ADVERSARIAL LEARNING OF LABEL DEPENDENCY: A NOVEL FRAMEWORK FOR MULTI-CLASS CLASSIFICATION

Che-Ping Tsai, Hung-Yi Lee

National Taiwan University

r06922039@ntu.edu.tw, tlkagkb93901106@gmail.com

ABSTRACT

Recent work has shown that exploiting relations between labels improves the performance of multi-label classification. We propose a novel framework based on generative adversarial networks (GANs) to model label dependency. The discriminator learns to model label dependency by discriminating real and generated label sets. To fool the discriminator, the classifier, or generator, learns to generate label sets with dependencies close to real data. Extensive experiments and comparisons on two large-scale image classification benchmark datasets (MS-COCO and NUS-WIDE) show that the discriminator improves generalization ability for different kinds of models.

Index Terms— Multi-label classification, Generative Adversarial Network

1. INTRODUCTION

Multi-label classification is a fundamental but challenging problem in machine learning with applications such as multi-object recognition [1, 2], image classification [3], text categorization [4], and music categorization [5]. In contrast to single-label classification, multilabel predictors must not only relate labels with the corresponding instances, but also exploit dependencies between labels due to label co-occurrences. Take for instance multi-label image categorization: *beach* and *sky* usually appear together in the same image, whereas *airplane* and *dog* do not often co-occur.

A simple approach to using deep networks – for example CNNs – for multi-label classification is to recast the problem as multiple disjoint binary classification by replacing cross-entropy loss with logistic loss or ranking loss [6]. To model label dependency, however, recent work has focused on capturing cross-label correlation using probabilistic graphical networks [7, 8], dependency networks [9], recurrent neural networks (RNNs) [10], and so on.

In this work, we propose a new framework under which to train a multi-label classifier. We use a generative adversarial network (GAN) to model the label distribution for multi-label classification. This framework is built upon a conditional GAN (cGAN). The classifier here plays the role of a conditional generator, whose input is an instance, and which outputs a set of labels as with a typical multi-label classifier. A discriminator is trained to model label dependency: it takes an object and a set of labels as input, and outputs a score. The set of labels comes either from the training data or from the output of the classifier, that is, the generator. The discriminator learns to discriminate the real and generated label sets. To tell the real label from the generated ones, the discriminator must model the correlation of the input instances and their corresponding label sets as well as the label dependency of the label sets in the training data. The classifier then learns to fool the discriminator by generating label sets with what seems to be the correct dependencies, given an input instance by the discriminator. The classifier and discriminator are learned iteratively as in a typical GAN.

As the proposed framework is general and independent of the network architecture of the classifier, we believe the discriminator can be easily appended to other models to help learn label dependencies. Evaluation on two public multi-label image classification datasets shows that the discriminator facilitates generalization ability among CNNs with different architectures. To the best of our knowledge, this is the first attempt to utilize GANs for multi-label classification.

2. RELATED WORK

Multi-label classification has been widely studied in image classification. A straightforward way to deal with multi-label classification is to decompose it into multiple binary classification tasks, such as binary relevance [11] using neural networks. To further improve performance, recent work has taken into account interdependency between labels. Gong et al. [6] evaluate various objectives with a CNN architecture, and find that weighted approximate ranking loss works best with CNNs. To better model the structure of label correlations, traditional graphical models have also been used for this task [7, 8]; latent space methods [12, 13] have also been proposed. Wang et al. [10] and Chen et al. [14] combine CNN and RNN to jointly embed images and semantic structure of labels in the same embedding space. Zhu et al. [15] further propose a spatial regularization network to capture both spatial and semantic relations.

The work mentioned above mainly considers the global representation of the whole image, ignoring the relationships between semantic labels and local image regions, which is difficult to decipher given complex backgrounds. To handle such cases, Wei et al. [13] propose a Hypothesis-CNN-Pooling framework to aggregate the label scores of each proposal using category-wise maxpooling. Yang et al. [16] transform the multi-label recognition problem into a multi-class, multi-instance learning problem and make use of label-view information of the proposals to enhance features. Newer work [17, 18] uses long short term memory (LSTM) units to iteratively discover a sequence of attentional and informative regions and further predict labeling scores.

The proposed approach is independent of the above approaches. It is possible to further improve the above approaches with a discriminator.



