

Network Cartography: Seeing the Forest and the Trees

Jia Wang, Kevin Chen-Chuan Chang, Hari Sundaram
Department of Computer Science
University of Illinois at Urbana-Champaign
Champaign, IL 61820
{jiawang4, kcchang, hs1}@illinois.edu

ABSTRACT

Real-world networks are often complex and large with millions of nodes, posing a great challenge for analysts to quickly see the big picture for more productive subsequent analysis. We aim at facilitating exploration of node-attributed networks by creating representations with conciseness, expressiveness, interpretability, and multi-resolution views. We develop such a representation as a *map*—among the first to explore principled network cartography for general networks. In parallel with common maps, ours has landmarks, which aggregate nodes homogeneous in their traits and interactions with nodes elsewhere, and roads, which represent the interactions between the landmarks. We capture such homogeneity by the similar roles the nodes played. Next, to concretely model the landmarks, we propose a probabilistic generative model of networks with roles as latent factors. Furthermore, to enable interactive zooming, we formulate novel model-based constrained optimization. Then, we design efficient linear-time algorithms for the optimizations. Experiments using real-world and synthetic networks show that our method produces more expressive maps than existing methods, with up to 10 times improvement in network reconstruction quality. We also show that our method extracts landmarks with more homogeneous nodes, with up to 90% improvement in the average attribute/link entropy among the nodes over each landmark. Sense-making of a real-world network using a map computed by our method qualitatively verify the effectiveness of our method.

1. INTRODUCTION

Over the last decades, the size of datasets for real-world networks has expanded by orders of magnitude—While this data deluge has blessed our building more robust analytics, it has challenged our ability to explore these large networks with intuition. Often, a network (e.g., Facebook) captures real-world information of entities as *nodes* on the network, with *attributes* (e.g., gender, age) describing the properties of and *edges* indicating the interactions (e.g., friend-link between users) between entities. Such attributed, unweighted networks proliferate in many domains, e.g., social (e.g., Facebook, Twitter) and biology (e.g., protein-protein

interaction) networks. We thus address the problem of *exploring large, real-world networks by creating compact representations that reveal the essential interactions in the data*, at desirable resolutions to see from the “trees” to the “forest” on the network.

We believe a compact representation for effectively exploring a network should satisfy the following requirements. *Conciseness*: The representation should preserve essential attribute and link patterns, while minimizing uninteresting detail, of the network. *Expressiveness*: It should be able to encapsulate a diversity of characteristic patterns that may exist in the network. *Multi-resolutions*: It can represent the network at various controllable levels of details. *Locatability*: It should capture not only the big picture but also every node on the network—i.e., we can locate each individual node on the representation.

Network cartography. The metaphor of a *map* satisfies our requirements. Maps, by definition of cartography [13], are “*graphic representations that facilitate a spatial understanding of things, concepts, conditions, processes, or events in the human world.*” In particular, for navigating the physical world, we are familiar with common geographic maps with “landmarks” and their spatial connections of “roads.” Such geographic maps are necessarily concise and expressive of all kinds of terrains, and allows zooming to various resolutions and locating every point of interest. Thus, for our objective, to compactly represent a large network, a *network map* should comprise *landmarks*, representing a group of individual nodes, and *roads*, abstracting the interactions between these landmarks—We aim to study this cartography of networks.

While the needs are pressing, unfortunately, there has been little prior work for network maps. Rosvall et al. [21] and Guimera et al. [11] made the early attempts to create maps to summarize complex networks. However, they either assume a specific domain (metabolic networks) or restrict to only link topologies. To our knowledge, our study is the first to create maps for general networks of both attribute and link information. We informally state our problem as follows.

Problem: For an attributed, unweighted network G , given the number of landmarks K , create a *map* for G as a graph \mathcal{M}_G with K nodes to represent G .

Challenge: *Discovering landmarks.* What is a landmark on a network? While the important landmarks in the familiar geographical map have emerged due to centuries of cartographic practice, the semantics of the landmarks in a map designed to compactly represent a network are unclear.

Insight: *From homogeneity to roles.* We observe that landmarks on a map signify *homogeneity* of points of interests—E.g., on a geographic map, a landmark (such as “Navy Pier”) is an aggregate representation of points (shops, restaurants, piers) sharing similar

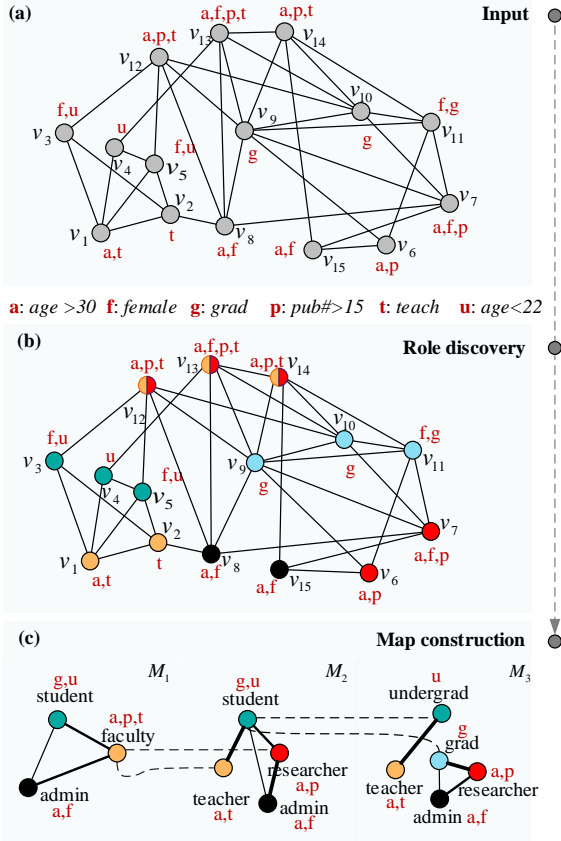


Figure 1: Example network G_{ex} with node attributes (in red) representing the people and their interaction in an academic department. The middle figure colors each node by its role(s) resulted from role discovery. The bottom set of figures show a sequence of three maps with increasing resolution constructed from the discovered roles, the relationships between the finer roles and those of the previous map are shown in dashed lines; each role is also labeled by the attributes it will likely have.

properties (at very close spatial coordinates) and similar interconnections with other landmarks. Thus, in a network, a landmark should aggregate nodes that are similar, in terms of their attributes and interactions. Such homogeneous nodes, in *sociology*, define *roles* [4] in an interconnected society, as individuals that share characteristic patterns of behaviors. We thus propose to discover roles in a network as landmarks on its map.

To illustrate, Figure 1 shows an example network G_{ex} in part (a) and the roles (as discovered via our algorithm) of the network in part (b). In G_{ex} , each node is a person in an academic department, with attributes such as *a*: “age over 30” and *t*: “teach” for node v_1 . While G_{ex} , a toy network of 15 nodes, already appears “complex,” its nodes share common patterns of behaviors. To reveal, we annotate the network with the implicit roles in part (b) by coloring the nodes— nodes with a color share a certain behavior pattern. Now, we can observe that people in orange-role $R_3 = \{v_1, v_2\}$ “teach and interact with green-role $R_4 = \{v_3, v_4, v_5\}$.” And, the latter group comprises young people. The red-role $R_1 = \{v_6, v_7\}$ “does not teach, has many publications, are aged over 30, and interacts with blue role $R_2 = \{v_9, v_{10}, v_{11}\}$ ”. A plausible theory is that R_4 is *undergrad* and R_2 *grad* students; then, R_1 and R_3 are perhaps *researcher* and *teacher*. Further, with roles as landmarks, in part

(c), we can construct maps of different resolutions, e.g., \mathcal{M}_1 with three roles and \mathcal{M}_3 five roles.

Challenge: Formulating roles for landmarks. The challenge of formulating roles for landmarks lies in properly capturing the desired homogeneity. While there exist diverse formulations [8, 1, 7], they fall short for landmark homogeneity, as Section 2 will further discuss. First, they are too restricted for real-world networks, as none of them captures overlapping roles and inter-role interactions. Second, they address only links but not attributes. We thus need to develop formulation for coherently integrating the dual-mode behaviors in potentially overlapping roles.

Insight: Probabilistic dual-mode regular equivalence. First, we conceptually capture link homogeneity by regular equivalence (RE) [5], defined as “two nodes are regularly equivalent if they are equally related to equivalent others”. This definition, however, is too restrictive for real-world networks. Our model for role extraction probabilistically captures the essence of RE to handle noisy real-world data, and incorporates overlapping roles and non-uniform inter-role interactions to flexibly express the variations in complex networks. Second, we inject attributes as the second mode of the role behavior, in addition to links. Thus, a role represents both link and attribute patterns, enforcing homogeneity in both modes.

