

Associating Neural Word Embeddings with Deep Image Representations using Fisher Vectors

Benjamin Klein, Guy Lev, Gil Sadeh, and Lior Wolf
The Blavatnik School of Computer Science
Tel Aviv University

beni.klein@gmail.com, levguy@gmail.com, gilsadeh@gmail.com, wolf@cs.tau.ac.il

Abstract

In recent years, the problem of associating a sentence with an image has gained a lot of attention. This work continues to push the envelope and makes further progress in the performance of image annotation and image search by a sentence tasks. In this work, we are using the Fisher Vector as a sentence representation by pooling the word2vec embedding of each word in the sentence. The Fisher Vector is typically taken as the gradients of the log-likelihood of descriptors, with respect to the parameters of a Gaussian Mixture Model (GMM). In this work we present two other Mixture Models and derive their Expectation-Maximization and Fisher Vector expressions. The first is a Laplacian Mixture Model (LMM), which is based on the Laplacian distribution. The second Mixture Model presented is a Hybrid Gaussian-Laplacian Mixture Model (HGLMM) which is based on a weighted geometric mean of the Gaussian and Laplacian distribution. Finally, by using the new Fisher Vectors derived from HGLMMs to represent sentences, we achieve state-of-the-art results for both the image annotation and the image search by a sentence tasks on four benchmarks: Pascal1K, Flickr8K, Flickr30K, and COCO.

1. Introduction

The Fisher Vector [33] is an advanced pooling technique, which provided state-of-the-art results on many different applications in computer vision [38, 32, 2, 34]. In this work, the Fisher Vector is used in order to represent sentences by pooling the word2vec embedding [28] of each word in the sentence. The Fisher Vector of a set of local descriptors is obtained as a concatenation of gradients of the log-likelihood of the descriptors in the set with respect to the parameters of a Gaussian Mixture Model that was fitted on a training set in an unsupervised manner. Many different improvements were suggested for the Fisher Vector [35, 45, 39] but all of them are in the context of the Gaus-

sian Mixture Model. In [13], Jia et al. showed empirically that the statistics of gradient based image descriptors, such as SIFT [25], often follow a heavy-tailed distribution which suggests that a Gaussian distribution does not capture well the descriptors' distribution, and that the Euclidean distance is not a suitable distance. They advocate for the selection of a distance measure according to the appropriate probabilistic model that fits the distribution of the empirical data, and show that for the application of SIFT feature matching, a significant improvement is obtained by using the Laplacian distribution and the L_1 distance instead of the Euclidean distance. Motivated by their findings, this paper presents and evaluates new variants of Fisher Vectors that are based on the Laplacian distribution.

By using the common assumption in the Fisher Vector that the covariance matrix is a diagonal one, we define the multivariate Laplacian distribution and the Laplacian Mixture Model (LMM). We explain how to fit a LMM by deriving the Expectation-Maximization (EM) equations and supply the Fisher Vector definition for this model. Similar to [33], we approximate the diagonal of the Fisher Information Matrix in order to normalize the dynamic range of the different dimensions in the Fisher Vector variant presented.

In order to gain the benefits of the two distributions in a single model, we define a new distribution, the Hybrid Gaussian-Laplacian distribution, which can be seen as a weighted geometric mean of the Gaussian and Laplacian distributions. As before, we define the Hybrid Gaussian-Laplacian Mixture Model (HGLMM), derive the EM equations for fitting a HGLMM model, derive the Fisher Vector definition and approximate the diagonal of the Fisher Information Matrix. Although the distribution of each dimension in each component is a weighted geometric mean of a Gaussian and a Laplacian distribution, we show that in the EM algorithm, there is a sharp binary choice between the Gaussian and the Laplacian distributions that results directly from the Maximization-Step derivations.

We employ the new variants of the Fisher Vectors for tasks that match texts with images. In our experiments,

the images are represented by the VGG [40] Convolutional Neural Network as a single vector. The text is represented as a set of vectors obtained by the word2vec [28] method. This set is converted to a Fisher Vector based on one of the distributions: GMM, LMM, or HGLMM. Text to image matching is done using the Canonical Correlations Analysis algorithm [12]. This combination of methods proves to be extremely potent.

2. Previous work

Fisher Vectors The Fisher Vector representation was first introduced in [33] and has been used successfully in various contexts. Since its introduction, many improvements have been suggested which have dramatically enhanced its performance. Some of the most widely used improvements were introduced by Perronnin et al. [34]. The first improvement is to apply an element-wise power normalization function, $f(z) = \text{sign}(z)|z|^\alpha$ where $0 \leq \alpha \leq 1$ is a parameter of the normalization. The second improvement is to apply a L2 normalization on the Fisher Vector after applying the power normalization function. By applying these two operations [34] achieved state-of-the-art accuracy on CalTech 256 and showed superiority over the traditional Bag of Words (BoW) model. Sydorov et. al [45] introduced Deep Fisher Kernels which is a successful attempt of combining two methods with a significant impact on object recognition - Fisher Kernels and Deep Learning. In the traditional pipeline of object recognition, the first two steps (extracting local descriptors and pooling) are done in an unsupervised manner, independently of the task. Only in the last step, when learning the classifier is the nature of the task taken into account. One of the advantages of Deep Learning is that the entire pipeline is optimized for the task - including the layers that are in charge of extracting descriptors and pooling. Sydorov et. al formulated the traditional pipeline of using Fisher kernel with a SVM classifier as a single multi-layer feed forward network. Therefore, both the GMM parameters and the weight vector of the classifiers are tuned according to the nature of the specific task. Simonyan et al. [39] were motivated by the recent success of Convolutional Neural Networks (CNN) [21] and proposed a version of the state-of-the-art Fisher Vector image encoding that can be stacked in multiple layers. Their version obtained competitive results with CNNs on the ILSVRC-2010 dataset. Furthermore, they demonstrated that their Fisher Vector version and CNNs representations are complementary and by combining the two, they achieved a significantly improved accuracy. Our method differs from previous work in that previous work concentrated on Fisher Vectors derived from the Gaussian Mixture Model and in our work, we are deriving the Fisher vector for other distributions.

