

# GC-Flow: A Graph-Based Flow Network for Effective Clustering

Tianchun Wang<sup>1</sup> Farzaneh Mirzazadeh<sup>2</sup> Xiang Zhang<sup>1</sup> Jie Chen<sup>2</sup>

<sup>1</sup>The Pennsylvania State University <sup>2</sup>MIT-IBM Watson AI Lab, IBM Research

## GCN's Limitation

Graph convolutional network (GCN) [1] is a seminal deep learning model for semi-supervised learning on graphs. It capitalizes on the graph structure underlying data to achieve effective classification.

However, GCN and other GNNs may not produce node representations with useful information for goals different from classification; e.g., the representations do not cluster well.

## GC-Flow

**Question:** Can we build a graph representation model that is effective for not only classification but also clustering?

**Answer:** Rather than construct a discriminative model  $p(y|\mathbf{x})$  as all GNNs do, we build a generative model  $p(\mathbf{x}|y)p(y)$  whose class conditional likelihood is defined by explicitly modeling the representation space.

**FlowGMM** [2]: FlowGMM is a normalizing flow that maps the distribution of input features to a Gaussian mixture, resulting in well-structured clusters.

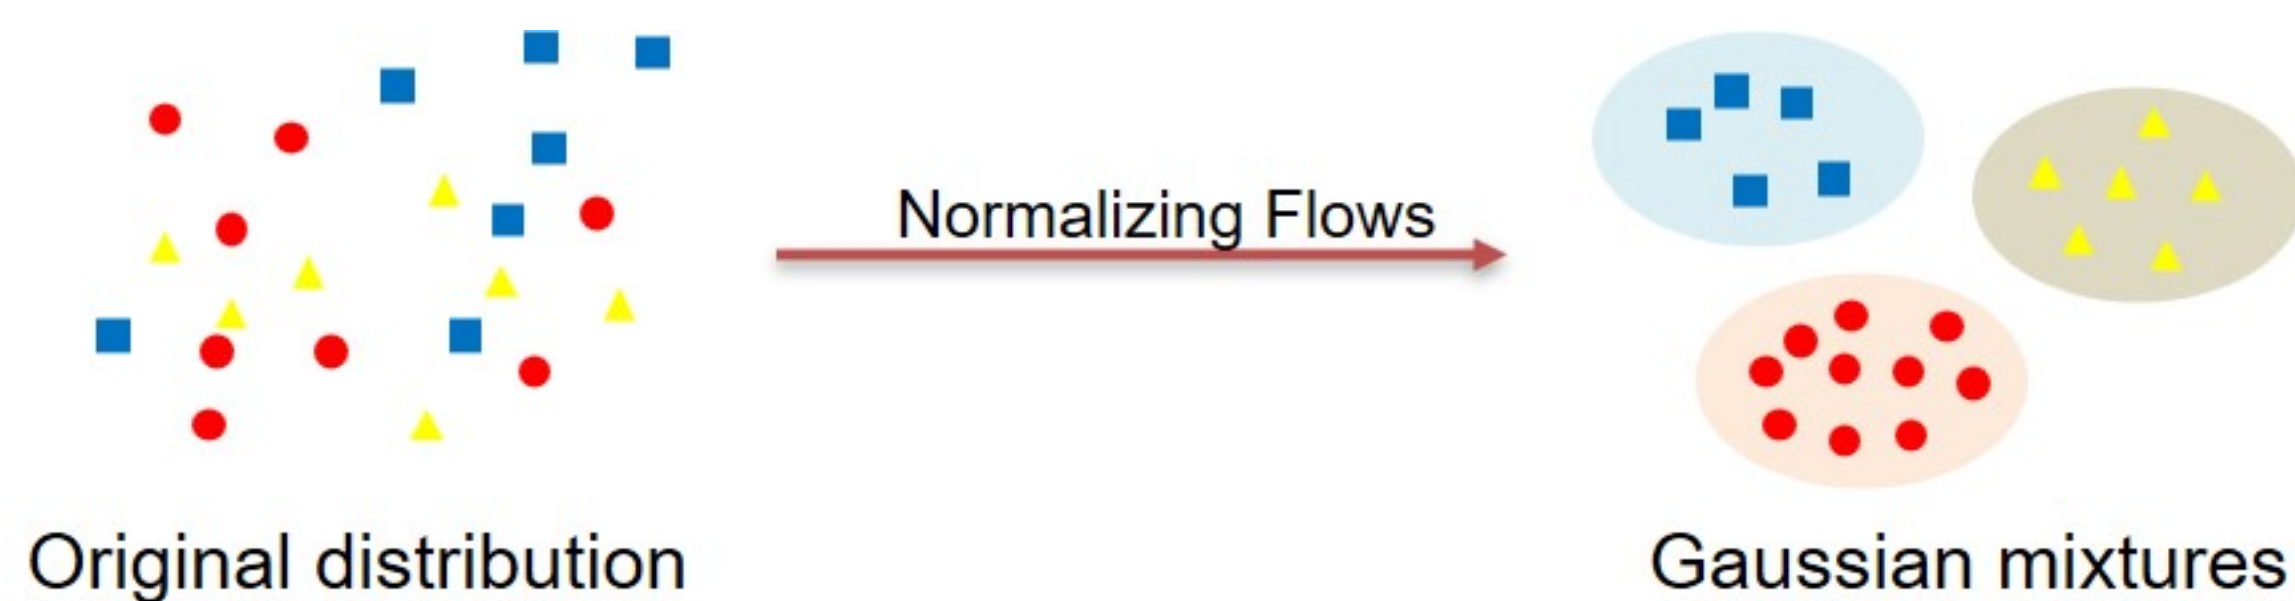


