Improving an SVD-based combination strategy of anomaly detectors for traffic labelling

Johan Mazel
NII/JFLI
johanmazel@nii.ac.jp

Romain Fontugne
The University of Tokyo/JFLI
romain@hongo.wide.ad.jp

Hiroshi Esaki
The University of Tokyo
hiroshi@wide.ad.jp

Kensuke Fukuda
NII
kensuke@nii.ac.jp

ABSTRACT

Network anomaly detection is a crucial task in traffic monitoring. Tools targeting this particular problematic need to be thoroughly evaluated to assess their efficiency. Such evaluation needs reliable ground truth data to be effective. The goal of the present article is to assist researchers in the evaluation of detectors by providing them with labelled anomaly traffic traces. One of the promising strategies to provide reliable ground truth data is to combine the output of the multiple anomaly detectors with different theoretical background. In this paper, we provide an in-depth analysis of the Singular Value Decomposition (SVD) based combination strategy that has been recently applied to anomaly detectors (MAWILab). This analysis highlights the key drawback of the method to efficiently discriminate the anomalous traffic from the harmless one. We then propose several techniques to overcome this drawback and improve the discrimination power of the combination strategy. Our evaluation using four anomaly detectors and four years of real backbone traffic traces (MAWI) emphasizes the accuracy gain of the proposed techniques over the original study.

Categories and Subject Descriptors

C.2.0 [**Computer-Communication Networks**]: General—*Security and protection (e.g., firewalls)*

General Terms

Measurement, Security

Keywords

network anomaly detection, combination strategy, anomaly labelling

AINTEC '12, November 14-16, 2012, Bangkok, Thailand

1. INTRODUCTION

Anomalies have a detrimental effect on legitimate users' access to Internet resources. Identifying anomalous events is a crucial network management task that requires automation. Consequently, anomaly detection has received a lot of attention in the last decade, and numerous detectors have been proposed. Operators, however, often disregard the alarms reported by anomaly detectors because of several drawbacks discrediting them [10, 21, 23, 24]. The key task for improving anomaly detectors is to thoroughly evaluate their output. However, identifying anomaly detectors vulnerabilities is particularly difficult due to a lack of ground truth data for real traffic.

Anomaly detectors evaluation usually rely on two type of dataset: simulated data and real traffic. Simulating anomalies is a common way to evaluate an anomaly detector [12, 18, 22, 25]. In this case, the parameters of anomalies are tunable (e.g. intensity and time duration), helping researchers to measure the sensitivity of their detectors to particular kinds of anomalies. However, simulating traffic as diverse as it is on the Internet is notoriously difficult [7], especially for anomalous traffic. Consequently, the evaluation of a detector with simulated anomalies is restricted to certain kinds of anomaly, and thus, is insufficient for measuring the detector performance [20]. With real anomalies, researchers evaluate anomaly detectors by manually checking the reported alarms [4, 6, 13, 14], or by comparing them to those reported by other anomaly detectors [9, 12–14]. Sometimes researchers also construct ground truth data by manually inspecting the analysed traffic [1]. However, these evaluations are hardly comparable, trustworthy, or reproducible, as they require significant human intervention and traffic traces are usually inaccessible due to privacy issues.

Ideally, an anomaly detector has to be evaluated using ground truth data containing real and non-specific traffic where a wide range of anomalies is located. This

ground truth data should be publicly available to allow all researchers to access the same data set and compare their results. Furthermore, the data set should follow the evolution of the Internet traffic to include traffic from emerging applications and anomalies. Such data exist in the MAWILab repository [8] regarding the MAWI archive [5].

The goal is to find and label anomalies in the traffic from the MAWI archive. The main advantages of the MAWI archive are that it is updated daily and it currently contains more than eleven years of real publicly available Internet traffic data. However, manually labelling anomalies in such a large data set is certainly impractical, and therefore, the challenge we face is to accurately locate anomalies in an automated and therefore unsupervised manner. The numerous anomaly detectors that have recently been proposed in literature are the main support that will help us reach our goal. Therefore, we are selecting diverse anomaly detectors and combining their results to accurately locate anomalies in the MAWI archive. The synergy between detectors with different theoretical backgrounds allows a more accurate level of detection to be achieved.

Our contribution is twofold. Firstly, we thoroughly analyse the output of the previously proposed SVD-based combination strategy of anomaly detectors [8]. We are thus able to identify the main drawback that deteriorate the discrimination performance. We here name discrimination the classification of anomalies as harmless or dangerous ones. Secondly, we propose two new discrimination methods that outperform the original one. We then carefully evaluate these methods and determine which one is the best. In the end, our enhancement of the combination strategy in crease the ground truth data reliability.