Challenge: Multi-resolution continuum. As maps should support multi-resolutions, maintaining an intuitive continuum between resolutions is challenging. On a geographic map, when we zoom into a region (e.g., Chicago), the changes are limited to *within* the region (we may now see Millennium Park), and the interconnections to the rest of the world should be stable (the whole region is still located in the state of Illinois). Such continuum is intuitive for network maps too. On map \mathcal{M}_1 (Figure 1c), when we zoom into landmark *faculty*, it gets refined into two finer landmarks in \mathcal{M}_2 : *teacher* and *researcher*, while the other landmarks are not changed; later, when we zoom into *student*, resulting in *undergrad* and *grad* in \mathcal{M}_3 , the continuum maintains. While desired, realizing such continuum is challenging— Since our landmarks are determined *inter-dependently*, changing one landmark will necessarily change others.

Insight: Multi-scale constraints. To maintain continuum in multi-resolution views, as we use roles for landmarks, we must enforce the semantics of continuum in role discovery. As we model role discovery as a likelihood maximization problem, we enforce the continuum as an additional constraint in the objective function, which thus enables a principled realization of multi-resolution stability.

Our solution. We summarize our framework of constructing a map for a given network G as follows, as Figure 1 shows.

- **Role discovery.** Compute θ as the *parameters of a model* that describes how the observed network G is generated probabilistically from the latent roles (Figure 1b).
- **Map construction.** Use θ to compute and construct the map \mathcal{M}_G as a *node-labeled weighted graph*: each node of \mathcal{M}_G represent a role, the weight of the edge between two roles represent how likely they will interact, and each node is labeled with how likely that role will have each attribute (Figure 1c).

Specifically, we propose a new probabilistic, generative model for networks to extract roles as landmarks of the map (Section 3). The model explicitly parameterizes each node-role affiliation pair and each role-role interaction pair. The same role affiliation parameters will be the explainers in the logistic regression models for the node attributes. We design an objective function to optimize

the likelihood of the observations under the model; our objective function formulation also encourages solution sparsity. Then, we process the model parameters into more intuitive quantities to construct the map, and achieve a multi-scale view by constraining the new solution to maintain the relationships and landmarks from the coarser scale, while determining the new landmarks at the finer scale and their relationship with the landmarks at the coarser scale (Section 4). We use an efficient block-coordinate gradient ascent algorithm to perform the optimization in linear time w.r.t the network size (Section 5).

In experiments, compared with baseline methods, ours shows significant improvement in model expressiveness measured by network reconstruction quality—up to 10 times for synthetic networks and 27% for real-world networks. In terms of homogeneity, our method outperforms existing methods by up to 90% in terms the attribute/link entropy among the nodes over each landmark. In a case study using real-world networks, a set of multi-resolution maps created by our method actually help us to easily see the salient structures of the network from complex observed interactions.

Contributions. We summarize our contributions as follows.

- *Network cartography through role discovery.* We are the first to propose the use of roles to serve as map landmarks for attributed networks; in this, we are inspired by earlier work on regular equivalence [5]. The landmarks along with their attribute distribution and the resulting interrelationships between landmarks result in meaningful maps.
- *A probabilistic model for the map.* We propose an expressive, probabilistic generative model of the network with roles as the latent factors. The expressive nature of the model allows us to model a wide variety of attributed networks. During the optimization process, we enforce sparsity constraints to achieve concise maps.
- *Simultaneous multiple resolution.* We are able to create maps that simultaneously show different parts of the map at different resolutions.
- *Scalable cartography.* Our algorithm runs in linear time (with respect to network size). Furthermore, it is easy to implement using standard optimization toolboxes.
- *Extensive experiments.* We not only synthesize networks with diverse structures to verify the effectiveness of our method in a wide range of networks, but also use real-world networks to show the practicality of our method, evaluated using a variety of measures. Case studies further evaluate our method’s usability in action.

2. RELATED WORK

Network cartography. Cartography [13] has a long history and was intensively used in visualizing geographic data, but their use in exploring large general networks has received attention only recently. Rosvall et al [21] creates map of sciences by revealing community structures through efficiently encoding of random walk paths. Guimera et al [11] used universal roles detected from optimizing network modularity to create maps for metabolic networks.

Attributed networks. We model node-attributed networks, for which alternative models with different goals have been proposed; e.g. BAGC [26] uses a Bayesian approach, CESNA [27] models overlapping communities, and MAG [16] models clusters as generated from attributes. Methods for graph clustering using both structure/attribute information have also been proposed [29, 22]. Issues of these methods will be discussed later in this section.

Role theory. We create maps with landmarks as roles. The theory of roles have been studied in sociology [4] where roles are defined as “a set of connected behaviors, rights, obligations, beliefs, and norms as conceptualized by people in a social situation” [10], and various formulations have been proposed [6, 8, 5, 18]. For the alternative formulation of universal roles, which does not fit in our picture, see survey [20].

Intergroup interactions. Our model accounts for intergroup interactions in contrast with only intra-group found in communities. Some other methods also accounts for this aspect but with limitations for creating maps. Mixed-Membership Stochastic Block-modeling (MMB) [1] models each observed edge of each node as generated from one of its roles, but it does not model the intensity of role affiliation, as also pointed out in [27, 1]. (BAGC [26] suffers from similar issues) as MMB. In fact, MMB cannot model hub and peripheral nodes, which in our model can be represented by nodes with strong affiliations to other roles. MMB does not model node attributes as key to information-enriched maps either. In terms of goals, MMB focus more on generalizability of model while we focus on summarization quality in computing maps. Generalized Blockmodeling [7] does not capture varying interaction intensities between blocks (roles). DEDICOM [2] uses matrix factorization for real-valued data, thus cannot model binary links; other NMF-based methods [25, 19] suffers from similar issues. The model by Han et al [12] uses separate parameters to explain node attributes and links.

Overlapping groups. Models for overlapping groups of nodes has less been studied. DEDICOM [2] and CESNA [27] are two examples. Note that mixed memberships [1, 26] do not sufficiently capture the concept of overlapping groups, because their models only model the probability, rather than the intensity, of the affiliation of nodes to the groups.

Multi-resolution views. Local multiple-resolution views implies a hierarchical structure of the data (zooming corresponds to branching in a hierarchy). Methods for hierarchical clustering [9] cannot be directly applied to our case in which groups are overlapping. For graph summarization at different levels of granularity, Tian et al [23] proposed k -SNAP, but this operation does not support local views of part of the map.

3. ROLE DISCOVERY

In this section, we develop our probabilistic generative model of networks, with roles as latent factors, to extract the landmarks of a map.

Probabilistic regular equivalence. We first identify RE (regular equivalence) as the most appropriate semantics of link homogeneity for maps, compared to ROLX [15] and structural equivalence [6] (SE) as representatives of alternative formulations. ROLX defines roles as *universal* in terms of structural functions, i.e, whether it is a bridge connecting otherwise separated groups, or a hub connecting to many nodes, etc.; it misses the interaction patterns of *specific* networks, thus falls short in expressness. SE considers nodes equivalent (with same role) only if they connects to the same other *nodes* — too restrictive since there can be too many nodes for SE to find a small number of useful roles. On the other hand, RE is free from both issues by considering nodes equivalent if they connects to the same roles rather than the same nodes.

We make significant extension of RE with a probabilistic generative model, regarding roles as the latent factors that drives the noisy observed networks. For extracting latent factors [9], although matrix factorizations (such as NMF [17]) are also widely studied and

used, they have not yet been well studied for effectively handling binary data matrices. Since binary data carry special semantics, thus methods for factorizing general real-valued data cannot be applied. From an efficiency perspective, NMF for a general $\mathbb{R}^{N \times N}$ data matrix (adjacency matrix in our case) needs $O(N^2)$ time [17], which is prohibitive for large networks; the NMF-based algorithm ASALSAN [2] specialized for graph processing has linear complexity but can only be used with weighted links. Probabilistic generative models, on the other hand, are a much more expressive class of tool which model various types of observations, including binary observations, e.g. by Bernoulli models. As shown later, an efficient algorithm can also be derived from a well designed probabilistic model.

