

A Framework for Deep Constrained Clustering - Algorithms and Advances

Hongjing Zhang¹ (✉), Sugato Basu², and Ian Davidson¹

¹ Department of Computer Science, University of California, Davis, CA 95616, USA
hjzzhang@ucdavis.edu, davidson@cs.ucdavis.edu

² Google Research, Mountain View, CA 94043, USA
sugato@google.com

Abstract. The area of constrained clustering has been extensively explored by researchers and used by practitioners. Constrained clustering formulations exist for popular algorithms such as k-means, mixture models, and spectral clustering but have several limitations. A fundamental strength of deep learning is its flexibility, and here we explore a deep learning framework for constrained clustering and in particular explore how it can extend the field of constrained clustering. We show that our framework can not only handle standard together/apart constraints (without the well documented negative effects reported earlier) generated from labeled side information but more complex constraints generated from new types of side information such as continuous values and high-level domain knowledge. (Source code available at: http://github.com/blueocean92/deep_constrained_clustering)

Keywords: Constrained Clustering · Deep Learning · Semi-supervised Clustering · Reproducible Research

1 Introduction

Constrained clustering has a long history in machine learning with many standard algorithms being adapted to be constrained [3] including EM [2], K-Means [25] and spectral methods [26]. The addition of constraints generated from ground truth labels allows a semi-supervised setting to increase accuracy [25] when measured against the ground truth labeling.

However, there are several limitations in these methods and one purpose of this paper is to explore how deep learning can make advances to the field beyond what other methods have. In particular, we find that existing non-deep formulations of constrained clustering have the following limitations:

- *Limited Constraints and Side Information.* Constraints are limited to simple together/apart constraints typically generated from labels. In some domains, experts may more naturally give guidance at the cluster level or generate constraints from continuous side-information.

- *Negative Effect of Constraints.* For some algorithms though constraints improve performance when *averaged* over many constraint sets, *individual* constraint sets produce results worse than using no constraints [8]. As practitioners typically have one constraint set their use can be “hit or miss”.
- *Intractability and Scalability Issues.* Iterative algorithms that directly solve for clustering assignments run into problems of intractability [7]. Relaxed formulations (i.e. spectral methods [17, 26]) require solving a full rank eigen-decomposition problem which takes $O(n^3)$.
- *Assumption of Good Features.* A core requirement is that good features or similarity function for complex data is already created.

Since deep learning is naturally scalable and able to find useful representations we focus on the first and second challenges but experimentally explore the third and fourth. Though deep clustering with constraints has many potential benefits to overcome these limitations it is not without its challenges. Our major contributions in this paper are summarized as follows:

- We propose a deep constrained clustering formulation that cannot only encode standard together/apart constraints but new triplet constraints (which can be generated from continuous side information), instance difficulty constraints, and cluster level balancing constraints (see section 3).
- Deep constrained clustering overcomes a long term issue we reported in PKDD earlier [8] with constrained clustering of profound practical implications: overcoming the negative effects of individual constraint sets.
- We show how the benefits of deep learning such as scalability and end-to-end learning translate to our deep constrained clustering formulation. We achieve better clustering results than traditional constrained clustering methods (with features generated from an auto-encoder) on challenging datasets (see Table 2).

Our paper is organized as follows. First, we introduce the related work in section 2. We then propose four forms of constraints in section 3 and introduce how to train the clustering network with these constraints in section 4. Then we compare our approach to previous baselines and demonstrate the effectiveness of new types of constraints in section 5. Finally, we discuss future work and conclude in section 6.

2 Related Work

Constrained Clustering. Constrained clustering is an important area and there is a large body of work that shows how *side information* can improve the clustering performance [24, 25, 28, 4, 26]. Here the side information is typically labeled data which is used to generate *pairwise* together/apart constraints used to partially reveal the ground truth clustering to help the clustering algorithm. Such constraints are easy to encode in matrices and enforce in procedural algorithms though not with its challenges. In particular, we showed [8] performance improves

with larger constraint sets when **averaged** over many constraint sets generated from the ground truth labeling. However, for a significant fraction (just not the majority) of these constraint sets performance is *worse* than using no constraint set. We recreated some of these results in Table 2.

Moreover, side information can exist in different forms beyond labels (i.e. continuous data), and domain experts can provide guidance beyond pairwise constraints. Some work in the supervised classification setting [14, 21, 20, 10] seek alternatives such as relative/triplet guidance, but to our knowledge, such information has not been explored in the non-hierarchical clustering setting. Complex constraints for hierarchical clustering have been explored [1, 5] but these are tightly limited to the hierarchical structure (i.e., x must be joined with y before z) and not directly translated to non-hierarchical (partitional) clustering.

Deep Clustering. Motivated by the success of deep neural networks in supervised learning, unsupervised deep learning approaches are now being explored [27, 13, 30, 11, 22]. There are approaches [30, 22] which learn an encoding that is suitable for a clustering objective first and then applied an external clustering method. Our work builds upon the most direct setting [27, 11] which encodes one self-training objective and finds the clustering allocations for all instances within one neural network.