The paper is structured as follows. Related work are presented in Section 2. The previously proposed method is briefly exposed in 3. Finally, our contribution is described and evaluated in Section 4. We conclude in Section 5.

2. RELATED WORK

Providing ground truth data to evaluate anomaly detectors is a challenge that has been addressed several times in the past. For example, the DARPA Intrusion Detection Evaluation Program [15] has been a great effort to provide labelled traffic to evaluate intrusion detection systems (IDS). It has been extensively studied, mainly through the KDD Cup 1999 data (KDD'99), and has been a profitable support for researchers. The main distinctions between this work and ours are the size of the network measured and the detectors to be evaluated. The DARPA Intrusion Detection Evaluation Program focuses on the evaluation of IDS and provides labelled LAN traffic where the packet payload is available and flows are complete. Whereas our work focuses on the evaluation of backbone traffic anomaly detectors and we provide labelled backbone traffic where the packet payload is not available, and the flows are incomplete and asymmetric. Furthermore, several critical drawbacks of the KDD'99 have been reported [16]. For example, the traffic data was captured in 1998 thus it contains no traffic from recent applications or anomalies. Therefore, this data must be carefully used as it is not representative of real traffic [27] and does not contain recent anomalies.

Closer to our work, Owezarski [19] recently proposed a data set containing real backbone traffic where anomalies are precisely located. In this work the traffic is captured at different points in the RENATER network, which is supposed to be anomaly free, and the researchers generate two kinds of anomalies (i.e. flash crowd and DDoS attack). Their experiment consist of different scenarios where the intensity of the anomalies varies. Thus, the sensitivity of the detectors to DDoS and flash crowd is easily identified. However, there are only a few kinds of anomalies in their data and they are not a realistic representation of the diverse anomalies found on the Internet. Due to privacy issues, their data is not downloadable and only accessible by visiting their laboratory.

3. COMBINING NETWORK ANOMALY DETECTORS

In order to provide reliable ground truth data, we take advantage of a combination of several anomaly detectors. The main advantages in combining anomaly detectors is that: (1) the detectors diversity allow us to broaden our analysis to a wide range of anomaly types and (2) the consensus view of independent detectors provides reliable results. The four selected anomaly detectors are based on different theoretical backgrounds.

They are presented in the next section. We will then explain how we combine their outputs.

3.1 Anomaly detectors

Principal component analysis.

Principal component analysis is dimension reduction method that has been extensively studied for detecting network traffic anomalies.

Principal component analysis (PCA) is an unsupervised technique highlighting the main features of the data. This is perhaps the most studied technique for anomaly detection in backbone traffic. It was first applied by Lakhina et al. [12], and it has received much attention in the last few years [13, 21, 22]. The key idea underlying a PCA-based anomaly detector is the extraction of the main features defining a normal traffic behaviour using PCA, then the distinct traffic is reported as anomalous.

An inherent problem with PCA-based detectors is the retrieval of the original anomalous traffic flows [21]. In

our experiments, we overcame this difficulty by using random projection techniques (sketches) [11, 14].

Gamma-law.

Dewaele et al. introduced an anomaly detection method based on sketching and multi-resolution gamma modelling [6]. In a nutshell, the traffic is split into sketches and modelled using Gamma distribution. Traffic that is distant from an adaptively computed reference is reported as anomalous.

Kullback-Leibler divergence.

The work presented in [4] detected the prominent changes in traffic by applying the Kullback-Leibler (KL) divergence to several kinds of histograms that monitor distinct traffic features. Furthermore, association rule mining allows for the extraction of the sets of traffic features that describes the anomalies detected by the histograms.

Hough transform.

The Hough transform is a pattern recognition technique that allows for the identification of a specific shape in a picture. This technique has been applied to several domains including anomaly detection of backbone traffic [9]. The approach proposed in [9] consists of first, monitoring the traffic in a 2-D scatter plot where the anomalous traffic appears as "lines", and second, identifies the anomalies with the Hough transform. The original data is retrieved from the identified plots, and the alarms reported by this method are aggregated sets of flows.

3.2 Alarms combination with SCANN

Once each detector has been run on traffic, we first use a graph-based similarity estimator to systematically uncovers the relations between the alarms reported by the detectors. We thus obtain several set of alarms where each set is associated with a single anomaly.

We then need to discern the sets of false positive alarms representing the harmless traffic from the sets of true positive alarms standing for the dangerous traffic. This classification is done using SCANN [17], an unsupervised combination strategy based on correspondence analysis.