Image Annotation and Image Search There has been a recent growing interest in methods that can bridge between the domains of vision and NLP. The works of [29, 19, 23] have focused on generating novel descriptive sentences for a query image. Kulkarni et al. [18] suggested a system that generates a descriptive text (not from a fixed set) for a given query image. Their pipeline contains the following steps; First, object detectors find candidate objects in an input image. The candidate objects are passed to a set of classifiers that assign attributes to each candidate. In parallel, each pair of candidate objects is processed by prepositional relationship functions which provide spatial relationships. A Conditional Random Field (CRF) is constructed that incorporates the unary image potentials computed in the previous steps and high order text based potentials computed from large document corpora. Finally, a labeling of the graph is predicted and a sentence is generated according to it. Other works [48, 10, 42] have focused on developing bi-directional mappings. Farhadi et al. [7] developed a method that can compute a score linking an image to a sentence. By using this score, a descriptive sentence from a fixed set can be given to a query image, or a relevant image from a fixed set can be found for a given query sentence. They suggested an intermediate space into which both the images and the sentences are mapped. In [15], Karpathy et al. introduce a model of bidirectional retrieval of images and sentences. Unlike previous works, they do not map images or sentences into a common space. Instead, their model works at a finer scale and embeds fragments of images and fragments of sentences into a common space. The sentence fragments are represented as dependency tree relations that are based on the dependence tree [43] of the sentence. The image fragments are represented by using a CNN [21]. First, objects in the image are detected using Region Convolutional Neural Network (RCNN) [8]. The top 19 detected locations and the entire image are used as image fragments. Each image fragment is embedded using a CNN which takes the image inside a given bounding box and returns the embedding. Finally, they suggest a similarity score for any image-sentence pair. Their method achieved state-of-the-art results on image-sentence retrieval tasks on Pascal1K [36], Flickr8K [10] and Flickr30K [11] datasets. In our method, we return to the paradigm in which the images and sentences are mapped into a common domain and show significant improvement over the state-of-the-art for these three datasets. Similar to the previous work, we are using a CNN that takes an image as input and embeds it into a single vector by taking the representation of the last layer. The sentences are treated in a different way than in the previous work. Specifically, we employ word2vec [28] and map every word in the sentence to a vector. All of the vectors that belong to a sentence are then pooled into a single vector by using Fisher Vector with LMM and HGLMM



Image	Correct Sentences	Our Five nearest neighbors
	<p>A brunette woman taking a big bite of some food.</p> <p>A young woman in black enjoys a bite to eat.</p> <p>A woman in a black top mugs for the camera.</p> <p>A woman wearing a black shirt is eating.</p> <p>A woman about to eat in a white room.</p>	<p>A woman wearing a black shirt is eating.</p> <p>A brunette woman taking a big bite of some food.</p> <p>Two girls holding drinks and looking at something on a cellphone.</p> <p>Girls are preparing food to eat.</p> <p>A young woman in black enjoys a bite to eat.</p>
	<p>Black dog walking on the beach after swimming in the ocean.</p> <p>A black dog walks on the beach near the rocks.</p> <p>A black dog walks along an ocean front.</p> <p>A black dog walks on the sand.</p> <p>A black dog on a rocky beach.</p>	<p>A black dog on a rocky beach.</p> <p>A black dog walks on the beach near the rocks.</p> <p>A black dog walks on the sand.</p> <p>A black dog walks along an ocean front.</p> <p>Three dogs run on beach, two playing with unknown object.</p>

Figure 1. Shown are two examples for our image annotation retrieval system. The sentences colored in green are correct results that were associated with the queried image.

distributions. Finally, the representations of the images and sentences are mapped into a common space by using the CCA algorithm [12].

Representing text as vectors Word2vec [27] is a recently developed technique for building a neural network (NN) that maps words to real-number vectors, with the desideratum that words with similar meanings will map to similar vectors. This technique belongs to the class of methods called “neural language models”. Using a scheme that is much simpler than previous work in this domain, where neural networks with many hidden units and several non-linear layers were normally constructed (e.g., [1]), word2vec [27] constructs a simple log-linear classification network [30]. Two such networks are proposed: the Skip-gram architecture and the Continuous Bag-of-words (CBOW) architecture. In our experiments, we employ the Skip-gram architecture, which is considered preferable.

Recently, the attention has shifted from single words into representing sentences and paragraphs. The classical method in this domain is BoW [41]. Clinchant et al. [3] also used Fisher Vector based on GMM for text. Socher et al. [44] have analyzed sentences using a recursive parse tree. A sentence is then represented by a matrix. In a recent contribution [20] the NN learns to predict the following word in a paragraph based on a representation that concatenates the vector representation of the previous text and the vector representations of a few words from the paragraph.

3. Laplacian Mixture Model (LMM)

The Laplacian Mixture Model (LMM) is a parametric probability density function represented as a weighted sum of multivariate Laplacians. The multivariate Laplacian itself represents a distribution over vectors in R^D . Note that unlike the multivariate Gaussian, it is not uniquely defined [6]. In our formulation, similar to the underlying GMM distributions of conventional Fisher Vectors [33], it is assumed that each multivariate Laplacian has a diagonal covariance matrix. Under this assumption, the probability density function of a single multivariate Laplacian is:

$$f(x; m_1, s_1) = \prod_{d=1}^D \frac{1}{2s_{1,d}} \exp\left(-\frac{|x_d - m_{1,d}|}{s_{1,d}}\right),$$

where $m_1 \in R^D$ and $s_1 \in R^D$ are called the location parameter vector and the scale parameter vector of the multivariate Laplacian, and the second index d is the index of the vector coordinates. The LMM is defined by a set of parameters $\lambda = \{\tau_k, m_k, s_k\}_{k=1 \dots K}$ where $\tau_k \in R$, $m_k \in R^D$, and $s_k \in R^D$ denote respectively the weight, location parameter vector and the scale parameter vector of the k^{th} component and where K is the number of components in the mixture.

