

Self-Supervised Visual Representations for Cross-Modal Retrieval

Yash Patel

The Robotics Institute, Carnegie Mellon University, PA, USA
yashp@andrew.cmu.edu

Lluís Gomez

Computer Vision Center, Universitat Autònoma de Barcelona, Spain
gomez@cvc.uab.es

Marçal Rusiñol

Computer Vision Center, Universitat Autònoma de Barcelona, Spain
marcal@cvc.uab.es

Dimosthenis Karatzas

Computer Vision Center, Universitat Autònoma de Barcelona, Spain
dimos@cvc.uab.es

C.V. Jawahar

CVIT, KCIS, IIT Hyderabad, India
jawahar@iiit.ac.in

ABSTRACT

Cross-modal retrieval methods have been significantly improved in last years with the use of deep neural networks and large-scale annotated datasets such as ImageNet and Places. However, collecting and annotating such datasets requires a tremendous amount of human effort and, besides, their annotations are limited to discrete sets of popular visual classes that may not be representative of the richer semantics found on large-scale cross-modal retrieval datasets. In this paper, we present a self-supervised cross-modal retrieval framework that leverages as training data the correlations between images and text on the entire set of Wikipedia articles. Our method consists in training a CNN to predict: (1) the semantic context of the article in which an image is more probable to appear as an illustration, and (2) the semantic context of its caption. Our experiments demonstrate that the proposed method is not only capable of learning discriminative visual representations for solving vision tasks like classification, but that the learned representations are better for cross-modal retrieval when compared to supervised pre-training of the network on the ImageNet dataset.

KEYWORDS

Self-Supervised Learning, Visual Representations, Cross-Modal Retrieval

ACM Reference Format:

Yash Patel, Lluís Gomez, Marçal Rusiñol, Dimosthenis Karatzas, and C.V. Jawahar. 2019. Self-Supervised Visual Representations for Cross-Modal Retrieval. In *International Conference on Multimedia Retrieval (ICMR '19)*, June 10–13, 2019, Ottawa, ON, Canada. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3323873.3325035>

1 INTRODUCTION

The emergence of large-scale annotated datasets such as ImageNet [4], Places [41] and MS-COCO [16] has undoubtedly been one of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICMR '19, June 10–13, 2019, Ottawa, ON, Canada

© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6765-3/19/06...\$15.00
<https://doi.org/10.1145/3323873.3325035>

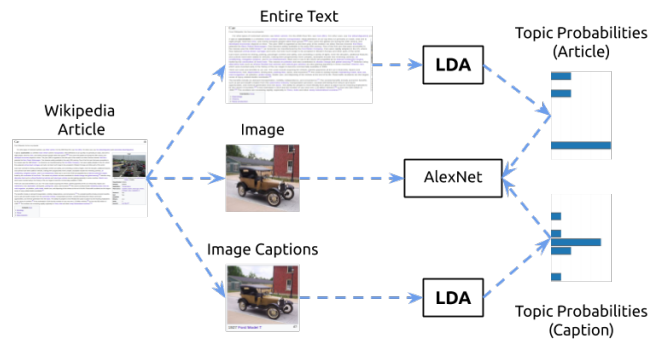


Figure 1: Method overview: Wikipedia articles contain textual description of a subject, these articles are also accompanied with illustrative images supporting the text. These images are often accompanied by captions. A Latent Dirichlet Allocation (LDA) [3] topic modeling framework generates a global contextual representation of the textual information from entire text article. The same LDA model generates a local contextual representation from the per-image caption. These two text representations are jointly used to supervised the training of deep CNN.

key ingredients for the tremendous impact of deep learning on almost every computer vision task. However, there is a major issue with the supervised learning setup in large scale datasets; collecting and manually annotating those datasets requires a great amount of human effort. On the other hand, the fact that the annotations on such datasets are usually limited to discrete sets of popular visual classes, may not necessarily be an optimal training setup for cross-modal retrieval datasets.