Correspondence analysis [2] is a multivariate statistical technique for analysing multiway tables. It represents a data set in a lower-dimensional space based on Singular Value Decomposition (SVD). Although its role is similar to the principal component analysis one, correspondence analysis is designed for categorical data.

SCANN stores all the decisions of the detectors in a table, so that each entry is a vector representing the decision of all detectors for a certain set of alarms. This table is reduced with correspondence analysis, thereby, the entries are then smaller vectors containing only the main features characterizing the detectors decisions. The benefit of this reduced table is to take into account only significant decisions. For instance, a particularly irrelevant detector is one constantly making the same decision; in the first table built by SCANN this detector decisions are constant values that are then ignored in the reduced table because they do not help for discriminating sets of alarms.

Thus, the reduced table contains the characteristics of each set of alarm in a low-dimensional space. The original SCANN algorithm aims at combining classifier for n-class problems. In order to classify instances into these n classes, SCANN uses n reference points in the new low-dimensional space. Each instance is then assigned to the class associated to the closest reference point. In our case, we have a 2-class problem: normal or anomalous. We thus use two reference points which are two representative alarms either unanimously detected or not detected. At the end, the class of each alarm is determined by the closest representative set of alarms in the low-dimensional space.

4. IMPROVING SCANN RESULTS DISCRIMINATION

As presented in section 3.2, the combination method used, SCANN, creates a new space where each set of similar alarm is projected. In our case, we use two theoretical reference points inside this new space. First, we build a purely theoretical "normal" point which is the origin of the space and that would not have been reported by any detectors. This point is theoretical because, in our case, a set of alarms has at least been reported once. Second, we create an "abnormal" point that is an anomaly unanimously reported by the detectors as anomalous. The current method to classify alarms projected into this new space is to compare the distance of the considered set of alarms to these two reference points. The term "distance" used here actually refers to the Euclidean distance in the space created by SCANN's SVD. At this point, it is interesting to note that we apply SCANN on each traffic trace independently, the Euclidean Distance is therefore used on spaces of different dimensions. If a set of alarms is closer to the "normal" reference point, it will be considered as normal, i.e. harmless. Corollary, if it is closer to the "abnormal" point, it is abnormal.

We will first analyse how alarms are located in the new space generated by SCANN's SVD. We will then devise several methods that rely on our analysis to improve the discrimination between harmless and dangerous anomalies. Finally, we will compare our methods with the original one used in MAWILab.

4.1 SCANN combination method analysis

In order to better understand the way sets of alarms are mapped into space created by SVD, we first build

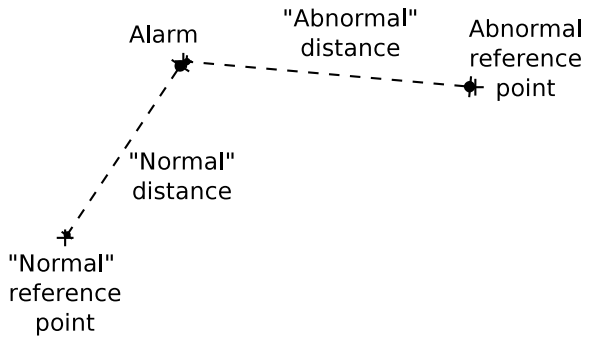

Figure 1: Example of anomaly location and distances in a 2-dimensional space yielded by SCANN's SVD.

a two-dimensional histogram regarding normalized distance to "normal" and "abnormal" reference points. We name these two distances "normal" distance and "abnormal" distance. These two distances are displayed on Figure 1 for an theoretic 2-dimensional space. Unless specified otherwise, the analysis and evaluation presented in this paper consider anomalies found in MAWI traces from 2003 to 2006. The distances are normalized regarding their respective maximum values inside each traffic trace.

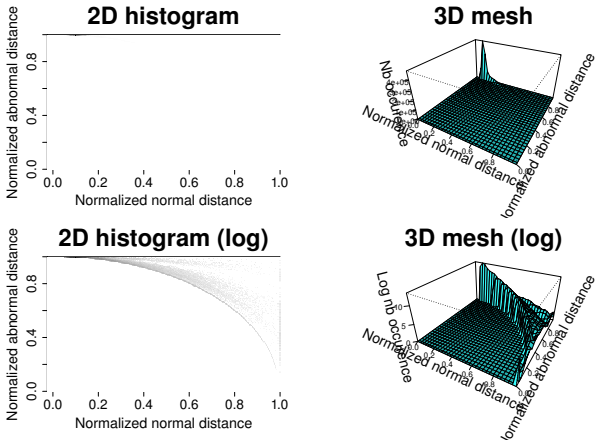

Figure 2: 2-dimensional histograms (the darker the point, the more alarms are located at the distance tuple) and 3D meshes (right) of the "normal" and "abnormal" distances of sets of alarms to the two corresponding reference point (bottom row uses log of occurrences).

