DyNMF: Role Analytics in Dynamic Social Networks

Yulong Pei, Jianpeng Zhang, George Fletcher, Mykola Pechenizkiy

Eindhoven University of Technology, the Netherlands

{y.pei.1,j.zhang.4,g.h.l.fletcher,m.pechenizkiy}@tue.nl

Abstract

Roles of nodes in a social network (SN) represent their functions, responsibilities or behaviors within the SN. Roles typically evolve over time, making role analytics a challenging problem. Previous studies either neglect role transition analysis or perform role discovery and role transition learning separately, leading to inefficiencies and limited transition analysis. We propose a novel dynamic non-negative matrix factorization (DyNMF) approach to simultaneously discover roles and learn role transitions. DyNMF explicitly models temporal information by introducing a role transition matrix and clusters nodes in SNs from two views: the current view and the historical view. The current view captures structural information from the current SN snapshot and the historical view captures role transitions by looking at roles in past SN snapshots. DyNMF efficiently provides more effective analytics capabilities, regularizing roles by temporal smoothness of role transitions and reducing uncertainties and inconsistencies between snapshots. Experiments on both synthetic and real-world SNs demonstrate the advantages of DyNMF in discovering and predicting roles and role transitions.

1 Introduction

Roles of nodes in a social network (SN) represent their functions in the network. For example, nodes may function as the centers of stars, members of cliques or peripheral nodes. An example is shown in Figure 1(a) where different colors indicate different roles, e.g., yellow nodes are centers of stars. In contrast with communities, which group together nodes that are well connected to each other [Fortunato, 2010], roles group nodes with similar structural properties [Henderson *et al.*, 2012]. Role discovery is the task of partitioning the nodes by their structural patterns in the SN [Rossi *et al.*, 2015]. Discovering roles in graphs can shed light on numerous graph mining tasks such as graph transferring [Henderson *et al.*, 2012], anomaly detection [Rossi *et al.*, 2013] and community detection [Ruan and Parthasarathy, 2014]. Therefore, the problem of role discovery has attracted an increasing amount of attention recently.

Existing role discovery methods focus predominantly on static SNs. For example, non-negative matrix factorization (NMF) based methods, such as RolX [Henderson *et al.*, 2012] and GLRD [Gilpin *et al.*, 2013], cluster a node-feature matrix to discover roles. Stochastic blockmodels employ Bayesian methods for role discovery [Airoldi *et al.*, 2009; Fu *et al.*, 2009]. In the real world, however, *dynamic* SNs are ubiquitous and structures of the networks will change over time. State-of-the-art static methods are not easy to extend to dynamic SNs directly. Few attempts have been made to discover roles and analyze role transitions in dynamic SNs. These attempts either neglect role transition analysis or perform role discovery and role transition learning separately. Evolutionary role clustering method [Choobdar *et al.*, 2015] integrates temporal information into a weighting function for user similarity and clustering. However, role transitions have not been analyzed in this work. DBMM [Rossi *et al.*, 2012; 2013] directly uses RolX to discover roles in each SN snapshot and then analyzes the role transition based on obtained node-role matrices. Although role transitions are analyzed in this model, role discovery and role transition analytics are two separate steps, i.e., role transition information can be learned only after the role discovery process, as shown in Figure 1(b). As we show in this paper, this strategy is inefficient and unstable in practice. These problems also remain in other studies for dynamic role discovery [Li *et al.*, 2013; Abnar *et al.*, 2015]. A summary comparison of the state of the art in role discovery methods can be found in Table 1.

To sum up, there are two issues in previous work: (1) *lack of role transition analysis*; and (2) *inefficiency in role transition analysis*. To tackle these issues, in this work we propose a new dynamic non-negative matrix factorization (DyNMF) approach. DyNMF is a unified model to discover role and role transition simultaneously in dynamic SNs. An illustration of DyNMF is shown in Figure 1(c) where we can simultaneously obtain the role matrix of snapshot $t+1$ and the role transition from snapshot t to $t+1$ by using information in snapshot $t+1$ and the role matrix of snapshot t. In particular, DyNMF can solve the two issues effectively and efficiently:

- For the issue of *lack of role transition analysis*, DyNMF explicitly introduces a *role transition matrix* for role transition, where roles and role transitions are modeled in a unified framework. Current and historical views are combined for role analytics. The current view follows

Methods	Unsupervised	Dynamic	Transition	Role Definition
RolX [Henderson *et al.*, 2012]	√	-	-	role explanation after role discovery
GLRD [Gilpin *et al.*, 2013]	√	-	-	role explanation after role discovery
DBMM [Rossi *et al.*, 2013]	√	√	√	role explanation after role discovery
LAP [Li *et al.*, 2013]	-	√	-	predefined roles based on data sets
ERM [Choobdar *et al.*, 2015]	√	√	-	predefined roles based on influence and blockage
SSRM [Abnar *et al.*, 2015]	√	√	-	predefined roles (e.g., leader, mediator)
DyNMF (our method)	√	√	√	role explanation after role discovery

Table 1: Comparison of role discovery methods.