Figure: The framework of FlowGMM.

**GC-Flow:** We propose **graph convolutional normalizing flow** (GC-Flow), which extends FlowGMM that acts on data points separately to one that acts on all graph nodes collectively.




Figure: Representation space of Cora. Left to right: FlowGMM; GCN; GC-Flow.

## Mathematical Definition

Let  $\mathbf{A} \in \mathbb{R}^{n \times n}$  be the adjacency matrix of a given graph with  $n$  nodes and let  $\widehat{\mathbf{A}}$  be a certain normalization of  $\mathbf{A}$  (e.g., the normalization defined by GCN). Let  $\mathbf{X}$  be the  $n \times D$  input feature matrix. Starting with  $\mathbf{X}^{(0)} \equiv \mathbf{X}$ , we define GC-Flow  $\mathbf{F}(\mathbf{X}) : \mathbb{R}^{n \times D} \rightarrow \mathbb{R}^{n \times D}$  as a composition of  $T$  constituent flows  $\mathbf{F} = \mathbf{F}_T \circ \mathbf{F}_{T-1} \circ \dots \circ \mathbf{F}_1$ , where each constituent flow  $\mathbf{F}_i$  has parameter  $\mathbf{W}^{(i-1)}$  and computes

$$\mathbf{X}^{(i)} = \mathbf{F}_i(\underbrace{\widehat{\mathbf{A}}\mathbf{X}^{(i-1)}}_{\tilde{\mathbf{x}}^{(i)}}; \mathbf{W}^{(i-1)}), \quad i = 1, \dots, T. \quad (1)$$

The final node representation is  $\mathbf{Z} \equiv \mathbf{X}^{(T)} \in \mathbb{R}^{n \times D}$ .

