



Multilayer Block Models for Exploratory Analysis of Computer Event Logs

Corentin Larroche^(✉)

French National Cybersecurity Agency (ANSSI), Paris, France
corentin.larroche@ssi.gouv.fr

Abstract. We investigate a graph-based approach to exploratory data analysis in the context of network security monitoring. Given a possibly large batch of event logs describing ongoing activity, we first represent these events as a bipartite multiplex graph. We then apply a model-based biclustering algorithm to extract relevant clusters of entities and interactions between these clusters, thereby providing a simplified situational picture. We illustrate this methodology through two case studies addressing network flow records and authentication logs, respectively. In both cases, the inferred clusters reveal the functional roles of entities as well as relevant behavioral patterns. Displaying interactions between these clusters also helps uncover malicious activity. Our code is available at <https://github.com/cl-anssi/MultilayerBlockModels>.

Keywords: Cybersecurity · Network modelling · Multiplex networks

1 Introduction

The ever-increasing prevalence of cyber threats compels network administrators and information security specialists to constantly monitor ongoing activity inside sensitive networks. Such monitoring often boils down to collecting and analyzing event logs, which can be generally defined as structured records reflecting all sorts of actions. These logs are continuously produced in massive amounts by nearly all devices within a monitored network, which is both a blessing and a curse: while wide and systematic log collection provides a detailed situational picture, extracting actionable insights from such volumes of data is highly challenging. Data mining algorithms are thus needed to efficiently analyze and summarize all the available information.

Interestingly, various types of event logs can be naturally represented as heterogeneous graphs. Consider for instance a network flow recorded at the border of a protected network, indicating that a connection has occurred between internal host u and external host v . A collection of such events can be abstracted into a bipartite graph whose top (resp. bottom) nodes are the internal (resp. external) hosts. In addition, each network flow is further characterized by additional categorical variables, such as the transport layer protocol (e.g. TCP or UDP) or the source and destination ports. Therefore, each internal–external host pair can

be linked by various kinds of edges. This suggests that a multiplex graph [17] might be an appropriate representation. However, when considering a real-world network, such a graph could easily have tens of thousands of nodes, making it hard to interpret without any further processing.

To alleviate this issue, we propose to use a model-based multiplex graph biclustering algorithm. This algorithm extends previous work on latent block modelling [10, 12] by factoring in the existence of multiple types of edges. It allows us to extract relevant clusters of entities along with their interactions across various network layers, thereby creating a simplified view of the information contained in the logs. While graph biclustering has already been used for event log analysis [20], to the best of our knowledge, the use of multilayer models to account for the heterogeneity of events is novel. In addition, we present detailed case studies on two publicly available datasets consisting of network flow records and authentication logs, respectively.

The rest of this paper is structured as follows. We first introduce the statistical model and inference procedure we use for multiplex graph biclustering in Sect. 2. The two case studies are then presented in Sects. 3 and 4, respectively. Finally, we review some related work in Sect. 5 and discuss potential areas of improvement in Sect. 6.

2 The Multilayer Latent Block Model

Before diving into the specifics of authentication logs and network flows, we first introduce the statistical tools we use to analyze them. Section 2.1 presents our generative model for multiplex bipartite graphs, and Sect. 2.2 describes the algorithms we use for model inference and selection.

2.1 Model Description

Key notations. Let $\mathcal{U} = \{u_1, \dots, u_I\}$, $\mathcal{V} = \{v_1, \dots, v_J\}$ be two node sets of size I and J , respectively. We consider a multiplex bipartite graph $\mathcal{G} = (\mathcal{U}, \mathcal{V}, \mathcal{E})$, where $\mathcal{E} \subset \mathcal{U} \times \mathcal{V} \times [L]$ denotes the edge set (with $[L] \stackrel{\text{def}}{=} \{1, \dots, L\}$). Specifically, each edge $(u, v, \ell) \in \mathcal{E}$ between a top node u and a bottom node v is further characterized by an edge type ℓ , and L denotes the number of possible edge types. The multiplex graph \mathcal{G} can then be characterized by L biadjacency matrices $\mathbf{B}^{(1)}, \dots, \mathbf{B}^{(L)}$, where for each $\ell \in [L]$, the matrix $\mathbf{B}^{(\ell)} = (b_{ij}^{(\ell)}) \in \mathbb{R}^{I \times J}$ is defined by

$$\forall (i, j) \in [I] \times [J], b_{ij}^{(\ell)} = \begin{cases} 1 & \text{if } (u_i, v_j, \ell) \in \mathcal{E} \\ 0 & \text{otherwise.} \end{cases}$$