3.1. Fitting a LMM

Given a set of data points $\{x_1, x_2, \dots, x_N\}$ one would like to estimate the parameters λ of a LMM. Similar to a

GMM, an expectation maximization [4] approach could be taken here. Specifically:

Expectation Step Let $Z_i = k$ be the event of x_i being associated with mixture component k . Let $T_{k,i}^{(t)} = P(Z_i = k | X = x_i; \lambda^{(t)})$ be the conditional probability of sample x_i being associated with component k , given the sample x_i and the current estimation at iteration t of the parameters $\lambda^{(t)}$. It is straightforward to show that:

$$T_{k,i}^{(t)} = \frac{\tau_k^{(t)} \cdot f(x_i; m_k^{(t)}, s_k^{(t)})}{\sum_{r=1}^K \tau_r^{(t+1)} \cdot f(x_i; m_r^{(t)}, s_r^{(t)})}$$

where $f(x_i; m_k^{(t)}, s_k^{(t)})$ is the pdf of the k^{th} multivariate Laplacian evaluated at point x_i according to the current estimation of parameters, $\lambda^{(t)}$.

Maximization Step Let $Q(\lambda | \lambda^{(t)})$ be the expected value of the log likelihood function, with respect to the conditional distribution of Z given X under the current estimate of the parameters $\lambda^{(t)}$.

$$Q(\lambda | \lambda^{(t)}) = E_{Z|X, \lambda^{(t)}} [\log(L(\lambda; X, Z))] .$$

By deriving $Q(\lambda | \lambda^{(t)})$ according to λ and solving the resulting equations (the full details are given in the supplementary material), one gets the following maximization step expressions:

$$\tau_k^{(t+1)} = \frac{\sum_{i=1}^N T_{k,i}^{(t)}}{\sum_{r=1}^K \sum_{i=1}^N T_{r,i}^{(t)}} \quad (1)$$

$$\sum_{m_{k,d}^{(t+1)} \leq x_{i,d}} T_{k,i}^{(t)} = \sum_{m_{k,d}^{(t+1)} > x_{i,d}} T_{k,i}^{(t)} \quad (2)$$

$$s_{k,d}^{(t+1)} = \frac{\sum_{i=1}^N T_{k,i}^{(t)} |x_{i,d} - m_{k,d}^{(t+1)}|}{\sum_{i=1}^N T_{k,i}^{(t)}} \quad (3)$$

Unlike the EM for GMMs, where the mean is provided explicitly, the location parameter for LMMs is not explicitly given. Rather it is defined as the weighted median of equation 2. An efficient solution to equation 2 would be to sort the values of $x_{i,d}$ for all the samples i and then to iterate over the sorted values and to choose a value for $m_{k,d}^{(t+1)}$ which minimizes the gap between the expressions at both sides of the equation $|\sum_{m_{k,d}^{(t+1)} \leq x_{i,d}} T_{k,i}^{(t)} - \sum_{m_{k,d}^{(t+1)} > x_{i,d}} T_{k,i}^{(t)}|$.

3.2. Fisher Vector of a LMM

The Fisher Vector that was introduced in [33] are the gradients of the log-likelihood of the data with respect to the parameters of the GMM. By following the same path, a variant of the Fisher Vector can be derived for the LMM. Let $X = \{x_1, x_2, \dots, x_N\}$ be a set of samples and λ are the parameters of a LMM. Denote the log-likelihood of the samples X with respect to the parameters λ by $\mathcal{L}(X|\lambda)$. Then the Fisher Vector of the LMM is (see supplementary material):

$$\frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}} = \sum_{i=1}^N \frac{T_{k,i}}{s_{k,d}} \cdot \begin{cases} 1 & \text{if } x_{i,d} > m_{k,d} \\ -1 & \text{otherwise} \end{cases} \quad (4)$$

$$\frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}} = \sum_{i=1}^N T_{k,i} \left(\frac{|x_{i,d} - m_{k,d}|}{s_{k,d}^2} - \frac{1}{s_{k,d}} \right) \quad (5)$$

As in [33], the diagonal of the Fisher Information Matrix F is approximated in order to normalize the dynamic range of the different dimensions of the gradient vectors. Denote by $f_{m_{k,d}}$ and $f_{s_{k,d}}$ the terms of the diagonal of F which correspond respectively to $\frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}}$ and $\frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}}$. Therefore, the normalized partial derivatives are $f_{m_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}}$ and $f_{s_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}}$. It is shown in the supplementary that $f_{m_{k,d}}$ and $f_{s_{k,d}}$ are approximately:

$$f_{m_{k,d}} = \frac{N\tau_k}{s_{k,d}^2}, f_{s_{k,d}} = \frac{N\tau_k}{s_{k,d}^2}$$

3.3. LMM and ICA

In [37], Sanchez et al. state that applying the Principal Components Analysis (PCA) on the data before fitting the GMM is key to make the Fisher Vector work. In experiments on PASCAL VOC 2007, they show that accuracy does not seem to be overly sensitive to the exact number of PCA components. The explanation is that transforming the descriptors by using PCA is a better fit to the diagonal covariance matrix assumption.

Following this observation, a transformation that will cause the transformed descriptors to be a better fit to the diagonal covariance matrix assumption is sought for the LMM. The optimal transformation will result with transformed descriptors that are dimension independent and are non-Gaussian signals. Therefore, and since PCA suffers from the implicit assumption of underlying Gaussian Distribution [16], the Independent Component Analysis (ICA) [22] is chosen in our experiments. It seems that, for the image-text matching experiments we run, ICA is not only preferable when using LMM, but also preferable when using GMM.

