate labels such as *playing*, *ball* and *grass*. The goal would be to modify the representation into $\mathbf{D}^{(\tilde{s})}$ (such that, for example, $\mathbf{D}^{(\tilde{s})} \subset c_{cat}$ and $\mathbf{D}^{(\tilde{s})} \not\subset c_{dog}$) and to not modify the information in the environments.

To achieve this, consider that by increasing the value of $\mathbf{U}_{c_+, :}$, one can increase the distance estimate with respect to class c_+ , thus expressing that $\mathbf{D}^{(s)} \not\subset c_+$. Practically, one can set the values of $\tilde{\mathbf{U}}_{c_+, :}$ to the mean of all rows in $\mathbf{U}$ corresponding to the classes $\bar{c}$ for which $\mathbf{D}^{(s)} \not\subset \bar{c}$. The opposite can be done for class c_- , *i.e.*, one can decrease the value of $\mathbf{U}_{c_-, :}$ such that $\mathbf{D}^{(\tilde{s})} \subset c_-$. To set the values of $\tilde{\mathbf{U}}_{c_-, :}$, one can perform a SVD on the matrix composed of all n_c template representations $\mathbf{T}$, thus obtaining $\mathbf{U}_T \mathbf{S}_T \mathbf{V}_T^T$. As the templates by definition contain estimated distances for samples of all classes, it is then easy to see that by setting $\tilde{\mathbf{U}}_{c_-, :} = \mathbf{U}_{T_{c_-}, :}$ we express that $\mathbf{D}^{(\tilde{s})} \subset c_-$ as desired. A valid representation can then be reconstructed with the outer product $\mathbf{D}^{(\tilde{s})} = \sum_k \sigma_k \tilde{\mathbf{U}}_{:, k} \otimes \mathbf{V}_{k, :}^T$ where σ_k are the eigenvalues of $\mathbf{D}^{(s)}$. In the next section this is illustrated by retrieving images after modifying the representations.

4 Experiments

We show how CoDiR compares to a (binary) cross-entropy baseline for multi-label image classification. Additionally, CoDiR’s qualities related to (de)compositions, compression and rank are examined.

4.1 Setup

The experiments are performed on the COCO dataset [4] which contains multiple labels and descriptive captions for each image. We use the 2014 train/val splits of this dataset as these sets contain the necessary labels for our experiment, where we split the validation set into two equal, arbitrary parts to have a validation and test set for the classification task. We set $n_c = 91$, *i.e.*, we use all available 91 class labels (which includes 11 supercategories that contain other labels, *e.g.*, ‘animal’ is the supercategory for ‘zebra’ and ‘cat’). An image can contain more than one class label. To construct environments we use either the class labels, CoDiR (class), or the captions, CoDiR (capt). For the latter, a vocabulary is built of the n_l most frequently occurring adjectives, nouns and verbs. For each

Table 1. F1 scores, precision (PREC) and recall (REC) for different models for the multi-label classification task. σ is the standard deviation of the F1 score over three runs. All results are the average of three runs.

MODEL	METHOD	n_e	n_l	R	F1	PREC	REC	σ
ResNet-18	BXENT (single)	—	—	—	0.566	0.579	0.614	$3.6e^{-3}$
ResNet-18	CoDiR (class)	300	91	40	0.601	0.650	0.613	$8.0e^{-3}$
ResNet-101	BXENT (single)	—	—	—	0.570	0.582	0.623	$1.3e^{-2}$
ResNet-101	CoDiR (class)	300	91	40	0.627	0.664	0.648	$2.5e^{-3}$
Inception-v3	BXENT (single)	—	—	—	0.638	0.663	0.669	$5.4e^{-3}$
Inception-v3	CoDiR (class)	300	91	40	0.617	0.648	0.646	$4.7e^{-3}$
ResNet-18	BXENT (joint)	—	300	—	0.611	0.631	0.654	$1.1e^{-3}$
ResNet-18	BXENT (joint)	—	1000	—	0.614	0.637	0.653	$9.3e^{-3}$
ResNet-18	CoDiR (capt)	300	300	40	0.629	0.680	0.641	$2.7e^{-3}$
ResNet-18	CoDiR (capt)	1000	1000	100	0.638	0.686	0.651	$1.9e^{-3}$
ResNet-101	BXENT (joint)	—	300	—	0.598	0.619	0.640	$1.1e^{-2}$
ResNet-101	BXENT (joint)	—	1000	—	0.592	0.611	0.638	$7.0e^{-3}$
ResNet-101	CoDiR (capt)	300	300	40	0.645	0.696	0.655	$2.8e^{-2}$
ResNet-101	CoDiR (capt)	1000	1000	100	0.657	0.702	0.666	$1.3e^{-2}$
Inception-v3	BXENT (joint)	—	300	—	0.644	0.671	0.675	$1.5e^{-2}$
Inception-v3	BXENT (joint)	—	1000	—	0.63	0.655	0.663	$3.0e^{-2}$
Inception-v3	CoDiR (capt)	300	300	40	0.660	0.699	0.675	$1.9e^{-3}$
Inception-v3	CoDiR (capt)	1000	1000	100	0.661	0.700	0.676	$6.5e^{-3}$