Fig. 1. Proposed framework for multi-label classification, shown in (a). Fig (b) illustrates three kinds of discriminator inputs. The upper one is the positive example (*real label set y, image x*), a matched pair sampled from real data distribution. The others are negative examples. The middle is the generated pair (*generated label set* \hat{y} , *image x*), where \hat{y} is sampled from the output of generator \tilde{y} . The bottom is the mismatched pair (*random sampled label set y', image x*). The discriminator learns to assign low scores to the negative pairs, and high scores to the positive pairs as genuine.

3. METHODOLOGY

The overview of the proposed framework is shown in Fig. 1(a). Here, multi-label image classification is considered an instance of multiclass classification. Let x denote an input image for which its corresponding ground-truth labels are $y \in \{0, 1\}^{|S|}$, where |S| is the number of labels (or classes), and y is a set of labels. The generator and discriminator in Fig. 1 are trained iteratively. That is, we fix one and update the other several times.

3.1. Classifier (Generator) and Discriminator

The generator G^1 is a classifier with sigmoid activation functions at the output layer. The classifier here can have a wide variety of architectures, for example, VGG-16 [19], Inception_v3 [20], Resnet-101 [21], or Resnet-152 [21]. *G* takes the input image *x* and predicts the probability of each label to generate a probability distribution \tilde{y} . During the training phase, the predicted label set \hat{y} is sampled from \tilde{y} by considering the sigmoid output as a probability. During testing, labels whose sigmoid outputs are larger than 0.5 are in the output label set.

The discriminator D receives a label set y (or \hat{y}) and an image x, and produces a score, D(y, x) (or $D(\hat{y}, x)$), which represents the measurement of how "real" the label distribution is and the degree of match between the image and the label set. D contains a feature extractor network f_{ext} which produces $z = f_{ext}(x)$ as the feature vector for image x. Then a feedforward network takes z and y (or \hat{y}) as input and outputs a single value D(y, x) (or $D(\hat{y}, x)$). During the testing phase, only the generator is used; the discriminator is left unused.

3.2. Classifier Training

G is pretrained as a typical multi-class classifier by minimizing the binary cross-entropy loss $\mathcal{L}_{logistic}$ using the ground truth labels $y = [y_1, y_2, ..., y_{|S|}]^T$ and the predicted probability of each label $\tilde{y} =$

 $[\tilde{y}_1, \tilde{y}_2, ..., \tilde{y}_{|S|}]^T$:

$$\mathcal{L}_{logistic} = \mathbb{E}_{(x,y)\sim data} [\sum_{i=1}^{|S|} y_i log \, \tilde{y}_i + (1-y_i) log (1-\tilde{y}_i)], \quad (1)$$

where (x, y) is an (image, label set) pair sampled from training data, and \tilde{y} is the output label distribution of G given x sampled from the training data. As all labels are considered independently in $\mathcal{L}_{logistic}$, the classifier learned by minimizing $\mathcal{L}_{logistic}$ is not guaranteed to model dependencies between labels.

During the iterative training, the loss function of generator G is

$$\mathcal{L}'_G = \mathcal{L}_G + \alpha \mathcal{L}_{logistic},\tag{2}$$

where α determines the scale of the logistic loss and

$$\mathcal{L}_G = -\mathbb{E}_{x \sim data, \hat{y} \sim G(x)}[D(\hat{y}, x)], \tag{3}$$

where x is an input image sampled from training data. \hat{y} is the output label set of G sampled from \tilde{y} . $D(\hat{y}, x)$ is the score assigned by the discriminator given an (image, label set) pair. With (3), the generator learns not only to minimize the logistic loss $\mathcal{L}_{logistic}$, but also to produce a reasonable combination of labels to fool the discriminator D by maximizing the value $D(\hat{y}, x)$. Because D takes the whole label set y or \hat{y} as input, it uses the dependency between the labels to discriminate real and generated label sets. Therefore, the G learned from D takes into account label dependency.

Equation (3) requires the composition of the generator and discriminator to be fully differentiable. Here, since the label set \hat{y} is a discrete multi-hot vector, we use the Gumbel-softmax trick for a Bernoulli distribution of each label [22, 23], which we here term *Gumbel sigmoid*, to reparameterize the sampling procedure to make it differentiable.