Setting. In this paper, we consider the extraction of K landmarks given an undirected and unweighted graph $G = (\mathcal{V}, \mathcal{E}, \mathcal{A})$ of N nodes and E links, where \mathcal{V} represents the node set and \mathcal{E} the edge set (directed graphs can be handled with slight modification of our method as described later; handling weighted graphs will be discussed till the end of the paper). Each node v is labeled by L binary attributes $a_{v,1}, \dots, a_{v,L} \in \{0, 1\}$, and we use \mathcal{A} to denote the set of observed node-attribute pairs $\{(v, i) : a_{v,i} = 1\}$.

Regarding the notation, we will use boldface capital letters like \mathbf{X} for matrices, calligraphic capital letters like \mathcal{X} for sets, boldface lowercase letters like \mathbf{x} for vectors, and lowercase letters for scalars. The roles will be indexed by k and the attributes by i .

3.1 Model Description

We model the landmarks in a map as roles. In the example network G_{ex} in Figure 1, we observe group $g_1 = \{v_6, v_7\}$ “often interact with ($g_2 = \{v_9, v_{10}, v_{11}\}$)” and “have more than 15 publications” but “do not teach courses”. A plausible theory would be that g_1 are research professors and g_2 graduates, and g_1 ’s behaviors can be explained by the *pure* role of “researcher”. We also observe “aged over 30” and “frequent interactions with group $g_3 = \{v_3, v_4, v_5\}$ ” for group $g_4 = \{v_1, v_2\}$; prior knowledge tells us g_4 are lecturers and g_3 are undergrads, and g_4 ’s behaviors are due to the pure role of “teacher”. Other roles include “admin”, “undergrad”, and “graduate”, as extracted in map \mathcal{M}_3 in Figure 1.

Plural and non-uniform role affiliation. To characterize the roles, we observe that each node can participate in *multiple* roles at different intensities. First, role affiliation is *plural*. In G_{ex} , group $g_5 = \{v_{12}, v_{13}, v_{14}\}$ (known as professors in prior) do both teaching and research and interact with both undergrad and graduate students; these behaviors are explained by their simultaneous roles “teacher” (abbrv. *tea*) and “researcher” (abbrv. *res*). Second, the affiliation is *non-uniform*. Note that professor v_{12} might not conduct role “teacher” so much as lecturers as indicated by less interactions with undergrads.

To capture the plurality and non-uniformity of role affiliation, we represent the affiliation of each node v to a set of K roles by vector $\mathbf{x}_v = (x_{v,1}, \dots, x_{v,K})$, where $x_{v,k}$ quantifies how strongly v is affiliated to the role k ; we also require $\mathbf{x}_v \geq 0$ for easier interpretation of the results. For instance, in G_{ex} , the lecturer v_1 will have dominantly large $x_{1,tea}$ compared to other components of \mathbf{x}_1 due to a pure role of “teacher”, while the professor v_{12} who does research but less teaching will have $x_{12,res} > x_{12,tea} > 0$. Stacking \mathbf{x}_v^T form the *node-role membership matrix* $\mathbf{X} \in \mathbb{R}_{\geq 0}^{N \times K}$.

Non-uniform inter-role interaction. We observe that inter-role interactions are also *non-uniform* in intensities. In G_{ex} , although $g_6 = \{v_8, v_{15}\}$ (married female aged over 30, possibly administrative staff) interact with both g_5 (professors) and g_2 (graduates),

administrative staff usually interact more with professors, because they often assist the professors in scheduling events.

To reflect the non-uniformity of inter-role interactions, we quantify the tendency for two nodes of roles k and l to interact by $r_{k,l} \geq 0$. We then define the *inter-role interaction matrix* as $\mathbf{R} = (r_{i,j}) \in \mathbb{R}_{\geq 0}^{K \times K}$. For instance, in G_{ex} , the strong interaction between professors and graduates (abbrv. *gra*) will be captured by $r_{res,gra}, r_{gra,res} > 0$ and possibly also $r_{gra,res} > r_{res,gra}$ as graduates usually more actively turn to professors.

Based on the above notions, we model our belief for two nodes u, v interact as follows. We know that if nodes u, v are affiliated to roles k and l respectively and roles k and l tends to interact, nodes u and v also tends to interact. In addition, our belief for (u, v) being formed will add up as we observe more such evidences from other pairs of roles k', l' . For simplicity, we assume that our belief grows linearly with the role affiliation and inter-role interaction intensities. Thus we formulate this belief, the *predictor* of link (u, v) , as

$$\rho_{u,v} := \sum_{k,l \in [K]} x_{u,k} \cdot r_{k,l} \cdot x_{v,l} = \mathbf{x}_u^T \mathbf{R} \mathbf{x}_v$$

Since our observations are binary, we need a *link function* that transform $\rho_{u,v}$ into probability values. Observing that the marginal potential of forming links normally decrease as more evidences are present, we model the probability of forming link (u, v) with the predictor $\rho_{u,v}$ as

$$p_{u,v} := \varphi(\rho_{u,v}) = 1 - \exp(-\mathbf{x}_u^T \mathbf{R} \mathbf{x}_v) \quad (1)$$

While another common choice for transforming predictors to have range $(0, 1)$ would be the sigmoid function $\sigma(\cdot)$, it has domain $(-\infty, \infty)$ but $\rho_{u,v}$ is restricted to be nonnegative, hence inappropriate for our case. We also support our use of $\varphi(\cdot)$ by efficiency considerations as we will discuss in Section 5.

Links and attributes as covariates. Finally, we observe that the links and attributes of each node are *covariates* with its roles as the underlying factor. In G_{ex} , the attribute “age > 30” and interactions of “being approached by graduate students” of the professor v_{12} can both be explained by a “researcher” role. Similarly, the attribute “female” and the interactions of “being contacted by professors” can be explained by a “admin” role.

To capture the role-based covariation of attributes and links, we will use an attribute predictor as the linear combination of the role affiliation vector \mathbf{x}_v and some parameter \mathbf{w}_i to model our belief of node v to have attribute i (for simplicity, we assume that the attributes of a node are conditionally independent given the roles of that node). For example, in G_{ex} , the role “researcher” will indicate a high weight $w_{res,(pub>15)}$ so that it would contributes positively and strongly to observing many publications for professors. Then, we transform the predictor to probability using a link function. For the above purpose, a standard logistic regression would suffice, by which we specify the attribute predictor as $\mu_{v,i} := \mathbf{w}_i^T \mathbf{x}_v$ and the probability of observing $a_{v,i} = 1$ as

$$q_{v,i} := \sigma(\mu_{v,i}) = \frac{1}{1 + \exp(-\mathbf{w}_i^T \mathbf{x}_v)} \quad (2)$$

Note the intercept term is ignored in the predictor to simplify our presentation. We denote $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_L)$.

Now, our problem of discovery roles as landmarks of the map for a given network is defined as follows.

DEFINITION 1 (ROLE DISCOVERY). *Given an undirected unweighted graph G with node attributes and the number of roles K ,*

we will find a set of roles represented in the following parameters: the matrix \mathbf{R} for inter-role interactions, a vector \mathbf{w}_i for each attribute i for predicting the attributes of any node v given the roles of v , and a vector \mathbf{x}_v for each node v for representing the roles affiliations of v .

Model extensibility. Although binary attributes are discussed here, the model can be easily adapted to categorical attributes (e.g. occupation) using a multinomial logistic model in place of the binary model described earlier. The model also directly applies to directed graphs which will result in an asymmetric matrix \mathbf{R} . To support weighted graphs, we can modify the model by replacing $\varphi(x)$ in (1) with an identity function, while the algorithm needs to be designed differently (e.g. the ASALSAN algorithm [2] with some adaptation), for which we will not discuss in more details here.

3.2 Objective Function

We extract the roles from network G by optimizing the parameters of our model of G in the following aspects. First, we optimize the parameters such that they *accurately express* the observations in the original network G using a set of roles. Since the number of landmarks K is usually small for concise maps, it implies that each landmark must be homogeneous for the observations to be represented compactly using the roles. In particular, we capture the accuracy of G by how well our model can be used to reconstruct G . Second, we optimize for *sparsity* of the parameters for conciseness of representation. That is, we wish G be represented using as few inter-role interactions as possible and each node can be described by their most important roles, such that users can focus on the most salient patterns of G .