As an alternative to the fully supervised setup, self-supervised learning methods aim at learning discriminative visual features by designing auxiliary tasks for which the target labels are free to obtain. These labels provide supervision for the training of computer vision models the same way as in supervised learning, but the supervisory signal can be directly obtained from the training data, either from the images themselves [5, 21] (uni-modal training) or from a complementary modality that is found naturally correlated with them [1, 9, 18] (multi-modal training). Unlike supervised

learning, where visual features are learnt from human generated labels, in self-supervised learning labels are automatically obtained.

In this paper we present a self-supervised cross-modal retrieval framework that leverages as supervisory signal the correlations found between images and text on a large collection of illustrated articles in order to learn discriminative visual features that could potentially transfer well to any general computer vision task, such as image classification or object detection. We hypothesise that the learned representations by using such approach can be used more naturally in a cross-modal retrieval framework than the representations learned from annotated datasets for image classification.

Gomez and Patel *et al.* [9, 19] have proposed in the past self-supervised representation learning using Wikipedia articles. Their method consists in learning a Latent Dirichlet Allocation (LDA) model from the entire corpus of text articles, and then training a CNN to predict the semantic context of images by using as training labels the semantic level representations (the probability distribution over semantic topics) of the articles in which they appear, as provided by the LDA model. An assumption made in their method is that all the images within a given text article have the same target semantic representation, which is obtained from the LDA model. However, images within a Wikipedia article can be drastically different in terms of appearance and semantic content. To overcome this, we create a new *Wikipedia dataset with captions* which is similar to the one used in the TextTopicNet [9, 19] method, but also containing the image captions from Wikipedia. Thus, as illustrated in Figure 1, the training data in our method comes in a triplet form (image, text article, image caption).

Our intuition is that adding another target representation based on image captions could provide more image specific training self-supervision. Furthermore, we experimentally show that our training procedure leads to significantly better results for both cross-modal retrieval and image classification.

Following are the major contributions made in this paper:

- We propose a multi-task learning framework to train a CNN that predicts text representations obtained from text articles (global context) and per-image captions (local context).
- We experimentally demonstrate that the self-supervisedly learned visual features are generic enough for other computer vision tasks and outperform other self-supervised and naturally supervised approaches on standard benchmarks.
- Without using any form of semantic information, our method outperforms both unsupervised and supervised approaches on cross-modal retrieval (image-to-text and text-to-image) benchmarks on Wikipedia [23] and Pascal sentences [7].
- The Wikipedia image-article dataset [19] consist of only images and text articles, and as an auxiliary contribution, we release a large scale dataset obtained from English Wikipedia consisting of images, per-image captions and text articles.

2 RELATED WORK

2.1 Self-Supervised Visual Representations

As an alternative to fully-supervised algorithms, there has recently been a growing interest in self-supervised or naturally-supervised approaches that make use of non-visual signals, intrinsically correlated to the image, as a form to supervise visual feature learning.

The objective of those methods is to learn visual representations (without human annotations) that are generic to work well across a wide range of object classes and at the same time are discriminating enough to be useful for classical computer vision tasks such as image classification, object detection, semantic segmentation etc.

Among various categories of such methods, our work is most similar to multi-modal self-supervised learning algorithms that attempt to utilize the structure in one modality to provide the training supervision for co-occurring modality.

Owens *et al.* [18] make use of sound as a modality to provide supervisory signal. They do so by training a deep CNN to predict a hand-crafted statistical summary of sound associated with a video frame. Gomez & Patel *et al.*[9] make use of Wikipedia documents which consist of text articles and co-occurring images. First, a Latent Dirichlet Allocation (LDA) [3] topic model is learned on the entire Wikipedia dataset. Second, text articles are represented in the form of topic-probabilities using learned LDA model. Finally, a convolutional neural network is trained on images in Wikipedia, where the target label is the representation of corresponding text article.