3.3. Discriminator Training

To train the discriminator D, (*real label set y, image x*) sampled from real data distribution serves as positive examples. For negative examples, we not only provide D with the generated pairs, (*generated label set*, \hat{y} , *image x*), but also the mismatched pairs, (*random sampled y'*, *image x*), where y' is a label set randomly sampled

¹In contrast to the typical cGAN setting, the conditional generator here is conditioned on the image, not the class information.

from the training data which does not correspond to image x [24]. Fig. 1(b) illustrates these three kinds of discriminator inputs. Here, we use a Wasserstein GAN with a gradient penalty (WGAN-gp). The loss function of discriminator \mathcal{L}_D is

$$\mathcal{L}_{D} = -\mathbb{E}_{(x,y)\sim data}[D(y,x)] + \frac{1}{2}\mathbb{E}_{x\sim data,\hat{y}\sim G(x)}[D(\hat{y},x)] + \frac{1}{2}\mathbb{E}_{x\sim data,y'\sim data}[D(y',x)] + \lambda\mathcal{L}_{gp}.$$
(4)

In the first term, an image x and its label set y is sampled from the training dataset, from which D learns to assign large values to genuine samples. In the second term, the image x and the label set \hat{y} predicted by the generator is generated, for which D is expected to assign a small value. To help the discriminator to learn how to discriminate mismatched image/label pairs, we use negative sampling as (the third term): we randomly select mismatched (image x, label set y') pairs from the training data, to which D also learns to assign low scores. That is, the discriminator learns to assign low scores to two kinds of errors: the generated label set from the classifier and realistic label sets with the wrong images that mismatch the condition information.

The last term \mathcal{L}_{gp} is the gradient penalty [25]; λ determines the gradient penalty scale.

$$\mathcal{L}_{gp} = \mathbb{E}_{x \sim data, \hat{y} \sim G(x), y^* \sim interp(\hat{y}, y^*)} [(\|\nabla D(y^*, x)\| - 1)^2]$$
(5)

We interpolate the real label y and the generated label \hat{y} with random weights between 0 and 1 to generate y^* , and apply the gradient penalty to y^* . This gradient penalty is essential to stabilize the training.

4. EXPERIMENT

We evaluated the proposed approach on the MS-COCO benchmark with 80 labels, and NUS-WIDE with 81 labels. As performance measures we used macro/micro precision, macro/micro recall, and macro/micro F1-measure. Note that macro/micro P/R/F1 scores are abbreviated as C/O-P/R/F1, respectively. Generally, C/O-F1 is more important [26].

In the experiments, we compared the proposed approach with WARP [6], CNN-RNN [18], order-free CNN-RNN with visual attention (Att-RNN) [14], and RLSD [27], which also models label dependency. To show that the proposed approach is general, we compared four kinds of generator architectures: CNNs, VGG-16 [19], Inception_v3 [20], Resnet-101 [21], and Resnet-152 [21].

4.1. Implementation details

The generator was pretrained on ImageNet with 1000 categories and fine-tuned on target datasets (MS-COCO with 80 labels and NUS-WIDE with 81 labels). We removed the output layer of the pretrained generator as the feature extractor network f_{ext} in the discriminator. Its parameters remained fixed during training. For the discriminator, the image features $z = f_{ext}(x)$ and label set were both linearly projected onto 256-dimension vectors, and then simply concatenated and fed into 8 fully connected layers of 512 dimensions with leaky relu activation functions. It is common to update the generator and discriminator with different numbers of steps. According to our experimental observation, as the pretrained generators were strong enough, we trained the discriminators 3 times per generator update. For each discriminator training iteration, we randomly sampled 3 batches for 3 kinds of inputs: matched pairs, generated pairs, and mismatched pairs, as indicated in Fig. 1(b); this is essential for training stability.

We based the whole optimization process on the Adam optimizer with a learning rate of 0.0001 for both the generator and the discriminator. The logistic loss weight of the generator α and the gradient penalty scale of the discriminator λ were both set to 10. In Gumbel sigmoid, the inverse temperature was set to 0.9 without annealing.

We followed [28] for our data augmentation strategies. Specifically, we first resized the image to 256×256 , after which we extracted five patches (four corner patches and the center patch) with a size from the set {256,224,192,168,128}. Finally, we resized the patches to 224×224 . For testing, we simply resized all images to 224×224 and conducted single-crop evaluation.

4.2. Experimental results

4.2.1. Microsoft COCO

Microsoft COCO (MS-COCO) is a large-scale dataset for object detection, segmentation, and image captioning. It has also been used for multi-label classification. It comprises a training set of 82,081 images, and a validation set of 40,137 images from 80 classes. Since the ground truth labels of the 2014 challenge are not available, we followed [14] and [18] in utilizing the validation set to evaluate our methods.

