

A probabilistic constrained clustering for transfer learning and image category discovery

Yen-Chang Hsu, Zhaoyang Lv
Georgia Institute of Technology, USA

Joel Schlosser, Phillip Odom, Zsolt Kira
Georgia Tech Research Institute, USA

Abstract

Neural network-based clustering has recently gained popularity, and in particular a constrained clustering formulation has been proposed to perform transfer learning and image category discovery using deep learning. The core idea is to formulate a clustering objective with pairwise constraints that can be used to train a deep clustering network; therefore the cluster assignments and their underlying feature representations are jointly optimized end-to-end. In this work, we provide a novel clustering formulation to address scalability issues of previous work in terms of optimizing deeper networks and larger amounts of categories. The proposed objective directly minimizes the negative log-likelihood of cluster assignment with respect to the pairwise constraints, has no hyper-parameters, and demonstrates improved scalability and performance on both supervised learning and unsupervised transfer learning.

1. Introduction

Unsupervised clustering algorithms have found a number of applications ranging from data analysis, visualization, and importantly transfer learning and unsupervised image category discovery [16, 14, 10]. Traditional methods that group data with a predefined metric, however, are very sensitive to the resulting choices or do not optimize the feature space jointly with clustering. If there is a complex nonlinear mapping between the input data and the desired semantic meaning of the clusters, the resulting clustering often cannot capture such meaning. In contrast, constrained clustering methods explicitly define the desired pairwise relationships as side information to the clustering algorithm. When paired with deep learning, the algorithm can learn the metric or feature representation to fit the desired relationships as well as optimize the assignments themselves.

Within this class of methods, the constraints often represent must-link/cannot-link or similar/dissimilar pairs (we use the latter in this paper). If a pair of data should be assigned to the same cluster, then it is a similar pair; otherwise it is dissimilar. The constraints could be regarded as an in-

terface to inject prior knowledge for clustering. Such information can be obtained via human labeling (supervised or semi-supervised setting), spatial or temporal relationships e.g. in videos (unsupervised setting) [15], or through a similarity metric learned on a different domain (transfer learning) [11]. Hence the constrained clustering formulation is a powerful concept that can bridge (semi-)supervised learning, unsupervised learning, and transfer learning.

A recent work [11] proposed a neural network-based constrained clustering objective in the form of a contrastive loss with KL-divergence. It demonstrated the transfer of learned pairwise similarity across tasks by clustering. Such a strategy can perform novel vision tasks, e.g. image category discovery with unlabeled examples of unseen categories. However, there are two critical challenges for this method: there is a degradation in the clustering performance when the chosen backbone network exceeds certain depth or when the number of clusters in the objective is larger (than a size of 10). These challenges prevent it from being applied to many real-world problems, which requires scalable solutions. Therefore, in this work we propose a new objective in the form of clustering likelihood to mitigate the mentioned issues.

2. The Constrained Clustering Likelihood

This section considers the constrained clustering problem. Suppose $X = \{x_1, \dots, x_n\}$ denotes a dataset of n instances, where each instance belongs to a set of unknown clusters $c \in \{1, 2, \dots, k\}$. The set of constraints $S = S^+ \cup S^-$ contains two types of pairwise relationships, where S^+ is the set of tuples (i, j) where (x_i, x_j) is a similar pair, and S^- contains the same for dissimilar pairs.

Similar to [9, 11], we define a neural network $f_\theta(x_i) = [p_{i,1}, p_{i,2}, \dots, p_{i,k}]^T$ that maps each sample x_i to a categorical distribution, i.e. $\sum_{c=1}^k p_{i,c} = 1$, which is the probability of data x_i belonging to cluster c . This is done by reinterpreting the softmax output of a neural network as outputting a probability distribution over cluster assignments, and a specialized loss function that utilizes the constraints is used to update θ , the set of parameters in the neural network.

We introduce y_i to represent x_i 's assigned cluster with probability $P(y_i = c|x_i) = p_{i,c}$. For a pair of data (x_i, x_j) , the probability of assigning both to the same cluster c is:

$$P(y_i = c, y_j = c|x_i, x_j) = P(y_i = c|x_i)P(y_j = c|x_j).$$

In the above formulation we impose an assumption that y_i and y_j are independent given their source data x_i and x_j . For the simplicity of the representation, we omit the notation of conditioning since all of the following probabilities are conditioned on X .

Now we describe the probability of (x_i, x_j) being a similar pair by marginalizing their joint distribution along the cluster index c :

$$P(y_i = y_j) = \sum_{c=1}^k P(y_i = c, y_j = c) = f_{\theta}(x_i)^T f_{\theta}(x_j). \quad (1)$$

Then the probability of (x_i, x_j) being a dissimilar pair becomes straightforward:

$$P(y_i \neq y_j) = 1 - P(y_i = y_j). \quad (2)$$

The constrained clustering likelihood \mathcal{L} now can be defined as the product of all the probabilities of similar and dissimilar pairs:

$$\mathcal{L}(\theta|S^+, S^-) = \prod_{(i,j) \in S^+} P(y_i = y_j) \prod_{(i,j) \in S^-} P(y_i \neq y_j). \quad (3)$$

For finding θ , we minimize the negative log-likelihood of equation (3), which is referred to as CCL and optimize this objective using stochastic gradient descent. For collecting the final cluster assignment, we only feed the data forward into f_{θ} again after the likelihood converged.