Deep Clustering with Pairwise Constraints. Most recently, the semi-supervised clustering networks with pairwise constraints have been explored: [12] uses pairwise constraints to enforce small divergence between similar pairs while increasing the divergence between dissimilar pairs assignment probability distributions. However, this approach did not leverage the unlabeled data, hence requires lot's of labeled data to achieve good results. Fogel et al. proposed an unsupervised clustering network [9] by self-generating pairwise constraints from mutual KNN graph and extends it to semi-supervised clustering by using labeled connections queried from the human. However, this method cannot make out-of-sample predictions and requires user-defined parameters for generating constraints from mutual KNN graph.

3 Deep Constrained Clustering Framework

Here we outline our proposed framework for deep constrained clustering. Our method of adding constraints to and training deep learning can be used for most deep clustering method (so long as the network has a k unit output indicating the degree of cluster membership) and here we choose the popular deep embedded clustering method (DEC [27]). We sketch this method first for completeness.

3.1 Deep Embedded Clustering

We choose to apply our constraints formulation to the deep embedded clustering method DEC [27] which starts with pre-training an autoencoder ($x_i = g(f(x_i))$) but then removes the decoder. The remaining encoder ($z_i = f(x_i)$) is then fine-tuned by optimizing an objective which takes first z_i and converts it to a soft

allocation vector of length k which we term $q_{i,j}$ indicating the degree of belief instance i belongs to cluster j . Then q is self-trained on to produce p a unimodal “hard” allocation vector which allocates the instance to primarily only one cluster. We now overview each step.

Conversion of z to Soft Cluster Allocation Vector q . Here DEC takes the similarity between an embedded point z_i and the cluster centroid u_j measured by Student’s t -distribution [18]. Note that v is a constant as $v = 1$ and q_{ij} is a soft assignment:

$$q_{ij} = \frac{(1 + \|z_i - \mu_j\|^2/v)^{-\frac{v+1}{2}}}{\sum_{j'} (1 + \|z_i - \mu_{j'}\|^2/v)^{-\frac{v+1}{2}}} \quad (1)$$

Conversion of Q To Hard Cluster Assignments P . The above normalized similarities between embedded points and centroids can be considered as soft cluster assignments Q . However, we desire a target distribution P that better resembles a hard allocation vector, p_{ij} is defined as:

$$p_{ij} = \frac{q_{ij}^2 / \sum_i q_{ij}}{\sum_{j'} (q_{ij'}^2 / \sum_i q_{ij'})} \quad (2)$$

Loss Function. Then the algorithm’s loss function is to minimize the distance between P and Q as follows. Note this is a form of self-training as we are trying to teach the network to produce unimodal cluster allocation vectors.

$$\ell_C = KL(P||Q) = \sum_i \sum_j p_{ij} \log \frac{p_{ij}}{q_{ij}} \quad (3)$$

The DEC method requires the initial centroids given (μ) to calculate Q are “representative”. The initial centroids are set using k-means clustering. However, there is no guarantee that the clustering results over an auto-encoders embedding yield a good clustering. We believe that constraints can help overcome this issue which we test later.

3.2 Different Types of Constraints

To enhance the clustering performance and allow for more types of interactions between human and clustering models we propose four types of guidance which are pairwise constraints, instance difficulty constraints, triplet constraints, and cardinality and give examples of each. As traditional constrained clustering methods put constraints on the final clustering assignments, our proposed approach constrains the q vector which is the soft assignment. A core challenge when adding constraints is to allow the resultant loss function to be differentiable so we can derive back propagation updates.

Pairwise Constraints Pairwise constraints (must-link and cannot-link) are well studied [3] and we showed they are capable of defining any ground truth set

partitions [7]. Here we show how these pairwise constraints can be added to a deep learning algorithm. We encode the loss for must-link constraints set ML as:

$$\ell_{ML} = - \sum_{(a,b) \in ML} \log \sum_j q_{aj} * q_{bj} \quad (4)$$

Similarly loss for cannot-link constraints set CL is:

$$\ell_{CL} = - \sum_{(a,b) \in CL} \log (1 - \sum_j q_{aj} * q_{bj}) \quad (5)$$

Intuitively speaking, the must-link loss prefers instances with same soft assignments and the cannot-link loss prefers the opposite cases.

Instance Difficulty Constraints A challenge with self-learning in deep learning is that if the initial centroids are incorrect, the self-training can lead to poor results. Here we use constraints to overcome this by allowing the user to specify which instances are easier to cluster (i.e., they belong strongly to only one cluster) and by ignoring difficult instances (i.e., those that belong to multiple clusters strongly).

