NUWC-NPT Technical Memo 22-057

Unmanned Aerial Vehicle (UAV) Cyber Attacks: Anomaly Detection Using Neural Network-Based Intrusion Detection

Sherida Jacob
Undersea Warfare Combat Systems Department

31 August 2022

Naval Undersea Warfare Center Division
Newport, Rhode Island

ADMINISTRATIVE INFORMATION

This technical memo was prepared under NUWC Division Newport Network Activity No. 300000168428/0010, principal investigator Sherida Jacob (Code 2511). The sponsoring activity is the NUWC Internal Investment Program, which is managed by Derek K. Potvin (Code 00X) of the NUWC Strategic Planning Office.

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
31-08-2022	Technical Memo	START DATE	END DATE 31-08-2022

4. TITLE AND SUBTITLE

Unmanned Aerial Vehicle (UAV) Cyber-Attacks: Anomaly Detection Using Neural Network-Based Intrusion Detection

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
N/A	N/A	N/A
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**
N/A	N/A	N/A

6. AUTHOR(S)

Sherida Jacob

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
NUWC Division Newport Code 25 – Undersea Warfare Combat Systems Department 1176 Howell, Newport, RI, 02841	NUWC-NPT Technical Memo 22-057

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)
NUWC Division Newport Code 00X – Internal Investmentst 1176 Howell, Newport, RI, 02841	NUWC Division Newport	NUWC-NPT Technical Memo 22-057

12. DISTRIBUTION/AVAILABILITY STATEMENT

Distribution Statement A. Approved for public release: distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The use of unmanned aerial vehicles (UAVs) in civilian, commercial, and military applications has grown and continues to grow. An emerging area of interest for UAVs is for several to form flying ad-hoc networks (FANETs). In both military and non-military settings, FANETs can be used for tasks such as extending communication networks, managing crops, or conducting surveillance. However, with the increased use of UAVs in the national air space, it is important to prevent or mitigate cyber-attacks such as spoofing or anomaly insertion.

Neural network-based intrusion detection techniques provide a way to potentially identify anomalous data that indicates a cyber-attack. Although well-trained neural networks can classify data and make predictions with a high degree of accuracy, they are not infallible.

This research explored the potential of using hierarchical clustering techniques to identify anomalous data points within neural network hidden layer activations, and identifying the nodes in which the activations adversely impacted the results.

15. SUBJECT TERMS

UAV, unmanned aerial vehicle, FANET, flying ad-hoc network, cyber, cybersecurity, cyber attack, spoofing, neural network, machine learning, intrusion detection

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
(U)	(U)	(U)	(U)	15

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Sherida Jacob (Code 2511)	401-832-4776

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

AUC	Area under the ROC
CHAODA	Clustered Hierarchical Anomaly and Outlier Detection Algorithms
CLAM	Clustered Learning of Approximate Manifolds
DoD	Department of Defense
FAA	Federal Aviation Administration
FANET	Flying ad-hoc network
GPS	Global positioning system
MNIST	Modified National Institute of Standards and Technology
NUWC	Naval Undersea Warfare Center

LIST OF ABBREVIATIONS AND ACRONYMS (Cont'd)

PCA	Principal component analysis
ROC	Receiver operating characteristic
UAV	Unmanned autonomous vehicle
UMAP	Uniform manifold approximation and projection

1. INTRODUCTION

In the United States, the use of unmanned autonomous vehicles (UAVs) is only growing in military, commercial, and civilian sectors. This increase is due to their declining costs and advances in technology.[1] The Department of Defense (DoD) has invested over an estimated $34 billion dollars in the research and development, procurement, and construction of drones between 2013 and 2018 [1]. For non-military drones, the Federal Aviation Administration (FAA) registered over 1.4 million drones as of November 2019 [2]. Of those 1.4 million, 1 million drones were categorized as hobby or recreational aircrafts, and over 400,000 drones were categorized as non-hobby (commercial, media, government, etc.) aircrafts. Based on the registration trends, overall market, and operational information, the FAA predicts that registration will surpass 1.39 million for recreational drones and 835,000 for non-hobby drones by 2023 [2].