3. Results

We evaluate the proposed CCL and compare it with KL-divergence based constrained clustering loss (KCL) [9, 11] in two settings using image datasets.

In the *supervised* setting, we construct S from ground-truth category labels. The motivation is to evaluate the upper-bound performance and explore scalability regarding the number of clusters and depth of the neural network, ablating the effect of noise. Tables 1 and 2 demonstrate the superiority of CCL over KCL in terms of scalability.

In the *unsupervised transfer learning* (i.e. image category discovery) setting, we follow the procedure for transferring across tasks [11], which uses a learned similarity function as the constraint provider. The predicted constraints are noisy and therefore demonstrate the performance of our algorithm in a real-world setting. The scenario of unknown number of clusters is also evaluated by

Table 1: The accuracy on CIFAR10 with different loss functions and different neural network architectures. All training configurations, except the loss functions, are the same. The performance for Cross Entropy (CE) is the classification accuracy on test set and is regarded as the upper bound performance of supervised learning. For KCL and CCL, we give the clustering accuracy [19] on the test set.

Networks	CE	KCL	CCL (ours)
LeNet	84.4	83.6	83.9
VGG8	89.4	89.6	89.5
VGG11	90.0	85.6	90.4
VGG16	91.1	21.9	91.3
PReActResNet-18	92.6	81.7	92.7
PReActResNet-101	93.0	21.9	93.3

Table 2: The accuracy of supervised clustering with different number of clusters in the image datasets.

Dataset	#class	Network	CE	KCL	CCL
MNIST	10	LeNet	99.4	99.5	99.3
CIFAR10	10	VGG8	89.4	89.6	89.5
CIFAR100	100	VGG8	64.1	44.3	64.0

Table 3: Image category discovery on ImageNet. The values are the average of three random subsets in ImageNet₁₁₈. Each subset has 30 classes. The "ACC" (clustering accuracy) has $K = 30$ while the "ACC (100)" sets $K = 100$. All methods use the features (outputs of average pooling) from Resnet-18 pre-trained with ImageNet₈₈₂ classification. NMI stands for normalized mutual information.

Method	ACC	ACC(100)	NMI	NMI(100)
K-means	71.9%	34.5%	0.713	0.671
LSC	73.3%	33.5%	0.733	0.655
LPNMF	43.0%	21.8%	0.526	0.500
CCN-KCL	73.8%	65.2%	0.750	0.715
CCN-CCL	74.4%	71.5%	0.762	0.765

setting a large k (e.g. 100) in the network output. Table 3 shows the results of discovering 30 held-out categories in ImageNet [7]. In the case of unknown number of clusters, CCL (71.5%) demonstrates a significant advantage over KCL (65.2%). A similar trend also happens in the experiments of discovering characters in the Omniglot dataset (supplementary table 4) which includes comprehensive comparisons. Therefore we empirically conclude that CCL is a better clustering objective for these types of category discovery tasks.

Acknowledgments

This work was supported by the National Science Foundation and National Robotics Initiative (grant # IIS-1426998).