4. Hybrid Gaussian-Laplacian Mixture Model

By combining the Gaussian and the Laplacian distributions into one hybrid distribution model one can hope to benefit from the properties of the two distributions. We define the Hybrid Gaussian-Laplacian unnormalized distribution, $h(x; \mu, \sigma, m, s, b)$, for the univariate case to be:

$$h(x; \mu, \sigma, m, s, b) = l(x; m, s)^b \cdot g(x; \mu, \sigma)^{1-b},$$

where $l(x; m, s)$ is the Laplacian distribution parameterized with location parameter m and scale parameter s and $g(x; \mu, \sigma)$ is the Gaussian distribution parameterized by the mean μ and the standard deviation σ . The parameter b is constrained to be in the range $[0, 1]$. According to these definitions, the Hybrid Gaussian-Laplacian Mixture Model is a weighted geometric mean of the Laplacian distribution and the Gaussian distribution. Geometric means of distributions are used, for example, to perform smooth transition between distributions in a simulated annealing framework [31]. More recently, geometric means have emerged as the distribution of the averaged predictions of multiple neural networks employing a softmax layer [9].

As before in the case of LMM, the Hybrid Gaussian-Laplacian can be defined for the multivariate case by assuming that the dimensions are independent. Under this assumption, the probability density function of a single multivariate Hybrid Gaussian-Laplacian is:

$$\prod_{d=1}^D h(x_d; \mu_d, \sigma_d, m_d, s_d, b_d)$$

The HGLMM is a mixture model of the Hybrid Gaussian-Laplacian multivariate distribution. It is defined by a set of parameters $\lambda = \{\tau_k, \mu_k, \sigma_k, m_k, s_k, b_k\}_{k=1 \dots K}$, where $\tau_k \in R$ is the weight of the k^{th} component and $\mu_k, \sigma_k, m_k, s_k, b_k \in R^D$ are the parameters of the Gaussian and Laplacian distributions and the weights of the geometric mean.

4.1. Fitting a HGLMM

Once again, the EM algorithm can be used in order to estimate the parameters λ of the HGLMM given a set of data points $\{x_1, x_2, \dots, x_N\}$.

Expectation Step Let $T_{k,i}^{(t)} = P(Z_i = k | X = x_i; \lambda^{(t)})$ then for the HGLMM the following equation holds:

$$T_{k,i}^{(t)} = \frac{\tau_k^{(t)} \cdot h(x; \mu_k^t, \sigma_k^t, m_k^t, s_k^t, b_k^t)}{\sum_{r=1}^K \tau_r^{(t)} \cdot h(x; \mu_r^t, \sigma_r^t, m_r^t, s_r^t, b_r^t)}$$

Note that, in essence, we sum multiple unnormalized distributions. However, as shown below, in practice the hybrid distribution h is a normalized one.

Maximization Step As before, the expected value of the log likelihood function, with respect to the conditional distribution of Z given X under then current estimate of the parameters $\lambda^{(t)}$ is computed:

$$Q(\lambda | \lambda^{(t)}) = E_{Z|X, \lambda^{(t)}} [\log(L(\lambda; X, Z))]$$

Deriving $Q(\lambda | \lambda^{(t)})$ according to the parameters $\{\tau, \mu, \sigma, m, s\}$ and solving the resulting equations yields the same expressions that one would get for the maximization step in the GMM and in the LMM, up to the values of $T_{k,i}^t$, which are different and are defined in the Expectation Step. Specifically, the equations for τ_k , m_k , and s_k are given above in equations 1, 2, 3, and the equations for μ_k and σ_k are as follows:

$$\mu_{k,d}^{(t+1)} = \frac{\sum_{i=1}^N T_{k,i}^{(t)} \cdot x_{i,d}}{\sum_{i=1}^N T_{k,i}^{(t)}} \quad (6)$$

$$(\sigma_{k,d}^{(t+1)})^2 = \frac{\sum_{i=1}^N T_{k,i}^{(t)} \left(x_{i,d} - \mu_{k,d}^{(t+1)} \right)^2}{\sum_{i=1}^N T_{k,i}^{(t)}} \quad (7)$$

In order to maximize the log-likelihood of the data according to the parameters b one can simply look at the contribution of the parameter $b_{k,d}$ to the log-likelihood. Specifically, omitting the iteration index, let:

$$L_{b_{k,d}} = \sum_{i=1}^N T_{k,i}^{(t)} \left(-\log(2s_{k,d}) - \frac{|x_{i,d} - m_{k,d}|}{s_{k,d}} \right)$$

$$G_{b_{k,d}} = \sum_{i=1}^N T_{k,i}^{(t)} \left(-\log(\sqrt{2\pi}\sigma_{k,d}) - \frac{(x_{i,d} - \mu_{k,d})^2}{2\sigma_{k,d}^2} \right)$$

Then the contribution of $b_{k,d}$ to the log-likelihood is:

$$b_{k,d} \cdot L_{b_{k,d}} + (1 - b_{k,d}) \cdot G_{b_{k,d}} \quad (8)$$

Therefore, under the constraint that $0 \leq b_{k,d} \leq 1$, the value of $b_{k,d}^{(t+1)}$ that maximizes the log-likelihood is:

$$b_{k,d}^{(t+1)} = \begin{cases} 1 & \text{if } L_{b_{k,d}} > G_{b_{k,d}} \\ 0 & \text{otherwise} \end{cases} \quad (9)$$

Due to equation 9, after the maximization step, the distribution of each dimension in each component of the HGLMM is either a Gaussian or Laplacian. Therefore, the final output of the EM algorithm also contains sharp selections between the two distributions, and the hybrid mixture model is a normalized probability model.