(a) Example of roles and role transitions in SNs.

(b) Role detection and role transition analysis in previous studies.

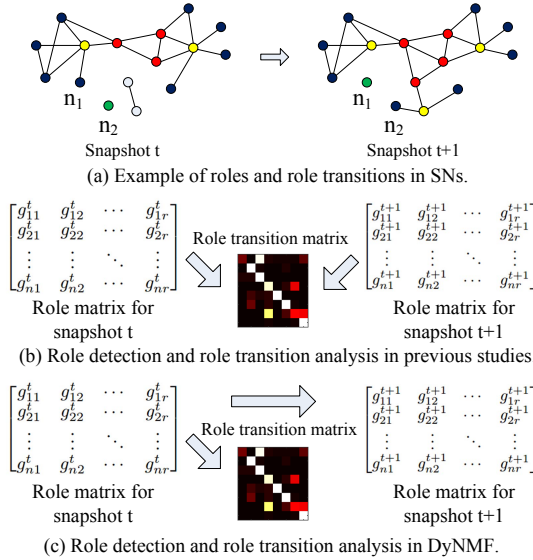(c) Role detection and role transition analysis in DyNMF.

Figure 1: Examples of roles and role analytics in previous methods and DyNMF.

RolX to discover roles in the current SN snapshot, while the historical view learns role transitions using past role information and the current SN snapshot.

- For the issue of *inefficiency in role transition analysis*, DyNMF, as a unified model, supports the simultaneous discovery of both roles and role transitions. In particular, it requires only one pass over the data to obtain roles and role transitions compared with previous studies.

DyNMF is also advantageous in a further aspect: by combining current and historical views, we can regularize the roles by capturing the temporal smoothness of roles and also reduce uncertainties and inconsistencies between snapshots. Thus, temporal information is better explored compared to [Choobdar *et al.*, 2015] and the discovered roles are more stable compared to DBMM.

All of these advantages of DyNMF are validated through extensive experiments on both synthetic and real-world SNs. The results validate DyNMF is advantageous in discovering roles, capturing role transitions, and predicting roles.

2 Role Analytics Using DyNMF

We first briefly revisit the original RolX approach to static role discovery using NMF [Henderson *et al.*, 2012]. We then introduce DyNMF, our new approach for dynamic role analytics in dynamic SNs.

2.1 NMF based Role Discovery

Non-negative matrix factorization (NMF) [Lee and Seung, 2001] is a popular model in multivariate analysis and linear algebra where a matrix is factorized into two matrices, with the property that all three matrices have no negative elements. There are several advantages in NMF including ease of implementing inference and ease of interpreting results. Hence, this model is widely used in text mining.

RolX [Henderson *et al.*, 2012] is the first method to discover roles using NMF. Given a node-feature matrix $V_{n \times l}$, where n and l is the number of nodes and features respectively, the idea of RolX is to generate a rank r approximation $GF \approx V$ where r is the number of roles, matrix $G_{n \times r}$ denotes the nodes' membership and matrix $F_{r \times l}$ represents the association of roles and features. Thus, the problem of role discovery is to seek two low rank matrices G and F to satisfy:

$$\min_{G,F} \|V - GF\|^2, \ \ s.t. \ G \geq 0, F \geq 0 \qquad (1)$$

where $\| \cdot \|$ is the Frobenius norm. The non-negativity constraint in Eq. (1) makes the representation of the original data easier to interpret and more semantically meaningful compared with other factorization methods, e.g., SVD and PCA [Lee and Seung, 2001]. Using multiplicative update rules, the solution for Eq. (1) is shown as follows:

$$G \leftarrow G \circ \frac{VF^T}{GFF^T}, \ \ F \leftarrow F \circ \frac{G^T V}{G^T GF} \qquad (2)$$

where $\circ$ denotes the element-wise product. By introducing notions of sparsity, diversity and alternativeness, RolX was extended with more guidance used as constraints in the matrix factorization for role discovery [Gilpin *et al.*, 2013]. However, these methods focus only on static SNs and it is nontrivial to extend them to dynamic SNs directly.

2.2 DyNMF Approach

To solve problems in previous studies on dynamic role discovery, i.e., (1) requiring separate analysis for role transition after discovering roles [Rossi *et al.*, 2012; 2013] and (2) no role transition analysis [Abnar *et al.*, 2015; Choobdar *et al.*, 2015], we propose a novel dynamic non-negative matrix factorization (DyNMF) approach for role discovery and role transition analysis simultaneously.