Our work is more closely related to [9, 19, 20], however, as previously mentioned, their approach makes use of same target representation for all images within a text article.

2.2 Cross-Modal Representation Learning

Two general categories of the representation learning methods for cross-modal retrieval can be: (a) *real-valued*, (b) *binary valued*. The binary methods are more focused on efficiency and aim to map the items from different modalities on a common binary hamming space [27, 29, 37, 42].

Our approach falls in the category of *real-valued* methods. Within this category of methods the training for cross-modal retrieval could be: unsupervised [2, 8, 11, 30, 38] or supervised [10, 33, 34, 39].

Zhang *et al.*[40] propose a multimodal hashing method, called semantic correlation maximization (SCM), which integrates semantic labels into the hashing learning procedure. Gong *et al.*[10] propose a novel three-view CCA (CCA-3V) framework, which explicitly incorporates the dependence of visual features and text.

Wang *et al.*[34] propose a novel regularization framework for the cross-modal matching problem, called LCFS (Learning Coupled Feature Spaces). It unifies coupled linear regressions, l_{21} -norm and trace norm into a generic minimization formulation so that subspace learning and coupled feature selection can be performed simultaneously. Furthermore, they extend this framework to more than two-modality case in [33], where the extension version is called JFSSL (Joint Feature Selection and Subspace Learning). Wang *et al.*[31] propose an adversarial learning approach for cross-modal retrieval.

While most of these supervised or unsupervised approaches attempts to learn a common embedding space for the purpose of cross-modal retrieval, they assume that the visual representations are provided by a pre-trained CNN (either AlexNet [14] or VGG-16 [28]) on ImageNet dataset [25]. The cost (human annotation effort) of this pre-training is not accounted by cross-modal retrieval methods. Further, the underlying assumption of is that ImageNet pre-trained features transfer well for cross-modal retrieval.

3 WIKIPEDIA DATASET WITH CAPTIONS

In order to obtain a training dataset for our method, we scrapped the entire English Wikipedia while considering only articles with at least 50 words and illustrated with at-least one image. Similarly to the preprocessing of ImageCLEF dataset we filtered small images (< 256 pixels) and images with formats other than JPG. Furthermore, we only considered the images for which captions are available. With these constraints our dataset consists of 1.8 million images with captions and the associated text article that they appear with.

4 METHOD

In this section, we first elaborate the Latent Dirichlet Allocation (LDA) [3], which is used for representing text articles and image captions, and thus for generating target representations for training the CNN. Then, we go over the training of the CNN.

4.1 Latent Dirichlet Allocation

LDA [3] is a generative statistical model of a text corpus where each document can be viewed as a mixture of various topics, and each topic is characterized by a probability distribution over words. LDA can be represented as a three level hierarchical Bayesian model. Given a text corpus consisting of M documents and a dictionary with N words, Blei *et al.* define the generative process [3] for a document d as follows:

- Choose $\theta \sim \text{Dirichlet}(\alpha)$.
- For each of the N words w_n in d :
 - Choose a topic $z_n \sim \text{Multinomial}(\theta)$.
 - Choose a word w_n from $P(w_n | z_n, \beta)$, a multinomial probability conditioned on the topic z_n .

where θ is the mixing proportion and is drawn from a Dirichlet prior with parameter α , and both α and β are corpus level parameters, sampled once in the process of generating a corpus. Each document is generated according to the topic proportions $z_{1:K}$ and word probabilities over β . The probability of a document d in a corpus is defined as:

$$P(d | \alpha, \beta) = \int_{\theta} P(\theta | \alpha) \left(\prod_{n=1}^N \sum_{z_K} P(z_K | \theta) P(w_n | z_K, \beta) \right) d\theta$$

Learning LDA [3] on a document corpus provides two set of parameters: word probabilities given topic $P(w | z_{1:K})$ and topic probabilities given document $P(z_{1:K} | d)$. Therefore each document is represented in terms of topic probabilities $z_{1:K}$ (K is the number of topics) and word probabilities over topics. Any unseen document can be represented in terms of a probability distribution over topics of the learned LDA model by projecting it into the topic space.