Figure 2 shows that points closer to normal reference point are far from the abnormal reference point and vice versa. One can also notice that there is a clear minimum value for "abnormal" distance when the "normal" one is fixed and vice versa. This means that points are close to the line between the two references point. Furthermore, the majority of points are close to the "normal" reference point, and thus, far from the "abnormal" one.

Table 1: Heuristics labelling the traffic corresponding to a set of alarms into three categories ("Attack", "Special", and "Unknown"). These are originated from the anomalies previously reported [3, 9] and the manual inspection of MAWI.

Label	Category	Details
Attack	Sasser	Traffic on ports 1023/tcp, 5554/tcp or 9898/tcp
Attack	RPC	Traffic on port 135/tcp
Attack	SMB	Traffic on port 445/tcp
Attack	Ping	High ICMP traffic
Attack	Other attacks	Traffic with more than 7 packets and: SYN, RST or FIN flag $\geq 50\%$ Or, http, ftp, ssh, dns traffic with SYN flag $\geq 30\%$
Attack	NetBIOS	Traffic on ports 137/udp or 139/tcp
Special	Http	Traffic on ports 80/tcp and 8080/tcp with less than 30% of SYN flag
Special	dns, ftp, ssh	Traffic on ports 20/tcp, 21/tcp, 22/tcp or 53/tcp&udp with less than 30% of SYN flag
Unknown	Unknown	Traffic that does not match other heuristics

In order to understand where anomalies are located in the space, we then perform a breakdown of MAWILab anomalies regarding a heuristic-based classification. The used heuristics have been already applied in several previous studies on MAWI archive [3, 8] (cf. table 1). They classify anomalies into three different classes: Attack, Special and Unknown. We here consider Attack anomalies as the dangerous ones. They are therefore our target and our goal is to separate them from the harmless anomalies, which are here the anomalies annotated as Special and Unknown.

Figure 3a displays such breakdown for the "abnormal" distance. This figure shows that anomalies considered as "attack" by heuristics are closer to the "abnormal" reference point than the rest of flagged anomalies. Figure 3b exposes the anomaly breakdown for the "normal" distance. This figure highlights that anomalies considered as "attack" by heuristics are farther from the "normal" reference point than the rest of anomalies. These two figures mean that the majority of events are reported by small number of detectors. Corollary, only a minority of these anomalies are reported by many detectors.

We then perform a breakdown of anomalies flagged as "Attack" by heuristics regarding anomaly type. Figures 4a and 4b displays such breakdown regarding three anomalies: Sasser anomalies, Ping flood and SYN scan. Anomalies classified as "Attack" are close to the "ab-