However, with the increased use of UAVs, it is important to minimize their vulnerability to cyber-attacks. Depending on the mission, UAVs can collect and transmit sensitive data. Neural networks are used to process this data to perform image classification, object detection, etc. In addition, UAVs range in size from 1 pound to over 4,000 pounds, and can reach an altitude higher than 50,000 feet. If a UAV falls from the sky or is used in an unintended manner, it can cause loss of life or severe property damage. In recent years, there have been a few examples of UAVs being a threat to public safety or national security. In 2009, Iraqi militants exploited an unsecured communications link and intercepted "live video feeds from U.S. Predator drones" [4]. In 2011, at Creech Air Force Base in Nevada, a keylogging virus was discovered on the ground control systems for Predator and Reaper drones [5]. In 2013, a UAV crashed into Manhattan, NY, narrowly missing a pedestrian [6]. In 2015, drug cartels spoofed and jammed Department of Homeland Security (DHS) drones near the border of Mexico [7].

1.1 NEURAL NETWORKS AND ADVERSARIAL INPUT ATTACKS

Deep neural networks are a series of algorithms designed to imitate the functions of a human brain. They can model non-linear, complex relationships (while building on previous knowledge) in order to process and interpret large amounts of complex data. The nodes in a neural network consist of weights and biases and act like neurons. The nodes serve as learning parameters that are modified as the network is trained. Neural networks "learn" when the nodes interact to modify their values to produce the desired output.

Research shows that—although neural networks are successful at pattern recognition, clustering, and classification—they do not understand the "underlying characteristics" of the data. Therefore, they struggle with tasks that require reasoning, and are susceptible to adversarial examples [8]. Adversarial examples are inputs that are nearly identical to data that was correctly categorized previously, but contain noise. This noise tricks the network into miscategorizing the input with a high level of confidence. Malicious adversarial examples trick the neural network into misclassifying the input as the attacker's desired output.

[1] This includes a longer battery life as well as technology improvements in cameras, sensors, software, and hardware.

Many neural networks are trained using the same datasets, so they learn the same way, and this makes them vulnerable to adversarial inputs. One way to mitigate adversarial input attacks is to use neural network-based intrusion detection techniques to identify the noise that causes the misclassification of data. The characteristics of the noise can then be used to train the network to identify anomalous inputs. Clustered Learning of Approximate Manifolds (CLAM) and Clustered Hierarchical Anomaly and Outlier Detection Algorithms (CHAODA) are techniques that work in conjunction to identify anomalies in data.

CLAM learns the general geometric and topological properties of the data, and then attempts to build graphs that map the underlying manifold. It works by divisively clustering the dataset until each cluster contains only one data point. CHAODA is an unsupervised algorithm that implements six anomaly detection algorithms to explore the properties of the graphs to find anomalies. The CHAODA algorithms assess relative cluster cardinality, relative component cardinality, graph neighborhood, the child-parent cardinality ratio, random walks, and stationary probability. In general, data with similar characteristics are clustered together. Small clusters and data points that are further away from other data points are considered anomalous. For each algorithm, an outlier score is assigned to each point in the data set. The mean of all the scores for each point create an AUC score. Note that a low AUC score is indicative of a model that cannot correctly classify data.

An unsupervised anomaly detection algorithm that implements clustering and manifold mapping utilizing the geometric and topological properties of data is an alternative way of identifying anomalous data in high dimensional spaces. The Navy has systems that process large amounts of complex data. Neural networks can assist the warfighter in processing the data efficiently and accurately. A neural network that uses unsupervised algorithms can process large amounts of complex, unlabeled data, and classify the data by identifying patterns that are not obvious to a human. Implementing appropriate clustering techniques can provide insight on how the data relates to each other and how new data relate to known data. The anomaly detection algorithm explored in this research can be used to verify the output of a neural network.