To encode the above requirements, we propose an objective function $f(\theta)$ with respect to parameters $\theta = \{\mathbf{X}, \mathbf{R}, \mathbf{W}\}$. θ captures model accuracy by the log-likelihood of observing G under a model with parameters θ , while the sparsity requirement is enforced by the ℓ_1 -norm regularizations of \mathbf{R} and \mathbf{X} . Formally,

$$f(\theta) = (1 - \alpha)\ell_{\mathcal{E}} + \alpha\ell_{\mathcal{A}} - \alpha_R\|\mathbf{R}\|_1 - \alpha_X\|\mathbf{X}\|_1 \quad (3)$$

where $\ell_{\mathcal{E}}$ and $\ell_{\mathcal{A}}$ are the log-likelihoods of links and attributes given by

$$\ell_{\mathcal{E}} = \sum_{(u,v) \in \mathcal{E}} \log \varphi(\rho_{u,v}) + \sum_{(u,v) \notin \mathcal{E}} \log(1 - \varphi(\rho_{u,v})) \quad (4)$$

$$\ell_{\mathcal{A}} = \sum_{(v,i) \in \mathcal{A}} \log \sigma(\mu_{v,i}) + \sum_{(v,i) \notin \mathcal{A}} \log(1 - \sigma(\mu_{v,i})) \quad (5)$$

and weight α is used to impose emphasis on links or attributes, and α_X, α_R to control the degree of representation sparsity.

In summary, the optimal parameters θ^* , which maximizes $f(\theta)$ subject to the constraints that \mathbf{X} and \mathbf{R} , will represent our extracted roles. That is,

$$\theta^* = \arg \max_{\mathbf{X} \geq 0, \mathbf{R} \geq 0, \mathbf{W}} f(\theta) \quad (6)$$

The detailed algorithm for optimizing $f(\theta)$ will be introduced later in Section 5.

4. MAP CONSTRUCTION

Thus far we have encoded the information needed to create a map for the given network in the model parameters θ . In this section, we will develop methods to render θ as an intuitively interpretable map. Further, we propose methods to allow for local multiple-resolution views of any part of the map.

4.1 From Model to Map

The computed model parameters θ^* in the stage of role discovery needs to be transformed to interpretable quantities to construct a map. For instance, while a larger $x_{v,k}$ apparently indicates the stronger affiliation of node v to role k , $x_{v,k}$ does not have an intuitive physical meaning, such as probability, length, or frequency which are natural for human interpretation. The same issue also holds for parameters \mathbf{x}_v and \mathbf{w}_i . This makes it non-intuitive, if not deceptive, to make sense of a map constructed directly from θ^* (e.g. using \mathbf{R} as the weighted adjacency matrix for the metagraph).

In order to construct maps as metagraphs with interpretable weights and labels, we introduce a *virtual node* v_k for each role k , which is an *imaginary* node affiliated purely to role k , i.e., $\mathbf{x}_{v_k} = c\hat{\mathbf{e}}_k$ for some $c > 0$, where $\hat{\mathbf{e}}_k$ is the unit vector along the k th dimension, and c controls the affiliation strength of v_k to role k . Intuitively, the virtual node v_k can be viewed as an idealized representative for role k so that its behaviors as expected from our model naturally characterize role k . We might choose the value of parameter c as follows. First, we may set c to some constant independent of the actual affiliations of the nodes to role k , thus in effect characterize role k itself by v_k . Second, we may set $c = \bar{x}_{\cdot,k} = (\sum_{v \in \mathcal{L}(k)} x_{v,k}) / |\mathcal{L}(k)|$, where $\mathcal{L}(k) = \{v : x_{v,k} > 0\}$, as the average affiliation of nodes with role k ; such created virtual node v_k will instead represent the role k as actually practiced by the real nodes in data.

Using the virtual node v_k as a representative for role k , we characterize the behaviors of role k intuitively by the probability of observing each attribute i on v_k , given by

$$\psi_{k,i} := \sigma(\mathbf{w}_i^T c\hat{\mathbf{e}}_k) = (1 + \exp(-c \cdot w_{k,i}))^{-1}$$

Similarly, to characterize the inter-role interaction between each pair of roles k and l , we use the probability of interaction between an virtual node v_k with purely role k and a virtual node v_l with purely role l , which we give by

$$\omega_{k,l} := \varphi(c\hat{\mathbf{e}}_k^T \mathbf{R} c\hat{\mathbf{e}}_l) = 1 - \exp(-c^2 r_{k,l})$$

For later reference, we denote $\mathbf{P}_{attrib} = (\psi_{k,i})_{(k,i) \in [K] \times [L]}$ and $\mathbf{P}_{link} = (\omega_{k,l})_{(k,l) \in [K]^2}$. Such transformed θ will be finally visualized as the output map.

In summary, we formally define the steps for creating a map as follows.

DEFINITION 2 (MAP CONSTRUCTION). *Given an unweighted network G with L binary node attributes, a map for G with K landmarks as a metagraph \mathcal{M}_G is constructed by*

1. Compute $\theta^* = \arg \max f(\theta)$ given G ;
2. Use θ^* to compute $\psi_{k,i}$ for each $(i, k) \in [L] \times [K]$ and $\omega_{k,l}$ for each $(k, l) \in [K]^2$;
3. Construct graph $\mathcal{M}_G = (\mathcal{R}, \mathcal{I}, \psi_{\cdot, \cdot}, \omega_{\cdot, \cdot})$: each node in \mathcal{R} represents a landmark, each edge $(k, l) \in \mathcal{I}$ is weighted by $\omega_{k,l}$, and each node $k \in \mathcal{R}$ is labeled with $\psi_{k,i}$ for each attribute i .

4.2 Local Multi-Resolution Views

With an overview of the network by a coarser map, we allow users to select *part* of the map for higher-resolution views in the same global context, as in map \mathcal{M}_2 in Figure 1 where we selected the landmark “student” to zoom in, and revealed the finer-granular landmarks “undergrad” and “graduate” in \mathcal{M}_3 , while other landmarks remains unchanged, and the finer-granular connections between the new landmarks and the existing landmarks are revealed too, such as that “graduate” interacts more with professors than “undergrad”.

As a first attempt to changing the resolution locally, one might extract a subgraph for the selected landmark and compute a map with that subgraph as input, but this method loses the connections between the newly created landmarks with those existing outside the subgraph. One might also attempt the alternative way of computing from scratch a new map with more landmarks, but it violates our locality requirement and destructs the existing non-selected landmarks. Although hierarchical clustering also produces finer-granular views of data, it does hard partitioning of the data points while our groups of nodes are overlapping; it also considers only intra-cluster connections but not the connections between the clusters.




Figure 2: Obtaining $\tilde{\theta}_p^*$ from θ_p^* to initialize θ_n^*

We achieve local multiple-resolution views for some role k' in a given map \mathcal{M}_p constructed with θ_n^* by constructing a new map \mathcal{M}_n with parameters θ_n^* , and constraining our optimization of θ_n^* such that the non-selected roles (those other than k') in \mathcal{M}_p remain little changed. We measure the change between maps \mathcal{M}_p and \mathcal{M}_n by the distance between their respective parameters θ_p^* and θ_n^* with respect to the non-split roles. Formally, we define the distance between models θ_n and θ_p using sum of squared errors as

$$d_{map}(\theta_p, \theta_n) = \|\text{reduce}_{k_1', k_2'}(\theta_n) - \text{reduce}_{k'}(\theta_p)\|_2^2$$

where $\text{reduce}_{k'}(\theta)$ (resp. $\text{reduce}_{k_1', k_2'}(\theta)$) transforms θ by removing its components associated with role k' (resp. k_1', k_2'), for which we illustrate the details in Figure 2. Using the operator $\text{reduce}_{k'}(\cdot)$, we actually measure the divergence between the maps only with respect to roles other than the selected k' .

In summary, we formally define the problem of computing local multi-resolution views on a map as follows.

DEFINITION 3 (LOCAL MULTI-RESOLUTION VIEWS). *Given a previously computed map \mathcal{M}_p constructed from θ_p^* with K roles, and a specified landmark k' to view at finer resolution, we compute a map \mathcal{M}_p with model θ_n^* of $K + 1$ roles by splitting role k' into the sub-roles k_1' and k_2' , such that θ_n^* satisfies our requirements for appropriate landmarks by maximizing $f(\theta_n^*)$, meanwhile keeps the roles other than k' relatively unchanged by minimizing $d_{map}(\theta_p^*, \theta_n^*)$*

Note that while we demonstrate the case of two sub-roles, the algorithm can be easily generalized to split any role k' into multiple sub-roles.