**Parameterization:** In addition to  $\mathbf{W}^{(i-1)}$ ,  $\widehat{\mathbf{A}}$  can also be parameterized.

**Constituent flows:** Any normalizing flow will do.

**Properties:** 1) GC-Flow is a normalizing flow; 2) GC-Flow is a GNN.

## Probability Model and Training Objective

A major distinction between GC-Flow and a usual GNN lies in the training objective (which requires a probability model). We use a maximum-likelihood type of objective, which is equivalent to maximizing the likelihood that  $\mathbf{Z}$  forms a Gaussian mixture:

$$\max_{\theta} \mathcal{L} := \frac{1-\lambda}{|\mathcal{D}_l|} \sum_{(\mathbf{x}, y=k) \in \mathcal{D}_l} \log p(\mathbf{x}, y=k) + \frac{\lambda}{|\mathcal{D}_u|} \sum_{\mathbf{x} \in \mathcal{D}_u} \log p(\mathbf{x}), \quad (2)$$

where  $\mathcal{D}_l$  and  $\mathcal{D}_u$  denote the set of labeled and unlabeled nodes, respectively.

We use  $p(\mathbf{X})$  and  $\pi(\mathbf{Z})$  to denote the joint distribution of the node feature vectors and that of the representations, respectively. The change-of-variable formula gives

$$p(\mathbf{X}) = \pi(\mathbf{Z}) |\det \nabla \mathbf{F}(\mathbf{X})|. \quad (3)$$

We assume the node features to be iid  $p(\mathbf{X}) = p(\mathbf{x}_1)p(\mathbf{x}_2) \dots p(\mathbf{x}_n)$  and also model the representations to be iid  $p(\mathbf{Z}) = p(\mathbf{z}_1)p(\mathbf{z}_2) \dots p(\mathbf{z}_n)$ .

## Determinant Lemma

The Jacobian determinant of the GC-Flow  $\mathbf{F}$  is

$$|\det \nabla \mathbf{F}(\mathbf{X})| = |\det \widehat{\mathbf{A}}|^{TD} \prod_{j=1}^T \prod_{i=1}^n |\det \nabla \mathbf{f}_j(\tilde{\mathbf{x}}_i^{(j)})|. \quad (4)$$

Based on the lemma, we model the class prior  $p(y)$  and the class conditional likelihood  $p(\mathbf{x}|y)$  as

$$p(\mathbf{x}_i|y_i=k) := \mathcal{N}(\mathbf{z}_i; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) |\det \widehat{\mathbf{A}}|^{TD/n} \times \prod_{j=1}^T |\det \nabla \mathbf{f}_j(\tilde{\mathbf{x}}_i^{(j)})|$$

$$p(y_i=k) := \phi_k. \quad (5)$$

Then, the marginal likelihood is

$$p(\mathbf{x}_i) = \pi(\mathbf{z}_i) |\det \widehat{\mathbf{A}}|^{TD/n} \prod_{j=1}^T |\det \nabla \mathbf{f}_j(\tilde{\mathbf{x}}_i^{(j)})|, \quad (6)$$

which is consistent with (3). We use (5) and (6) to compute the labeled part and the unlabeled part of the loss (2), respectively.

## Clustering & Classification Performance

Table: Comparison of GMM-based generative models, GNN-based discriminative models, and GC-Flow for semi-supervised clustering and classification.