4.2 Network Architecture

Throughout our experiments, we make use of AlexNet architecture [14]. The choice of AlexNet is justified because most of the existing self-supervised methods make use of this same architecture [1, 9, 18, 19, 21, 32]. Further, we compare to cross-modal retrieval methods with reported performance using AlexNet [10, 11, 23, 26, 33, 34]. Thus the use of AlexNet architecture is for fair comparisons. Till the $fc7$ layer the architecture is same as standard AlexNet [14], which is followed by two fully-connected branches one prediction

caption level topic probabilities and other predicting article level topic probabilities.

4.3 Learning Self-Supervised Representations

The multimodal document from Wikipedia can be thought as a training sample, x_i . This multimodal document consists of text article x_i^A , image captions x_i^C and images x_i^I . Let $\Phi(x_i^A)$ and $\Phi(x_i^C)$ be the text topic probability distributions given by LDA 4.1 for the document text and the image captions accordingly. The deep CNN is trained to predict the above topic distributions given the corresponding article image, and producing as outputs: $f_A(x_i^I, \Theta)$ (for article) and $f_C(x_i^I, \Theta)$ (for caption).

The loss is computed as the cross entropy between the LDA topic distribution and the predicted distribution. The overall risk term on the training data will be:

$$R = \sum_{i=1}^{i=N} \left[\sum_{topic=1}^{topic=K} \Phi(x_i^A)_{topic} \log(f_A(x_i^I, \Theta)_{topic}) + \sum_{topic=1}^{topic=K} \Phi(x_i^C)_{topic} \log(f_C(x_i^I, \Theta)_{topic}) \right] \quad (1)$$

where N is total number of samples in the training data, K is the number of topics in the LDA model [3] and Θ maps to the learnt CNN parameters. Note that K is a hyper-parameter and we fix $K = 40$ throughout the experiments.

Learning to predict the target topic probability distributions we minimize a sigmoid cross-entropy loss as shown in the overall risk term Eq. 1. We use a Stochastic Gradient Descent (SGD) optimizer, with base learning rate of 0.001, with a step decay after every 200,000 iterations by a factor of 0.1, and momentum of 0.9. The batch size is set to 128. With these settings the network converges after 500,000 iterations of training.

5 EXPERIMENTS

We will first compare the learnt visual representations with other self-supervised methods on the task on image classification on a standard benchmark datasets (Section 5.1). Next, we compare our method with various cross-modal retrieval methods (Section 5.2).

5.1 Self-Supervised Features for Image Classification

Self-supervised learned features are tested for image classification on PASCAL VOC 2007 [6] dataset. In total there are 9,963 images, and 20 semantic classes. The data has been split into 50% for training/validation and 50% for testing. The setup here is multi-label, that is, each image can be classified into multiple classes.

We extract features from the top layers of the CNN (fc7, fc6, pool5) for each image of the dataset. Then, for each class we perform a grid search over the parameter space of an one-vs-all Linear SVM classifier¹ to optimize its validation accuracy.

In Table 1, we compare our results on the PASCAL VOC2007 test set with different state-of-the-art self-supervised learning algorithms using features from different top layers and SVM classifiers.

¹Liblinear implementation from <http://scikit-learn.org/>

Method	max5	pool5	fc6	fc7
Ours	-	53.8	54.9	56.8
TextTopicNet (Wikipedia) [19]	-	51.9	54.2	55.8
TextTopicNet (ImageCLEF) [9]	-	47.4	48.1	48.5
Sound [18]	39.4	46.7	47.1	47.4
Texton-CNN	28.9	37.5	35.3	32.5
K-means [13]	27.5	34.8	33.9	32.1
Tracking [36]	33.5	42.2	42.4	40.2
Patch pos. [5]	26.8	46.1	-	-
Egomotion [1]	22.7	31.1	-	-
ImageNet [14]	63.6	65.6	69.6	73.6
Places [41]	59.0	63.2	65.3	66.2

Table 1: PASCAL VOC2007 mAP comparison for image classification with supervised (bottom), and self-supervised (middle) methods.