In order to find θ_n^* that both maximizes $f(\theta_n^*)$ for quality landmarks and minimizes $d_{map}(\theta_p^*, \theta_n^*)$ for locality of the finer view, we optimize θ_n^* under a revised objective function

$$f_z(\theta_n; \theta_p^*) := f(\theta_n) - \beta \cdot d_{map}(\theta_p^*, \theta_n) \quad (7)$$

where β is the hyperparameter controlling the tradeoff between the landmark quality and the locality of the resolution change. For the revised optimization objective, we will design an efficient algorithm in Section 5.

5. ALGORITHMS

In this section, we develop efficient algorithms to optimize the objective functions formulated in Section 3 for role discovery and Section 4 for computing local multi-resolution views in selected landmarks.

Algorithm for landmark extraction. Given the optimization objective function $f(\theta)$ formulated in (3), we derive the gradients of $f(\theta)$ with respect to variables $\mathbf{X}, \mathbf{R}, \mathbf{W}$ as follows. For easier derivation, we split the matrix variables \mathbf{X} and \mathbf{W} into the vector variables \mathbf{x}_v and \mathbf{w}_b . Then

$$\frac{\partial f}{\partial \mathbf{R}} = (1 - \alpha) \left(\underbrace{\sum_{(u,v) \in \mathcal{E}} \frac{1 - \varphi(\rho_{u,v})}{\varphi(\rho_{u,v})} \mathbf{x}_u \mathbf{x}_v^T}_{\text{present links}} - \underbrace{\sum_{(u',v') \notin \mathcal{E}} \mathbf{x}_{u'} \mathbf{x}_{v'}^T}_{\text{absent links}} \right) - \alpha_R \cdot \text{sign}(\mathbf{R}) \quad (8)$$

$$\frac{\partial f}{\partial \mathbf{x}_v} = (1 - \alpha) \left(\underbrace{\sum_{(v,u) \in \mathcal{E}} \frac{1 - \varphi(\rho_{u,v})}{\varphi(\rho_{u,v})} \mathbf{R} \mathbf{x}_u}_{\text{neighbors of } v} - \underbrace{\sum_{(v,u') \notin \mathcal{E}} \mathbf{R} \mathbf{x}_{u'}}_{\text{non-neighbors of } v} \right) + \alpha \left(\sum_{(v,i) \in \mathcal{A}} (1 - \sigma(\mu_{v,i})) \mathbf{w}_i - \sum_{(v,i') \notin \mathcal{A}} \sigma(\mu_{v,i'}) \mathbf{w}_{i'} \right) - \alpha_X \text{sign}(\mathbf{x}_v) \quad (9)$$

$$\frac{\partial f}{\partial \mathbf{w}_i} = \alpha \left(\sum_{(v,i) \in \mathcal{A}} (1 - \sigma(\mu_{v,i})) \mathbf{x}_v - \sum_{(v,i') \notin \mathcal{A}} \sigma(\mu_{v,i'}) \mathbf{x}_v \right) \quad (10)$$

Note that the optimization problem (6) is *nonconvex*. The problem of optimizing the variable \mathbf{R} with all other variable fixed, however, is convex. This is also true for variables \mathbf{x}_v and \mathbf{w}_i . Thus we resort to the *block-coordinate gradient ascent* algorithm to optimize $f(\theta)$, where we cycle through the variables and optimize one of them at a time until convergence to a local minima. To optimize each variable, we employ the gradient ascent algorithm with the learning rate set proportional to the normalized gradient. After each gradient ascent update of any variable z subject to the non-negativity constraint, we need an additional step to ensure the updated z' still satisfies the constraint. To this end, we project the negative entries of z' to 0 after each iteration, denoted by $z' \leftarrow \text{proj}^+(z)$. The steps are summarized in Algorithm 1.

On Initialization. At the beginning of the optimization, we need to initialize each variable properly. Initializing the variables to small uniformly distributed random values around 0.1 turns out to work well in practice. An initialization with large random values, however, will typically lead to useless results with the algorithm stopping at an early stage. In addition, special attention must be paid to the \mathbf{R} when the given network G is undirected. In this case, \mathbf{R} needs to be initialized as some *symmetric* matrix. It can be easily shown that the symmetry property will then continue to hold throughout the optimization. In practice, setting \mathbf{R} to $(\mathbf{M} + \mathbf{M}^T)/2$ for some random matrix \mathbf{M} will do the work, while failures in confirming to symmetry will result in unexpected nonsense results in most cases.

Linear-time computation. The cost of Algorithm 1 concentrates on the evaluation of the log-likelihoods (4), (5) for evaluating the objective function and the gradients (8)–(10). To analyze the time for objective function evaluation, it amounts to analyzing the evaluation of $\rho_{u,v}$ for each link $(u, v) \in \mathcal{E}$ individually and the sum of

Algorithm 1: NETLAND

Input : Network G , number of landmarks K
Output: Landmarks $\theta^* = (\mathbf{X}^*, \mathbf{R}^*, \mathbf{W}^*)$

```

// initialization
1 Initialize entries of  $\mathbf{X}, \mathbf{W}$  with small random values;
2 Initialize  $\mathbf{R}$  with  $(\mathbf{M} + \mathbf{M}^T)/2$  for some matrix  $\mathbf{M}$  with
  small random entries;
3 repeat
  // evaluate gradients with
  // precomputed  $\tilde{\mathbf{x}}$ 
4  $\tilde{\mathbf{x}} \leftarrow \sum_v \mathbf{x}_v$ ;
5 Repeat  $\mathbf{R}^{(t+1)} \leftarrow \text{proj}^+(\mathbf{R}^{(t)} + \beta^{(t+1)} \frac{\partial f}{\partial \mathbf{R}^{(t)}})$ ;
  // enforcing symmetry of  $\mathbf{R}$  by
  // projection
6 if  $G$  is undirected then
7    $\mathbf{R} \leftarrow (\mathbf{R} + \mathbf{R}^T)/2$ ;
8 for each node  $v$  do
9   Repeat  $\mathbf{x}_v^{(t+1)} \leftarrow \text{proj}^+(\mathbf{x}_v^{(t)} + \beta^{(t+1)} \frac{\partial f}{\partial \mathbf{x}_v^{(t)}})$ ;
10  Update  $\tilde{\mathbf{x}}$  incrementally due to new  $\mathbf{x}_v$ ;
11 for each attribute  $i$  do
12   Repeat  $\mathbf{w}_i^{(t+1)} \leftarrow \mathbf{w}_i^{(t)} + \beta^{(t+1)} \frac{\partial f}{\partial \mathbf{w}_i^{(t)}}$ ;
13 until  $f$  not improved;

```

$\rho_{u',v'}$ for all absent links $(u', v') \notin \mathcal{E}$.

Simply evaluating $\rho_{u,v}$ for each pair of vertices is prohibitive. While for each observed link (u, v) , the term $\rho_{u,v}$ takes $O(K^2)$ time for evaluating $\mathbf{x}_u^T \mathbf{R} \mathbf{x}_v$, there are $O(N^2)$ separated pairs of nodes, for which computing $\rho_{u',v'}$ once for each such pair is impossible.

Fortunately, the sum of $\rho_{u',v'}$ for missing links can be improved by first rewriting it into

$$\sum_{(u',v') \notin \mathcal{E}} \rho_{u',v'} = \sum_{(u,v) \in \mathcal{V}^2} \mathbf{x}_u^T \mathbf{R} \mathbf{x}_v - \sum_{(u,v) \in \mathcal{E}} \rho_{u,v}$$

We observe that the first term on the right side can be simplified into $\tilde{\mathbf{x}}^T \mathbf{R} \tilde{\mathbf{x}}$ where $\tilde{\mathbf{x}} = \sum_{v \in \mathcal{V}} \mathbf{x}_v$. The value of $\tilde{\mathbf{x}}^T \mathbf{R} \tilde{\mathbf{x}}$ can be precomputed in $O(NK + K^2)$ time and cached before optimizing \mathbf{R} in each outer iteration. The total time for evaluating $f(\theta)$ is then $O((N + E)K^2)$. The costly last term in (8) can be similarly computed.

With the precomputation technique, the time for a full (outer) iteration of the algorithm scales *linearly* w.r.t the size of G . In details, each gradient ascent step for \mathbf{R} takes $O((N + E)K^2)$ time, each \mathbf{x}_v takes $O(d_v K^2 + LK)$ time where d_v is the degree of node v , and each \mathbf{w}_i takes $O(NLK)$ time. Summing up, the overall time complexity of Algorithm 1 is $O((N + E)K^2 + NLK)$; the memory needed is $O(NK + K^2 + KL)$ proportional to the size of θ . The performance in practice can be further boosted using optimized matrix operation libraries.