	Image search					Image annotation					Sentence
	r@1	r@5	r@10	median rank	mean rank	r@1	r@5	r@10	median rank	mean rank	mean rank
GCS [43]	6.1	18.5	29.0	29.0	NA	4.5	18.0	28.6	32.0	NA	NA
SDT-RNN [43]	6.6	21.6	31.7	25.0	NA	6.0	22.7	34.0	23.0	NA	NA
DFE [15]	9.7	29.6	42.5	15.0	NA	12.6	32.9	44.0	14.0	NA	NA
BRNN [14]	11.8	32.1	44.7	12.4	NA	16.5	40.6	54.2	7.6	NA	NA
SC-NLM [17]	12.5	37.0	51.5	10.0	NA	18.0	40.9	55.0	8.0	NA	NA
NIC [47]	19.0	NA	64.0	5.0	NA	20.0	NA	61.0	6.0	NA	NA
m-RNN [26]	11.5	31.0	42.4	15.0	NA	14.5	37.2	48.5	11.0	NA	NA
Mean Vec	19.1	45.3	60.4	7.0	27.1	22.6	48.8	61.2	6.0	28.8	12.5
GMM	20.6	48.5	64.1	6.0	21.9	28.4	57.7	70.1	4.0	20.1	10.8
LMM	19.8	47.6	62.7	6.0	23.2	27.7	56.6	69.0	4.0	21.1	11.5
HGLMM	20.6	49.4	64.0	6.0	22.1	28.5	58.4	71.7	4.0	18.1	11.5
GMM+HGLMM	21.2	50.0	64.8	5.0	21.1	31.0	59.3	73.7	4.0	18.4	10.6

Table 1. Results on the Flickr8K benchmark [10]. Shown are the recall rates at 1, 5 and 10 retrieval results (higher is better). Also shown, the mean and median rank of the first ground truth (lower is better). There are three tasks: image annotation, image search, and sentence similarity. We compare the results of [43, 15, 14, 17, 47, 47, 26] to the mean vector baseline and to Fisher Vectors based on GMM, LMM and HGLMM. In addition we report results for the combination of the GMM and HGLMM Fisher Vectors.

	Image search					Image annotation					Sentence
	r@1	r@5	r@10	median rank	mean rank	r@1	r@5	r@10	median rank	mean rank	mean rank
DFE [15]	10.3	31.4	44.5	13.0	NA	16.4	40.2	54.7	8.0	NA	NA
BRNN [14]	15.2	37.7	50.5	9.2	NA	22.2	48.2	61.4	4.8	NA	NA
SC-NLM [17]	16.8	42.0	56.5	8.0	NA	23.0	50.7	62.9	5.0	NA	NA
NIC [47]	17.0	NA	57.0	7.0	NA	17.0	NA	56.0	7.0	NA	NA
m-RNN [26]	12.6	31.2	41.5	16.0	NA	18.4	40.2	50.9	10.0	NA	NA
LRCN [5]	14.0	34.9	47.0	11.0	NA	NA	NA	NA	NA	NA	NA
Mean Vec	20.5	46.3	59.3	6.8	32.4	24.8	52.5	64.3	5.0	27.3	16.3
GMM	23.9	51.6	64.9	5.0	24.8	33.0	60.7	71.9	3.0	19.0	15.1
LMM	23.6	51.2	64.4	5.0	25.2	32.5	59.9	71.5	3.2	19.2	15.6
HGLMM	24.4	52.1	65.6	5.0	24.5	34.4	61.0	72.3	3.0	18.1	14.7
GMM+HGLMM	25.0	52.7	66.0	5.0	23.7	35.0	62.0	73.8	3.0	17.4	14.2

Table 2. Mean results on the Flickr30K benchmark [11]. For details, see Table 1. GCS and SDT-RNN results [43] are not available for this specific benchmark.

4.2. Fisher Vector of a HGLMM

A Fisher Vector variant can also be defined for the HGLMM. Although the HGLMM has more parameters than the GMM, the Fisher Vector of the HGLMM has the same length as the Fisher Vector of the GMM which is $2KD$. The reason is that according to equation 9, each dimension d in each component k is either a Laplacian univariate distribution or a Gaussian univariate distribution and therefore the contribution to the Fisher Vector is $\left\{ \frac{\partial \mathcal{L}(X|\lambda)}{\partial \mu_{k,d}}, \frac{\partial \mathcal{L}(X|\lambda)}{\partial \sigma_{k,d}} \right\}$ if $b_{k,d} = 0$ and $\left\{ \frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}}, \frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}} \right\}$ if $b_{k,d} = 1$. The values of the Fisher Vector for the HGLMM are computed by taking the gradients of the log-likelihood of the data with respect to the parameters of the HGLMM. The resulting equations yields the same expres-

sions that one would get for the Fisher Vector of the GMM and of the LMM, up to the values of $T_{k,i}$ which are different and have the same expressions as in the expectation step of the HGLMM. Therefore the values of the Fisher Vector for the HGLMM when $b_{k,d} = 0$ are:

$$\frac{\partial \mathcal{L}(X|\lambda)}{\partial \mu_{k,d}} = \sum_{i=1}^N T_{k,i} \cdot \frac{x_{i,d} - \mu_{k,d}}{\sigma_{k,d}^2} \quad (10)$$

$$\frac{\partial \mathcal{L}(X|\lambda)}{\partial \sigma_{k,d}} = \sum_{i=1}^N T_{k,i} \left(\frac{(x_{i,d} - \mu_{k,d})^2}{\sigma_{k,d}^3} - \frac{1}{\sigma_{j,d}} \right) \quad (11)$$

When $b_{k,d} = 1$, equations 4 and 5 provide the relevant Fisher Vector coordinates.