We encode user supervision with an $n \times 1$ constraint vector M . Let $M_i \in [-1, 1]$ be an instance difficulty indicator, $M_i > 0$ means the instance i is easy to cluster, $M_i = 0$ means no difficulty information is provided and $M_i < 0$ means instance i is hard to cluster. The loss function is formulated as:

$$\ell_I = \sum_{t \in \{M_t < 0\}} M_t \sum_j q_{tj}^2 - \sum_{s \in \{M_s > 0\}} M_s \sum_j q_{sj}^2 \quad (6)$$

The instance difficulty loss function aims to encourage the easier instances to have sparse clustering assignments but prevents the difficult instances having sparse clustering assignments. The absolute value of M_i indicates the degree of confidence in difficulty estimation. This loss will help the model training process converge faster on easier instances and increase our model's robustness towards difficult instances.

Triplet Constraints Although pairwise constraints are capable of defining any ground truth set partitions from labeled data [7], in many domains no labeled side information exists or strong pairwise guidance is not available. Thus we seek triplet constraints, which are weaker constraints that indicate the relationship within a triple of instances. Given an anchor instance a , positive instance p and negative instance n we say that instance a is more similar to p than to n . The loss function for all triplets $(a, p, n) \in T$ can be represented as:

$$\ell_T = \sum_{(a,p,n) \in T} \max(d(q_a, q_n) - d(q_a, q_p) + \theta, 0) \quad (7)$$

where $d(q_a, q_b) = \sum_j q_{aj} * q_{bj}$ and $\theta > 0$. The larger value of $d(q_a, q_b)$ represents larger similarity between a and b . The variable θ controls the gap distance between

positive and negative instances. ℓ_T works by pushing the positive instance’s assignment closer to anchor’s assignment and preventing negative instance’s assignment being closer to anchor’s assignment.

Global Size Constraints Experts may more naturally give guidance at a cluster level. Here we explore clustering size constraints, which means each cluster should be approximately the same size. Denote the total number of clusters as k , total training instances number as n , the global size constraints loss function is:

$$\ell_G = \sum_{c \in \{1, \dots, k\}} \left(\sum_{i=1}^n q_{ic} / n - \frac{1}{k} \right)^2 \quad (8)$$

Our global constraints loss function works by minimizing the distance between the expected cluster size and the actual cluster size. The actual cluster size is calculated by averaging the soft-assignments. To guarantee the effectiveness of global size constraints, we need to assume that during our mini-batch training the batch size should be large enough to calculate the cluster sizes. A similar loss function can be used (see section 3.4) to enforce other cardinality constraints on the cluster composition such as upper and lower bounds on the number of people with a certain property.

3.3 Preventing Trivial Solution

In our framework the proposed must-link constraints we mentioned before can lead to trivial solution that all the instances are mapped to the same cluster. Previous deep clustering method [30] have also met this problem. To mitigate this problem, we combine the reconstruction loss with the must-link loss to learn together. Denote the encoding network as $f(x)$ and decoding network as $g(x)$, the reconstruction loss for instance x_i is:

$$\ell_R = \ell(g(f(x_i)), x_i) \quad (9)$$

where ℓ is the least-square loss: $\ell(x, y) = \|x - y\|^2$.

3.4 Extensions to High-level Domain Knowledge-Based Constraints

Although most of our proposed constraints are generated based on instance labels or comparisons. Our framework can be extended to high-level domain knowledge-based constraints with minor modifications.

Cardinality Constraints. For example, cardinality constraints [6] allow expressing requirements on the number of instances that satisfy some conditions in each cluster. Assume we have n people and want to split them into k dinner party groups. An example cardinality constraint is to enforce each party should have the same number of males and females. We split the n people into two

groups as M (males) and F (females) in which $|M| + |F| = n$ and $M \cap F = \emptyset$. Then the cardinality constraints can be formulated as:

$$\ell_{Cardinality} = \sum_{c \in \{1, \dots, k\}} \left(\sum_{i \in M} q_{ic}/n - \sum_{j \in F} q_{jc}/n \right)^2 \quad (10)$$

For upper-bound and lower-bound based cardinality constraints [6], we use the same setting as previously described, now the constraint changes as for each party group we need the number of males to range from L to U . Then we can formulate it as:

$$\ell_{CardinalityBound} = \sum_{c \in \{1, \dots, k\}} \left(\min(0, \sum_{i \in M} q_{ic} - L)^2 + \max(0, \sum_{i \in M} q_{ic} - U)^2 \right) \quad (11)$$

Logical Combinations of Constraints. Apart from cardinality constraints, complex logic constraints can also be used to enhance the expressivity power of representing knowledge. For example, if two instances x_1 and x_2 are in the same cluster then instances x_3 and x_4 must be in different clusters. This can be achieved in our framework as we can dynamically add cannot-link constraint $CL(x_3, x_4)$ once we check the soft assignment q of x_1 and x_2 .

