Managing Scope Changes for Cellular Network-level Anomaly Detection

Gabriela F. Ciocarlie*, Chih-Chieh Cheng*, Christopher Connolly*,
Ulf Lindqvist*, Szabolcs Nováczki†, Henning Sanneck†, Muhammad Naseer-ul-Islam†
*SRI International
{gabriela.ciocarlie,chih-chieh.cheng,christopher.connolly,ulf.lindqvist}@sri.com
†NSN Research
{szabolcs.novaczki,henning.sanneck,muhammad.naseer-ul-islam}@nsn.com

***Abstract*—The Self-Organizing Networks (SON) concept is increasingly being used as an approach for managing complex, dynamic mobile radio networks. In this paper we focus on the verification component of SON, which is the ability to automatically detect problems such as performance degradation or network instability stemming from configuration management changes. In previous work, we have shown how Key Performance Indicators (KPIs) that are continuously collected from network cells can be used in an anomaly detection framework to characterize the state of the network. In this study, we introduce new methods designed to handle scope changes. Such changes can include the addition of new KPIs or cells in the network, or even re-scoping the analysis from the level of a cell or group of cells to the network level. Our results, generated using real cellular network data, suggest that the proposed network-level anomaly detection can adapt to such changes in scope and accurately identify different network states based on all types of available KPIs.**

***Index Terms*—network automation, self-organized networks (SON), SON verification, anomaly detection, diagnosis**

I. Introduction

Modern radio networks for mobile broadband (voice and data) are complex and dynamic, not only in terms of behavior and mobility of users and their devices, but also in terms of the many elements that make up the network infrastructure. Effective management of such networks requires some form of automated detection of problems such as performance degradation or network instability. In order to prevent network-level degradation, either the actions that change network-element configurations must be coordinated *a priori*, or their effects must be verified by a SON verification framework.

A. SON Verification

In previous work [5], we proposed a novel SON verification framework that uses anomaly detection and diagnosis techniques and operates within a certain spatial scope larger than an individual cell, e.g., a small group of cells or cell cluster being in scope for a SON optimization, an existing administrative network domain, etc. Key Performance Indicators (KPIs) are continuously collected from network cells and used as indicators (e.g., call-drop statistics, channel-quality-indicator statistics, handover statistics, throughput, etc.) for network performance.

The first component of the SON verification framework [5] was the network-level anomaly detection component, which aimed to detect anomalies within a given scope of analysis (at the level of either a network or group of cells). The anomaly detection framework used a multi-level clustering approach in which KPI data was first clustered to form different states in which the network could be, and to express the overall state of the network as a combination of states that correspond to different parts of the network. Our framework used topic modeling [1], a type of statistical model that provides an efficient framework for training of clusters and performing inference of the mixing proportion of the clusters (Figure 1). Given the raw KPI data, topic models learn clusters of network states, and output the codebook and the mixing proportions of the clusters at each timestamp. For each cluster, a codebook records a profile: the average of the cluster, called a centroid. Clusters are also known as the topics in topic modeling. The mixing proportion of clusters at each timestamp represents the overall state of the cells in scope.

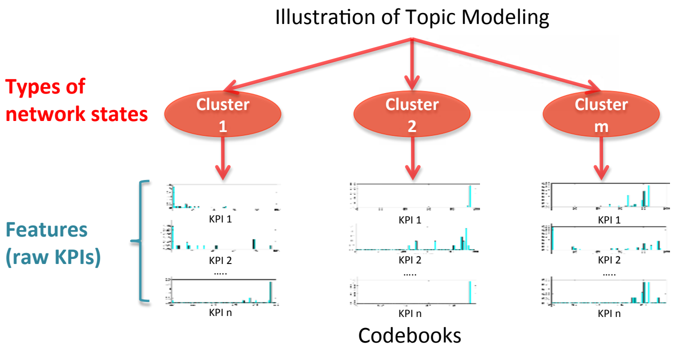

Fig. 1. Illustration of topic modeling on clustering KPI data in a multi-level clustering framework.