Note that the precomputation technique is valid for our choice of $\varphi(x)$ as the link function but *not* for the logistic function $\sigma(x)$. With $\sigma(x)$, term $\rho_{u,v}$ has to be evaluated *once for each* such link. In fact, with $\sigma(x)$, term $\sum_{(u',v') \notin \mathcal{E}} \mathbf{x}_{u'} \mathbf{x}_{v'}^T$ in (8) will be replaced by $\sum_{(u',v') \notin \mathcal{E}} \sigma(\rho_{u',v'}) \cdot \mathbf{x}_{u'} \mathbf{x}_{v'}^T$ which cannot be decomposed as the product of two cheap sums as in case of $\varphi(x)$ mean function. This results in prohibitive $O(N^2)$ complexity.

Algorithm for zooming. We create locally refined view θ_n^* with

Algorithm 2: Zoom

Input : Network G , landmarks θ_p^* , selected landmark k'
Output: Landmarks θ_n^* with refined k'

```

1  $\tilde{\theta}_p^* \leftarrow \text{reduce}_{k'}(\theta_p^*)$ ;
2 Append parameters for  $k'_1, k'_2$  to  $\tilde{\theta}_p^*$  and initialize them with
  small random numbers;
3 Initialize  $\theta_n$  by  $\tilde{\theta}_p^*$ ;
4 Compute  $\theta_n^* = \arg \max_{\theta_n} f_z(\theta_n; \theta_p^*)$  using Algorithm 1
  with adapted gradients

```

respect to the selected landmark k' in a map with parameters θ_p^* by optimizing $f_z(\theta_n; \theta_p^*)$ as formulated in (7). Let the sub-roles to be created be k'_1 and k'_2 . We derive the gradients of $f_z(\theta_n; \theta_p^*)$ as follows.

$$\frac{\partial f_z}{\partial \theta_n} = \frac{\partial f}{\partial \theta_n} - 2\beta \cdot \text{mask}_{k'_1, k'_2}(\theta_n - \tilde{\theta}_p^*)$$

where $\tilde{\theta}_p^*$ is the extended θ_p^* which will be used to initialize θ_n . In particular, we create $\tilde{\theta}_p^*$ from θ_p^* by removing its parameters for the role k' and appending new (randomly initialized) parameters for the sub-roles k'_1 and k'_2 ; see Figure 2 for an illustration of this process. The operator $\text{mask}_{k'_1, k'_2}(\theta_n)$ sets the parameters in θ_n corresponding to the sub-roles k'_1, k'_2 to 0. Intuitively, $\text{mask}(\cdot)$ keeps the newly created sub-roles from being penalized by its distance from $\tilde{\theta}_p^*$. The weight β controls the degree of the closeness enforcement.

We optimize θ_n^* using Algorithm 1 with the adapted gradients and also a special initialization of θ_n for *faster convergence*. In particular, we initialize θ_n to $\tilde{\theta}_p^*$ which inherits the roles other than k' , so that the algorithm's overhead will be limited to the optimization of the sub-roles. The algorithm for zooming is summarized in Algorithm 2.

6. EXPERIMENTAL EVALUATION

We will both quantitatively verify if the maps computed by our method have the desired properties of maps we identified in Section 1. We will also demonstrate the usability of our method by creating multi-resolution maps for real-world networks.

Datasets. The following real-world and synthetic datasets will be used in our tests. We chose the real networks so that they exhibit interesting strong intergroup interactions instead of only intra-group interactions found in communities.

- **Hospital**¹. This dataset [24] contains the directed network of contacts among 75 people in a hospital ward in France during five days. Each person is labeled with one of {patient, admin, nurse, doctor}. An arc is created for each pair of people if they interacted at least 5 times during the period.
- **Enron**². This dataset represents the directed network of email communications between 184 users. Each user is labeled by one or more the 32 topics (manually extracted in [3]) they ever used. We create an arc for each (ordered) pair of users if they communicated at least 15 emails.
- **SYN-X**. We synthesize θ as parameters of our proposed model with $K = 5$ landmarks, which in turn generates random undirected networks, such that the network corresponds

¹<http://www.sociopatterns.org/datasets>

²<http://cis.jhu.edu/~parky/Enron>




Figure 3: Convergence of NETLAND (Algorithm 1).

to one of $\{\text{bip}, \text{star}, \text{comm}, \text{rand}\}$ representing bipartite, star (wheel), community, and random structures of G_{map} (constructed from \mathbf{R}). Thus, a variety of common graph structures will be captured in the experiments. We then synthesize networks according our probabilistic generative model with parameter θ .

Methods. We will evaluate and compare the following methods.


- **NETLAND**. Our model of networks for creating maps. By default, we set the weight for attribute likelihoods as $\alpha = 0.5$ and sparsity regularization weights α_R, α_X to 0 when compared with other method and 0.2 otherwise.
- **CESNA** [28]. This the state-of-the-art model for large node-attributed networks. It aims at overlapping groups, models attributes as generated from group memberships, as well allows flexible community affiliation. By default, $\alpha = 0.5$ which has the same role as ours. In CESNA, K represents the number of communities, analogous to our landmarks.
- **BINDEDI**. The original DEDICOM [2, 14] model uses a similar form $\mathbf{A} \approx \mathbf{X}\mathbf{R}\mathbf{X}^T$ to model links, which allows arbitrary membership affiliation and pairwise intergroup interactions, but handles only weighted graphs and does not handle attributes. We non-trivially adapt it for binary links by applying the link function $\varphi(\cdot)$ to $\mathbf{X}\mathbf{R}\mathbf{X}^T$ and $\mathbf{A} \sim \text{Bernoulli}(\varphi(\mathbf{X}\mathbf{R}\mathbf{X}^T))$. The modified DEDICOM is equivalent to our model with $\alpha = 0$.

6.1 Comparing with Baselines

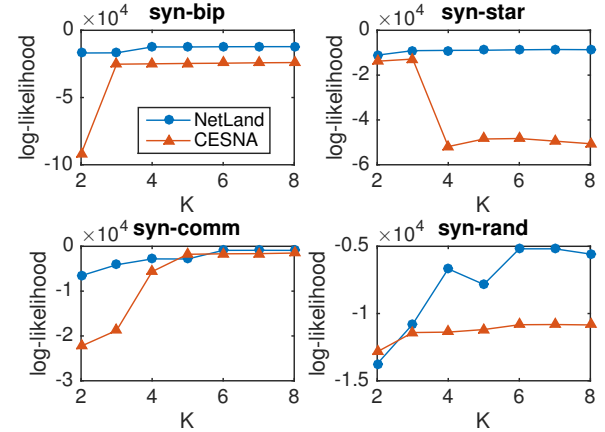
Convergence. We evaluate the convergence performance of Algorithm 1 as indicator for efficiency. Figure 3 shows how the value of the objective function $f(\theta)$ changes with the number of full iterations. We observe that the algorithm improves significantly at the first few iterations and converges quickly. Thus our algorithm is scalable given also the linear time complexity of each iteration.

Expressiveness. We measure a map’s expressiveness by quality of network reconstruction using the map θ^* , in terms of the log-likelihood of observing the network given the model θ^* underlying our map. In Figure 4(a), we show the log-likelihoods of the real-world networks at different K computed by NETLAND and CESNA (method BINDEDI is not considered here due to its inability to handle node attributes; it corresponds to NETLAND with $\alpha = 0$ and the effect of α will be evaluated later). We observe that our method outperforms CESNA by up to $\sim 27\%$ on *hospital* and $\sim 13\%$ on *enron* in reconstruction quality. For both methods, the log-likelihoods improve with larger K , as expected due to more parameters.

To evaluate the reconstruction quality of NETLAND on networks with diverse structural characteristics, we further use four synthetic networks with structures often found in real-world networks, as shown in Figure 4(b). For *syn-star*, we observe that NETLAND drastically outperforms CENSA by orders of magnitude. Obvious advantage of our method is also observed for *syn-bip* and



(a) Real-world networks



(b) Synthetic networks

Figure 4: Expressiveness: log-likelihood of network at varying K using CESNA and NETLAND on real-world and synthetic datasets.