Consider a horn form constraint like $r \wedge s \wedge t \rightarrow u$. Denote $r = ML(x_1, x_2)$, $s = ML(x_3, x_4)$, $t = ML(x_5, x_6)$ and $u = CL(x_7, x_8)$. By forward passing the instances within r, s, t to our deep constrained clustering model, we can get the soft assignment values of these instances. By checking the satisfying results based on $r \wedge s \wedge t$, we can decide whether to enforce cannot-link loss $CL(x_7, x_8)$.

4 Putting It All Together - Efficient Training Strategy

Our training strategy consists of two training branches and effectively has two ways of creating mini-batches for training. For instance-difficulty or global-size constraints, we treat their loss functions as additive losses so that no extra branch needs to be created. For pairwise or triplet constraints we build another output branch for them and train the whole network in an alternative way.

Loss Branch for Instance Constraints. In deep learning it is common to add loss functions defined over the same output units. In the Improved DEC method [11] the clustering loss ℓ_C and reconstruction loss ℓ_R were added together. To this we add the instance difficulty loss ℓ_I . This effectively adds guidance to speed up training convergence by identifying “easy” instances and increase the model’s robustness by ignoring “difficult” instances. Similarly we treat the global size constraints loss ℓ_G as an additional additive loss. All instances whether or not they are part of triplet or pairwise constraints are trained through this branch and the mini-batches are created randomly.

Loss Branch For Complex Constraints. Our framework uses more complex loss functions as they define constraints on pairs and even triples of instances. Thus we create another loss branch that contains pairwise loss ℓ_P or triplet loss ℓ_T to help the network tune the embedding which satisfy these stronger constraints.

For each constraint *type* we create a mini-batch consisting of only those instances having that type of constraint. For each *example* of a constraint type, we feed the constrained instances through the network, calculate the loss, calculate the change in weights but do not adjust the weights. We sum the weight adjustments for all constraint examples in the mini-batch and then adjust the weights. Hence our method is an example of batch weight updating as is standard in DL for stability reasons. The whole training procedure is summarized in Algorithm 1.

Algorithm 1 Deep Constrained Clustering Framework

Input: X : data, m : maximum epochs, k : number of clusters, N : total number of batches and N_C : total number of constraints batches.

Output: latent embeddings Z , cluster assignment S .

Train the stacked denosing autoencoder to obtain Z

Initialize centroids μ via k-means on embedding Z .

for $epoch = 1$ **to** m **do**

for $batch = 1$ **to** N **do**

 Calculate ℓ_C via Eqn (3), ℓ_R via Eqn (9).

 Calculate ℓ_I via Eqn (6) or ℓ_G via Eqn (8).

 Calculate total loss as $\ell_C + \ell_R + \{\ell_I|\ell_G\}$.

 Update network parameters based on total loss.

end for

for $batch = 1$ **to** N_C **do**

 Calculate ℓ_P via Eqn (4, 5) or ℓ_T via Eqn (7).

 Update network parameters based on $\{\ell_P|\ell_T\}$.

end for

 Forward pass to compute Z and $S_i = \operatorname{argmax}_j q_{ij}$.

end for

5 Experiments

All data and code used to perform these experiments are available online (http://github.com/blueocean92/deep_constrained_clustering) to help with reproducibility. In our experiments we aim to address the following questions:

- How does our end-to-end deep clustering approach using traditional pairwise constraints compare with traditional constrained clustering methods? The latter is given the same auto-encoding representation Z used to initialize our method.
- Are the new types of constraints we create for deep clustering method useful in practice?
- Is our end-to-end deep constrained clustering method more robust to the well known negative effects of constraints we published earlier [8]?

5.1 Datasets

To study the performance and generality of different algorithms, we evaluate the proposed method on two image datasets and one text dataset:

MNIST: Consists of 70000 handwritten digits of 28-by-28 pixel size. The digits are centered and size-normalized in our experiments [15].

FASHION-MNIST: A Zalando’s article images-consisting of a training set of 60000 examples and a test set of 10000 examples. Each example is a 28-by-28 grayscale image, associated with a label from 10 classes.

REUTERS-10K: This dataset contains English news stories labeled with a category tree [16]. To be comparable with the previous baselines, we used 4 root categories: `corporate/industrial`, `government/social`, `markets` and `economics` as labels and excluded all documents with multiple labels. We randomly sampled a subset of 10000 examples and computed TF-IDF features on the 2000 most common words.

5.2 Implementation Details

Basic Deep Clustering Implementation. To be comparable with deep clustering baselines, we set the encoder network as a fully connected multilayer perceptron with dimensions $d - 500 - 500 - 2000 - 10$ for all datasets, where d is the dimension of input data(features). The decoder network is a mirror of the encoder. All the internal layers are activated by the ReLU [19] nonlinearity function. For a fair comparison with baseline methods, we used the same greedy layer-wise pre-training strategy to calculate the auto-encoders embedding. To initialize clustering centroids, we run k-means with 20 restarts and select the best solution. We choose Adam optimizer with an initial learning rate of 0.001 for all the experiments. We adopt standard metrics for evaluating clustering performance which measure how close the clustering found is to the ground truth result. Specifically, we employ the following two metrics: normalized mutual information(**NMI**)[23, 29] and clustering accuracy(**Acc**)[29]. In our baseline comparisons we use IDEC [11], a non-constrained improved version of DEC published recently.