A second component of the anomaly detection framework [5] was an interpretation module that could characterize topic model clusters as either normal or abnormal. This simple classifier used the semantics of the KPI to generate labels; consequently, classification was achieved only for KPIs that were supposed to maintain a low level (e.g., drop call rate) or a high level (e.g., call success rate).

B. Contributions

The key challenge for deploying a SON verification in a real environment is to cope with changes in the networks. This paper proposes novel methods for managing scope changes which are prevalent in real cellular networks. Our implementation and experiments focus on the problem of adapting to changes in the types of KPIs and cells, and also changes in the analysis scope that span smaller to larger group of cells. The work described here provides several contributions by:

- proposing a novel approach that uses cell-level degradation information as features for the topic modeling to adapt to all KPIs,
- applying an incremental topic modeling approach that can adapt to changes in the analysis scope, and operates across all KPIs,
- building a system to implement the algorithms, applying the system to a real KPI dataset, and analyzing the performance of the proposed framework.

II. Network-Level Anomaly Detection and Scope Changes

Due to the dynamic nature of cellular networks, the scope of anomaly detection analysis can change with the addition of new KPIs or cells in the network, or from the level of a cell or group of cells up to the network. The first step towards addressing scope changes was to consider the neutral KPIs (e.g., throughput), which were not included in the original framework. Neutral KPIs exhibit an operating area, which is deployment-specific with both a lower and upper bound. Hence, they cannot be directly used by the topic modeling component because the interpretation module cannot interpret them. To address this, we propose to combine anomaly detection capabilities applied at different scopes: cell level and group/network levels.

Another important aspect of our SON verification framework is adaptability to different cell scopes (increased or decreased number of cells in the scope). With the previous approach [5], clustering is first applied for the largest scope (the largest number of cells) and then, the state of the network can be determined for subsets of the largest number of cells. However, in a real deployment, if the scope needs to be enlarged (e.g., when new cells are added to the network), we would require an incremental approach for topic modeling, which will gradually update the clusters with information from the larger scope. Consequently, we propose an incremental approach for topic modeling.

A. Topic Modeling Using Neutral KPIs

To cope with all types of KPIs, we propose to extend the topic-modeling framework (Figure 2) to include output from our previous work on cell-level anomaly detection [3], [4]. Hence, for the neutral KPIs we propose to use the KPI degradation level generated by our ensemble method when applied to multiple individual univariate and multivariate methods. The KPI degradation level computes a numerical measure to indicate the severity of a degradation. Ideally a normal cell would exhibit a KPI level of 0; thus, the interpretation module would know how to label the neutral KPIs. For the non-neutral KPIs, we continue to use the raw KPI values in the topic modeling.

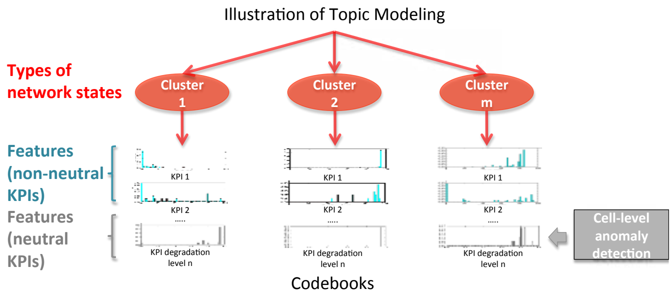

Fig. 2. Illustration of topic modeling on clustering both non-neutral and neutral KPI data in a multi-level clustering framework.