References

- [1] E. Amid, A. Gionis, and A. Ukkonen. Semi-supervised kernel metric learning using relative comparisons. *arXiv preprint arXiv:1612.00086*, 2016. 4
- [2] S. Anand, S. Mittal, O. Tuzel, and P. Meer. Semi-supervised kernel mean shift clustering. *IEEE transactions on pattern analysis and machine intelligence*, 36(6):1201–1215, 2014. 4
- [3] M. Bilenko, S. Basu, and R. J. Mooney. Integrating constraints and metric learning in semi-supervised clustering. In *Proceedings of the twenty-first international conference on Machine learning*, page 11. ACM, 2004. 4
- [4] D. Cai, X. He, X. Wang, H. Bao, and J. Han. Locality preserving nonnegative matrix factorization. In *IJCAI*, volume 9, pages 1010–1015, 2009. 4
- [5] X. Chen and D. Cai. Large scale spectral clustering with landmark-based representation. In *AAAI*, volume 5, page 14, 2011. 4
- [6] J. V. Davis, B. Kulis, P. Jain, S. Sra, and I. S. Dhillon. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pages 209–216. ACM, 2007. 4
- [7] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009. 2
- [8] K. He, X. Zhang, S. Ren, and J. Sun. Identity mappings in deep residual networks. In *European Conference on Computer Vision*, pages 630–645. Springer, 2016. 3
- [9] Y.-C. Hsu and Z. Kira. Neural network-based clustering using pairwise constraints. *ICLR workshop*, 2016. 1, 2
- [10] Y.-C. Hsu, Z. Lv, and Z. Kira. Deep image category discovery using a transferred similarity function. *arXiv preprint arXiv:1612.01253*, 2016. 1
- [11] Y.-C. Hsu, Z. Lv, and Z. Kira. Learning to cluster in order to transfer across domains and tasks. In *International Conference on Learning Representations (ICLR)*, 2018. 1, 2, 3, 4
- [12] 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. 3
- [13] J. MacQueen et al. Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 281–297. Oakland, CA, USA., 1967. 4
- [14] A. Ošep, P. Voigtlaender, J. Luiten, S. Breuers, and B. Leibe. Large-scale object discovery and detector adaptation from unlabeled video. *arXiv preprint arXiv:1712.08832*, 2017. 1
- [15] A. Ošep, P. Voigtlaender, J. Luiten, S. Breuers, and B. Leibe. Large-scale object discovery and detector adaptation from unlabeled video. *CoRR*, abs/1712.08832, 2017. 1
- [16] L. Shu, H. Xu, and B. Liu. Unseen class discovery in open-world classification. *arXiv preprint arXiv:1801.05609*, 2018. 1
- [17] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3
- [18] X. Wang, B. Qian, and I. Davidson. On constrained spectral clustering and its applications. *Data Mining and Knowledge Discovery*, pages 1–30, 2014. 4
- [19] Y. Yang, D. Xu, F. Nie, S. Yan, and Y. Zhuang. Image clustering using local discriminant models and global integration. *IEEE Transactions on Image Processing*, 19(10):2761–2773, 2010. 2

Supplementary

A. Setup for supervised clustering

Network Architecture: We use convolution neural networks with a varied number of layers: LeNet [12]; VGG [17] and the PreActResNet [8]. We also add a VGG8, which only has one convolution layer before each pooling layer, as the supplement between LeNet and VGG11. The number of output nodes in the last fully connected layer is the maximum number of clusters. We set it to the true number of categories for this section. Since the clustering objectives KCL and CCL both work on pairs of inputs, we insert a pairwise enumeration layer [11] between the network outputs and the loss function. Therefore the dense pairs in a mini-batch are all subject to the clustering loss.

Training configurations: All networks in this section are trained from scratch with randomly initialized weights. On the MNIST dataset, we use mini-batch size 100 with initial learning rate 0.1, which was dropped every 10 epochs by a factor of 0.1. We trained 30 epochs in total. On CIFAR10 we use the same setting except that the learning rate is dropped every 30 epochs and it is trained for 70 epochs in total. We use SGD for the optimization on MNIST and CIFAR10. For CIFAR100, the mini-batch size was 1000 and Adam was used as the optimizer with initial learning rate 0.001, which dropped every 70 epochs by a factor of 0.1. The training lasted 160 epochs.

B. Setup for image category discovery

We follow the unsupervised transfer learning procedures described in [11] to evaluate the clustering performance with noisy constraints. We use the original implementation and replace KCL by our CCL. We use the same training configurations, including the same similarity prediction function to make sure performance gain only by adopting CCL. We follow [11] to use CCN for Constrained Clustering Network, and use CCN-KCL and CCN-CCL to differentiate the two approaches in the tables. Noted that the clustering accuracy is directly calculated on the target dataset

Table 4: Unsupervised cross-task transfer from $Omniglot_{bg}$ to $Omniglot_{eval}$ for discovering the characters in $Omniglot_{eval}$. The performance is averaged across 20 alphabets which have 20 to 47 letters. The ACC and NMI without brackets have the number of clusters equal to ground-truth. The "(100)" means the algorithms use $k = 100$, *i.e.* one hundred outputs from the network. The characteristics of how each algorithm utilizes the pairwise constraints are marked in the "Constraints in" column, where metric stands for the metric learning of feature representation.

Method	Constraints in		ACC	ACC (100)	NMI	NMI (100)
	Metric	Clustering				
K-means [13]			21.7%	18.9%	0.353	0.464
LPNMF [4]			22.2%	16.3%	0.372	0.498
LSC [5]			23.6%	18.0%	0.376	0.500
ITML [6]	o		56.7%	47.2%	0.674	0.727
SKMS [2]	o		-	45.5%	-	0.693
SKKm [2]	o		62.4%	46.9%	0.770	0.781
SKLR [1]	o		66.9%	46.8%	0.791	0.760
CSP [18]		o	62.5%	65.4%	0.812	0.812
MPCK-means [3]	o	o	81.9%	53.9%	0.871	0.816
CCN-KCL [11]	o	o	82.4%	78.1%	0.889	0.874
CCN-CCL (ours)	o	o	83.3%	80.2%	0.897	0.893

instead of a holdout set since it is an unsupervised clustering setting.