Pairwise Constraints Experiments. We randomly select pairs of instances and generate the corresponding pairwise constraints between them. To ensure transitivity we calculate the transitive closure over all must-linked instances and then generate entailed constraints from the cannot-link constraints [7]. Since our loss function for must-link constraints is combined with reconstruction loss, we use grid search and set the penalty weight for must-link as 0.1.

Instance Difficulty Constraints Experiments. To simulate human-guided instance difficulty constraints, we use k-means as a base learner and mark all the incorrectly clustered instances as difficult with confidence 0.1, we also mark the correctly classified instances as easy instances with confidence 1. In Figure 1 we give some example difficulty constraints found using this method.

Triplet Constraints Experiments. Triplet constraints can state that instance i is more similar to instance j than instance k . To simulate human guidance




Fig. 1. Example of instance difficulty constraints. Top row shows the “easy” instances and second row shows the “difficult” instances.

on triplet constraints, we randomly select n instances as anchors (i), for each anchor we randomly select two instances (j and k) based on the similarity between the anchor. The similarity is calculated as the euclidian distance d between two instances pre-trained embedding. The pre-trained embedding is extracted from our deep clustering network trained with 100000 pairwise constraints. Figure 2 shows the generated triplets constraints. Through grid search we set the triplet loss margin $\theta = 0.1$.



Fig. 2. Examples of the generated triplet constraints for MNIST and Fashion. The three rows for each plot shows the anchor instances, positive instances and negative instances correspondingly.

Global Size Constraints Experiments. We apply global size constraints to MNIST and Fashion datasets since they satisfy the balanced size assumptions. The total number of clusters is set to 10.

5.3 Experimental Results

Experiments on Instance Difficulty. In Table 1, we report the average test performance of deep clustering framework without any constraints in the left. In comparison, we report the average test performance of deep clustering framework with instance difficulty constraints in the right and we find the model learned with instance difficulty constraints outperforms the baseline method in all datasets. This is to be expected as we have given the algorithm more information than the baseline method, but it demonstrates our method can make good use of this extra information. What is unexpected is the effectiveness of speeding up the learning process and will be the focus of future work.

	MNIST	Fashion	Reuters	MNIST	Fashion	Reuters
Acc(%)	88.29 ± 0.05	58.74 ± 0.08	75.20 ± 0.07	91.02 ± 0.34	62.17 ± 0.06	78.01 ± 0.13
NMI(%)	86.12 ± 0.09	63.27 ± 0.11	54.16 ± 0.11	88.08 ± 0.14	64.95 ± 0.04	56.02 ± 0.21
Epoch	87.60 ± 12.53	77.20 ± 11.28	12.90 ± 2.03	29.70 ± 4.25	47.60 ± 6.98	9.50 ± 1.80

Table 1. Left table shows baseline results for Improved DEC [11] averaged over 20 trials. Right table lists experiments using instance difficulty constraints (mean ± std) averaged over 20 trials.

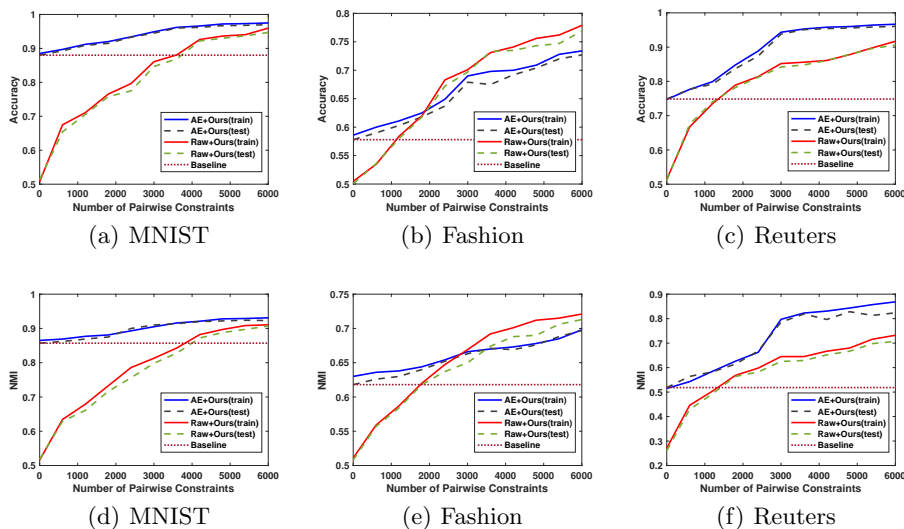


Fig. 3. Clustering accuracy and NMI on training test sets for different number of pairwise constraints. AE means an autoencoder was used to seed our method. The horizontal maroon colored baseline shows the IDEC’s [11] test set performance.