B. Incremental Topic Modeling

We follow an online learning version of the Hierarchical Dirichlet Processes (HDP) proposed in [9], which utilizes a stochastic gradient optimization to allow the training process to evolve incrementally over time, and we adapted it to use all KPIs as multivariate input. The idea behind the stochastic gradient optimization is that we subsample data, compute an approximation of the gradient based on the subsample and follow the gradient with a decreasing step-size. Instead of running over the whole dataset, the algorithm uses each data sample on the fly, and updates the model parameters by the gradients of that single data sample. The ideal assumption behind stochastic gradient learning is that there exists infinite data. Theoretically, with infinite data, stochastic gradient learning will eventually converge to an optimum point, as does training with all available data.

There are two main considerations for implementing an incremental learning algorithm for topic modeling on a real-world system: convergence and speed. Since the incremental learning algorithm sees one data point at a time and updates the model parameters according to one single data sample, the convergence will be more treacherous than its batch-learning counterpart and could require more time [2].

Depending on the size of the data, the speed of convergence of online learning may be slower or faster than its batch-learning counterpart. When batch-learning algorithms cannot or can barely handle the amount of data, online learning algorithms may converge faster [7]. However, the speed of each iteration is supposed to be fast, since online learning algorithms consider only one single data sample at a time.

III. Experimental Evaluation

Our experimental corpus consisted of real KPI data from a 3G network for the period 01/2013 through 03/2013, containing information for approximately 2,000 cells.

A. Topic Modeling with Neutral KPIs Results

Throughout our work, we assume that the type of KPI (neutral, increasing or decreasing) is given by the semantics of the KPIs. The set of eleven non-neutral KPIs consists of call-setup success rates and drop call rates. This was the original

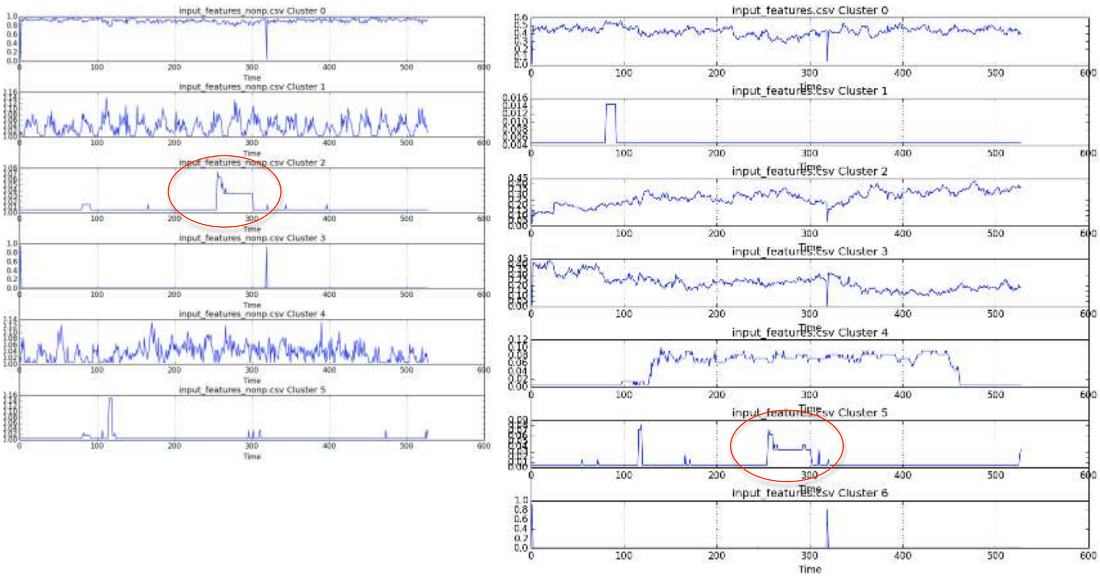

Fig. 3. Comparison between the cluster weights when using only the non-neutral KPIs (left) and when using all KPIs (right). Each row represents the weights for one cluster over time, where the x-axis denotes the time, and the y-axis denotes the weight. There are different numbers of clusters (6, respectively 7) and the order in which the clusters were number does not imply a one-to-one mapping. The mid-February anomaly is circled in red.

set of KPIs that was used in the previous work [5]. The neutral KPIs set includes throughput KPIs. These are the new KPIs for which topic modeling uses the KPI degradation level based on the ensemble method for cell anomaly detection [3], [4] instead of the raw KPI values.