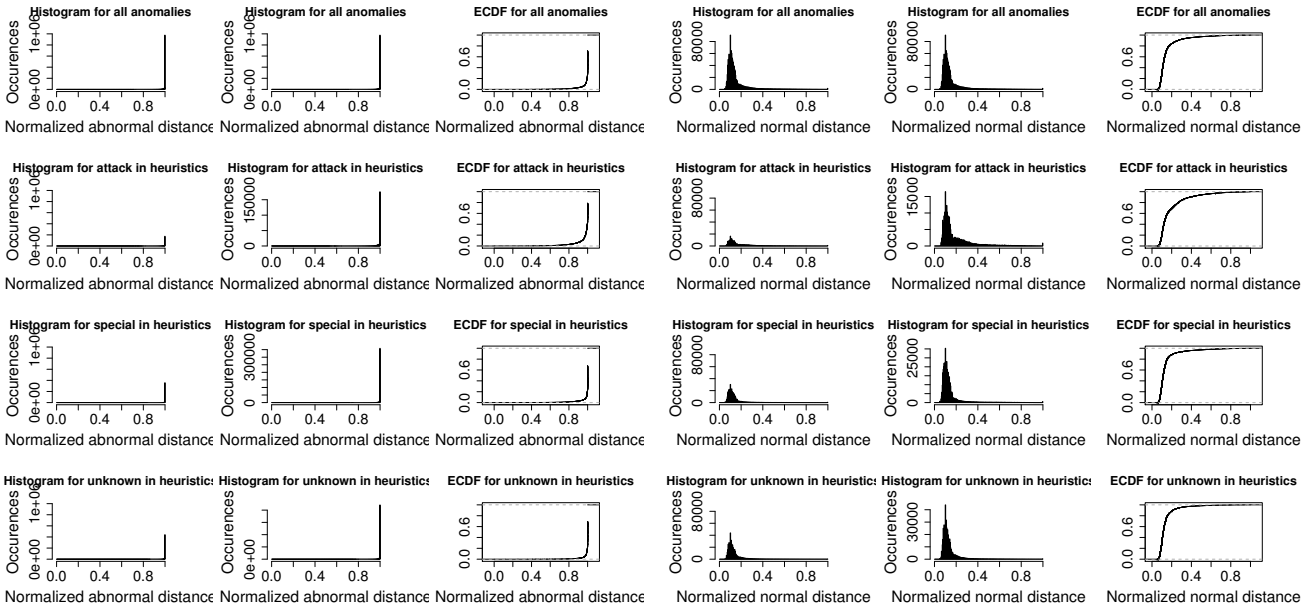

(a) Normalized "abnormal" distance

(b) Normalized "normal" distance

Figure 3: Histograms (first column with fixed common scale and second column with zoomed scale) and Empirical Cumulative Distribution Functions (ECDF) of distances regarding heuristics classification (first row, all anomalies, second row, attack, third row, special, fourth row, unknown).

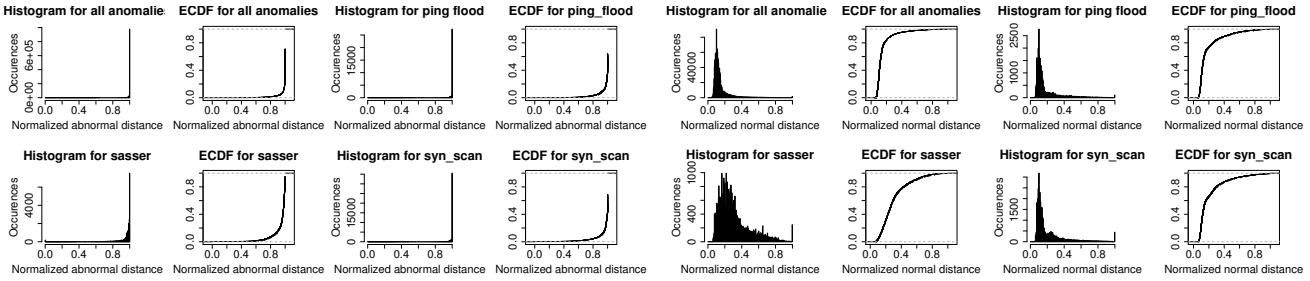

(a) Normalized "abnormal" distances

(b) Normalized "normal" distances

Figure 4: Histograms and ECDFs of distances regarding attacks (Sasser, Ping flood and SYN scan) in heuristics classification.

normal" reference point and far from the "normal" reference point. These results are consistent with the previously found trend.

This detailed study allow us to say that anomalies are close to the line between the two reference points. Moreover, anomalies flagged as "Attack" by heuristics are both, closer to the "abnormal" reference point, and, farther from the "normal" reference point, than the harmless anomalies.

This detailed study of the spatial distribution of alarms in the low-dimensional space tells us that this distribution is predictable. The strict adaptation of SCANN's class assignation technique (through comparison of distance to reference points) to our particular use case may be inadequate. In fact, in an n-class problem, SCANN makes no hypothesis on point locations in the low-dimensional space. While in our case, the study conducted clearly show that point distributions are predictable. We thus introduce two methods that intend to use this predictability to build a more reliable discrimination criteria.

4.2 Proposed discrimination methods

We then designed two methods to discriminate harmless alarms from dangerous one. The first of these methods uses the "normal" distance: every point farther

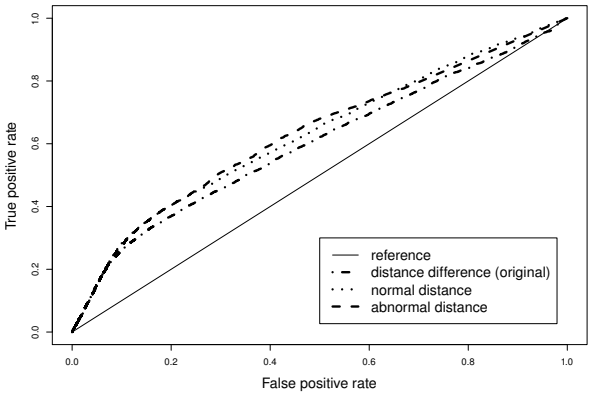

Figure 5: ROC curves for MAWI traces between 2003 and 2006 for discrimination methods.