DyNMF factorizes the current node-feature matrix from two different views: *current view* and *historical view*. Figure 2 is to illustrate the proposed DyNMF approach. The first
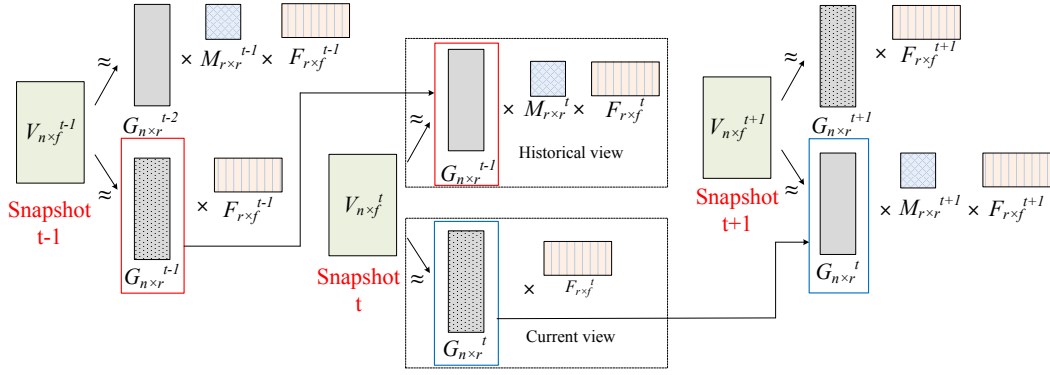
Figure 2: Graphical representation of DyNMF approach which consists of current view and historical view. To learn the node-role matrix $G^{(t)}$, role transition $M^{(t)}$ and associate matrix $F^{(t)}$ in snapshot t, we take node-feature matrix $V^{(t)}$ in snapshot t and node-role matrix $G^{(t-1)}$ from snapshot $t-1$ as the input.

view is derived from the current SN snapshot same to RolX introduced in Section 2.1. Formally, the current view based factorization component for snapshot t is:

$$\min_{G^{(t)}, F^{(t)}} \|V^{(t)} - G^{(t)} F^{(t)}\|^2 \qquad (3)$$

where $V^{(t)}$, $G^{(t)}$ and $F^{(t)}$ are the node-feature, node-role, and role-feature matrix in snapshot t, respectively.

The second view captures historical information by looking at the roles in the previous SN snapshot. We explicitly introduce a role transition matrix M to capture the transitions between previous roles and current roles. Formally, the historical view based factorization component for snapshot t is:

$$\min_{F^{(t)}, M^{(t)}} \|V^{(t)} - G^{(t-1)} M^{(t)} F^{(t)}\|^2 \qquad (4)$$

where $V^{(t)}$ and $F^{(t)}$ have the same meaning in Eq. (3), $M^{(t)}$ denotes the role transition matrix from snapshot $t-1$ to t and $G^{(t-1)}$ denote the node-role matrix from snapshot $t-1$.

By combining these two views of factorization components, we can make good use of the temporal information to discover roles and learn role transitions simultaneously. We formulate the objective function of DyNMF as follows:

$$\min_{G^{(t)}, F^{(t)}, M^{(t)}} L = \min_{G^{(t)}, F^{(t)}, M^{(t)}} \|V^{(t)} - G^{(t)} F^{(t)}\|^2 \qquad (5)$$
$$+ \|V^{(t)} - G^{(t-1)} M^{(t)} F^{(t)}\|^2$$
$$s.t.\ G^{(t)} \geq 0, F^{(t)} \geq 0, M^{(t)} \geq 0.$$

Note that we use only the first-order assumption in Eq. (5). One can extend it to any higher-order version by setting $k > 1$, i.e., the roles depend on more than one previous snapshots:

$$\min_{G^{(t)}, F^{(t)}, M_i^{(t)}, ..., M_k^{(t)}} \|V^{(t)} - G^{(t)} F^{(t)}\|^2 \qquad (6)$$
$$+ \|V^{(t)} - G^{(t-1)} M_1^{(t)} F^{(t)}\|^2 + ...$$
$$+ \|V^{(t)} - G^{(t-k)} M_k^{(t)} F^{(t)}\|^2$$
$$s.t.\ G^{(t)} \geq 0, F^{(t)} \geq 0, M_i^{(t)} \geq 0, 1 \leq i \leq k$$

where $M_k^{(t)}$ is the transition matrix to capture the role transition from snapshot $t-k$ to t. However, higher-order extension suffers from two limitations: (1) from the *computational*

perspective it is more complex since there are more parameters to learn, i.e., $M_k^{(t)}$; (2) from the *empirical* perspective, we have observed that higher-order extensions have not improved performance compared to DyNMF (in Section 3.2). Therefore, unless explicitly stated otherwise, we only focus on the first-order version, i.e., Eq. (5), in following sections.