For our initial experiments, we considered 100 handpicked cells and analyzed them for the period January 2013 to March 2013. The goal was to first determine if the KPI levels provide useful information on the network status before expending to a larger number of cells.

Figure 3 presents a comparison between the cluster weights that were generated using only the non-neutral KPIs (left) and using all KPIs (right). Each row represents the weights for one cluster over time, where the x-axis denotes the time, and the y-axis denotes the weight. Note that there are different numbers of clusters for the two cases (6, respectively 7) and the order in which clusters were numbered does not imply a one-to-one mapping, i.e., if Cluster 2 corresponds to a normal state in one case, Cluster 2 might not correspond to the same state for the other case. However, in both cases, we notice the mid-February anomaly that we referred to in the previous work (for Cluster 2, respectively Cluster 5) [5]. We recall that the number of clusters can be different because the topic modeling implementation uses an HDP [8] approach, which determines the number of topics automatically.

Table I summarizes the analysis of the 6 and 7 clusters, respectively. The analysis included a manual investigation of the histograms of each cluster (which lead to detailed characteristics), and has been confirmed by the automated cluster interpretation module (which provided either normal or abnormal labels for each cluster). We notice that Cluster 2 for the baseline case and Cluster 5 for the case with all KPIs correspond to strong anomalies across all 11 raw KPIs

	Baseline: 11 raw KPIs	11 raw KPIs + KPI degradation levels of new KPIs
Cluster 0	Normal	Normal
Cluster 1	Periodicity of some KPIs	Minor anomaly
Cluster 2	Strong anomaly	Normal, shows some periodicity
Cluster 3	Bad DCR_cs_voice	Normal, shows some periodicity
Cluster 4	Minor anomaly, periodicity of some KPIs	Anomaly conditions in the new 9 KPIs (11 raw KPIs are normal)
Cluster 5	Bad cell availability	Strong anomaly in 11 raw KPIs
Cluster 6	N/A	Bad DCR_cs_voice

TABLE I
SUMMARY OF THE ANALYSIS FOR EACH CLUSTER AND FOR EACH CASE. NOTE THAT THERE IS NO ONE-TO-ONE MAPPING BETWEEN THE CLUSTERS GENERATED FOR THE TWO CASES.

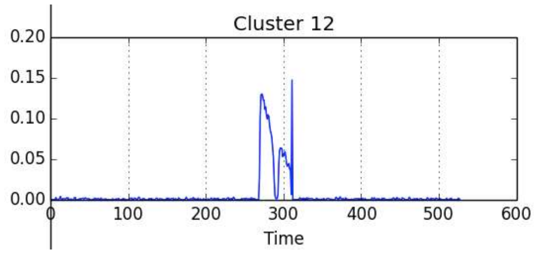

Fig. 4. Cluster weights for Cluster 12 for February. We notice the predominant anomaly in mid February.

exhibited in mid-February. Overall, the KPI levels seem to provide some discrimination on the network status, given the appearance of Cluster 4, which is characterized by abnormalities in the neutral KPIs. The rest of the clusters correspond to very similar states without a one-to-one mapping (i.e., different types of clusters were generated in different order for the two cases).