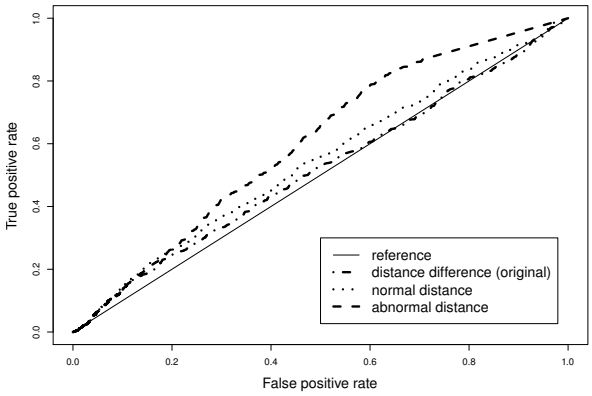

Figure 7: ROC curves for detection of Netbios anomalies in MAWI traffic of 2003 for all three discrimination methods.

from a defined threshold is considered as a real anomaly. The second method uses the "abnormal" distance: every alarm whose "abnormal" distance is smaller than a defined threshold is considered as abnormal. The next section evaluates the performance of these two methods.

4.3 Evaluation

We then evaluate our discrimination methods while considering the heuristics previously exposed as the ground-truth.

Figure 5 displays the results for anomalies found in MAWI traces from 2003 to 2006. We here use the ROC curves as introduced in the network anomaly detection field by [26] for three discrimination methods: the original one and the two method presented in section 4.2. ROC curves display the trade-off between true positive rate (proportion of anomalous instances detected as anomalous) and the false positive rate (proportion of normal instances detected as anomalous). A perfect curve would exhibit a perfect true positive rate with a null false positive rate. This would be represented on the plot as a step-shaped curved with a point located at the top-left corner. Figure 5 shows that both discrimination methods, either relying on either "abnormal" or "normal" distance, exhibit better performance than the original one that uses distance difference. The relatively low true positive rate and high false negative rate are explained by the fact that special and unknwon events are relatively close to attack in terms of distances.

Figure 6 displays a breakdown of performance regarding year. One can here notice that the discrimination method relying on "normal" distance exhibits better performance than the one using "abnormal" distance for all years except 2003.

In order to understand this discrepancy, we carefully analyse the anomaly landscape along the years. This analysis reveals that the majority of anomalies occurring in 2003 are Netbios anomalies due to the Blaster worm outbreak. This is not the case for the remaining years: 2004, 2005 and 2006.

Figure 7 shows ROC curves for Netbios, special and unknown anomalies only. It highlights the good performance of the "abnormal" distance discrimination method compared to the other methods to extract Netbios anomalies.

The discrepancy of "abnormal" distance-based discrimination method performance between 2003 and the others years can thus be explained by both the high number of Netbios anomaly in 2003 and the good performance of this discrimination method for that particular type of anomalies.

In the end, we can thus consider that the overall best performance is provided by the discrimination method based on "normal" distance.

5. CONCLUSION AND FUTURE WORK

We presented a detailed analysis of the combination step of a previously published anomaly detection results combination. We thus highlighted a drawback in this method: it fails to separate harmless anomalies from dangerous ones in the new space where alarms are projected by SVD. We thus proposed two new discrimination methods that cope with this limitation.

We then compared the performance of the original and proposed discrimination methods over four years of traffic from the MAWI repository. We are thus able to select one of the proposed method. It provides a substantial accuracy improvement in terms of anomaly detection. This improves the documentation provided by the MAWILab repository [1] concerning anomalies present in MAWI traces.