Finally, we aim to partition \mathcal{U} (resp. \mathcal{V}) into a fixed number H (resp. K) of clusters, which we denote $\mathbb{U} = \{\mathcal{U}_1, \dots, \mathcal{U}_H\}$ (resp. $\mathbb{V} = \{\mathcal{V}_1, \dots, \mathcal{V}_K\}$). For each top node u_i (resp. bottom node v_j), the unique index h such that $u_i \in \mathcal{U}_h$ (resp. k such that $v_j \in \mathcal{V}_k$) is denoted U_i (resp. V_j). In order to find the optimal partition, we adopt a model-based approach relying on the generative model described in the next paragraph.

Generative model. We propose a multilayer extension of the Poisson latent block model introduced by Govaert and Nadif [12]. This model relies on the fundamental assumption that the probability of an edge between two nodes depends on the cluster assignments of these nodes. The biadjacency matrices $\mathbf{B}^{(1)}, \dots, \mathbf{B}^{(L)}$ are then generated by the following hierarchical model:

- (i) for each $i \in [I]$, sample $U_i \sim \text{Multinomial}(\boldsymbol{\pi})$;
- (ii) for each $j \in [J]$, sample $V_j \sim \text{Multinomial}(\boldsymbol{\rho})$;
- (iii) for each $(i, j, \ell) \in [I] \times [J] \times [L]$, sample $b_{ij}^{(\ell)} \sim \text{Poisson} \left(\mu_i \nu_j \theta_{U_i V_j}^{(\ell)} \right)$.

In words, the proportion of top (resp. bottom) nodes falling into each cluster is controlled by a parameter $\boldsymbol{\pi} \in (0, 1)^H$ (resp. $\boldsymbol{\rho} \in (0, 1)^K$). The probability of an edge of type ℓ linking a top node u_i and a bottom node v_j then depends on three parameters: the node-specific factors μ_i and ν_j represent the overall propensity of each node to form edges across all layers, and the rate $\theta_{U_i V_j}^{(\ell)}$ controls the number of edges between clusters \mathcal{U}_{U_i} and \mathcal{V}_{V_j} in the ℓ -th layer. Note that the Poisson distribution is only used here to facilitate calculations: in practice, the biadjacency matrices only contain zeros and ones.

2.2 Model Inference and Selection

Given an observed graph $\mathcal{G} = (\mathcal{U}, \mathcal{V}, \mathcal{E})$, we now aim to find the optimal partitions \mathbb{U} and \mathbb{V} using the model introduced above. To that end, we need to infer the parameters $\boldsymbol{\pi}$, $\boldsymbol{\rho}$, $\boldsymbol{\mu} = (\mu_i) \in \mathbb{R}_+^I$, $\boldsymbol{\nu} = (\nu_j) \in \mathbb{R}_+^J$, and $\boldsymbol{\Theta} = (\boldsymbol{\Theta}^{(1)}, \dots, \boldsymbol{\Theta}^{(L)})$ (where $\boldsymbol{\Theta}^{(\ell)} = (\theta_{hk}^{(\ell)}) \in \mathbb{R}_+^{H \times K}$ for each $\ell \in [L]$). The number of top clusters H and bottom clusters K must also be set in a principled manner. These two problems are addressed in the next paragraphs.