image, each of the n_l labels is then assigned if the corresponding vocabulary word occurs in any of the captions. For the retrieval experiment we select a set of 400 images from the test set and construct their queries.³ All images are randomly cropped and rescaled to 224×224 pixels. We use three types of recent state-of-the-art classification models to compare performance: ResNet-18, ResNet-101 [2] and Inception-v3 [13]. For all runs, an Adam optimizer is used with learning rate $5.e-3$. ρ for the Fisher IPM loss is set to $1e^{-6}$. Parameters are found empirically based on performance on the validation set.

4.2 Results

Multi-label Image Classification. In this experiment the objects in the image are recognized. For each experiment images are fed through a neural network where the only difference between the baseline and our approach is the last layer. For the baseline, which we call ‘BXENT’, the classification model is trained with a binary cross-entropy loss over the outputs and optimal decision thresholds are

³ All dataset splits and queries available at <https://github.com/GR4HAM/CoDiR>.

selected based on the validation set F1 score. For CoDiR, classification is performed on the learned representations as explained in Sect. 3. We then conduct two types of experiments: (1) **BXENT (single)** vs **CoDiR (class)**: An experiment where only class labels are used. For BXENT (single), classification is performed on the output with dimension n_c . For CoDiR (class), environments are built with class labels, such that $n_l = n_c$. (2) **BXENT (joint)** vs **CoDiR (capt)**: An experiment where n_l additional contextual labels from image captions are used. The total amount of labels is $n_c + n_l$. For BXENT (joint) this means joint classification is performed on all $n_c + n_l$ outputs. For CoDiR (capt), there are n_c classes whereas environments are built with the selected n_l caption words. For all models, scores are computed over the n_c class labels.

With the same underlying architecture, Table 1 shows that the CoDiR method compares favorably to the baselines in terms of F1 score.⁴ When adding more detailed contextual information in the environments, as is the case for CoDiR (capt), our model outperforms the baseline in all cases.⁵

The performance of CoDiR depends on the parameters n_e and R . To measure this influence the multi-label classification task is performed for different n_e values. Increasing n_e or the amount of environments leads in general to better performance, although it plateaus after a certain level. For R , the max amount of labels per environment, an optimal value can also be found empirically between 0 and n_l . The reason is that combining a large amount of labels in any environment creates a unique subset to compare samples with. When R is too large, however, subsets with unique features are no longer created and performance deteriorates. Also, even when n_e and R are small, the outcome is not sensitive with regard to the choice of environments, suggesting that the amount and diversity are more important than the composition of the environments.

Retrieval. The experiments here are designed to show interpretability, composability and compressibility of the CoDiR representations. All models and baselines in these sections are pre-trained on the classification task above. We perform two types of retrieval experiments: (1) NN: the most similar sample to a reference sample is retrieved; (2) M-NN: a sample is retrieved with modified class membership while contextual information in the environments is retained. Specifically: “Given an input s_r that belongs to class c_+ but not c_- , retrieve the instance in the dataset that is most similar to s_r that belongs to c_- and not c_+ ”, where c_+ and c_- are class labels (see Fig. 2). We will show that CoDiR is well suited for such a task, as its structure can be exploited to create modified representations $\mathbf{D}^{(\tilde{s}_r)}$ through decomposition as explained in Sect. 3.

This task is evaluated as shown in Table 2a where the goal is to achieve a good combination of M-NN PREC and F1% (for the latter, higher percentages are better). We use the highly structured sigmoid outputs of the BXENT (single) and BXENT (joint) models as baselines, denoted as **SEM (single)** and **SEM (joint)** respectively. With SEM (joint) it is possible to directly modify class

⁴ Multi-label scores as defined by [12].