We intend to study more sophisticated methods that

[1]http://www.fukuda-lab.org/mawilab/

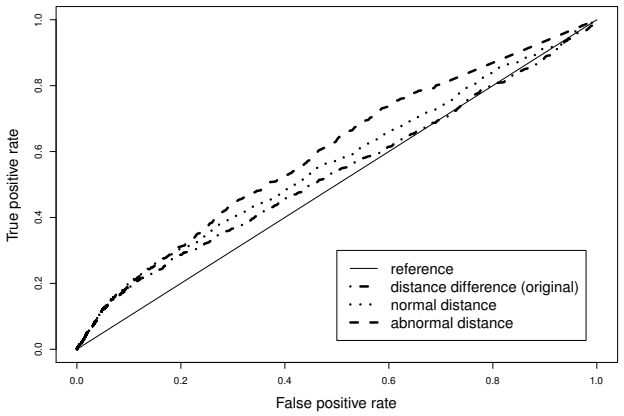

(a) ROC curves for 2003

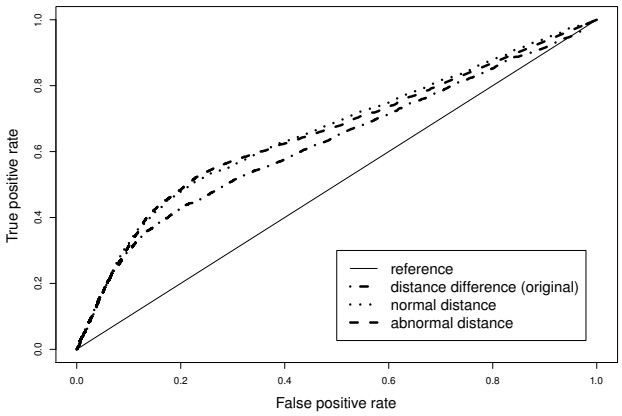

(b) ROC curves for 2004

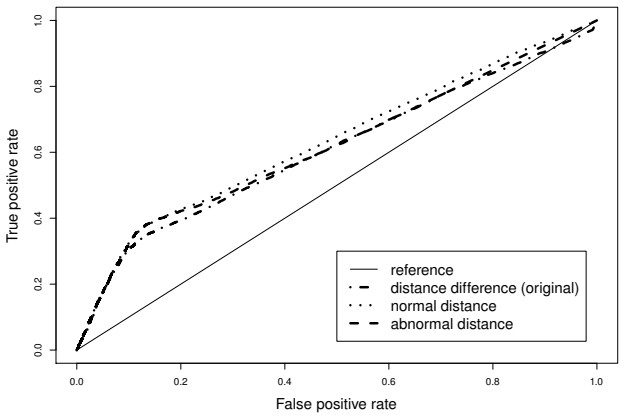

(c) ROC curves for 2005

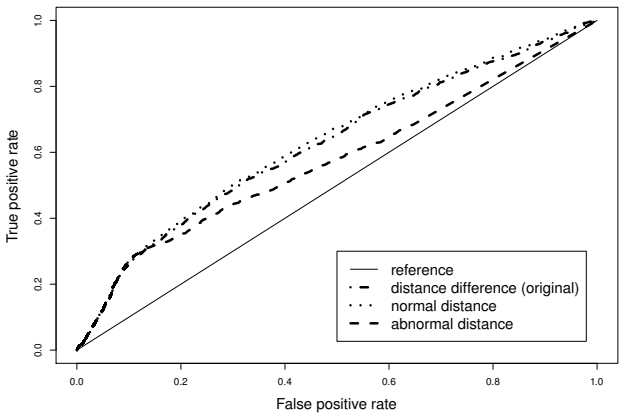

(d) ROC curves for 2006

Figure 6: Breakdown regarding year of ROC curves for discrimination methods.

will take into account more information about the detectors and the reported alarms. Thereby, we will further increase the discrimination power of the combination strategy and provide more reliable results for the MAWILab repository.

6. ACKNOWLEDGMENTS

This research has been funded by the National Institute of Informatics.

7. REFERENCES

[1] P. Barford, J. Kline, D. Plonka, and A. Ron. A signal analysis of network traffic anomalies. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurment*, IMW '02, pages 71–82, New York, NY, USA, 2002. ACM.

[2] J.-P. Benzécri. *Correspondence Analysis Handbook*. Marcel Dekker, New York, 1992.

[3] P. Borgnat, G. Dewaele, K. Fukuda, P. Abry, and K. Cho. Seven years and one day: Sketching the evolution of internet traffic. In *INFOCOM 2009, IEEE*, pages 711 –719, april 2009.

[4] D. Brauckhoff, X. Dimitropoulos, A. Wagner, and K. Salamatian. Anomaly extraction in backbone networks using association rules. In *Proceedings of the 9th ACM SIGCOMM conference on Internet Measurement Conference*, IMC '09, pages 28–34, New York, NY, USA, 2009. ACM.

[5] K. Cho, K. Mitsuya, and A. Kato. Traffic data repository at the WIDE project. In *USENIX 2000 Annual Technical Conference: FREENIX Track*, pages 263–270, 2000.

[6] G. Dewaele, K. Fukuda, P. Borgnat, P. Abry, and K. Cho. Extracting hidden anomalies using sketch and non gaussian multiresolution statistical detection procedures. In *Proceedings of the 2007 workshop on Large scale attack defense*, LSAD '07,

pages 145–152, New York, NY, USA, 2007. ACM.

[7] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM Transaction on Networking*, 9(4):392–403, 2001.

[8] R. Fontugne, P. Borgnat, P. Abry, and K. Fukuda. Mawilab: combining diverse anomaly detectors for automated anomaly labeling and performance benchmarking. In *Proceedings of the 6th International COnference*, Co-NEXT '10, pages 1–12, New York, NY, USA, 2010. ACM.