Model inference. The parameters of the model and the cluster assignments are obtained by maximizing the complete data log-likelihood

$$L_C(\mathbb{T}, \mathbb{U}, \mathbb{V}) = \sum_{i=1}^I \log \pi_{U_i} + \sum_{j=1}^J \log \rho_{V_j} + \sum_{i=1}^I \sum_{j=1}^J \sum_{\ell=1}^L \left\{ b_{ij}^{(\ell)} \log \left(\mu_i \nu_j \theta_{U_i V_j}^{(\ell)} \right) - \mu_i \nu_j \theta_{U_i V_j}^{(\ell)} \right\},$$

where $\mathbb{T} = \{\boldsymbol{\pi}, \boldsymbol{\rho}, \boldsymbol{\mu}, \boldsymbol{\nu}, \boldsymbol{\Theta}\}$ denotes the complete set of parameters. To that end, we use the block expectation-maximization algorithm described in [12], with minor adjustments to factor in the multilayer nature of the data. This algorithm first introduces soft cluster assignment matrices $\mathbf{U} = (\tilde{u}_{ih}) \in [0, 1]^{I \times H}$ and $\mathbf{V} = (\tilde{v}_{jk}) \in [0, 1]^{J \times K}$, with $\sum_{h=1}^H \tilde{u}_{ih} = \sum_{k=1}^K \tilde{v}_{jk} = 1$ for all $i \in [I]$ and $j \in [J]$. It then maximizes the fuzzy criterion

$$G(\mathbb{T}, \mathbf{U}, \mathbf{V}) = L_S(\mathbb{T}, \mathbf{U}, \mathbf{V}) + H(\mathbf{U}) + H(\mathbf{V}),$$

where $H(\mathbf{U}) = -\sum_{i=1}^I \sum_{h=1}^H \tilde{u}_{ih} \log \tilde{u}_{ih}$ denotes the total entropy of the soft cluster assignments of the top nodes, $H(\mathbf{V})$ denotes the total entropy for the bottom nodes, and

$$L_S(\mathbb{T}, \mathbf{U}, \mathbf{V}) = \sum_{i=1}^I \sum_{h=1}^H \tilde{u}_{ih} \log \pi_h + \sum_{j=1}^J \sum_{k=1}^K \tilde{v}_{jk} \log \rho_k \\ + \sum_{h=1}^H \sum_{k=1}^K \sum_{i=1}^I \sum_{j=1}^J \sum_{\ell=1}^L \tilde{u}_{ih} \tilde{v}_{jk} \left[b_{ij}^{(\ell)} \log \left(\theta_{hk}^{(\ell)} \right) - \mu_i \nu_j \theta_{hk}^{(\ell)} \right]$$

is the fuzzy likelihood function. The criterion $G(\mathbb{T}, \mathbf{U}, \mathbf{V})$ is maximized by alternatively performing two steps: optimizing \mathbf{U} and \mathbf{V} with Θ fixed (E-step), and optimizing Θ with \mathbf{U} and \mathbf{V} fixed (M-step). These two steps are iterated until G stabilizes. The node partitions \mathbb{U} and \mathbb{V} can then be obtained by assigning each node to the most probable cluster according to \mathbf{U} and \mathbf{V} . Algorithm 1 describes the detailed inference procedure. Note that since the result depends on the random initialization of the parameters, we run the whole procedure 50 times with different initializations and return the model with the highest likelihood.

Model selection. In order to run the aforementioned inference procedure, we must first set the number of top clusters H and bottom clusters K . In the absence of any prior knowledge, we use the integrated completed likelihood (ICL [3]) to pick the best values out of a predefined set of candidates. More specifically, for each $(H, K) \in \{2, \dots, 16\}^2$, we run the inference procedure to obtain the optimal parameter set \mathbb{T} and node partitions \mathbb{U}, \mathbb{V} . We then compute

$$\text{ICL}(\mathbb{U}, \mathbb{V}; H, K) = L_C(\mathbb{T}, \mathbf{U}, \mathbf{V}) - \frac{H-1}{2} \log I - \frac{K-1}{2} \log J - \frac{LHK}{2} \log(LIJ)$$

for each candidate model, and the highest-scoring model is selected. The ICL penalizes the complete data log-likelihood by subtracting terms proportional to the number of free parameters, thus enabling a trade-off between goodness-of-fit and model complexity. Note that it relies on several approximations, which makes it an arguably imperfect evaluation criterion. This has motivated many subsequent contributions on model selection for stochastic and latent block models (see e.g. [4, 19, 27]). However, for the sake of simplicity, we leave the use of more sophisticated criteria for future work.

3 First Case Study—Network Flows

Having described our modelling tools, we now move on to our first case study. Section 3.1 describes the dataset we consider and the preprocessing steps we apply to turn it into a multiplex network. We then discuss the results obtained by applying our methodology in Sect. 3.2.