⁵ For reference: a k-Nearest Neighbors ($k = 3$) on pre-trained ImageNet features of a ResNet-18 achieves a F1 of 0.221.

Table 2. Methods are used in combination with three different base models: ResNet-18/ResNet-101/Inception-v3. All results are the average of three runs.

Method	NN F1	M-NN PREC	F1%	Method	F1
SEM(single)	.64/.66/.70	.53/.55/.55	93/87/89	SEM (single)	0.00/0.00/0.00
SEM(joint)	.71/.70/.73	.29/.28/.31	97/100/96	CoDiR (class)	0.10/0.06/0.07
CNN(joint)	.71/.70/.70	.37/.26/.33	92/90/92	C-CoDiR(5)(class)	0.06/0.08/0.09
CM	.72/.74/.74	.19/.15/.18	100/100/100	SEM (joint)	0.00/0.10/0.00
CoDiR	.70/.72/.72	.30/.30/.27	97/97/95	CoDiR (capt)	0.08/0.15/0.20
C-CoDiR(5)	.70/.72/.72	.30/.29/.26	97/94/93	C-CoDiR(5)(capt)	0.10/0.14/0.19

(a) For the NN and M-NN retrieval, the F1 score of class labels and the precision (PREC) of the modified labels are shown for the first retrieved sample. The proportion of the F1 score of M-NN over NN for the caption words is shown as F1%.

(b) F1 score for a simple logistic regression on pre-trained representations to classify a previously unseen label ("panting dogs"). For the last three models, $n_l = 300$.

Fig. 2. Example of a retrieval result for both NN and M-NN. For NN, based on the representation $D^{(s_r)}$, the most similar instance is retrieved. For M-NN $D^{(s_r)}$ is modified into $D^{(\bar{s}_r)}$ before retrieving the most similar instance.

labels while maintaining all other information. It is thus a ‘best-case scenario’-baseline for which one can strive, as it combines a good M-NN precision and F1% score. SEM (single) on the other hand only contains class information and thus presents a best-case scenario for the M-NN precision score yet a worst-case scenario for the F1% score. Additionally we compare with a simple baseline consisting of **CNN** features from the penultimate layer of the BXENT (joint) models with $n_l = 300$. We also use those features in a **Correlation Matching (CM)** baseline, that combines different modalities (CNN features and word caption labels) into the same representation space [10]. The representations of these baseline models cannot be composed directly. In order to compare them to the ‘M-NN’ method, therefore, we define templates as the average feature vector for a particular class. We then modify the representation for a sample s by subtracting the template of c_+ and adding the template of c_- . All representations except SEM (single) are built from the BXENT (joint) models with $n_l = 300$. For CoDiR they are built from CoDiR (capt) with $n_l = 300$.

For all baselines similarity is computed with the cosine similarity, whereas for CoDiR we exploit its structure as: $\text{similarity} = \text{mean_cos}(D^{(\bar{s}_r)}, D^{(s)})$ over all classes c for which $\text{cos}(D_{c,:}^{(\bar{s}_r)}, T_c) > 0.75 \times t_c$. Here, notations are taken from Sect. 3 and $D^{(\bar{s}_r)}$ is the modified representation of the reference sample. $\text{mean_cos}(D^{(\bar{s}_r)}, D^{(s)})$ is the mean cosine similarity between $D^{(\bar{s}_r)}$ and $D^{(s)}$ with the mean calculated over class dimensions. The similarity is thus calculated over class dimensions where classes with low relevance, *i.e.*, those that have a low similarity with the templates, are not taken into account.

The advantages of the composability of the representations can be seen in Table 2a where CoDiR (capt) has comparable performance to the fully semantic SEM (joint) representations. CNN (joint) manages to obtain a decent M-NN precision score, thus changing class information well, but at the cost of losing contextual information (low F1%), performing almost as poorly as SEM (single). Whereas CM performs well on the NN task, it doesn't change the class information accurately and thus (inadvertently) retains most contextual information.

Rank. While the previous section shows that the structure of CoDiR representations provides access to semantic information derived from the labels on which they were trained, we hypothesize that the representations contain additional information beyond those labels, reflecting local, continuous features in the images. To investigate this hypothesis, we perform an experiment, similar to [15], to determine the rank of a matrix composed of 1000 instance representations of the test set. To maintain stability we take only the first 3 rows (corresponding to 3 classes) and all 300 environments of each representation. Each of these is flattened into a 1D vector of size 900 to construct a matrix of size 1000 * 900. Small singular values are thresholded as set by [8]. The used model is the CoDiR (capt) ResNet-18 model with $n_l = 300$. We obtain a rank of 499, which exceeds the amount of class and environment labels (3 + 300) within, suggesting that the representations contain additional structure beyond the original labels.