The objective function in Eq. (5) (or Eq. (6)) is not convex for all parameters $G^{(t)}$, $F^{(t)}$ and $M^{(t)}$ simultaneously. We use the multiplicative update rules to solve this optimization problem due to its good compromise between speed and ease of implementation [Lee and Seung, 2001]. The optimization is done by iterating the three following steps until the convergence (or the number of iteration exceeds a given threshold): (1) fix matrices $F^{(t)}$ and $M^{(t)}$ to update $G^{(t)}$ (2) fix matrices $G^{(t)}$ and $F^{(t)}$ to update $M^{(t)}$ and (3) fix matrices $G^{(t)}$ and $M^{(t)}$ to update $F^{(t)}$ (Algorithm 1). Formally, using the Karush-Kuhn-Tucker (KKT) conditions to solve the objective function in Eq. (5), we get the update rules:

$$G^{(t)} \leftarrow G^{(t)} \circ \frac{V^{(t)} F^{(t)T}}{G^{(t)} F^{(t)} F^{(t)T}} \qquad (7)$$

$$M^{(t)} \leftarrow M^{(t)} \circ \frac{G^{(t-1)T} V^{(t)} F^{(t)T}}{G^{(t-1)T} G^{(t-1)} M^{(t)} F^{(t)} F^{(t)T}} \qquad (8)$$

$$F^{(t)} \leftarrow F^{(t)} \circ \qquad (9)$$

$$\frac{G^{(t)T} V^{(t)} + M^{(t)T} G^{(t-1)T} V^{(t)}}{G^{(t)T} G^{(t)} F^{(t)} + M^{(t)T} G^{(t-1)T} G^{(t-1)} M^{(t)} F^{(t)}}$$

Algorithm 1 Optimization Algorithm

Input: previous node-role matrix $G^{(t-1)}$, node-feature matrix $V^{(t)}$
Output: node-role matrix $G^{(t)}$, transition matrix $M^{(t)}$
  Initialize $G^{(t)}$, $M^{(t)}$ and $F^{(t)}$ with random non-negative values
  while not converge **do**
    Update $G^{(t)}$ according to Eq. (7)
    Update $M^{(t)}$ according to Eq. (8)
    Update $F^{(t)}$ according to Eq. (9)
  end while

Complexity analysis of DyNMF. For simplicity, given two matrices $M_{n \times r}$ and $N_{r \times l}$, the computational complexity of the

multiplication of M and N is $O(nrl)$. The complexity of Algorithm 1 is $O(nrl + nr^2 + lr^2)$ and therefore by considering the number of iteration i the number of snapshots t the complexity is $O(ti(nrl + nr^2 + lr^2))$. The complexity of feature extraction and model selection, discussed below, can be referred to [Henderson *et al.*, 2012].

2.3 Model Selection

In this paper, we view role as a latent class and use a model selection method to determine the number of roles in the discovery process because this method has better generalization in SNs [Henderson *et al.*, 2012]. We adopt the model selection method proposed in [Henderson *et al.*, 2012]. It uses the Minimum Description Length (MDL) [Rissanen, 1978] to decide on the number of roles. The appropriate number of roles r is the one that minimizes the description length $\mathcal{L}$, defined as sum of the coding cost $\mathcal{M}$ and the model description cost $\mathcal{E}$. $\mathcal{M}$ is defined as $br(n + l)$, where b is the bits used for each element. $\mathcal{E}$ is defined as sum of the KL divergence based errors from current view $\mathcal{E}_1$ and historical view $\mathcal{E}_2$:

$$\mathcal{E}_1 = \sum_t \sum_{i,j} \left(V_{i,j}^{(t)} \log \frac{V_{i,j}^{(t)}}{U_{i,j}^{(t)}} - V_{i,j}^{(t)} + U_{i,j}^{(t)} \right) \quad (10)$$

$$\mathcal{E}_2 = \sum_t \sum_{i,j} (V_{i,j}^{(t)} \log \frac{V_{i,j}^{(t)}}{W_{i,j}^{(t)}} - V_{i,j}^{(t)} + W_{i,j}^{(t)}) \quad (11)$$

where $U^{(t)} = G^{(t)} F^{(t)}$ and $W^{(t)} = G^{(t-1)} M^{(t)} F^{(t)}$. Note that to calculate $\mathcal{E}_1$ and $\mathcal{E}_2$, we sum the errors of all snapshots. We choose the optimal number of roles with the minimal description length $\mathcal{L}$ empirically.