Algorithm 1: Block expectation-maximization algorithm for maximum likelihood inference of the node clusters and model parameters.

Data: Biadjacency matrices $\mathbf{B}^{(1)}, \dots, \mathbf{B}^{(L)}$; number of top clusters H ; number of bottom clusters K ; stopping criterion ϵ .

Result: Set of estimated parameters \mathbb{T} ; node partitions \mathbb{U} and \mathbb{V} .

$M \leftarrow \sum_{\ell=1}^L \sum_{i=1}^I \sum_{j=1}^J b_{ij}^{(\ell)}$;

foreach $i \in [I], j \in [J]$ **do**

$\mu_i \leftarrow \frac{1}{\sqrt{M}} \sum_{\ell=1}^L \sum_{j=1}^J b_{ij}^{(\ell)}$; $\nu_j \leftarrow \frac{1}{\sqrt{M}} \sum_{\ell=1}^L \sum_{i=1}^I b_{ij}^{(\ell)}$;

end

Randomly initialize $\Theta, \pi, \rho, \mathbf{U}, \mathbf{V}$; $\Delta \leftarrow \infty$;

while $\Delta > \epsilon$ **do**

$G_{\text{old}} \leftarrow G(\mathbb{T}, \mathbf{U}, \mathbf{V})$;

foreach $i \in [I], h \in [H]$ **do**

$s_{ih} \leftarrow \log \pi_h + \sum_{k=1}^K \sum_{\ell=1}^L \sum_{j=1}^J \tilde{v}_{jk} \left(b_{ij}^{(\ell)} \log \theta_{hk}^{(\ell)} - \mu_i \nu_j \theta_{hk}^{(\ell)} \right)$;

$\tilde{u}_{ih} \leftarrow \frac{\exp(s_{ih})}{\sum_{h'=1}^H \exp(s_{ih'})}$;

end

foreach $j \in [J], k \in [K]$ **do**

$t_{jk} \leftarrow \log \rho_k + \sum_{h=1}^H \sum_{\ell=1}^L \sum_{i=1}^I \tilde{u}_{ih} \left(b_{ij}^{(\ell)} \log \theta_{hk}^{(\ell)} - \mu_i \nu_j \theta_{hk}^{(\ell)} \right)$;

$\tilde{v}_{jk} \leftarrow \frac{\exp(t_{jk})}{\sum_{k'=1}^K \exp(t_{jk'})}$;

end

foreach $h \in [H], k \in [K]$ **do**

$\pi_h \leftarrow \frac{1}{I} \sum_{i=1}^I \tilde{u}_{ih}$; $\rho_k \leftarrow \frac{1}{J} \sum_{j=1}^J \tilde{v}_{jk}$;

foreach $\ell \in [L]$ **do**

$\theta_{hk}^{(\ell)} \leftarrow \frac{\sum_{i=1}^I \sum_{j=1}^J \tilde{u}_{ih} \tilde{v}_{jk} b_{ij}^{(\ell)}}{(\sum_{i=1}^I \tilde{u}_{ih} \mu_i)(\sum_{j=1}^J \tilde{v}_{jk} \nu_j)}$;

end

end

$G_{\text{new}} \leftarrow G(\mathbb{T}, \mathbf{U}, \mathbf{V})$; $\Delta \leftarrow |1 - G_{\text{new}}/G_{\text{old}}|$;

end

return \mathbb{T} , ROUND(\mathbf{U}), ROUND(\mathbf{V})

3.1 Data Description

The dataset we use was originally created for the Mini-Challenge 3 of the VAST 2013 competition [26]. It represents two weeks of simulated network traffic between approximately 1400 hosts, most of which belong to an enterprise network. We refer to these hosts as internal ones, while hosts that are not part of the enterprise network are called external hosts. The dataset contains benign traffic as well as various kinds of attacks. Some of these attacks are rather noisy—e.g. distributed denials of service (DDoS), port scans—while others are more subtle (e.g. data exfiltration). Overall, 53 out of 200 external hosts are involved in malicious traffic, which is a rather high ratio. This, along with the small number of hosts and synthetic nature of the data, makes this dataset a somewhat easy

test case for our methodology. Moreover, a thorough description of the hosts and their respective roles is available, as well as a complete list of attack-related events. This allows us to evaluate the relevance of our results.