Figure 5: Impact of α on log-likelihoods

syn-rand. In these networks, non-community patterns dominate or mix with other inter-group interactions — structures which CESNA cannot model. On the other hand, for *syn-comm*, our improvement is marginal as K increases since both methods are able to handle community structures (corresponding to an \mathbf{R} with only diagonal nonzero elements).

We also evaluate the impact of weight α which controls the trade-off between link and attribute reconstruction qualities. The log-likelihoods at different α for methods NETLAND and CESNA are shown in Figure 5 (CENSA also has α as the weight for attribute log-likelihoods). We observe that log-likelihoods peaks at $\alpha = 0.7$ for NETLAND on both networks, suggesting 0.7 as a possible starting point to set α generally. On both datasets, the performance drops quickly as α approaches 1 with an overemphasis on attributes. On the other hand, a small α has less negative impact on performance. This is possibly due to the more observations (N^2) for links than for attributes (NL observations) so that the decrease in attribute likelihoods less affects the overall. A higher weight α for attributes would impact more on the overall likelihood and make the attributes better expressed.

Homogeneity. The key property our method must satisfy is the

homogeneity of links and attributes among the nodes in each landmark. We evaluate the homogeneity by the *entropy* of the link and attribute distributions in each landmark.

Since each node can belong to multiple landmarks, we approximately measure the entropy of the overlapping groups by first grouping the nodes by their main roles $\mathcal{G}(k) = \{v : (\arg \max_l x_{v,l}) = k\}$. Let $\mathbf{p}_{k,i} = (f_{k,i}/n_k, 1 - f_{k,i}/n_k)$ be the distribution of attribute i in group k , where $f_{k,i} = |\{v \in \mathcal{G}(k) : a_{v,i} = 1\}|$ and $n_k = |\mathcal{G}(k)|$, and $\mathbf{q}_v = (f'_{v,l}/n_k)_{l=1 \dots K}$ be the link distribution of node v with respect to the K roles, where $f'_{v,l} = |\{(v, u) \in \mathcal{E} : u \in \mathcal{G}(l)\}|$. The mean link distribution (with respect to the roles) of group k is thus $\bar{\mathbf{q}}_{v \in \mathcal{G}(k)} = (\sum_{v \in \mathcal{G}(k)} \mathbf{q}_v)/n_k$.

Given the grouping of nodes, we measure the homogeneity of attributes in $g(k)$ by summing up the entropy of each attribute over the nodes in $g(k)$, and measure the homogeneity of links in $g(k)$ by the summed deviation of each node’s link distribution among the roles from the mean of $g(k)$ summed entropy over each attribute and links of map θ^* . Finally, we give the homogeneity measures of map θ^* as follows.

$$h_{attrib}(\theta^*) = \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^L H(\mathbf{p}_{k,i})$$

$$h_{link}(\theta^*) = \frac{1}{K} \sum_{k=1}^K \sum_{v \in \mathcal{G}(k)} D_{KL}(\mathbf{q}_v \| \bar{\mathbf{q}}_{v \in \mathcal{G}(k)}),$$

where $H(\cdot)$ is (binary) entropy function and $D_{KL}(\cdot \| \cdot)$ is the KL-divergence function (relative entropy). h_{link} measures how much each node’s link distribution diverges from the mean distribution of its group.

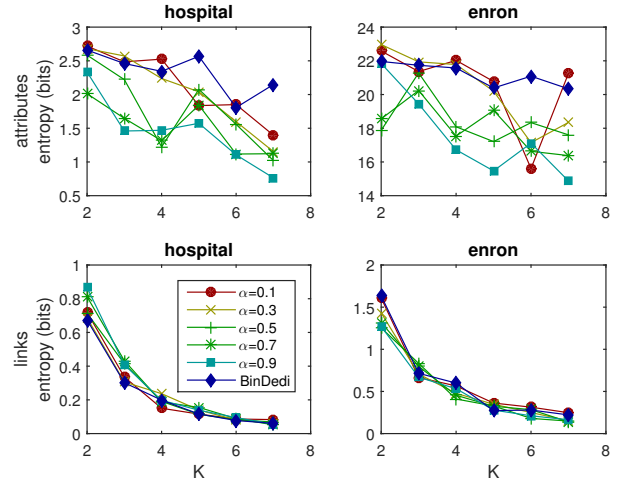
We show the results for homogeneity of both links and attributes in Figure 6. For our method, as K increases, it shows an overall trend of increasing homogeneity (decreasing entropy) on both attributes and links, as expected due to unmixed roles. On the other hand, a smaller α generally results in higher homogeneity for attributes and slightly lower homogeneity for links, as more weight is posed on attributes in optimization.

Next, we compare our method to BINDEDI in Figure 6(a), We observe that NETLAND consistently outperforms BINDEDI which only considers only the topology aspect of networks. For link homogeneity, although our methods model the additional attribute observations, the resulted link homogeneity is still quite comparable to that of BINDEDI which focuses entirely on links.


We also compare our method to CESNA which handles attributes but does not model intergroup interactions, with results given in Figure 6(b). Our method outperforms CESNA in most cases by a significant margin (up to $\sim 40\%$ for *enron* at $K = 7, \alpha = 0.5, 0.9$). We explain this advantage by the greater expressiveness of our model which accounts for intergroup interactions, beyond CESNA’s state-of-the-art modeling of node-attributed networks. Moreover, we observe improvement also in attribute homogeneity over CESNA in most cases (up to $\sim 30\%$ for *hospital* at $K = 7, \alpha = 0.5$), as roles turns out a better concept than communities in explaining attributes.

6.2 Case Study

In Figure 7, we show a sequence of maps \mathcal{M}_K by Algorithm 2 computed by Algorithm 2 with $\beta = 0.002$ for the *hospital* network. While we allow users to select any landmark to split, for this example, we split the largest role (in term of how many nodes have it as the main role) in each map. We used $c = 1$ for \mathbf{P}_{link} and \mathbf{P}_{attrib} , as defined in Section 4. For reference, we index the roles as R_1, \dots, R_6 .



(a) NETLAND v.s BINDEDI



(b) NETLAND v.s CESNA

Figure 6: Homogeneity: average attribute entropy H_{attrib} of landmarks and link entropy H_{link} of nodes.

Multi-resolution maps. First of all, we can intuitively see the localized incremental changes between consecutive maps from the coloring of the non-split landmarks and the heatmaps of \mathbf{P}_{link} , and \mathbf{P}_{attrib} . The interactions between non-split landmarks are not disturbed by the zooming, too.

To start, we see in map \mathcal{M}_1 that nurses has the largest population in the hospital (see corresponding \mathbf{P}_{attrib}). In the subsequent map \mathcal{M}_2 , role R_2 is characterized by mainly involving nurses and R_1 by patients. With help of \mathbf{P}_{link} indicating strong interaction from R_2 to R_1 , we see that the major interactions in this hospital are those from nurses to patients (possibly explained by that nurses actively perform regular checks for patients).

Increasing K reveals finer-resolution structures. In the map \mathcal{M}_5 , role R_1 are administrative staff who actively approach R_4 which involves mainly doctors and some patients; R_1 possibly corresponds to “administrative staff members assigned with specific doctors or patients to assist”. R_2 mainly involves patients and R_3 mainly nurses. R_5 involves mainly administrative staff and some other nurses who are passively approached by patients and doctors; they are possibly “on-duty staff”.