2.4 Feature Extraction

In previous studies on role discovery, a variety of feature extraction methods have been employed. In [Henderson *et al.*, 2012], the features consist of local and egonet properties based on counts of links and the egonet-based properties generated in a recursive fashion [Henderson *et al.*, 2011]. In [Zhao *et al.*, 2013], five types of network properties, i.e., homophily, triadic closure, reachability, embeddedness and structural holes, have been used as the features to infer the social roles in the SNs. The structural properties of nodes in the graph such as clustering coefficient and the locality index, have been considered as node features in [Choobdar *et al.*, 2015]. In order to compare our proposed method with the original NMF-based role discovery method, we use the same feature extraction method as RolX [Henderson *et al.*, 2012] and DBMM [Rossi *et al.*, 2013] in this work.

3 Experimental Study

3.1 Settings

To validate the advantages of our proposed DyNMF for role discovery and role transition learning, we conduct experiments on one synthetic data set and four real-world data sets[1]. A summary of these data sets is shown in Table 2.

For the synthetic data set, we generate a series of graph snapshots with four roles, i.e., star-center nodes, star-edge

---

[1] networkrepository.com/index.php

Data set	# Nodes	# Edges	# Roles	# Snapshots
Synthetic	100	1729	4	8
Enron	147	1666	8	9
Reality	6809	9467	11	10
Facebook	44416	196414	12	12
Slashdot	51068	130324	11	12

Table 2: Summary of data sets used in the experiments.

	Methods	NMI
Feature-based methods in social science	K-means	0.4831
	Agglomerative	0.5116
Graph-based methods	Spectral	0.6233
	MMSB	0.6382
NMF-based methods in data mining	RolX	0.7220
	DyNMF	**0.7816**

Table 3: Comparison of role discovery performance using *NMI*. The highest value is in bold.

nodes, bridges and cliques. In each snapshot, based on a fixed change rate, part of nodes will change to other roles randomly. As we have ground-truth labels of roles, we can use traditional clustering evaluation metrics, e.g., Normalized Mutual Information (NMI), to verify the experimental results. To further examine the sensitivity of DyNMF towards role changes in SNs, the change rate is set to vary from 0% to 25% and the performance will be compared. For the real-world data sets, there are no ground-truth labels of roles, so the numbers of roles shown in Table 2 are the "best" numbers of roles determined by the method introduced in Section 2.3. To evaluate the performance, *goodness-of-fit index* [Wasserman and Faust, 1994] is applied.

Our experimental study is aimed to analyze the performance of DyNMF on three analytics tasks:

- *Role discovery analysis*. This task aims to analyze the performance of role discovery using DyNMF quantitatively. We first use *NMI* to evaluate DyNMF on the synthetic data set. Then we examine the sensitivity of DyNMF towards role changes. We also use *goodness-of-fit index* to measure DyNMF on real-world SNs. Furthermore we compare models with different orders.
- *Role transition analysis*. The goal of this task is to verify the effectiveness and stability of the role transition learned by DyNMF. In particular, we calculate the traces of the normalized role transition matrices learned by DyNMF and DBMM respectively for comparison.
- *Role prediction analysis*. The evaluation of the prediction ability is another way to validate the effectiveness of the role transition. In this experiment, we propose two strategies to predict nodes' roles using the transition matrices learned by DyNMF.

3.2 Role Discovery Analysis

NMI on Synthetic Data Set To quantitatively evaluate the performance of DyNMF, we use Normalized Mutual Information (NMI) as the measurement on the synthetic data set. NMI is obtained by dividing the mutual information by the arithmetic average of the entropy of obtained cluster $\mathcal{C}$ and ground-truth cluster $\mathcal{D}$. We compare three types of methods:

Methods	Average *goodness-of-fit index*			
	Enron	Reality	Facebook	Slashdot
K-means	1.7722	1.5950	0.1802	0.0496
Agglomerative	0.8750	0.7875	0.0461	0.0888
Spectral	1.4235	0.9945	0.0864	0.1025
MMSB	1.0425	0.7596	0.1174	0.0950
RolX	1.1873	1.0685	0.1168	**0.0314**
DyNMF	**0.5788**	**0.6051**	**0.0393**	0.0345

Table 4: Comparison of role discovery performance using *goodness-of-fit index*. The lowest error in each data set is in bold.

Method	*goodness-of-fit index*		Running time	
	Enron	Reality	Enron	Reality
first-order	0.5788	0.6051	1.6ms	307.1s
second-order	0.5758	0.6344	1.8ms	370.5s
third-order	0.5818	0.6129	2.3ms	414.5s

Table 5: *Goodness-of-fit index* and running time vs. orders.

feature-based methods in social science, graph-based methods and feature-based methods in data mining. Note that all the feature-based methods use the same features extracted by the method introduced in Section 2.4 and graph-based methods use the adjacency matrix from the graph as the input.