Note that while the dataset encompasses several data sources, we only keep the network flow records between internal hosts and external hosts. We turn these records into a bipartite multiplex graph as follows. First, each internal (resp. external) host is represented by a top (resp. bottom) node. For each flow between an internal host IH and an external host EH, we then build an edge between IH and EH, of type (*Protocol, Destination Port, Direction*). Note that due to the large number of possible destination ports, we only keep 10 distinct values corresponding to well-known protocols¹ and represent the other values as a single **Other** token. As for the direction of the flow, it is either inbound (the internal host is the destination) or outbound (the internal host is the source). Note that each observed (IH, EH, Type) triple is represented by one single edge, regardless of its number of occurrences. See Table 1 for some descriptive statistics about the dataset and the obtained graph.

Table 1. Description of the datasets: number of top nodes I , number of bottom nodes J , number of layers L , number of distinct edges M , and total number of events N .

Dataset	I	J	L	M	N
VAST	1220	200	18	26,597	68,793,510
LANL	74,049	16,119	44	869,547	842,282,832

3.2 Results

Our methodology yields three clusters of internal hosts and three clusters of external hosts. These clusters, as well as the aggregated edges between them, are shown in Fig. 1. Note that for the sake of readability, only edges representing at least 40 events are displayed. The size of each cluster is proportional to the logarithm of the number of hosts it contains. Similarly, the width of each edge grows logarithmically with the number of underlying events.

Starting with internal clusters, we observe that cluster 1 initiates HTTP and SSH connections, and that it has more outbound flows than inbound ones. Similarly, cluster 3 primarily interacts through outbound flows, including HTTP and FTP traffic. In contrast, cluster 2 receives many inbound connections, with HTTP, SMTP and DNS among the most represented protocols. It also sends SMTP traffic to external cluster 3. Overall, we can thus safely assume that clusters 1 and 3 contain workstations while cluster 2 contains servers, which is

¹ TCP/20 (FTP-Data), TCP/21 (FTP), TCP/22 (SSH), TCP/23 (Telnet), TCP/25 (SMTP), TCP/53 (DNS), TCP/80 (HTTP), TCP/443 (HTTPS), TCP/465 (SMTPS), and TCP/587 (SMTP message submission).

exfiltrations from regular FTP traffic, or DDoS attacks from regular Web traffic. However, these high-volume attacks can be detected using other methods.

4 Second Case Study—Authentication Logs

Our second case study focuses on a different data source, namely authentication logs collected within a real-world enterprise network. The dataset, which we describe in Sect. 4.1, is significantly larger and more complex than the one studied in the previous section. It is thus more challenging to extract meaningful insights from it. However, as we show in Sect. 4.2, our methodology still allows us to infer some functional roles and uncover malicious behavior.

4.1 Data Description

We use the “Comprehensive, Multi-Source Cyber-Security Events” [14, 15] dataset released by the Los Alamos National Laboratory (LANL). This dataset represents 58 days of activity within the LANL’s enterprise network, recorded through several kinds of event logs. A red team exercise took place during this time span, meaning that penetration testers tried to breach the network in order to assess its security. The remote authentications performed by the red team are labelled, providing an interesting example of an advanced intrusion within an enterprise network. In particular, this attack is stealthier than those considered in the previous section, which makes it a more challenging test case for our method.

We focus on Windows authentication logs, more specifically on successful *LogOn* events. Each one of these events is described by a source user SU, a destination user DU, a source host SH, a destination host DH, a logon type LT and an authentication package AP. Note that SU and DU can be identical, and the same goes for SH and DH. As for LT and AP, they describe the type of session being created and the protocol used to authenticate the user, respectively. Examples of logon types include Interactive (local session with graphical interface) or Network (used for remote file accesses or remote procedure calls, for instance), and the most frequent authentication packages are Kerberos and NTLM. Note that these two additional fields are especially meaningful as they allow to distinguish different behaviors (e.g., remote administrative session versus simple file access). They can also help characterize users and hosts: for instance, since NTLM is considered less secure than Kerberos, NTLM authentication can sometimes be disabled for highly privileged accounts. Conversely, some legacy applications may not support Kerberos authentication, thus servers hosting such applications should frequently use NTLM.