From Table 1, we see that the baselines of Inception_v3, Resnet-101, and Resnet-152 without the discriminator outperform other methods in C-F1 and O-F1 due to the advanced deep neural network structures. Moreover, all four models trained with WGAN-gp achieve higher C/O-F1 scores than the baselines without discriminator, which further suggests that modeling label correlation improves multi-label classification. The performance gain is less obvious when the baseline model is stronger. For example, there is a 4.4%/2.7% performance gain in C/O-F1 for VGG-16 but only 0.6%/0% for Resnet-152. For deeper networks such as Resnet-101 and Resnet-152, they may implicitly learn label dependencies due the huge number of hidden layers; this limits the usefulness of WGAN-gp. We also note that models which use WGAN-gp achieve higher recall but lower precision. We find that baseline models predict 2.09 labels per instance on average, whereas models which use WGAN-gp predict 2.61 labels, which is about 25% higher than the previous and results in higher F1 scores.

Methods	C-P	C-R	C-F1	O-P	O-R	O-F1
WARP	59.3	52.5	55.7	59.8	61.4	60.7
CNN-RNN	66.0	55.6	60.4	69.2	66.4	67.8
Att-RNN	71.6	54.8	62.1	74.2	62.2	67.7
RLSD	67.6	57.2	62.0	70.1	63.4	66.5
VGG-16	74.2	44.8	56.0	77.6	52.5	62.6
+ WGAN-gp	62.6	58.3	60.4	67.5	63.3	65.3
Inception_v3	76.4	52.8	62.4	80.0	58.8	67.8
+ WGAN-gp	70.5	58.2	63.8	73.2	63.8	68.2
Resnet-101	76.2	53.4	62.8	80.8	58.9	68.1
+ WGAN-gp	70.5	58.7	64.0	72.3	64.6	68.2
Resnet-152	76.6	53.9	63.3	80.6	59.6	68.6
+ WGAN-gp	71.4	57.9	63.9	73.6	64.2	68.6

Table 1. Multi-label classification results on MS-COCO with 80 labels. Results of WARP, CNN-RNN, and RLSD are reported with the top 3 labels.

This shows that with WGAN-gp, the classifier better models label dependencies and thus extracts more labels that are not detected by the original classifier. Examples of multi-label classification results are shown in Fig. 2.

4.2.2. NUS-WIDE

NUS-WIDE is a web image dataset which contains 269,648 images and associated tags from Flickr. The images are further manually annotated into 81 concepts. Following the experimental settings of WARP [6] and Att-RNN [14], we removed images without annotations and used 150,000 images for training and 59,347 images for testing. The results are reported in Table 2. MS-COCO and NUS-WIDE show similar trends.

Methods	C-P	C-R	C-F1	O-P	O-R	0-F1
WARP	31.7	35.6	33.5	48.6	60.5	53.9
CNN-RNN	40.5	30.4	34.7	49.9	61.7	55.2
Att-RNN	59.4	50.7	54.7	69.0	71.4	70.2
RLSD	44.4	49.6	46.9	54.4	67.6	60.3
VGG-16	53.3	24.9	33.9	73.9	59.6	66.0
+ WGAN-gp	51.6	34.3	41.2	68.8	67.3	68.1
Inception_v3	67.9	44.1	53.5	74.7	64.8	70.3
+ WGAN-gp	62.4	50.5	55.8	71.4	70.9	71.2
Resnet-101	67.0	44.0	53.1	76.3	65.0	70.2
+ WGAN-gp	59.6	51.8	55.4	68.9	72.8	70.8
Resnet-152	69.1	41.8	52.1	75.9	65.1	70.1
+ WGAN-gp	65.2	46.2	54.1	71.3	70.8	71.1

Table 2. Multi-label classification results on NUS-WIDE with 81 labels. Results of WARP, CNN-RNN, and RLSD are reported with the top 3 labels.

4.2.3. Ablation study

In this section, we show all the mechanisms described in Sect. 3. Table 3 reports the macro/micro F1 scores of different type of models. In this experiment, we used Resnet-101 as the generator and performed classification on MS-COCO.