Given that the results on the small sample of cells indicated that the neutral KPIs could discriminate even better the network state, we further performed our analysis on the 2,000 cells. The total number of timestamps was

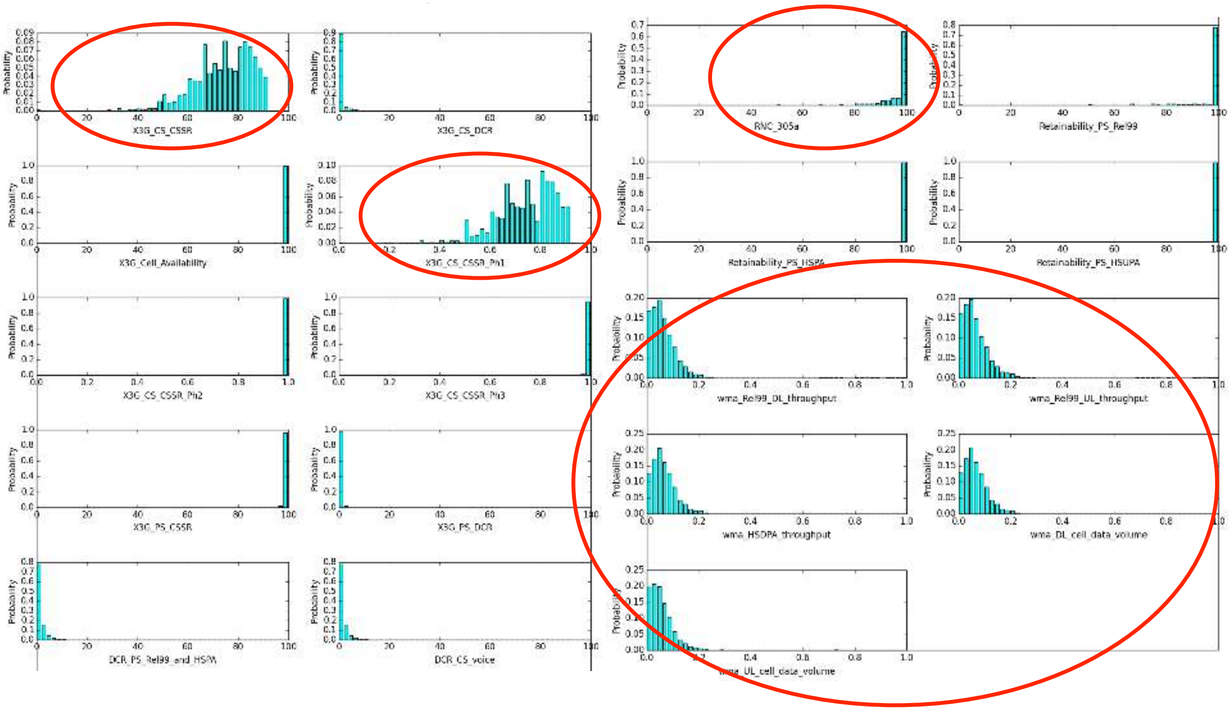

Fig. 5. KPI Histograms for Cluster 12. Anomaly conditions are circled.

1,834. The topic modeling method generated 20 clusters for the 2,000 cells, which were very similar to the original experiments presented in the previous work [5]. One interesting cluster was Cluster 12 for which we could observe the mid-February anomaly (Figure 4). Figure 5 presents the KPI histograms for Cluster 12. We can observe that the anomalous KPIs are the non-neutral ones: 3G_CS_CSSR, 3G_CS_CSSR_Ph2, RNC_305a and the neutral ones: Rell99_DL_throughput, Rell99_DL_throughput, UL_cell_data_volume, HSDPA_throughput and DL_cell_data_volume (for which we used the KPI levels computed using the ensemble method, hence they are prefixed by *wma*). Cluster 9 is an example of a normal cluster (Figure 6). After applying the interpretation method, which automatically classifies each cluster as either normal or abnormal based on KPI characteristics, we obtained 9 normal clusters and 11 abnormal clusters.