	Image search					Image annotation					Sentence
	r@1	r@5	r@10	median rank	mean rank	r@1	r@5	r@10	median rank	mean rank	mean rank
GCS [43]	16.4	46.6	65.6	NA	12.5	23.0	45.0	63.0	NA	16.9	10.5
SDT-RNN [43]	25.4	65.2	84.4	NA	7.0	25.0	56.0	70.0	NA	13.4	NA
DFE [15]	23.6	65.2	79.8	NA	7.6	39.0	68.0	79.0	NA	10.5	NA
Mean Vec	44.9	84.9	94.1	2.0	3.6	52.5	83.2	92.3	1.3	4.0	2.1
GMM	43.5	85.4	94.3	2.0	3.6	55.6	86.0	93.1	1.1	4.2	2.3
LMM	43.1	84.6	94.1	2.0	3.6	54.6	85.1	93.4	1.2	3.6	2.3
HGLMM	43.6	84.7	94.2	2.0	3.7	55.6	85.8	93.3	1.2	3.8	2.3
GMM+HGLMM	44.0	85.6	94.6	2.0	3.6	55.9	86.2	93.3	1.2	3.9	2.2

Table 3. Mean results on the Pascal1K benchmark [36]. For details, see Table 1. This dataset is rather small and the contribution of high dimensional representations is less obvious than in the Flickr datasets.

As in the LMM, the diagonal of the Fisher Information Matrix F is approximated in order to normalize the dynamic range of the different dimensions of the gradient vectors. Let $f_{\mu_{k,d}}$, $f_{\sigma_{k,d}}$, $f_{m_{k,d}}$ and $f_{s_{k,d}}$ be the terms of the diagonal of F that correspond respectively to $\frac{\partial \mathcal{L}(X|\lambda)}{\partial \mu_{k,d}}$, $\frac{\partial \mathcal{L}(X|\lambda)}{\partial \sigma_{k,d}}$, $\frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}}$ and $\frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}}$. Then, the normalized partial derivatives are $f_{m_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial \mu_{k,d}}$, $f_{s_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial \sigma_{k,d}}$, $f_{m_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial m_{k,d}}$ and $f_{s_{k,d}}^{-1/2} \cdot \frac{\partial \mathcal{L}(X|\lambda)}{\partial s_{k,d}}$. It can be shown that the approximated values of the terms of the diagonal are the same as the approximated values of the terms of the diagonal in the GMM and LMM:

$$f_{\mu_{k,d}} = \frac{N\tau_k}{\sigma_{k,d}^2}, f_{\sigma_{k,d}} = \frac{2N\tau_k}{\sigma_{k,d}^2}, f_{m_{k,d}} = \frac{N\tau_k}{s_{k,d}^2}, f_{s_{k,d}} = \frac{N\tau_k}{s_{k,d}^2}$$

5. Results

We perform our image annotation experiments on four benchmarks: Pascal1K [36], Flickr8K [10], Flickr30K [11], and COCO [24]. The datasets contain 1,000, 8,000, 30,000, and 123,000 images respectively. The annotation of the images was done using crowdsourcing via Amazon Mechanical Turk, with five independent sentences provided by five users to each image.

The Flickr8k dataset is provided with a training split of size 6091, a validation split of 1000 images, and a test split of size 1000. We use the same split. For Pascal1K, no training splits are given, and we use 20 splits of the same size used to report results in previous work: 800 train, 100 validation and 100 test images. The situation for Flickr30K is similar, and following previous work, we use 5 random splits of 1000 images for test, 1000 for validation, and the rest for training. For COCO, we follow Karpathy et al. [14] and use 5000 images for both validation and testing, and also report results on a subset of 1000 testing images. 5 random splits were used here as well. The reduced number of repeats on Flickr30K and COCO compared to Pascal1K

stems from the datasets size and the resulting computational burden.

The word2vec [28] vectors used to represent words were obtained from code.google.com/p/word2vec/. We found that performing ICA without reducing the dimensionality of the 300D word2vec vectors helps performance in all methods, including the vanilla GMM based Fisher Vectors. The original vectors and PCA provided somewhat lower results. We therefore apply ICA to the word2vec representation in all of our experiments below. All images were resized to a fixed size of 221 by 221 pixels and encoded as a single vector using the VGG [40] software. When employing the VGG [40] representation, the recommended pipeline is used. Namely, the original image is cropped in ten different ways into 224 by 224 pixel images: the four corners, the center, and their x-axis mirror image. The mean intensity is then subtracted in each color channel and the resulted images are encoded by the network. The average of the resulting 10 feature vectors is used as the single image representation.

Our system has two parameters: the number of components in the mixture models used to describe the sentence, and the regularization parameters of the regularized CCA algorithm [46] that matches between the image representation and the sentence representation. The first parameter was fixed to 30 throughout the experiments. This value was selected once using the Flickr8K validation split. The regularization parameter was selected in each repetition of each experiment based on the validation data. We use linear CCA, which is presumably sub-optimal compared to the Kernel CCA that [43] used as one of the baseline methods.

There are three tasks: image annotation, in which the goal is to retrieve, given a query image, the five ground truth sentences; image search, in which, given a query sentence, the goal is to retrieve the ground truth image; and sentence similarity, in which the goal is to retrieve given a query sentence the other four sentences associated with the same image. For the first two tasks, the results are reported as the

	Image search					Image annotation					Sentence
	r@1	r@5	r@10	median rank	mean rank	r@1	r@5	r@10	median rank	mean rank	mean rank
1K test images:											
BRNN [14]	20.9	52.8	69.2	4.0	NA	29.4	62.0	75.9	2.5	NA	NA
Mean Vec	24.2	56.4	72.4	4.0	14.7	33.2	61.8	75.1	3.0	14.5	14.4
GMM	24.2	59.2	76.0	4.0	11.3	39.0	67.0	80.3	3.0	11.2	12.4
LMM	25.0	59.5	76.1	4.0	11.1	38.6	67.8	79.8	3.0	12.1	12.2
HGLMM	24.9	58.8	76.5	4.0	11.1	37.7	66.6	79.1	3.0	10.8	13.7
GMM+HGLMM	25.1	59.8	76.6	4.0	11.1	39.4	67.9	80.9	2.0	10.4	12.9
5K test images:											
BRNN [14]	8.9	24.9	36.3	19.5	NA	11.8	32.5	45.4	12.2	NA	NA
Mean Vec	10.2	27.2	38.3	18.0	64.8	12.7	32.0	44.6	14.0	62.3	63.7
GMM	10.2	27.7	39.4	17.0	51.1	16.7	37.8	49.6	11.0	47.4	52.1
LMM	10.5	28.0	40.0	17.0	51.3	16.4	37.8	50.3	10.0	48.1	52.5
HGLMM	10.4	27.7	39.7	17.0	50.7	16.6	37.0	49.5	11.0	47.6	55.0
GMM+HGLMM	10.8	28.3	40.1	17.0	49.3	17.3	39.0	50.2	10.0	46.4	52.4