Experiments on pairwise constraints We randomly generate 6000 pairs of constraints which are a small fractions of possible pairwise constraints for MNIST (0.0002%), Fashion (0.0002%) and Reuters (0.006%).

Recall the DEC method is initialized with auto-encoder features. To better understand the contribution of pairwise constraints, we have tested our method with both auto-encoders features and raw data. As can be seen from Figure 3: the clustering performance improves consistently as the number of constraints increases in both settings. Moreover, with just 6000 pairwise constraints the performance on Reuters and MNIST increased significantly especially for the setup with raw data. We also notice that learning with raw data in Fashion achieves a better result than using autoencoder’s features. This shows that the autoencoder’s features may not always be suitable for DEC’s clustering objective.

Overall our results show pairwise constraints can help reshape the representation and improve the clustering results.

We also compare the results with recent work [12]: our approach(autoencoders features) outperforms the best clustering accuracy reported for MNIST by a margin of 16.08%, 2.16% and 0.13% respectively for 6, 60 and 600 samples/class. Unfortunately, we can't make a comparison with Fogel's algorithm [9] due to an issue in their code repository.

	Flexible CSP*	COP-KMeans	MPCKMeans	Ours
MNIST Acc	0.628 ± 0.07	0.816 ± 0.06	0.846 ± 0.04	0.963 ± 0.01
MNIST NMI	0.587 ± 0.06	0.773 ± 0.02	0.808 ± 0.04	0.918 ± 0.01
Negative Ratio	19%	45%	11%	0 %
Fashion Acc	0.417 ± 0.05	0.548 ± 0.04	0.589 ± 0.05	0.681 ± 0.03
Fashion NMI	0.462 ± 0.03	0.589 ± 0.02	0.613 ± 0.04	0.667 ± 0.02
Negative Ratio	23%	27%	37%	6 %
Reuters Acc	0.554 ± 0.07	0.712 ± 0.0424	0.763 ± 0.05	0.950 ± 0.02
Reuters NMI	0.410 ± 0.05	0.478 ± 0.0346	0.544 ± 0.04	0.815 ± 0.02
Negative Ratio	28%	73%	80%	0 %

Table 2. Pairwise constrained clustering performance (mean ± std) averaged over 100 constraints sets. Due to the scalability issues we apply flexible CSP with downsampled data(3000 instances and 180 constraints). Negative ratio is the fraction of times using constraints resulted in poorer results than not using constraints. See Figure 4 and text for an explanation why our method performs well.

Negative Effects of Constraints. Our earlier work [8] showed that for traditional constrained clustering algorithms, that the addition of constraints *on average* helps clustering but many individual constraint sets can hurt performance in that performance is worse than using **no** constraints. Here we recreate these results even when these classic methods use auto-encoded representations. In Table 2, we report the average performance with 3600 randomly generated pairwise constraints. For each dataset, we randomly generated 100 sets of constraints to test the negative effects of constraints[8]. In each run, we fixed the random seed and the initial centroids for k-means based methods, for each method we compare its performance between constrained version to unconstrained version. We calculate the negative ratio which is the fraction of times that unconstrained version produced better results than the constrained version. As can be seen from the table, our proposed method achieves significant improvements than traditional non-deep constrained clustering algorithms [25, 4, 26].

To understand why our method was robust to variations in constraint sets we visualized the embeddings learnt. Figure 4 shows the embedded representation of a random subset of instances and its corresponding pairwise constraints using t-SNE and the learned embedding z . Based on Figure 4, we can see the autoencoders