	Cora		Citeseer		Pubmed	
	Silhouette	Micro-F1	Silhouette	Micro-F1	Silhouette	Micro-F1
FlowGMM	<b>0.739 ± 0.015</b>	0.504 ± 0.021	<b>0.609 ± 0.034</b>	0.512 ± 0.044	<b>0.653 ± 0.031</b>	0.734 ± 0.014
GMM on $\mathbf{X}$	0.162 ± 0.000	0.163 ± 0.000	0.071 ± 0.000	0.085 ± 0.000	0.062 ± 0.000	0.581 ± 0.000
GMM on $\widehat{\mathbf{A}}\mathbf{X}$	0.144 ± 0.000	0.173 ± 0.000	0.089 ± 0.000	0.182 ± 0.000	0.183 ± 0.000	0.411 ± 0.000
GCN	0.340 ± 0.003	0.813 ± 0.007	0.314 ± 0.016	0.700 ± 0.025	0.453 ± 0.006	<b>0.791 ± 0.004</b>
GraphSAGE	0.346 ± 0.004	0.801 ± 0.005	0.278 ± 0.007	0.697 ± 0.007	0.440 ± 0.018	0.769 ± 0.011
GAT	0.383 ± 0.003	<b>0.825 ± 0.005</b>	0.304 ± 0.003	<b>0.702 ± 0.007</b>	0.435 ± 0.010	0.774 ± 0.005
GC-Flow	<b>0.734 ± 0.006</b>	<b>0.815 ± 0.011</b>	<b>0.538 ± 0.022</b>	<b>0.714 ± 0.011</b>	<b>0.669 ± 0.021</b>	<b>0.791 ± 0.009</b>
	Computers		Photo		Wiki-CS	
	Silhouette	Micro-F1	Silhouette	Micro-F1	Silhouette	Micro-F1
FlowGMM	<b>0.540 ± 0.024</b>	0.614 ± 0.026	<b>0.704 ± 0.027</b>	0.599 ± 0.089	<b>0.677 ± 0.011</b>	0.671 ± 0.011
GMM on $\mathbf{X}$	-0.018 ± 0.00	0.102 ± 0.000	-0.024 ± 0.00	0.120 ± 0.000	0.088 ± 0.000	0.124 ± 0.000
GMM on $\widehat{\mathbf{A}}\mathbf{X}$	-0.021 ± 0.00	0.062 ± 0.000	-0.041 ± 0.00	0.098 ± 0.000	0.026 ± 0.000	0.188 ± 0.000
GCN	0.357 ± 0.026	0.812 ± 0.016	0.388 ± 0.003	0.891 ± 0.012	0.264 ± 0.005	<b>0.775 ± 0.005</b>
GraphSAGE	0.434 ± 0.030	0.761 ± 0.024	0.386 ± 0.007	0.839 ± 0.020	0.233 ± 0.009	0.771 ± 0.003
GAT	0.431 ± 0.015	<b>0.814 ± 0.023</b>	0.425 ± 0.020	<b>0.900 ± 0.009</b>	0.278 ± 0.008	0.773 ± 0.003
GC-Flow	<b>0.487 ± 0.012</b>	<b>0.847 ± 0.007</b>	<b>0.655 ± 0.013</b>	<b>0.917 ± 0.004</b>	<b>0.717 ± 0.010</b>	<b>0.775 ± 0.002</b>

## More Results

Comparison with more clustering methods:

Table: Clustering performance of various GNN methods (Cora).

	NMI	ARI	Silhouette
DGI	0.592 ± 0.001	0.570 ± 0.002	0.330 ± 0.000
GRACE	0.475 ± 0.028	0.394 ± 0.047	0.153 ± 0.011
GCA	0.418 ± 0.053	0.259 ± 0.058	0.301 ± 0.005
GraphCL	0.577 ± 0.002	0.482 ± 0.003	0.297 ± 0.002
MVGRL	<b>0.612 ± 0.026</b>	<b>0.576 ± 0.062</b>	0.369 ± 0.013
R-GMM-VGAE	0.559 ± 0.006	0.557 ± 0.009	<b>0.430 ± 0.013</b>
GC-Flow	<b>0.621 ± 0.013</b>	<b>0.631 ± 0.008</b>	<b>0.734 ± 0.006</b>

Training behavior:

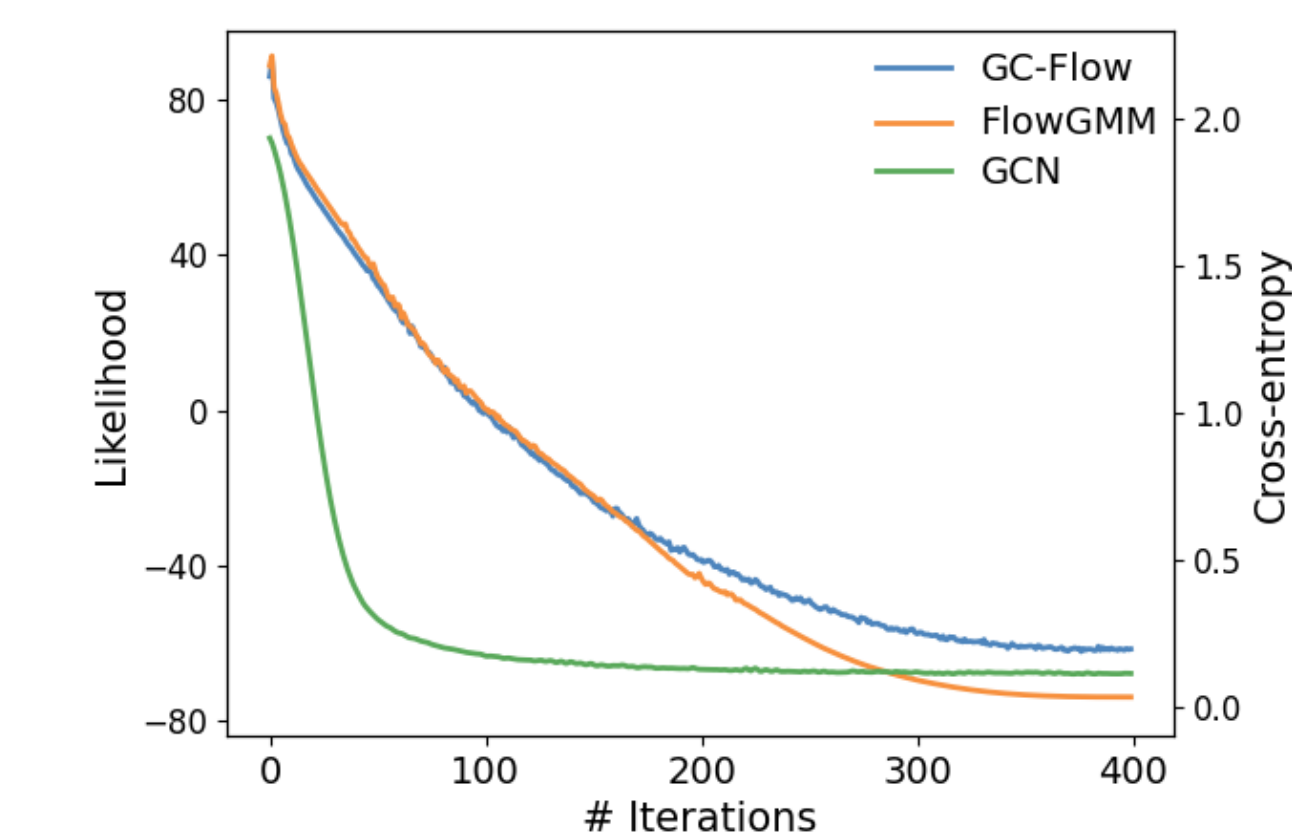


Figure: Convergence of the training loss (Cora).

Analysis on labeling rate:

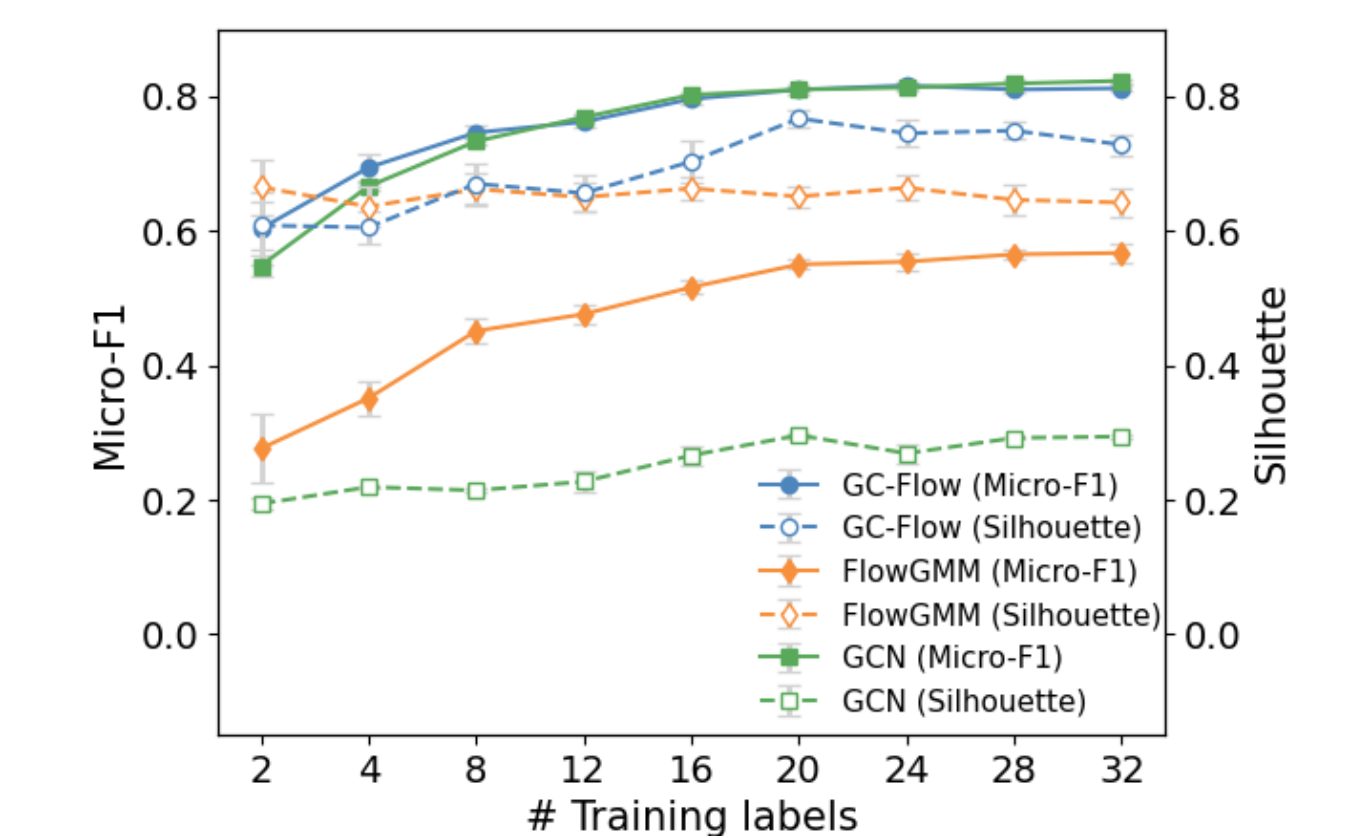


Figure: Performance variation w.r.t. the labeling rate (Cora).

Analysis on depth:

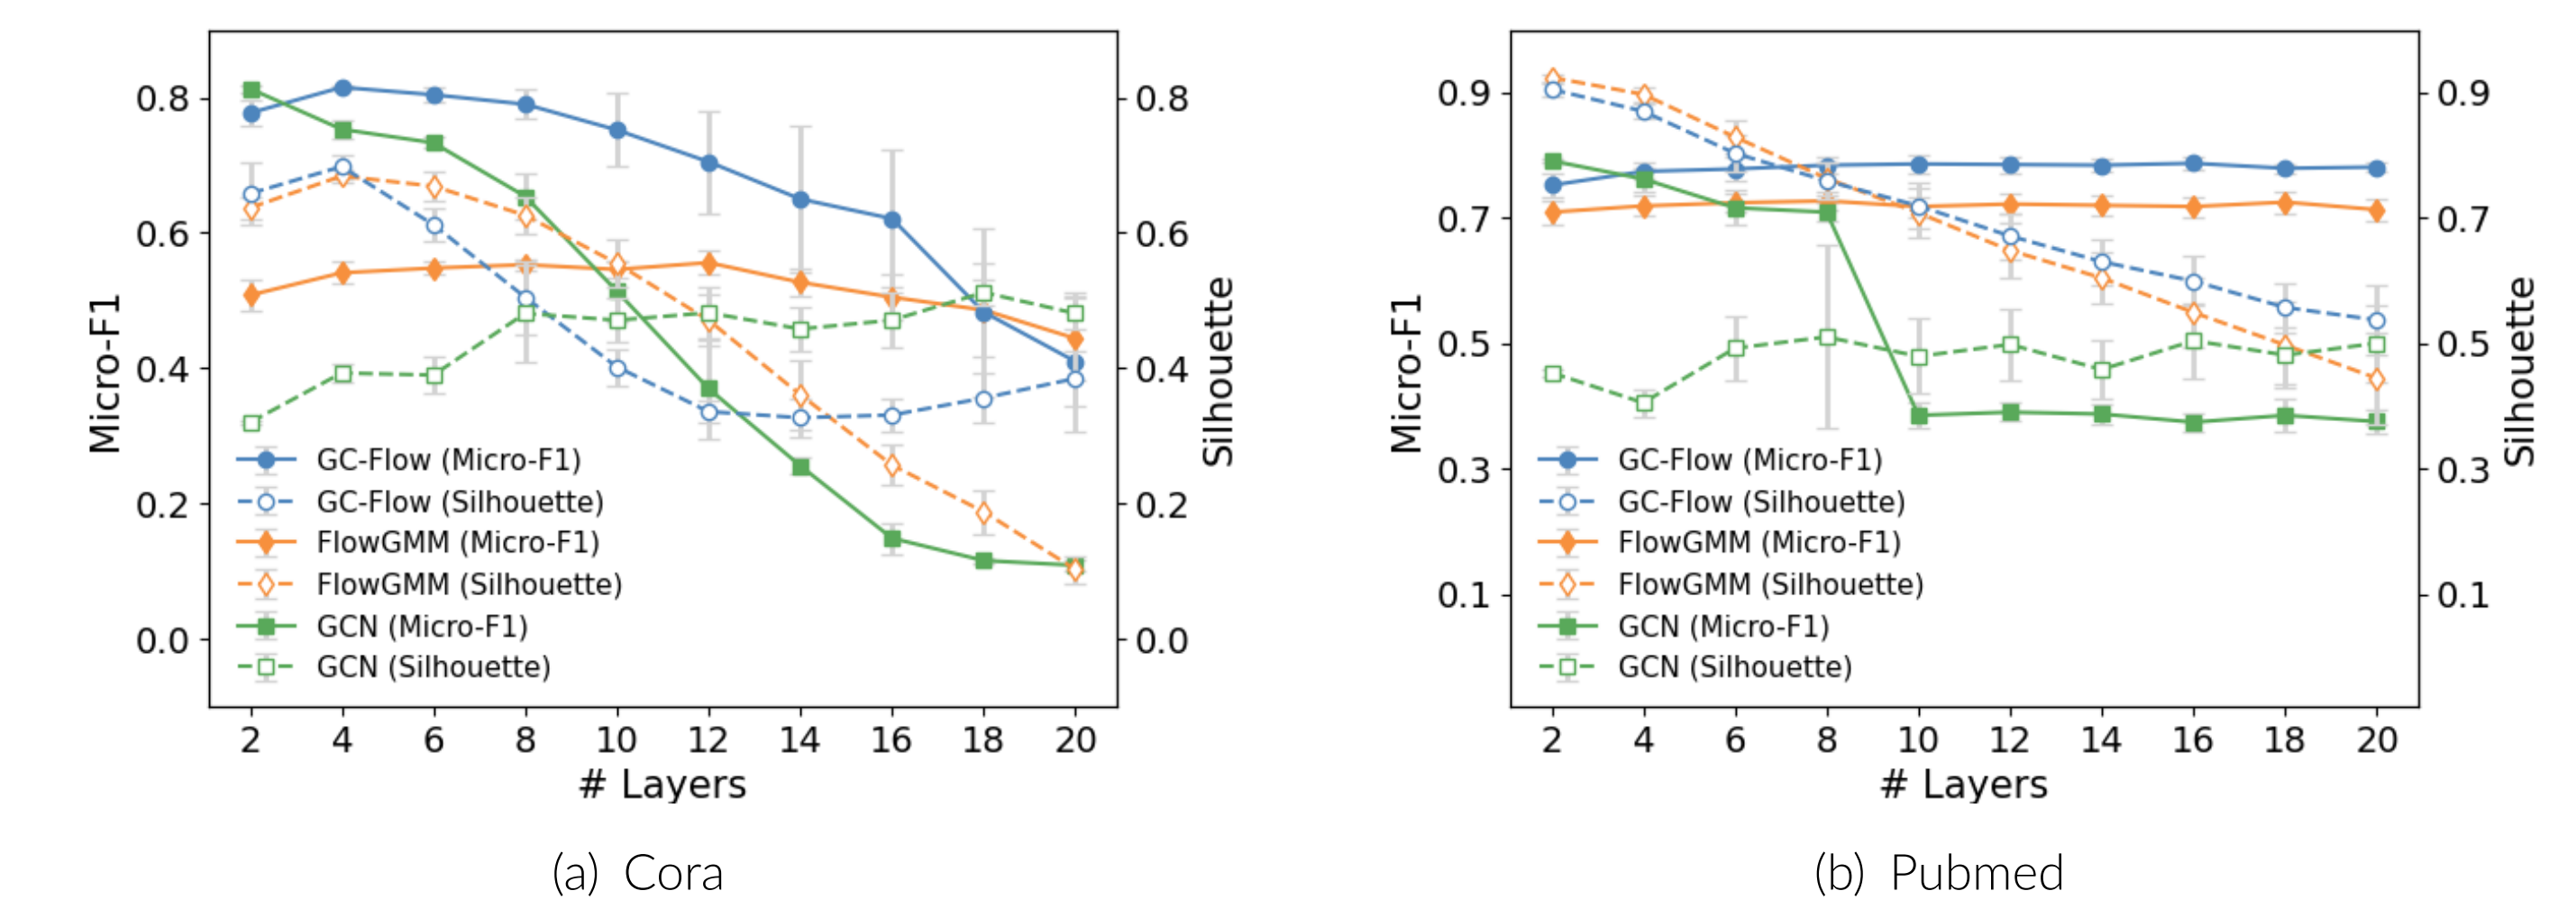


Figure: Performance variation w.r.t. the network depth (i.e., number of constituent flows).

Improving performance by parameterizing  $\widehat{\mathbf{A}}$ :

Table: Effect of modeling  $\widehat{\mathbf{A}}$ . Boldfaced numbers indicate improvement over GC-Flow.

	Pubmed		Computers		Photo	
	Silhouette	Micro-F1	Silhouette	Micro-F1	Silhouette	Micro-F1
GC-Flow	0.669 ± 0.021	0.791 ± 0.009	0.487 ± 0.012	0.847 ± 0.007	0.655 ± 0.013	0.917 ± 0.004
GC-Flow-p	<b>0.804 ± 0.010</b>	0.790 ± 0.007	<b>0.706 ± 0.019</b>	0.841 ± 0.006	<b>0.874 ± 0.011</b>	0.914 ± 0.008
GC-Flow-l	<b>0.856 ± 0.029</b>	0.783 ± 0.009	<b>0.582 ± 0.020</b>	<b>0.851 ± 0.009</b>	<b>0.842 ± 0.006</b>	0.911 ± 0.005

## References

- [1] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *Preprint arXiv:1609.02907*, 2016.
- [2] Pavel Izmailov, Polina Kirichenko, Marc Finzi, and Andrew Gordon Wilson. Semi-supervised learning with normalizing flows. In *ICML*, pages 4615–4630. PMLR, 2020.