Table 4. Mean results on the COCO benchmark [24]. For details, see Table 1.

recall rate at one result, at 5 results, or at the 10 first results. Also reported is the mean rank and the median rank of the first ground truth result. In the sentence similarity task, only the mean rank of the first ground truth result is reported.

In all tasks, CCA is trained on the training set, its parameter is tuned on the validation set, and testing is performed on the test set. For sentence similarity, the representations of the test samples are used after projection to the CCA space, however, the test images are not used. Note that all the parameters of the sentence representations, prior to the CCA computation, are learned in an unsupervised fashion on the corpus of word2vec vectors, without employing the three benchmark datasets.

Given a sentence, it is mapped to the set of word2vec vectors that are associated with the sentence’s words. In our experiments, we test the average vector of this set as a simple baseline representation, three types of Fisher Vector pooling (GMM, LMM, and HGLMM), and fusion of GMM with HGLMM. This fusion is done by concatenating the two Fisher Vector representations into a single vector. All Fisher Vectors were normalized by the L2 norm after applying the power normalization function with $\alpha = 0.5$, as described in Sec. 2.

We compare our results to the current state of the art, namely to the Grounded Compositional Semantics (GCS) [43], Semantic Dependency Tree Recursive Neural Network (SDT-RNN) [43], Deep Fragment Embedding Multi Instance Learning (DFE-MIL) [15], and the five concurrent technical reports described in Sec. 2: BRNN [14], SC-NLM [17], NIC [47], m-RNN [26], LRCN [5].

The results are reported in Table 1, 2, 3, 4 for the Flickr8K, Flickr30K, Pascal1K, and COCO benchmarks respectively. As can be seen, the Fisher Vector methods out-

perform the existing state of the art and the mean vector baseline. As a general trend, in all three benchmarks, Fisher Vectors based on GMM outperform those based on LMM. HGLMM based Fisher Vectors perform better than both, and combining GMM with HGLMM outperforms the other methods. Sample results can be seen in Figure 1.

6. Discussion

The Fisher Vector has proven to be useful in a variety of object recognition applications. In all those applications, it was derived according the Gaussian Mixture Model. In this paper we show that using a Fisher Vector derived from other distributions, namely, LMM and HGLMM, one can obtain improved accuracy in central computer vision tasks. We believe that these improvements would carry on to other tasks as well.

The normalization improvements that were suggested by Perronnin et al. [34] dramatically increased the performance achieved by the Fisher Vector technique and contributed to its success. It remains to be explored whether there are specific normalization techniques that are most suitable for the Fisher Vectors that we derive from the LMM and HGLMM distributions.

The Hybrid Gaussian-Laplacian Mixture Model (HGLMM) that we presented, allowed us to gain benefits from both underlying distributions by having the flexibility that each dimension in each component would be modeled according to the most suitable distribution. Such geometric-mean mixtures could be generalized to any two distributions and fit real world data of any distribution shape. It is also not limited to just two parametric distributions.

Acknowledgments

This research is supported by the Intel Collaborative Research Institute for Computational Intelligence (ICRI-CI).