We represent users as top nodes and hosts as bottom nodes, and we turn logon events into typed edges as follows. If SH and DH are identical, we create one edge between DU and DH with type (LT, AP, Local). Otherwise, we create one edge between SU and SH with type (LT, AP, From), and one edge between DU and DH with type (LT, AP, To). While this construction breaks the

link between the source and destination of an authentication, it still highlights the difference between hosts receiving many remote authentications (typically servers) and hosts from which these authentications originate (typically workstations). See Table 1 for some descriptive statistics about the dataset.

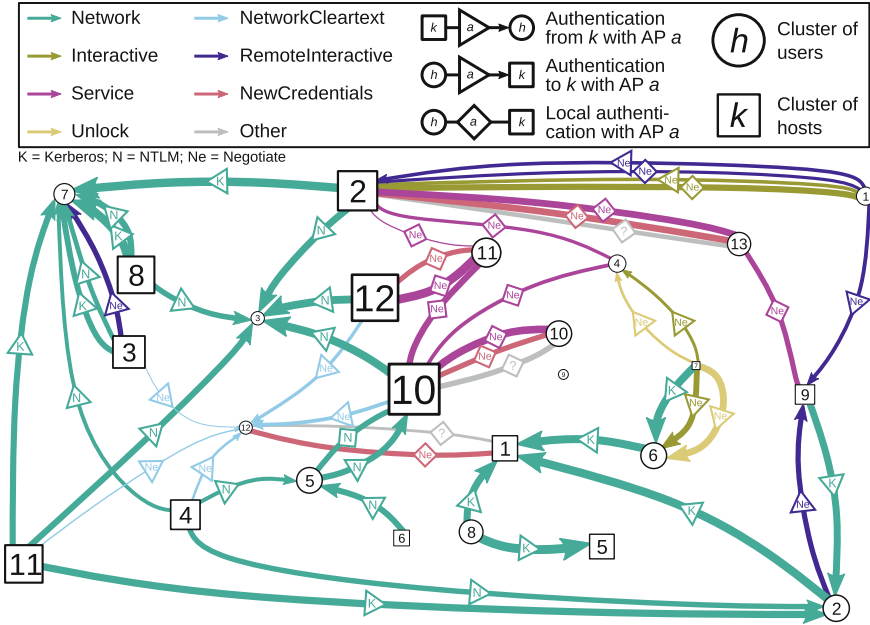


Fig. 2. Clusters found in the LANL dataset and interactions between them.

4.2 Results

The optimal model has thirteen user clusters and twelve host clusters. We display them in Fig. 2, along with the main interactions between them. Specifically, we display all edges (h, k, ℓ) such that $\theta_{hk}^{(\ell)} \geq 0.7$. Note that we select edges using the set of rate matrices Θ rather than the number of underlying events because the behaviors we seek to detect generate relatively few events. Thus cluster–cluster edges with high rate, which represent many user–host edges occurring at least once, are more relevant than those representing many individual events.

The global picture is expectedly more complex than in the previous section. In addition, little information is available on the true functional roles of users and hosts. However, we can still infer the meaning of some clusters and confirm their relevance. For instance, user clusters 10, 11 and 13 mainly interact through local authentications with logon type Service, indicating that they contain service accounts. User cluster 1 opens remote interactive sessions on many hosts, which

suggests the presence of administrator accounts. As for user clusters 3 and 5, they mostly consist of anonymous user credentials, which are typically used to access shared resources which do not require any access control (such as internal Web pages). This is consistent with the NTLM-authenticated network logons associated with these clusters. Interestingly, a significant proportion of the non-anonymous user names from these two clusters were used for red team activity. This suggests that looking for inconsistencies between the expected functional role of an entity and the cluster it falls into can help uncover malicious behaviors.

Inferring the functional roles of host clusters is more difficult, although it seems reasonable to assume that cluster 1 contains domain controllers and applicative servers. Indeed, this cluster is the one receiving the most remote authentications. Conversely, clusters 2, 3, 8, 9, and 11 are sources of remote authentications, suggesting that they mostly contain workstations. Finally, cluster 4 stands out as the source of remote NTLM authentications involving three user clusters. Further investigation reveals that most of these authentications were performed by the red team and originate from host C17693, which also happens to be the main source of red team activity. Once again, the simplified situational picture we generate thus preserves important clues that can lead to the detection of malicious behavior.