[9] R. Fontugne and K. Fukuda. A hough-transform-based anomaly detector with an adaptive time interval. In *Proceedings of the 2011 ACM Symposium on Applied Computing*, SAC '11, pages 471–477, New York, NY, USA, 2011. ACM.

[10] Y. Himura, K. Fukuda, K. Cho, and H. Esaki. An automatic and dynamic parameter tuning of a statistics-based anomaly detection algorithm. ICC '09, page 6, 2009.

[11] Y. Kanda, K. Fukuda, and T. Sugawara. Evaluation of anomaly detection based on sketch and pca. In *GLOBECOM 2010, 2010 IEEE Global Telecommunications Conference*, pages 1 –5, dec. 2010.

[12] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '04, pages 219–230, New York, NY, USA, 2004. ACM.

[13] A. Lakhina, M. Crovella, and C. Diot. Mining anomalies using traffic feature distributions. In *Proceedings of the ACM SIGCOMM 2005 conference*, SIGCOMM '05, pages 217–228, New York, NY, USA, 2005. ACM.

[14] X. Li, F. Bian, M. Crovella, C. Diot, R. Govindan, G. Iannaccone, and A. Lakhina. Detection and identification of network anomalies using sketch subspaces. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, IMC '06, pages 147–152, New York, NY, USA, 2006. ACM.

[15] R. Lippmann, J. W. Haines, D. J. Fried, J. Korba, and K. Das. The 1999 darpa off-line intrusion detection evaluation. *Computer Networks*, 34(4):579 – 595, 2000.

[16] J. Mchugh. Testing intrusion detection systems: a critique of the 1998 and 1999 darpa intrusion detection system evaluations as performed by lincoln laboratory. *ACM Transactions on Information and System Security*, 3(4):262–294, 2000.

[17] C. J. Merz. Using correspondence analysis to combine classifiers. *Machine Learning*, 36(1-2):33–58, 1999.

[18] G. Nychis, V. Sekar, D. G. Andersen, H. Kim, and H. Zhang. An empirical evaluation of entropy-based traffic anomaly detection. *IMC '08*, pages 151–156, 2008.

[19] P. Owezarski. A database of anomalous traffic for assessing profile based ids. In *International Workshop on Traffic Monitoring and Analysis (TMA '10)*, pages 59–72, 2010.

[20] H. Ringberg, M. Roughan, and J. Rexford. The need for simulation in evaluating anomaly detectors. *SIGCOMM Comput. Commun. Rev.*, 38(1):55–59, 2008.

[21] H. Ringberg, A. Soule, J. Rexford, and C. Diot. Sensitivity of pca for traffic anomaly detection. In *Proceedings of the 2007 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, SIGMETRICS '07, pages 109–120, New York, NY, USA, 2007. ACM.

[22] B. I. Rubinstein, B. Nelson, L. Huang, A. D. Joseph, S.-h. Lau, S. Rao, N. Taft, and J. D. Tygar. Antidote: understanding and defending against poisoning of anomaly detectors. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, pages 1–14, New York, NY, USA, 2009. ACM.

[23] B. I. Rubinstein, B. Nelson, L. Huang, A. D. Joseph, S.-H. Lau, N. Taft, and J. D. Tygar. Evading anomaly detection through variance injection attacks on pca. In *Proceedings of the 11th international symposium on Recent Advances in Intrusion Detection*, RAID '08, pages 394–395, Berlin, Heidelberg, 2008. Springer-Verlag.

[24] B. I. P. Rubinstein, B. Nelson, L. Huang, A. D. Joseph, S.-h. Lau, N. Taft, and D. Tygar. Compromising pca-based anomaly detectors for network-wide traffic. Technical Report UCB/EECS-2008-73, EECS Department, University of California, Berkeley, May 2008.

[25] A. Scherrer, N. Larrieu, P. Owezarski, P. Borgnat, and P. Abry. Non-Gaussian and Long Memory Statistical Characterisations for Internet Traffic with Anomalies. *IEEE Transaction on Dependable and Secure Computing*, 4(1):56–70, 02 2007.

[26] A. Soule, K. Salamatian, and N. Taft. Combining filtering and statistical methods for anomaly detection. In *Proceedings of the 5th ACM SIGCOMM conference on Internet Measurement*, IMC '05, pages 331–344, Berkeley, CA, USA, 2005. USENIX Association.

[27] M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani. A detailed analysis of the kdd cup 99 data set. *IEEE international conference on Computational intelligence for security and defense applications (CISDA '09)*, pages 53–58, 2009.