References

- [1] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin. A neural probabilistic language model. *J. Mach. Learn. Res.*, 3:1137–1155, Mar. 2003.
- [2] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman. The devil is in the details: an evaluation of recent feature encoding methods. In *British Machine Vision Conference*, 2011.
- [3] S. Clinchant and F. Perronnin. Textual similarity with a bag-of-embedded-words model. In *Proceedings of the 2013 Conference on the Theory of Information Retrieval*, page 25. ACM, 2013.
- [4] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.
- [5] J. Donahue, L. A. Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. *arXiv preprint arXiv:1411.4389v2*, 2014.
- [6] T. Eltoft, T. Kim, and T.-W. Lee. On the multivariate laplace distribution. *Signal Processing Letters, IEEE*, 13(5):300–303, 2006.
- [7] A. Farhadi, M. Hejrati, M. A. Sadeghi, P. Young, C. Rashtchian, J. Hockenmaier, and D. Forsyth. Every picture tells a story: Generating sentences from images. In *Computer Vision—ECCV 2010*, pages 15–29. Springer, 2010.
- [8] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *arXiv preprint arXiv:1311.2524*, 2013.
- [9] G. E. Hinton, N. Srivastava, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, abs/1207.0580, 2012.
- [10] M. Hodosh, P. Young, and J. Hockenmaier. Framing image description as a ranking task: Data, models and evaluation metrics. *J. Artif. Intell. Res. (JAIR)*, 47:853–899, 2013.
- [11] P. Y. A. L. M. Hodosh and J. Hockenmaier. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2014.
- [12] H. Hotelling. Relations between two sets of variates. *Biometrika*, pages 321–377, 1936.
- [13] Y. Jia and T. Darrell. Heavy-tailed distances for gradient based image descriptors. In *Advances in Neural Information Processing Systems*, pages 397–405, 2011.
- [14] A. Karpathy and L. Fei-Fei. Deep visual-semantic alignments for generating image descriptions. Technical report, Computer Science Department, Stanford University, 2014.
- [15] A. Karpathy, A. Joulin, and L. Fei-Fei. Deep fragment embeddings for bidirectional image sentence mapping. *arXiv preprint arXiv:1406.5679*, 2014.
- [16] Y. Ke and R. Sukthankar. Pca-sift: A more distinctive representation for local image descriptors. In *Computer Vision and Pattern Recognition, 2004. CVPR 2004. Proceedings of the 2004 IEEE Computer Society Conference on*, volume 2, pages II–506. IEEE, 2004.
- [17] R. Kiros, R. Salakhutdinov, and R. S. Zemel. Unifying visual-semantic embeddings with multimodal neural language models. *arXiv preprint arXiv:1411.2539*, 2014.
- [18] G. Kulkarni, V. Premraj, S. Dhar, S. Li, Y. Choi, A. C. Berg, and T. L. Berg. Baby talk: Understanding and generating simple image descriptions. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1601–1608. IEEE, 2011.
- [19] P. Kuznetsova, V. Ordonez, A. C. Berg, T. L. Berg, and Y. Choi. Collective generation of natural image descriptions. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pages 359–368. Association for Computational Linguistics, 2012.
- [20] Q. V. Le and T. Mikolov. Distributed representations of sentences and documents. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, volume 32 of *JMLR Proceedings*, pages 1188–1196. JMLR.org, 2014.
- [21] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [22] T.-W. Lee. Independent component analysis: theory and applications [book review]. *IEEE Transactions on Neural Networks*, 10(4):982–982, 1999.
- [23] S. Li, G. Kulkarni, T. L. Berg, A. C. Berg, and Y. Choi. Composing simple image descriptions using web-scale n-grams. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 220–228. Association for Computational Linguistics, 2011.
- [24] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollr, and C. Zitnick. Microsoft coco: Common objects in context. In D. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, editors, *Computer Vision ECCV 2014*, volume 8693 of *Lecture Notes in Computer Science*, pages 740–755. Springer International Publishing, 2014.
- [25] D. G. Lowe. Object recognition from local scale-invariant features. In *Computer vision, 1999. The proceedings of the seventh IEEE international conference on*, volume 2, pages 1150–1157. Ieee, 1999.
- [26] J. Mao, W. Xu, Y. Yang, J. Wang, and A. L. Yuille. Explain images with multimodal recurrent neural networks. *arXiv preprint arXiv:1410.1090*, 2014.
- [27] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.
- [28] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, pages 3111–3119, 2013.
- [29] M. Mitchell, X. Han, J. Dodge, A. Mensch, A. Goyal, A. Berg, K. Yamaguchi, T. Berg, K. Stratos, and

- H. Daumé III. Midge: Generating image descriptions from computer vision detections. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 747–756. Association for Computational Linguistics, 2012.
- [30] A. Mnih and G. Hinton. Three new graphical models for statistical language modelling. In *Proceedings of the 24th International Conference on Machine Learning, ICML '07*, pages 641–648, New York, NY, USA, 2007. ACM.
- [31] R. Neal. Annealed importance sampling. *Statistics and Computing*, 11(2):125–139, 2001.
- [32] X. Peng, C. Zou, Y. Qiao, and Q. Peng. Action recognition with stacked fisher vectors. In *Computer Vision–ECCV 2014*, pages 581–595. Springer, 2014.
- [33] F. Perronnin and C. Dance. Fisher kernels on visual vocabularies for image categorization. In *Computer Vision and Pattern Recognition, 2007. CVPR'07. IEEE Conference on*, pages 1–8. IEEE, 2007.
- [34] F. Perronnin, Y. Liu, J. Sánchez, and H. Poirier. Large-scale image retrieval with compressed fisher vectors. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 3384–3391. IEEE, 2010.
- [35] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *Computer Vision–ECCV 2010*, pages 143–156. Springer, 2010.
- [36] C. Rashtchian, P. Young, M. Hodosh, and J. Hockenmaier. Collecting image annotations using amazon’s mechanical turk. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon’s Mechanical Turk*, pages 139–147. Association for Computational Linguistics, 2010.
- [37] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek. Image classification with the fisher vector: Theory and practice. *International journal of computer vision*, 105(3):222–245, 2013.
- [38] K. Simonyan, O. M. Parkhi, A. Vedaldi, and A. Zisserman. Fisher vector faces in the wild. In *Proc. BMVC*, volume 1, page 7, 2013.
- [39] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep fisher networks for large-scale image classification. In *Advances in neural information processing systems*, pages 163–171, 2013.
- [40] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.
- [41] J. Sivic, B. C. Russell, A. A. Efros, A. Zisserman, and W. T. Freeman. Discovering objects and their location in images. In *Computer Vision, 2005. ICCV 2005. Tenth IEEE International Conference on*, volume 1, pages 370–377. IEEE, 2005.
- [42] R. Socher and L. Fei-Fei. Connecting modalities: Semi-supervised segmentation and annotation of images using unaligned text corpora. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 966–973. IEEE, 2010.
- [43] R. Socher, Q. Le, C. Manning, and A. Ng. Grounded compositional semantics for finding and describing images with sentences. In *NIPS Deep Learning Workshop*, 2013.
- [44] R. Socher, C. C. Lin, A. Y. Ng, and C. D. Manning. Parsing Natural Scenes and Natural Language with Recursive Neural Networks. In *Proceedings of the 26th International Conference on Machine Learning (ICML)*, 2011.
- [45] V. Sydorov, M. Sakurada, and C. H. Lampert. Deep fisher kernels—end to end learning of the fisher kernel gmm parameters. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2014.
- [46] H. Vinod. Canonical ridge and econometrics of joint production. *Journal of Econometrics*, 4(2):147 – 166, 1976.
- [47] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan. Show and tell: A neural image caption generator. *arXiv preprint arXiv:1411.4555*, 2014.
- [48] C. L. Zitnick, D. Parikh, and L. Vanderwende. Learning the visual interpretation of sentences. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1681–1688. IEEE, 2013.