Our method which leverages global and local contexts for self-supervised training achieves state-of-the-art performance as seen in Table 1. This demonstrates that a network that identifies global and local semantic contexts in which an image is more probable to appear gives better visual representations.

5.2 Cross-Modal Retrieval

As seen in Fig. 1, the final layer of the network projects the images on same representation as text as obtained by the LDA model (Section 4.1). Therefore, cross-modal retrieval can be directly done by making use of LDA topic probabilities for text and network final predictions for image. We use KL-divergence as a distance metric to short the samples of the target modality, since both the LDA encoding and the CNN output represent probability distributions.

5.2.1 Wikipedia. We use the Wikipedia retrieval dataset [23], which consists of 2,866 image-article pairs split into train and test set of 2,173 and 693 pairs respectively. Further, each image-document pair is labeled with one of the ten semantic classes [23]. In Table 2 we compare our results with supervised and unsupervised multi-modal retrieval methods discussed in [35] and [12]. Supervised methods make use of class or categorical information associated with each image-document pair, whereas unsupervised methods do not. All of these methods use LDA for text representation and CNN features from pre-trained CaffeNet, which is trained on ImageNet dataset in a supervised setting. We observe that our method leads to state-of-the-art performance, even when compared to fully supervised approaches. This demonstrates that training a network to predict both the global and local semantic contexts in which it is more probable to appear leads to better learning for retrieval task. Further note that, except ours and TextTopicNet [9, 19] all the other methods use ImageNet pre-trained network.

5.2.2 Pascal Sentences. We also evaluate our method on pascal sentences dataset [7] which is a subset of pascal VOC dataset. It contains 1000 pairs of an image along with several sentences from 20 categories. While, the other methods randomly split the dataset into 600 training and 400 testing samples, we test on all 1000 samples. This is due to the fact that we do not make use of this dataset for training at any point. Table 3 provides an extensive comparison with

Method	Image Query	Text Query	Average
Ours	39.10	43.40	41.25
TextTopicNet (Wikipedia) [19]	37.63	40.25	38.94
TextTopicNet (ImageCLEF) [9]	39.58	38.16	38.87
CCA [11, 23]	19.70	17.84	18.77
PLS [24]	30.55	28.03	29.29
SCM* [23]	37.13	28.23	32.68
GMMFA* [26]	38.74	31.09	34.91
CCA-3V* [10]	40.49	36.51	38.50
GMLDA* [26]	40.84	36.93	38.88
LCFS* [34]	41.32	38.45	39.88
JFSSL* [33]	42.79	39.57	41.18

Table 2: Mean average precision (MAP) comparison on Wikipedia dataset [23] with supervised (bottom), unsupervised (middle) and self-supervised (top) methods. Methods marked with asterisk make use of document (image-text) class category information.

Method	Image Query	Text Query	Average
Ours	32.6	36.0	34.3
TextTopicNet (Wikipedia)[19]	30.1	35.2	32.7
TextTopicNet (ImageCLEF) [9]	26.4	31.6	29.0
CCA [11, 23]	9.90	9.7	9.8
CFA [15]	18.7	21.6	20.2
KCCA (Poly) [11]	20.7	19.1	19.9
KCCA (RBF) [11]	23.3	24.9	24.1
Bimodal AE [17]	24.5	25.6	25.1
Multimodal DBN [30]	19.7	18.3	19.0
Corr-AE [8]	26.8	27.3	27.1
JRL [39]	30.0	28.6	29.3
CMDN [22]	33.4	33.3	33.4

Table 3: Mean average precision (MAP) comparison on pascal sentences dataset [7] with supervised image representations (bottom) and self-supervised image representations (top) methods.

existing methods. Compared to other retrieval methods that use self-supervised visual representations [9, 19], our method achieves 1.6% higher MAP with $\frac{1}{4^{th}}$ the size of training data. This demonstrates the efficacy of jointly using global and local self-supervision signals.