The NMI results of different methods are shown in Table 3. Note that these results are based on 5% role change rate. From the results, it can be observed that: (1) DyNMF outperforms all the other methods. It indicates the advantage of DyNMF in role discovery by explicitly leveraging the temporal smoothness in role transition. (2) Feature-based methods in data mining, i.e., RolX and DyNMF, perform better than other methods. This result demonstrate that latent models like NMF are better choices in role discovery.

Smoothness Experiment on Synthetic Data Set We investigate the influence of smoothness in role transition and sensitivity of DyNMF towards role changes in SNs. The change rate between two consecutive snapshots is set to vary from 0% to 25% and the NMIs of DyNMF and RolX are shown in Figure 3. This result shows the limitation of our proposed DyNMF in the condition that the assumption of smoothness in role transition does not hold, i.e., the change rate becomes larger. However, this limitation does not influence the performance of DyNMF on real-world data sets for role discovery and role transition analysis. More results and discussion are shown in following experiments.

Goodness-of-fit indices **on Real-world Data Sets** Due to lack of ground-truth roles in real-world data sets, evaluation of role discovery is challenging in practice. Based on structural equivalence [Lorrain and White, 1971], researchers in sociology have applied *goodness-of-fit indices* to measure how well the representation of roles and the relations among these roles fit a given SN [Wasserman and Faust, 1994]. In *goodness-of-fit indices*, it is assumed that the output of a role discovery method is an optimal model, and nodes belonging to the same role are predicted to be perfectly structurally equivalent. For more details about *goodness-of-fit indices*,
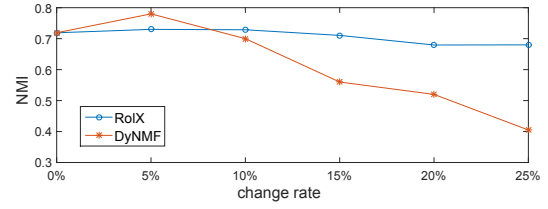

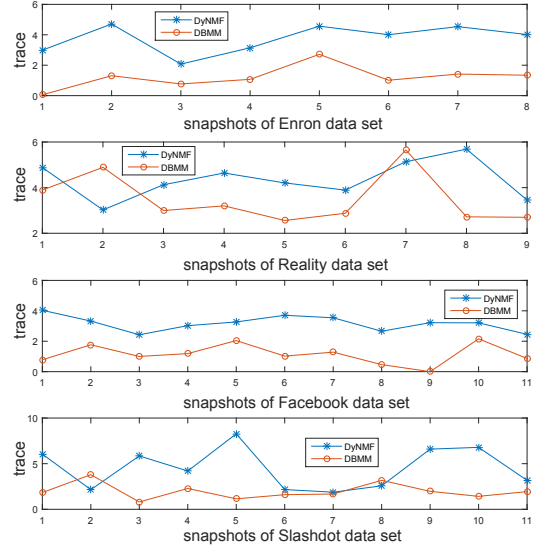Figure 3: The influence of change rates on DyNMF and RolX.

Figure 4: Traces of transition matrices using DBMM and DyNMF.

please refer to [Wasserman and Faust, 1994].

To validate the effectiveness of DyNMF, we compare it with some baselines including K-means, Agglomerative clustering, Spectral clustering and RolX. The comparison of average *goodness-of-fit index* over all snapshots is shown in Table 4. From the results, it can be observed that DyNMF outperforms other methods except on Slashdot. This may result from the extreme sparsity of Slashdot.

Higher-order Comparison on Real-world Data Set We validate the rationality and effectiveness of the first-order assumption empirically introduced in Eq. (6). We compare the *goodness-of-fit index* of DyNMF on Enron and Reality with second-order and third-order versions. The second-order DyNMF looks two steps back and uses role information from snapshot $t-1$ and $t-2$ to detect roles in snapshot t, i.e., $k=2$ in Eq. (6). Similarly, in third-order version, $k=3$. The results are shown in Table 5. It can be observed that higher order models have similar performance with first-order model, i.e., DyNMF, but increase the computational complexity due to more parameters, i.e., more transition matrices, to learn.

3.3 Role Transition Analysis

We use DBMM as a competitor approach to DyNMF in our analysis of the role transition stability. Since there is no constraint in the transition matrix M in DBMM, in order to compare DyNMF with DBMM, we also follow that work not

adding constraint to M. One could easily add different types of constraint in M, e.g., orthogonal. As the values of elements in transition matrices between different snapshots may be in different ranges, we normalize each M by row to make it be a Markov matrix $\tilde{M}$, i.e., $\forall i$, $\sum_j \tilde{M}_{ij} = 1$.