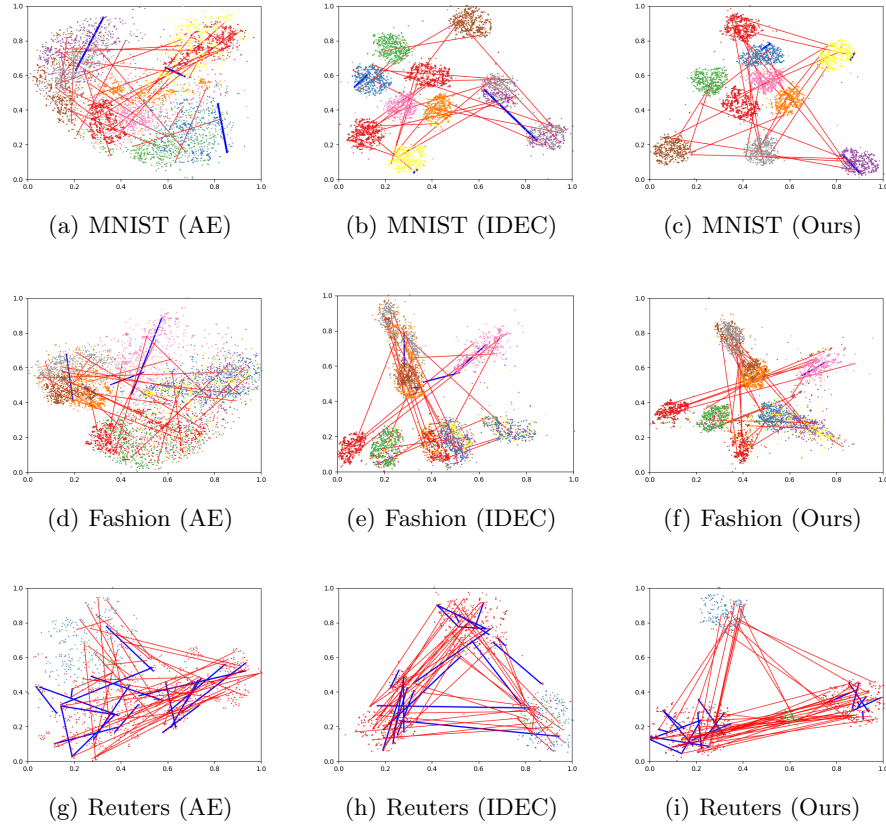


Fig. 4. We visualize (using t-SNE) the latent representation for a subset of instances and pairwise constraints, we visualize the same instances and constraints for each row. The red lines are cannot-links and blue lines are must-links.

embedding is noisy and lot's of constraints are inconsistent based on our earlier definition [8]. Further, we visualize the IDEC's latent embedding and find out the clusters are better separated. However, the inconsistent constraints still exist (blue lines across different clusters and redlines within a cluster); these constraints tend to have negative effects on traditional constrained clustering methods. Finally, for our method's results we can see the clusters are well separated, the must-links are well satisfied (blue lines are within the same cluster) and cannot-links are well satisfied (red lines are across different clusters). Hence we can conclude that end-to-end-learning can address these negative effects of constraints by simultaneously learning a representation that is consistent with the constraints and clustering the data. This result has profound practical significance as practitioners typically only have one constraint set to work with.

Experiments on triplet constraints We experimented on MNIST and FASHION datasets. Figure 2 visualizes example triplet constraints (based on embedding similarity), note the positive instances are closer to anchors than negative instances. In Figure 5, we show the clustering Acc/NMI improves consistently as the number of constraints increasing. Comparing with Figure 3 we can find the pairwise constraints can bring slightly better improvements, that’s because our triplets constraints are generated from a continuous domain and there is no exact together/apart information encoded in the constraints. Triplet constraints can be seen as a weaker but more general type of constraints.

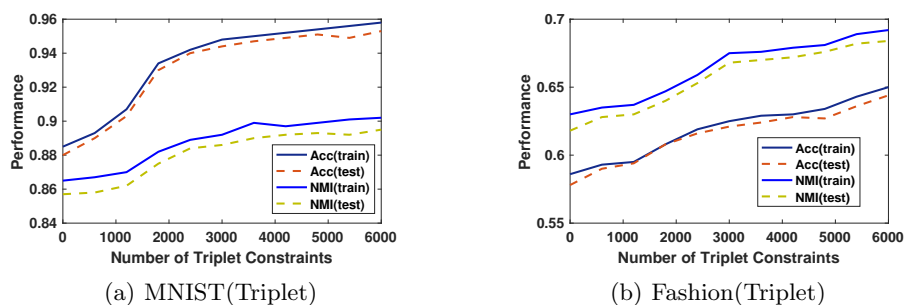


Fig. 5. Evaluation of the effectiveness of triplet constraints in terms of Acc/NMI.

Experiments on global size constraints To test the effectiveness of our proposed global size constraints, we have experimented on MNIST and Fashion training set since they both have balanced cluster sizes (see Figure 6). Note that the ideal size for each cluster is 6000 (each data set has 10 classes), we can see that blue bars are more evenly distributed and closer to the ideal size.

We also evaluate the clustering performance with global constraints on MNIST (Acc:0.91, NMI:0.86) and Fashion (Acc:0.57, NMI:0.59). Comparing to the baselines in table 1, interestingly, we find the performance improved slightly on MNIST but dropped slightly on Fashion.

6 Conclusion and Future Work

The area of constrained partitional clustering has a long history and is widely used. Constrained partitional clustering typically is mostly limited to simple pairwise together and apart constraints. In this paper, we show that deep clustering can be extended to a variety of fundamentally different constraint types including instance-level (specifying hardness), cluster level (specifying cluster sizes) and triplet-level.




Fig. 6. Evaluation of the global size constraints. This plot shows each cluster’s size before/after adding global size constraints.

Our deep learning formulation was shown to advance the general field of constrained clustering in several ways. Firstly, it achieves better experimental performance than well-known k-means, mixture-model and spectral constrained clustering in both an academic setting and a practical setting (see Table 2).

Importantly, our approach does not suffer from the negative effects of constraints [8] as it learns a representation that simultaneously satisfies the constraints and finds a good clustering. This result is quite useful as a practitioner typically has just one constraint set and our method is far more likely to perform better than using no constraints.