6 CONCLUSION

In this paper, we put forward a self-supervised method that takes advantage of the natural correlation between an article’s text and the images used to illustrate it, in order to learn useful visual representations.

We demonstrated that the learned visual features can transfer well to any general computer vision task such as image classification, while they can be directly used in a cross-modal retrieval framework yielding state of the art results both on the Wikipedia retrieval and the Pascal Sentences datasets.²

²Trained models and dataset are publicly available at: <https://github.com/yash0307/Self-Supervised-Retrieval>

REFERENCES

- [1] Pulkit Agrawal, Joao Carreira, and Jitendra Malik. 2015. Learning to see by moving. In *ICCV*.
- [2] Galen Andrew, Raman Arora, Jeff Bilmes, and Karen Livescu. 2013. Deep canonical correlation analysis. In *ICML*.
- [3] David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent Dirichlet allocation. *Journal of Machine Learning Research* (2003).
- [4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*.
- [5] Carl Doersch, Abhinav Gupta, and Alexei A Efros. 2015. Unsupervised visual representation learning by context prediction. In *ICCV*.
- [6] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. 2010. The pascal visual object classes (VOC) challenge. *IJCV* (2010).
- [7] Ali Farhadi, Mohsen Hejrati, Mohammad Amin Sadeghi, Peter Young, Cyrus Rashtchian, Julia Hockenmaier, and David Forsyth. 2010. Every picture tells a story: Generating sentences from images. In *ECCV*.
- [8] Fangxiang Feng, Xiaojie Wang, and Ruifan Li. 2014. Cross-modal retrieval with correspondence autoencoder. In *ACM-MM*.
- [9] Lluís Gomez, Yash Patel, Marçal Rusiñol, Dimosthenis Karatzas, and CV Jawahar. 2017. Self-supervised learning of visual features through embedding images into text topic spaces. In *CVPR*.
- [10] Yunchao Gong, Qifa Ke, Michael Isard, and Svetlana Lazebnik. 2014. A multi-view embedding space for modeling internet images, tags, and their semantics. *International Journal of Computer Vision* (2014).
- [11] David R Hardoon, Sandor Szedmak, and John Shawe-Taylor. 2004. Canonical correlation analysis: An overview with application to learning methods. *Neural Computation* (2004).
- [12] Cuicui Kang, Shengcai Liao, Yonghao He, Jian Wang, Wenjia Niu, Shiming Xiang, and Chunhong Pan. 2015. Cross-modal similarity learning: A low rank bilinear formulation. In *CIKM*.
- [13] Philipp Krähenbühl, Carl Doersch, Jeff Donahue, and Trevor Darrell. 2015. Data-dependent initializations of convolutional neural networks. In *ICLR*.
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *NIPS*.
- [15] Dongge Li, Nevenka Dimitrova, Mingkun Li, and Ishwar K Sethi. 2003. Multimedia content processing through cross-modal association. In *ACM-MM*.
- [16] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *ECCV*.
- [17] Jiquan Ngiam, Aditya Khosla, Mingyu Kim, Juhan Nam, Honglak Lee, and Andrew Y Ng. 2011. Multimodal deep learning. In *ICML*.
- [18] Andrew Owens, Jiajun Wu, Josh H McDermott, William T Freeman, and Antonio Torralba. 2016. Ambient sound provides supervision for visual learning. In *ECCV*.
- [19] Yash Patel, Lluís Gomez, Raul Gomez, Marçal Rusiñol, Dimosthenis Karatzas, and CV Jawahar. 2018. TextTopicNet-Self-Supervised Learning of Visual Features Through Embedding Images on Semantic Text Spaces. *arXiv preprint arXiv:1807.02110* (2018).