B. Incremental Topic Modeling Results

We have implemented the incremental learning algorithm proposed by Wang et al. [9]. This algorithm has previously only been used for experimentation on academic data sets and is still an ongoing research topic [10], [6]. Indeed, the algorithm requires many parameters that need to be tuned beforehand with background knowledge; it lacks an automated methodology for self-tuning parameters. The nature of stochastic learning behind this algorithm makes the convergence very unstable; significant effort is required to smooth the process. Therefore we demonstrate what we have investigated so far about the possibility of extending topic modeling for ever-increasing data with some initial results. To the best of our knowledge, this is the first time that incremental topic modeling is applied to real cellular network data.

We experimented with the incremental topic modeling on the full 3G dataset and both non-neutral and neutral KPIs. The ideal assumption behind stochastic gradient learning is that there exists infinite data. While we only have a finite set of data, we can loop through the dataset to simulate an infinite data environment. We ran the incremental algorithm for HDP by randomly choosing timestamps from the 3G dataset, and updated the model parameters accordingly. After running through 500 timestamps (out of 1,834 timestamps in total), one cluster that represented a normal condition was generated, while the other clusters included random noise. For the remaining of the timestamps, nothing changed. The profile of the normal cluster is almost identical to the one represented in Figure 6 (due to lack of space we do not show it). The reason we only got one normal cluster is that the 500 random timestamps are mostly normal, and the model is still far from convergence. This is in agreement with the fact that this algorithm did not produce meaningful clusters until it processed 200K documents from the Nature corpus [9]. With one normal cluster generated, our current results are not exhaustive. However, we believe they illustrate the initial feasibility of the incremental topic modeling approach in the context of cellular network data, and pave the way for future investigation using larger datasets.

In terms of speed, for 500 timestamps, the incremental learning algorithm for HDP took about 15 minutes on a single core of a server-level Linux machine, without extensive code