We calculate the traces of the normalized role transition matrices, i.e., $Tr(\tilde{M}) = \sum_i \tilde{M}_{ii}$, to further validate this statement. Figure 4 shows the traces of role transition matrices using DyNMF and DBMM on the four data sets, respectively. Higher trace value means less change of roles. From the results, some conclusions can be drawn:

(1) DyNMF can obtain higher trace in almost all the snapshots of the SNs. This result demonstrates the stability of the role transitions obtained by DyNMF. (2) The stability of transition is more obvious in Enron and Facebook. The reason may be the characteristics of these SNs. Intuitively, the roles (positions) of employees in a company, e.g., Enron, rarely change abruptly and few users in a SN, e.g., Facebook, can become cyberstars suddenly. Therefore, the changes of roles are more smooth in these data sets. In addition, the stability also depends on the temporal granularity of snapshots. (3) The finding of role stability is consistent with the discovery in [Revelle *et al.*, 2016] that some roles are persistent in SNs.

3.4 Role Prediction Analysis

We formulate the goal to predict the roles $\hat{G}^{(t+1)}$ in $t + 1$ snapshot using the role membership matrix $G^{(t)}$ in t snapshot and the transition matrices $M^{(i)}$ ($0 < i \leq t$). Note that these transition matrices are also normalized. To analyze the prediction ability, we design two strategies:

- *average strategy (AvgDyNMF)*. This strategy is based on the *global consistency* and assumes that the role transitions are smooth in all the snapshots, i.e., $\hat{G}^{(t+1)} = G^{(t)} \sum_i^t M^{(i)}/t$.
- *previous strategy (PrevDyNMF)*. This strategy is based on the *local consistency* and assumes that the role transition is smooth between two consecutive snapshots, i.e., $\hat{G}^{(t+1)} = G^{(t)} M^{(t)}$.

We report the squared Frobenius norm of the prediction error, i.e., $\|G^{(t+1)} - \hat{G}^{(t+1)}\|^2$. Note that there is no ground-truth role indicator $G^{(t+1)}$ and therefore we use the node-role matrices detected by RolX in each snapshot as the golden standard similar to the experiment in [Rossi *et al.*, 2013]. The prediction performance results are shown in Figure 5. To validate the effectiveness of DyNMF, we compare its performance with DBMM also using these two strategies.

In fact, DBMM learns the transition directly from the discovered roles, i.e., obtaining $M^{(t)}$ by minimizing $\|G^{(t+1)} - G^{(t)} M^{(t)}\|_F^2$ and intuitively calculating $G^{(t)} M^{(t)}$ to predict $\hat{G}^{(t+1)}$ will be a good approximation of $G^{(t+1)}$. Thus, this role prediction analysis does not aim to beat DBMM but to demonstrate that by learning role and transition simultaneously we can still obtain satisfactory results. Some conclusions can be drawn from these results:

(1) Our proposed DyNMF can effectively predict the roles based on the learned transition matrices since both the errors and the trends of curves for DyNMF are similar to those of
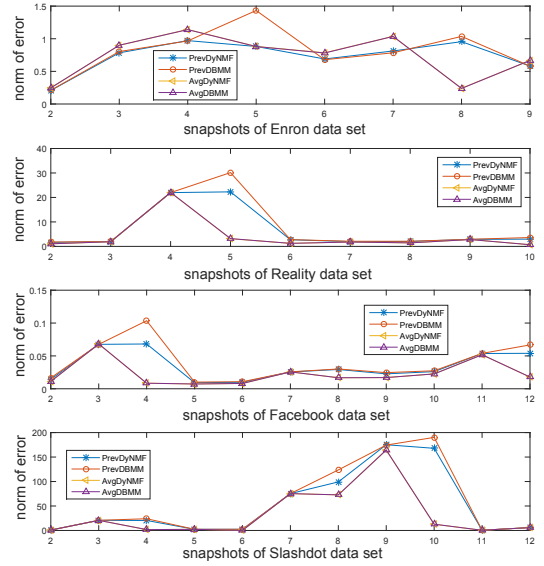

Figure 5: Role prediction using DBMM and DyNMF with average and previous strategies.

DBMM on all the data sets. (2) The *average strategy* performs better than the *previous strategy* on all data sets. It indicates that the global consistency plays a more important role in the dynamics of SNs and all nodes exhibit the average behaviors. This conclusion is consistent with [Rossi *et al.*, 2013]. (3) The gaps between *average strategy* and *previous strategy*, e.g., snapshot 8 in Enron, snapshot 5 in Reality, snapshot 4 in Facebook and snapshot 10 in Slashdot, reflect that there are more role changes happened. DyNMF can effectively handle such changes because in *average strategy* DyNMF performs equally well as DBMM and in *previous strategy* DyNMF outperforms DBMM on all data sets.