The representations can thus be compressed. Table 2a shows that C-CoDiR with $k = 5$, denoted as **C-CoDiR(5)**, approaches CoDiR's performance across all defined retrieval tasks. To show that the CoDiR representations contain information beyond the pre-trained labels, we also use cross-validation to perform a binary classification task with a simple logistic regression. A subset of 400 images of dogs is taken from the validation and test sets, of which 24 and 17 respectively are positive samples of the previously unseen label: *panting dogs*. The outcome in Table 2b shows that CoDiR and C-CoDiR(5) representations outperform the purely semantic representations of the SEM model, which shows that the additional continuous information is valuable.

5 Conclusion

CoDiR is a novel deep learning method to learn representations that can combine different modalities. The instance representations are obtained from images with a convolutional neural network and are structured along class and environment

dimensions. Templates are derived from the instance representations that generalize the class-specific information. In a classification task it is shown that this generalization improves as richer contextual information is added to the environments. When environments are built with labels from image captions, the CoDiR representations consistently outperform their respective baselines. The representations are continuous and have a high rank, as demonstrated by their ability to classify a label that was not seen during pre-training with a simple logistic regression. At the same time, they contain a clear structure which allows for a semantic interpretation of the content. It is shown in a retrieval task that the representations can be decomposed, modified and recomposed to reflect the modified information, while conserving existing information.

CoDiR opens an interesting path for deep learning applications to explore uses of structured representations, similar to how such structured matrices played a central role in many language processing approaches in the past. In zero-shot settings the structure might be exploited, for example, to make compositions of classes and environments that were not seen before. Additionally, further research might explore unsupervised learning or how the method can be applied to other tasks and modalities with alternative building blocks for the environments. While we demonstrate the method with a Wasserstein-based distance, other distance or similarity metrics could be examined in future work.

Acknowledgements. This work was partly supported by the FWO and SNSF, grants G078618N and #176004 as well as an ERC Advanced Grant, #788506. We thank Nvidia for granting us two TITAN Xp GPUs.

References

1. Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein generative adversarial networks. In: International Conference on Machine Learning, pp. 214–223 (2017)
2. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on CVPR, pp. 770–778 (2016)
3. Kusner, M., Sun, Y., Kolkin, N., Weinberger, K.: From word embeddings to document distances. In: ICML, pp. 957–966 (2015)
4. Lin, T.-Y., et al.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
5. Mroueh, Y., Sercu, T.: Fisher GAN. In: Advances in NeurIPS, pp. 2513–2523 (2017)
6. Murdock, B.B.: A theory for the storage and retrieval of item and associative information. *Psychol. Rev.* **89**(6), 609 (1982)
7. Nairne, J.S.: The myth of the encoding-retrieval match. *Memory* **10**(5–6), 389–395 (2002)
8. Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes. The Art of Scientific Computing, 3rd edn. Cambridge University Press, Cambridge (2007)
9. Rahimi, A., Recht, B.: Random features for large-scale kernel machines. In: Advances in Neural Information Processing Systems, pp. 1177–1184 (2008)

10. Rasiwasia, N., et al.: A new approach to cross-modal multimedia retrieval. In: Proceedings of the 18th ACM International Conference on Multimedia, pp. 251–260. ACM (2010)
11. Singh, S.P., Hug, A., Dieuleveut, A., Jaggi, M.: Context mover’s distance & barycenters: optimal transport of contexts for building representations. In: Proceedings of ICLR Workshop on Deep Generative Models (2019)
12. Sorower, M.S.: A literature survey on algorithms for multi-label learning. *Or. State Univ. Corvallis* **18**, 1–25 (2010)
13. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: CVPR, pp. 2818–2826 (2016)
14. Wu, L., Yen, I.E.H., Xu, K., Xu, F., Balakrishnan, A., Chen, P.Y., et al.: Word mover’s embedding: from Word2Vec to document embedding. In: Proceedings of the 2018 Conference on EMNLP, pp. 4524–4534 (2018)
15. Yang, Z., Dai, Z., Salakhutdinov, R., Cohen, W.W.: Breaking the softmax bottleneck: a high-rank RNN language model. In: ICLR (2018)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.