Rows (a) and row (b) show the results of Resnet-101 with all the mechanisms and the baseline model trained with logistic loss, as reported in Table 1. The classifier models in rows (c), (d), and (e) have the same network architecture as in rows (b), but we remove some GAN training tricks. In rows (c) and (d), we do not perform negative sampling. That is, the third term of \mathcal{L}_D in Eq. 4 is removed, and the weight of the second term becomes 1. In row (d), we replace the conditional discriminator with an unconditional one. Therefore, the only discriminator input is y or \hat{y} . The discriminators here need only distinguish between real and generated label sets. The scores in row (c) and (d) are both less than the baseline. In row (e), as we directly feed the generator continuous output distribution \tilde{y} to the discriminator, Gumbel sigmoid is not needed. However, since the real data y is discrete, the generator must sharpen the distribution \tilde{y} , which reduces performance.

5. CONCLUSION

In this paper, we propose a novel framework for multi-class classification.² Inspired by GAN, the discriminator learns to model la-

Methods	C-F1	O-F1
(a): Resnet-101	62.8	68.1
(b): Resnet-101 + WGAN-gp	64.0	68.3
(c): (b) w/o negative sampling	62.2	67.5
(d): (b) w/o conditional discriminator	62.5	67.6
(e): (b) w/o Gumbel sigmoid	62.3	67.1

 Table 3. Macro/micro F1 scores with/without specific modules. Results are evaluated on MS-COCO with the Resnet-101 generator.

		(B)
Constant.	person, sports ball,	wine glass, cup, fork,
Ground truth:	baseball bat, baseball glove	knife, pizza, dining table
Resnet-101	person, baseball bat	fork, knife, pizza, dining table
Resnet-101 +	person, sports ball,	wine glass, cup, fork,
WGAN-gp	baseball bat, baseball glove	knife, pizza, dining table
	(C)	(D)
Ground truth:	chair, couch, bed, book	person, laptop
Resnet-101	couch, tv, book	person,laptop
Resnet-101 +	chair, couch, tv,	laptop, mouse,
WGAN-gp	laptop, book	keyboard

Fig. 2. Multi-label classification results from MS-COCO. With WGAN-gp, classifiers better predict smaller-sized image objects. For example, (A): Resnet-101 + WGAN-gp correctly predicts *base-ball glove* and *sports ball* based on observations of *person* and *base-ball bat*. However, in (D), it incorrectly relates *mouse* and *keyboard* to *laptop*.

bel dependency by discriminating real and generated label sets. To fool the discriminator, the classifier learns to generate label sets with dependencies close to real data. Extensive experiments and comparisons on two large-scale image classification benchmark datasets show that with this discriminator, F1 scores are improved across different classifier models. In future work, because the proposed idea is a general framework for multi-class classification, we will apply the proposed approach on multi-class classification tasks other than image classification.

6. REFERENCES

- Kai Kang, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang, "Object detection from video tubelets with convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 817–825.
- [2] Kai Kang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Bin Yang, Tong Xiao, Cong Zhang, Zhe Wang, Ruohui Wang, Xiaogang Wang, et al., "T-cnn: Tubelets with convolutional neural networks for object detection from videos," *IEEE Transactions* on Circuits and Systems for Video Technology, 2017.
- [3] Matthew R Boutell, Jiebo Luo, Xipeng Shen, and Christo-

²This work was financially supported by the Ministry of Science and Technology of Taiwan.

pher M Brown, "Learning multi-label scene classification," *Pattern recognition*, vol. 37, no. 9, pp. 1757–1771, 2004.