Most importantly, we were able to show that our method achieves all of the above but still retains the benefits of deep learning such as scalability, out-of-sample predictions and end-to-end learning. We found that even though standard non-deep learning methods were given the same representations of the data used to initialize our methods the deep constrained clustering was able to adapt these representations even further. Future work will explore new types of constraints, using multiple constraints at once and extensions to the clustering setting.

Acknowledgements

We acknowledge support for this work from a Google Gift entitled: “Combining Symbolic Reasoning and Deep Learning”.

References

1. Bade, K., Nürnberger, A.: Creating a cluster hierarchy under constraints of a partially known hierarchy. In: SIAM, 2008.
2. Basu, S., Bilenko, M., Mooney, R.J.: A probabilistic framework for semi-supervised clustering. In: KDD, 2004.
3. Basu, S., Davidson, I., Wagstaff, K.: Constrained Clustering: Advances in algorithms, theory, and applications. CRC Press (2008)

4. Bilenko, M., Basu, S., Mooney, R.J.: Integrating constraints and metric learning in semi-supervised clustering. In: ICML, 2004.
5. Chatziafratis, V., Niazadeh, R., Charikar, M.: Hierarchical clustering with structural constraints. arXiv preprint arXiv:1805.09476 (2018)
6. Dao, T.B.H., Vrain, C., Duong, K.C., Davidson, I.: A framework for actionable clustering using constraint programming. In: ECAI 2016.
7. Davidson, I., Ravi, S.: Intractability and clustering with constraints. In: ICML, 2007.
8. Davidson, I., Wagstaff, K.L., Basu, S.: Measuring constraint-set utility for partitional clustering algorithms. In: PKDD, 2006.
9. Fogel, S., Averbuch-Elor, H., Goldberger, J., Cohen-Or, D.: Clustering-driven deep embedding with pairwise constraints. arXiv preprint arXiv:1803.08457 (2018)
10. Gress, A., Davidson, I.: Probabilistic formulations of regression with mixed guidance. In: ICDM, 2016.
11. Guo, X., Gao, L., Liu, X., Yin, J.: Improved deep embedded clustering with local structure preservation. In: IJCAI, 2017.
12. Hsu, Y.C., Kira, Z.: Neural network-based clustering using pairwise constraints. arXiv preprint arXiv:1511.06321 (2015)
13. Jiang, Z., Zheng, Y., Tan, H., Tang, B., Zhou, H.: Variational deep embedding: An unsupervised and generative approach to clustering. arXiv preprint arXiv:1611.05148 (2016)
14. Joachims, T.: Optimizing search engines using clickthrough data. In: KDD, 2002.
15. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* **86**(11), 2278–2324 (1998)
16. Lewis, D.D., Yang, Y., Rose, T.G., Li, F.: Rcv1: A new benchmark collection for text categorization research. In: JMLR, 2004.
17. Lu, Z., Carreira-Perpinan, M.A.: Constrained spectral clustering through affinity propagation. In: CVPR, 2008.
18. Maaten, L.v.d., Hinton, G.: Visualizing data using t-sne. *Journal of machine learning research* **9**(Nov), 2579–2605 (2008)
19. Nair, V., Hinton, G.E.: Rectified linear units improve restricted boltzmann machines. In: ICML, 2010.
20. Schroff, F., Kalenichenko, D., Philbin, J.: Facenet: A unified embedding for face recognition and clustering. In: CVPR, 2015.
21. Schultz, M., Joachims, T.: Learning a distance metric from relative comparisons. In: NIPS, 2004.
22. Shaham, U., Stanton, K., Li, H., Nadler, B., Basri, R., Kluger, Y.: Spectralnet: Spectral clustering using deep neural networks. arXiv preprint arXiv:1801.01587 (2018)
23. Strehl, A., Ghosh, J., Mooney, R.: Impact of similarity measures on web-page clustering. In: Workshop on artificial intelligence for web search (AAAI 2000). vol. 58, p. 64 (2000)
24. Wagstaff, K., Cardie, C.: Clustering with instance-level constraints. In: AAAI, 2000.
25. Wagstaff, K., Cardie, C., Rogers, S., Schrödl, S., et al.: Constrained k-means clustering with background knowledge. In: ICML, 2001.
26. Wang, X., Davidson, I.: Flexible constrained spectral clustering. In: KDD, 2010.
27. Xie, J., Girshick, R., Farhadi, A.: Unsupervised deep embedding for clustering analysis. In: ICML, 2016.
28. Xing, E.P., Jordan, M.I., Russell, S.J., Ng, A.Y.: Distance metric learning with application to clustering with side-information. In: NIPS, 2003.
29. Xu, W., Liu, X., Gong, Y.: Document clustering based on non-negative matrix factorization. In: SIGIR, 2003.

30. Yang, B., Fu, X., Sidiropoulos, N.D., Hong, M.: Towards k-means-friendly spaces: Simultaneous deep learning and clustering. arXiv preprint arXiv:1610.04794 (2016)