4 Conclusions

In this work we proposed DyNMF, a novel dynamic non-negative matrix factorization approach to discover roles and role transitions simultaneously in dynamic SNs. Current and historical views have been combined for the node-feature matrix factorization. The current view is based on structural information in the current snapshot and the historical view captures the correlation between previous roles and current roles using role transition matrices. We conducted comprehensive experiments on both synthetic and real-world data sets to validate the performance of DyNMF in role discovery and role transition learning. We analyzed the experimental results from three aspects including role discovery, role transition, and role prediction. The results indicate the effectiveness of our proposed method for the challenging problem of dynamic role analytics.

For further study, first towards improving performance, it is interesting to study the impact of constraints on role indicator matrices. A second area for improvement is to utilize faster and more scalable matrix factorization methods. A third area for future exploration is to exploit other feature extraction methods which can capture temporal information.

References

[Abnar *et al.*, 2015] Afra Abnar, Mansoureh Takaffoli, Reihaneh Rabbany, and Osmar R Zaïane. Ssrm: structural social role mining for dynamic social networks. *Social Network Analysis and Mining*, 5(1):1–18, 2015.

[Airoldi *et al.*, 2009] Edo M Airoldi, David M Blei, Stephen E Fienberg, and Eric P Xing. Mixed membership stochastic blockmodels. In *Advances in Neural Information Processing Systems*, pages 33–40, 2009.

[Choobdar *et al.*, 2015] Sarvenaz Choobdar, Pedro Ribeiro, Srinivasan Parthasarathy, and Fernando Silva. Dynamic inference of social roles in information cascades. *Data Mining and Knowledge Discovery*, 29(5):1152–1177, 2015.

[Fortunato, 2010] Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3):75–174, 2010.

[Fu *et al.*, 2009] Wenjie Fu, Le Song, and Eric P Xing. Dynamic mixed membership blockmodel for evolving networks. In *Proceedings of the 26th annual international conference on machine learning*, pages 329–336. ACM, 2009.

[Gilpin *et al.*, 2013] Sean Gilpin, Tina Eliassi-Rad, and Ian Davidson. Guided learning for role discovery (glrd): framework, algorithms, and applications. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 113–121. ACM, 2013.

[Henderson *et al.*, 2011] Keith Henderson, Brian Gallagher, Lei Li, Leman Akoglu, Tina Eliassi-Rad, Hanghang Tong, and Christos Faloutsos. It's who you know: graph mining using recursive structural features. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 663–671. ACM, 2011.

[Henderson *et al.*, 2012] Keith Henderson, Brian Gallagher, Tina Eliassi-Rad, Hanghang Tong, Sugato Basu, Leman Akoglu, Danai Koutra, Christos Faloutsos, and Lei Li. Rolx: structural role extraction & mining in large graphs. In *Proceedings of the 18th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 1231–1239. ACM, 2012.

[Lee and Seung, 2001] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, pages 556–562, 2001.

[Li *et al.*, 2013] Kang Li, Suxin Guo, Nan Du, Jing Gao, and Aidong Zhang. Learning, analyzing and predicting object roles on dynamic networks. In *Data Mining (ICDM), 2013 IEEE 13th International Conference on*, pages 428–437. IEEE, 2013.

[Lorrain and White, 1971] Francois Lorrain and Harrison C White. Structural equivalence of individuals in social networks. *The Journal of mathematical sociology*, 1(1):49–80, 1971.

[Revelle *et al.*, 2016] Matt Revelle, Carlotta Domeniconi, and Aditya Johri. Persistent roles in online social networks. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 47–62. Springer, 2016.

[Rissanen, 1978] Jorma Rissanen. Modeling by shortest data description. *Automatica*, 14(5):465–471, 1978.

[Rossi *et al.*, 2012] Ryan Rossi, Brian Gallagher, Jennifer Neville, and Keith Henderson. Role-dynamics: fast mining of large dynamic networks. In *Proceedings of the 21st international conference companion on World Wide Web*, pages 997–1006. ACM, 2012.

[Rossi *et al.*, 2013] Ryan A Rossi, Brian Gallagher, Jennifer Neville, and Keith Henderson. Modeling dynamic behavior in large evolving graphs. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 667–676. ACM, 2013.

[Rossi *et al.*, 2015] Ryan Rossi, Nesreen K Ahmed, et al. Role discovery in networks. *Knowledge and Data Engineering, IEEE Transactions on*, 27(4):1112–1131, 2015.

[Ruan and Parthasarathy, 2014] Yiye Ruan and Srinivasan Parthasarathy. Simultaneous detection of communities and roles from large networks. In *Proceedings of the second edition of the ACM conference on Online social networks*, pages 203–214. ACM, 2014.

[Wasserman and Faust, 1994] Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*, volume 8. Cambridge university press, 1994.

[Zhao *et al.*, 2013] Yuchen Zhao, Guan Wang, Philip S Yu, Shaobo Liu, and Simon Zhang. Inferring social roles and statuses in social networks. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 695–703. ACM, 2013.