To study R_3 in map \mathcal{M}_5 in more details, we computed map \mathcal{M}_6 by splitting the R_3 of \mathcal{M}_5 into the sub-roles R'_3 and R_6 . Role R'_3 still corresponds to nurses but now those inactive ones, R_6 captures active nurses, and now also some administrative staff absent in R_3 . This is because while we enforce stability of affiliations of nodes to the non-selected roles, our formulation still allows a small number of nodes to change their role compositions (controlled by β in (7)); particularly, some administrative staff become affiliated to the new

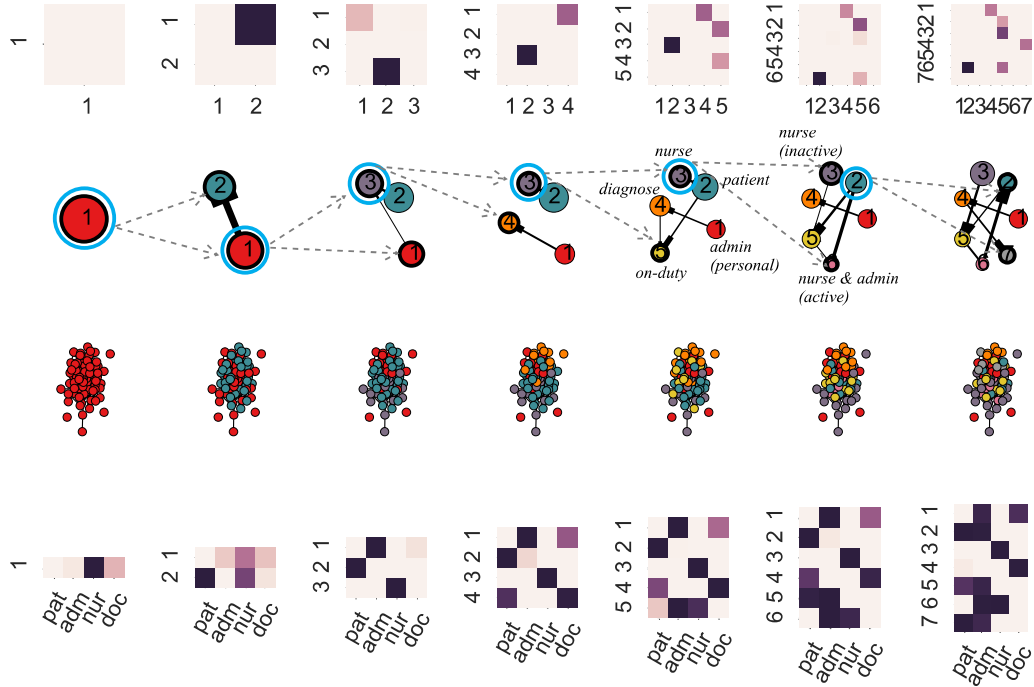


Figure 7: Cartography of the directed hospital network (enlarge for details): a continuous sequence of maps \mathcal{M}_K for $K = 1 \dots 7$ with increasing levels of resolution. The rows from top to bottom are: P_{link} of inter-landmark interaction probabilities, the map (landmarks selected for zooming are circled in bold blue and created sub-roles in bold black), input network (colored by main role), and P_{attrib} of probability of each attribute for each landmark. Landmarks in \mathcal{M}_5 and \mathcal{M}_6 are labeled by hypothesized names.




Figure 8: Locations of nodes: coordinates of nodes (x_v for node v) with respect to map \mathcal{M}_5 in Figure 7, in the space defined by each pair of landmarks. In each subplot, the set of points represent the set of all nodes.

R_6 in \mathcal{M}_6 , resulting in the added adm attribute of R_6 . We interpret R_6 as “active hospital staff (in approaching patients)”.

Locatability. Recall that each node v can be located using affiliation vector x_v . In Figure 8, we show a pairwise plot of \mathbf{X} ’s columns for map \mathcal{M}_5 in Figure 7. Each subplot shows the overall distribution of nodes in the corresponding landmarks. For instance, we observe in subplot-(5,1) that no nodes simultaneously takes roles R_1 (personal admins) and R_5 (on-calls), probably because a personal assistant does not often perform general on-call duties.

7. CONCLUSION

The task of extracting a big picture from complex real-world networks is challenging. In this paper, we presented effective cartographic methods to help with this task. We created intuitively understandable visual maps. To extract landmarks of the map, we recognized roles as the appropriate concept to model landmarks, and proposed a novel expressive model for role discovery. Furthermore, we allow for local multiple-resolution views for knowledge acquisition from networks at different granularities. We also propose efficient algorithms for computing the maps. In addition to an intuitive representation, experiments show our method outperforms state-of-the-art network models in representation accuracy by up to 27% for real networks and up to 10 times on synthetic networks. Case study also shows our created maps actually help users to quickly capture the big picture of a network.

8. REFERENCES

- [1] E. M. Airoldi, D. M. Blei, S. E. Fienberg, and E. P. Xing. Mixed membership stochastic blockmodels. *Advances in Neural Information Processing Systems*, pages 33–40, 2009.
- [2] B. W. Bader, R. Harshman, and T. G. Kolda. Temporal analysis of semantic graphs using asalsan. In *Data Mining, 2007. ICDM 2007. Seventh IEEE International Conference on*, pages 33–42. IEEE, 2007.
- [3] M. W. Berry, M. Browne, and B. Signer. topic annotated enron email data set. *Philadelphia: Linguistic Data Consortium*, 2001.
- [4] B. J. Biddle. *Role theory: Expectations, identities, and behaviors*. Academic Press, 2013.

- [5] S. P. Borgatti and M. G. Everett. The class of all regular equivalences: Algebraic structure and computation. *Social Networks*, 11(1):65–88, 1989.
- [6] R. S. Burt. Positions in networks. *Social forces*, 55(1):93–122, 1976.
- [7] P. Doreian, V. Batagelj, and A. Ferligoj. Generalized blockmodeling. 2005.
- [8] M. G. Everett and S. P. Borgatti. Regular equivalence: General theory. *Journal of mathematical sociology*, 19(1):29–52, 1994.
- [9] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics Springer, Berlin, 2001.
- [10] A. Giddens, M. Duneier, R. P. Appelbaum, and D. Carr. *Introduction to sociology*. WW Norton New York, 2000.
- [11] R. Guimera and L. A. N. Amaral. Functional cartography of complex metabolic networks. *Nature*, 433(7028):895–900, 2005.
- [12] Y. Han and J. Tang. Probabilistic community and role model for social networks. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 407–416. ACM, 2015.
- [13] J. B. Harley, D. Woodward, et al. *Cartography in prehistoric, ancient, and medieval Europe and the Mediterranean*, volume 1. University of Chicago Press Chicago, 1987.
- [14] R. A. Harshman, P. E. Green, Y. Wind, and M. E. Lundy. A model for the analysis of asymmetric data in marketing research. *Marketing Science*, 1(2):205–242, 1982.
- [15] K. Henderson, B. Gallagher, T. Eliassi-Rad, H. Tong, S. Basu, L. Akoglu, D. Koutra, C. Faloutsos, and L. 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.
- [16] M. Kim and J. Leskovec. Latent multi-group membership graph model. *arXiv preprint arXiv:1205.4546*, 2012.
- [17] C.-b. Lin. Projected gradient methods for nonnegative matrix factorization. *Neural computation*, 19(10):2756–2779, 2007.
- [18] M. J. Mandel. Local roles and social networks. *American sociological review*, pages 376–386, 1983.
- [19] M. Nickel, V. Tresp, and H.-P. Kriegel. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 809–816, 2011.
- [20] R. Rossi, N. K. Ahmed, et al. Role discovery in networks. *Knowledge and Data Engineering, IEEE Transactions on*, 27(4):1112–1131, 2015.
- [21] M. Rosvall and C. T. Bergstrom. Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123, 2008.
- [22] A. Silva, W. Meira Jr, and M. J. Zaki. Mining attribute-structure correlated patterns in large attributed graphs. *Proceedings of the VLDB Endowment*, 5(5):466–477, 2012.
- [23] Y. Tian, R. A. Hankins, and J. M. Patel. Efficient aggregation for graph summarization. In *Proceedings of the 2008 ACM SIGMOD international conference on Management of data*, pages 567–580. ACM, 2008.
- [24] P. Vanhems, A. Barrat, C. Cattuto, J.-F. Pinton, N. Khanafer, C. Régis, B.-a. Kim, B. Comte, and N. Voirin. Estimating potential infection transmission routes in hospital wards using wearable proximity sensors. *PLoS one*, 8(9):e73970, 2013.
- [25] F. Wang, T. Li, X. Wang, S. Zhu, and C. Ding. Community discovery using nonnegative matrix factorization. *Data Mining and Knowledge Discovery*, 22(3):493–521, 2011.
- [26] Z. Xu, Y. Ke, Y. Wang, H. Cheng, and J. Cheng. A model-based approach to attributed graph clustering. In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data*, pages 505–516. ACM, 2012.
- [27] J. Yang and J. Leskovec. Overlapping community detection at scale: a nonnegative matrix factorization approach. In *Proceedings of the sixth ACM international conference on Web search and data mining*, pages 587–596. ACM, 2013.
- [28] J. Yang, J. McAuley, and J. Leskovec. Community detection in networks with node attributes. In *Data Mining (ICDM), 2013 IEEE 13th international conference on*, pages 1151–1156. IEEE, 2013.
- [29] Y. Zhou, H. Cheng, and J. X. Yu. Graph clustering based on structural/attribute similarities. *Proceedings of the VLDB Endowment*, 2(1):718–729, 2009.