5 Related Work

Biclustering and latent block models. Biclustering, i.e., the idea of simultaneously partitioning the rows and columns of a matrix so as to form blocks of coefficients with similar values, can be traced back to the early 1970s [13]. Spectral methods [7, 18] are among the most popular approaches to this task. However, model-based biclustering has gained traction in the last two decades, following the introduction of the latent block model [10]. Various inference procedures have been proposed, most of which rely on the expectation-maximization (EM) algorithm. While we apply the simple variational approximation proposed in [10], other approaches include the addition of intermediary classification steps [11] or the use of Bayesian inference [16].

Multilayer stochastic block models. Even though we are not aware of any previous work on multilayer generalizations of the latent block model, a straightforward connection can be made with multiplex stochastic block models (SBMs). More specifically, our model is analogous to the degree-corrected SBM with independent layers described in [22]. Note, however, that many other multilayer generalizations of the SBM can be found in the literature. In particular, the existence of diversely complex statistical dependencies between layers has been addressed in various ways [2, 6, 21, 24].

Data exploration and visualization for cybersecurity. Finally, our work contributes to a vast research effort aiming to ease network security monitoring by providing insightful visualizations. In particular, displaying network flows

between internal and external hosts is a recurring challenge. VISUAL [1], VisFlowConnect [28], NFlowVis [8], and FloVis [25] are some of the existing tools designed to perform this task. However, none of them focuses on statistical modelling to build a more condensed view. Note that FloVis and NFlowVis use the hierarchical nature of IP addresses to aggregate hosts by subnetwork, and NFlowVis also uses k -medoids clustering to find external hosts with similar communication patterns. Even so, the use of behavior-based biclustering makes our method more effective at reducing the number of elements to display while preserving essential information. As for authentication logs, APTHunter [23] also focuses on interactivity rather than automated data reduction, only letting the user manually apply filters to make the authentication graph legible. Finally, the work of Glatz et al. [9] adopts an approach somewhat similar to ours: they first extract frequent itemsets from network flow records, then display them as a bipartite itemset–item graph. While this does provide a simple summary of large volumes of logs, the extracted itemsets still cover a small minority of the total number of events, leading to significant information loss.

6 Conclusion and Perspectives

We propose a graph-oriented approach to event log exploration for network security monitoring. Through the use of model-based multiplex graph biclustering, we aim to extract meaningful clusters of entities, such as groups of users or hosts sharing functional roles or behavioral patterns. Our case studies demonstrate that such meaningful clusters can indeed be uncovered. In addition, displaying interactions between these clusters can facilitate malicious behavior detection.

Aside from investigating better model selection criteria (see Sect. 2.2), extending our model to factor in the temporal dimension could be an interesting lead for future work. Previous contributions on temporal latent block models [5] provide foundations for such an extension. Finally, looking for meaningful groups of edge types in addition to entity clusters, similarly to the strata multilayer SBM [24], is another promising direction. Indeed, it could lead to an even simpler situational picture, especially for large and complex datasets such as the one studied in Sect. 4.

References

1. Ball, R., Fink, G.A., North, C.: Home-centric visualization of network traffic for security administration. In: VizSec/DMSec (2004)
2. Barbillon, P., Donnet, S., Lazega, E., Bar-Hen, A.: Stochastic block models for multiplex networks: an application to a multilevel network of researchers. *J. R. Stat. Soc. A Stat.* **180**(1), 295–314 (2017)
3. Biernacki, C., Celeux, G., Govaert, G.: Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Trans. Pattern Anal. Mach. Intell.* **22**(7), 719–725 (2000)