- [4] Pengcheng Yang, Xu Sun, Wei Li, Shuming Ma, Wei Wu, and Houfeng Wang, "Sgm: sequence generation model for multi-label classification," in *Proceedings of the 27th International Conference on Computational Linguistics*, 2018, pp. 3915–3926.
- [5] George Tzanetakis and Perry Cook, "Musical genre classification of audio signals," *IEEE Transactions on speech and audio processing*, vol. 10, no. 5, pp. 293–302, 2002.
- [6] Yunchao Gong, Yangqing Jia, Thomas Leung, Alexander Toshev, and Sergey Ioffe, "Deep convolutional ranking for multilabel image annotation," arXiv preprint arXiv:1312.4894, 2013.
- [7] Qiang Li, Maoying Qiao, Wei Bian, and Dacheng Tao, "Conditional graphical lasso for multi-label image classification," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 2977–2986.
- [8] Xin Li, Feipeng Zhao, and Yuhong Guo, "Multi-label image classification with a probabilistic label enhancement model.," in UAI, 2014, vol. 1, p. 3.
- [9] Yuhong Guo and Suicheng Gu, "Multi-label classification using conditional dependency networks," in *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, 2011, vol. 22, p. 1300.
- [10] Jiang Wang, Yi Yang, Junhua Mao, Zhiheng Huang, Chang Huang, and Wei Xu, "Cnn-rnn: A unified framework for multilabel image classification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2285–2294.
- [11] Grigorios Tsoumakas, Anastasios Dimou, Eleftherios Spyromitros, Vasileios Mezaris, Ioannis Kompatsiaris, and Ioannis Vlahavas, "Correlation-based pruning of stacked binary relevance models for multi-label learning," in *Proceedings of the 1st International Workshop on Learning from Multi-label Data*, 2009, pp. 101–116.
- [12] Chih-Kuan Yeh, Wei-Chieh Wu, Wei-Jen Ko, and Yu-Chiang Frank Wang, "Learning deep latent space for multilabel classification.," in AAAI, 2017, pp. 2838–2844.
- [13] Kush Bhatia, Himanshu Jain, Purushottam Kar, Manik Varma, and Prateek Jain, "Sparse local embeddings for extreme multilabel classification," in *Advances in Neural Information Processing Systems*, 2015, pp. 730–738.
- [14] Shang-Fu Chen, Yi-Chen Chen, Chih-Kuan Yeh, and Yu-Chiang Frank Wang, "Order-free rnn with visual attention for multi-label classification," *arXiv preprint arXiv:1707.05495*, 2017.
- [15] Feng Zhu, Hongsheng Li, Wanli Ouyang, Nenghai Yu, and Xiaogang Wang, "Learning spatial regularization with imagelevel supervisions for multi-label image classification," *arXiv* preprint arXiv:1702.05891, 2017.
- [16] Hao Yang, Joey Tianyi Zhou, Yu Zhang, Bin-Bin Gao, Jianxin Wu, and Jianfei Cai, "Exploit bounding box annotations for multi-label object recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 280–288.
- [17] Tianshui Chen, Zhouxia Wang, Guanbin Li, and Liang Lin, "Recurrent attentional reinforcement learning for multi-label image recognition," *arXiv preprint arXiv*:1712.07465, 2017.

- [18] Zhouxia Wang, Tianshui Chen, Guanbin Li, Ruijia Xu, and Liang Lin, "Multi-label image recognition by recurrently discovering attentional regions," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 464–472.
- [19] Karen Simonyan and Andrew Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv preprint arXiv:1409.1556, 2014.
- [20] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna, "Rethinking the inception architecture for computer vision," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2818–2826.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in *Proceed*ings of the IEEE conference on computer vision and pattern recognition, 2016, pp. 770–778.
- [22] Eric Jang, Shixiang Gu, and Ben Poole, "Categorical reparameterization with gumbel-softmax," *arXiv preprint arXiv:1611.01144*, 2016.
- [23] Zhuohan Li, Di He, Fei Tian, Wei Chen, Tao Qin, Liwei Wang, and Tie-Yan Liu, "Towards binary-valued gates for robust lstm training," arXiv preprint arXiv:1806.02988, 2018.
- [24] Scott Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee, "Generative adversarial text to image synthesis," arXiv preprint arXiv:1605.05396, 2016.
- [25] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville, "Improved training of wasserstein gans," in Advances in Neural Information Processing Systems, 2017, pp. 5767–5777.
- [26] Mohammad S Sorower, "A literature survey on algorithms for multi-label learning," .
- [27] Junjie Zhang, Qi Wu, Chunhua Shen, Jian Zhang, and Jianfeng Lu, "Multi-label image classification with regional latent semantic dependencies," *IEEE Transactions on Multimedia*, 2018.
- [28] Limin Wang, Yuanjun Xiong, Zhe Wang, and Yu Qiao, "Towards good practices for very deep two-stream convnets," *arXiv preprint arXiv:1507.02159*, 2015.
- [29] Xiangyang Xue, Wei Zhang, Jie Zhang, Bin Wu, Jianping Fan, and Yao Lu, "Correlative multi-label multi-instance image annotation," in *Computer Vision (ICCV), 2011 IEEE International Conference on.* IEEE, 2011, pp. 651–658.
- [30] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in Advances in neural information processing systems, 2012, pp. 1097–1105.
- [31] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition*, 2009. *CVPR 2009. IEEE Conference on.* Ieee, 2009, pp. 248–255.