- [20] Yash Patel, Lluís Gomez, Marçal Rusiñol, and Dimosthenis Karatzas. 2016. Dynamic Lexicon Generation for Natural Scene Images. In *ECCV*.
- [21] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. 2016. Context encoders: Feature learning by inpainting. In *CVPR*.
- [22] Yuxin Peng, Xin Huang, and Jinwei Qi. 2016. Cross-Media Shared Representation by Hierarchical Learning with Multiple Deep Networks. In *IJCAI*.
- [23] Nikhil Rasiwasia, Jose Costa Pereira, Emanuele Coviello, Gabriel Doyle, Gert RG Lanckriet, Roger Levy, and Nuno Vasconcelos. 2010. A new approach to cross-modal multimedia retrieval. In *ACM-MM*.
- [24] Roman Rosipal and Nicole Krämer. 2006. Overview and recent advances in partial least squares. In *Subspace, latent structure and feature selection*.
- [25] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision* (2015).
- [26] Abhishek Sharma, Abhishek Kumar, Hal Daume, and David W Jacobs. 2012. Generalized multiview analysis: A discriminative latent space. In *CVPR*.
- [27] Fumin Shen, Xiang Zhou, Yang Yang, Jingkuan Song, Heng Tao Shen, and Dacheng Tao. 2016. A Fast Optimization Method for General Binary Code Learning. *IEEE Transactions on Image Processing* (2016).
- [28] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [29] Jingkuan Song, Yang Yang, Yi Yang, Zi Huang, and Heng Tao Shen. 2013. Inter-media hashing for large-scale retrieval from heterogeneous data sources. In *ACM-SIGMOD*.
- [30] Nitish Srivastava and Ruslan R Salakhutdinov. 2012. Multimodal learning with deep boltzmann machines. In *NIPS*.
- [31] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. 2017. Adversarial cross-modal retrieval. In *ACM-MM*.
- [32] Dong Wang and Xiaoyang Tan. 2016. Unsupervised feature learning with c-svddnet. *Pattern Recognition* (2016).
- [33] Kaiye Wang, Ran He, Liang Wang, Wei Wang, and Tieniu Tan. 2016. Joint feature selection and subspace learning for cross-modal retrieval. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2016).
- [34] Kaiye Wang, Ran He, Wei Wang, Liang Wang, and Tieniu Tan. 2013. Learning coupled feature spaces for cross-modal matching. In *ICCV*.
- [35] Kaiye Wang, Qiyue Yin, Wei Wang, Shu Wu, and Liang Wang. 2016. A Comprehensive Survey on Cross-modal Retrieval. *CoRR* (2016).
- [36] Xiaolong Wang and Abhinav Gupta. 2015. Unsupervised learning of visual representations using videos. In *CVPR*.
- [37] Xing Xu, Fumin Shen, Yang Yang, Heng Tao Shen, and Xuelong Li. 2017. Learning discriminative binary codes for large-scale cross-modal retrieval. *IEEE Transactions on Image Processing* (2017).
- [38] Fei Yan and Krystian Mikolajczyk. 2015. Deep correlation for matching images and text. In *CVPR*.
- [39] Xiaohua Zhai, Yuxin Peng, and Jianguo Xiao. 2014. Learning cross-media joint representation with sparse and semisupervised regularization. *IEEE Transactions on Circuits and Systems for Video Technology* (2014).
- [40] Junbo Zhao, Michael Mathieu, Ross Goroshin, and Yann Lecun. 2016. Stacked what-where auto-encoders. In *ICLR*.
- [41] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. 2014. Learning deep features for scene recognition using places database. In *NIPS*.
- [42] Xiaofeng Zhu, Zi Huang, Heng Tao Shen, and Xin Zhao. 2013. Linear cross-modal hashing for efficient multimedia search. In *ACM-MM*.