4. Côme, E., Latouche, P.: Model selection and clustering in stochastic block models based on the exact integrated complete data likelihood. *Stat. Model.* **15**(6), 564–589 (2015)
5. Corneli, M., Latouche, P., Rossi, F.: Exact ICL maximization in a non-stationary time extension of the latent block model for dynamic networks. In: *ESANN* (2015)
6. De Bacco, C., Power, E.A., Larremore, D.B., Moore, C.: Community detection, link prediction, and layer interdependence in multilayer networks. *Phys. Rev. E* **95**(4), 042317 (2017)
7. Dhillon, I.S.: Co-clustering documents and words using bipartite spectral graph partitioning. In: *KDD* (2001)
8. Fischer, F., Mansmann, F., Keim, D.A., Pietzko, S., Waldvogel, M.: Large-scale network monitoring for visual analysis of attacks. In: *VizSec* (2008)
9. Glatz, E., Mavromatidis, S., Ager, B., Dimitropoulos, X.: Visualizing big network traffic data using frequent pattern mining and hypergraphs. *Computing* **96**(1), 27–38 (2014)
10. Govaert, G., Nadif, M.: Clustering with block mixture models. *Pattern Recognit.* **36**(2), 463–473 (2003)
11. Govaert, G., Nadif, M.: Block clustering with Bernoulli mixture models: comparison of different approaches. *Comput. Stat. Data Anal.* **52**(6), 3233–3245 (2008)
12. Govaert, G., Nadif, M.: Latent block model for contingency table. *Commun. Stat. Theory Methods* **39**(3), 416–425 (2010)
13. Hartigan, J.A.: Direct clustering of a data matrix. *J. Am. Stat. Assoc.* **67**(337), 123–129 (1972)
14. Kent, A.D.: Comprehensive, multi-source cyber-security events. Los Alamos National Laboratory (2015)
15. Kent, A.D.: Cybersecurity data sources for dynamic network research. In: *Dynamic Networks in Cybersecurity*. Imperial College Press (2015)
16. Keribin, C., Brault, V., Celeux, G., Govaert, G.: Estimation and selection for the latent block model on categorical data. *Stat. Comput.* **25**(6), 1201–1216 (2015)
17. Kivelä, M., Arenas, A., Barthélemy, M., Gleeson, J.P., Moreno, Y., Porter, M.A.: Multilayer networks *J. Complex Netw.* **2**(3), 203–271 (2014)
18. Kluger, Y., Basri, R., Chang, J.T., Gerstein, M.: Spectral biclustering of microarray data: coclustering genes and conditions. *Genome Res.* **13**(4), 703–716 (2003)
19. Lomet, A., Govaert, G., Grandvalet, Y.: Model selection in block clustering by the integrated classification likelihood. In: *COMPSTAT* (2012)
20. Metelli, S., Heard, N.: On Bayesian new edge prediction and anomaly detection in computer networks. *Ann. Appl. Stat.* **13**(4), 2586–2610 (2019)
21. Paul, S., Chen, Y.: Consistent community detection in multi-relational data through restricted multi-layer stochastic blockmodel. *Electron. J. Stat.* **10**(2), 3807–3870 (2016)
22. Peixoto, T.P.: Inferring the mesoscale structure of layered, edge-valued, and time-varying networks. *Phys. Rev. E* **92**(4), 042807 (2015)
23. Siadati, H., Saket, B., Memon, N.: Detecting malicious logins in enterprise networks using visualization. In: *VizSec* (2016)
24. Stanley, N., Shai, S., Taylor, D., Mucha, P.J.: Clustering network layers with the strata multilayer stochastic block model. *IEEE Trans. Netw. Sci. Eng.* **3**(2), 95–105 (2016)
25. Taylor, T., Paterson, D., Glanfield, J., Gates, C., Brooks, S., McHugh, J.: Flowvis: Flow visualization system. In: *CATCH* (2009)

26. Whiting, M., Cook, K., Paul, C., Whitley, K., Grinstein, G., Nebesh, B., Liggett, K., Cooper, M., Fallon, J.: Vast challenge 2013: Situation awareness and prospective analysis. In: VAST (2013)
27. Wyse, J., Friel, N., Latouche, P.: Inferring structure in bipartite networks using the latent blockmodel and exact ICL. *Netw. Sci.* **5**(1), 45–69 (2017)
28. Yin, X., Yurcik, W., Treaster, M., Li, Y., Lakkaraju, K.: Visflowconnect: net-flow visualizations of link relationships for security situational awareness. In: *VizSec/DMSec* (2004)