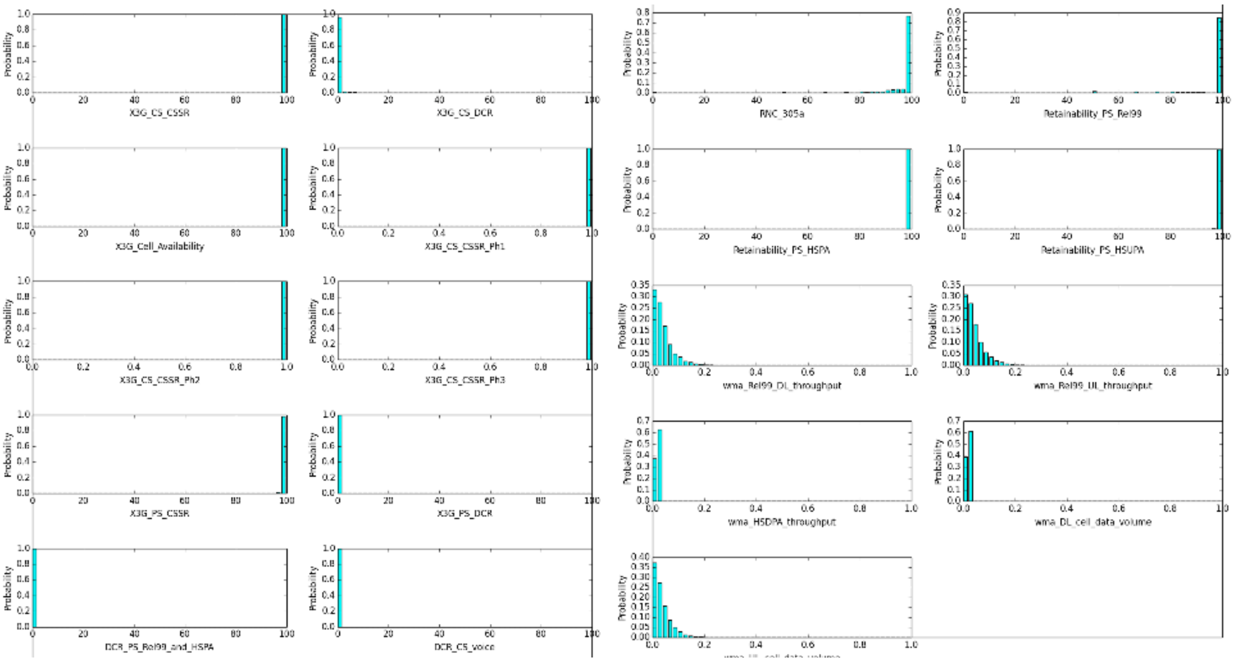

Fig. 6. Example of a normal cluster: Cluster 9

optimization. As a comparison, when the original HDP is trained, it takes 1-2 hours for one sweep of the same dataset on the same machine; convergence is usually achieved in about 20 iterations (sweeps of data). On a real deployment system, the incremental algorithm would update the model incrementally, whenever a new set of KPIs is collected from the cells. This action will take only a few seconds.

IV. Conclusions

This paper proposes novel extensions to a SON verification framework, designed to manage scenarios where the scope changes with the addition of new KPIs or cells in the network, or from the level of a cell or group of cells to the network level. The design was implemented and applied to a dataset consisting of KPI data collected from a real operational cell network. The experimental results suggest that the augmented network-level anomaly detection accurately identifies different network states based on all types of KPIs and can adapt to changes in scope. We are currently planning to apply the incremental topic modeling to a larger dataset (when available) and to explore techniques for tuning relevant parameters such as model priors, learning rate and annealing temperature, to achieve convergence.

Acknowledgment

We thank Lauri Oksanen, Kari Aaltonen, Michael Freed and Kenneth Nitz for valuable feedback.

References

[1] David Blei, Lawrence Carin, and David Dunson. Probabilistic Topic Models. IEEE Signal Processing Magazine, 27(6): 5565. 2010.

[2] Léon Bottou, Online Algorithms and Stochastic Approximations, Online Learning and Neural Networks, Edited by David Saad, Cambridge University Press, Cambridge, UK, 1998.

[3] Gabriela F. Ciocarlie, Ulf Lindqvist, Szabolcs Novaczki, and Henning Sanneck, Detecting Anomalies in Cellular Networks Using an Ensemble Method, 9th International Conference on Network and Service Management (CNSM), 2013.

[4] Gabriela F. Ciocarlie, Ulf Lindqvist, Kenneth Nitz, Szabolcs Nováczki and Henning Sanneck, On the Feasibility of Deploying Cell Anomaly Detection in Operational Cellular Networks , IEEE/IFIP Network Operations and Management Symposium (NOMS), Experience Session, 2014.

[5] Gabriela F. Ciocarlie, Christopher Connolly, Chih-Chieh Cheng, Ulf Lindqvist, Szabolcs Nováczki, Henning Sanneck, MuhammadNaseer-ul-Islam, Anomaly Detection and Diagnosis for Automatic Radio Network Verification , Submitted to the 6th International Conference on Mobile Networks and Managemen (MONAMI), 2014.

[6] D. Kim, P. Gopalan, D. Blei, and E. Sudderth, Efficient online inference for Bayesian nonparametric relational models, Neural Information Processing Systems, 2013.

[7] Yann Le Cun, Léon Bottou, Genevieve B. Orr and Klaus-Robert Müller, Efficient Backprop, Neural Networks, Tricks of the Trade, Lecture Notes in Computer Science LNCS 1524, Springer Verlag, 1998.

[8] Yee Whye Teh, Michael I. Jordan, Matthew J. Beal and David M. Blei, Hierarchical Dirichlet Processes., Journal of the American Statistical Association 101, 476: 1566-1581, 2006.

[9] Chong Wang, John Paisley and David M. Blei. Online variational inference for the hierarchical Dirichlet process, In AISTATS 2011.

[10] C. Wang and D. Blei, Truncation-free stochastic variational inference for Bayesian nonparametric models, Neural Information Processing Systems, 2012.