When creating a neural network that can accurately classify data, researchers must ensure that the model is not overfitting the data. Overfitting occurs when the neural network memorizes the data it was trained on instead of generalizing the data. When this occurs, the neural network cannot make accurate predictions on new data. Implementing the CHAODA algorithm on data can assist in verifying that a neural network is making correct classifications. If the neural network does not detect anomalies detected by CHAODA, the researcher can explore the causes for the discrepancy and adjust the neural network to improve the accuracy of the results.

2. RESEARCH LIMITATIONS

The original intent of this research was to investigate the potential to use hierarchical clustering techniques and neural networks to identify anomalous global positioning system (GPS) data in Flying Ad-Hoc Networks (FANETs). Due to the lack of available open-source UAV GPS data, the research focus shifted to gaining a better understanding of how hidden layer activations can provide insight into how neural networks make decisions.

3. EXPERIMENT

To gain insight into the neural network decision-making process, this research utilized Python to implement a feed-forward neural network using back propagation. Uniform manifold approximation and projection (UMAP) was used to plot the resulting activations at each layer. Activations represent meaningful features of the data. The Modified National Institute of Standards and Technology (MNIST) dataset was used as the input data. The MNIST dataset is a well-established dataset of handwritten digits. Miscategorizations of the input images can easily be reviewed for clarity to determine why it was miscategorized. Initial attempts to create the network from scratch were unsuccessful. The plots did not properly reflect the way the network interpreted activations for different classes within a layer (see Figure 1).

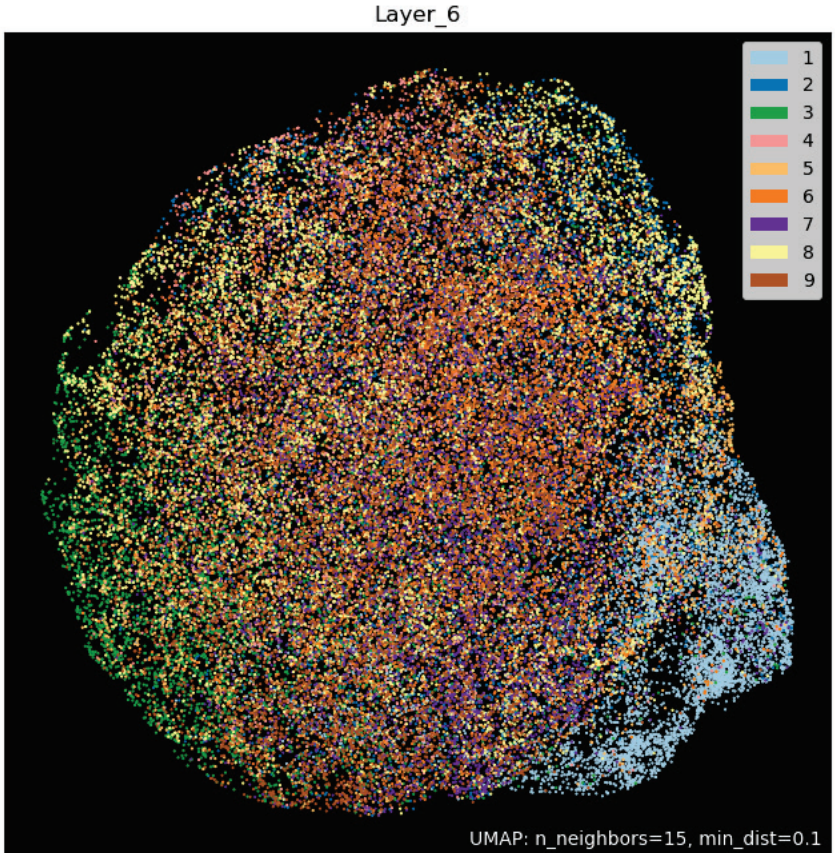

Figure 1. Incorrect Neural Network Activation Clusters

To correct this error, established neural network functions in tensorflow were used to build a multi-layered network. The get_activations function in the Python keract module was used to extract the hidden-layer activations. The results of these implementations are in Figures 2 to 6.

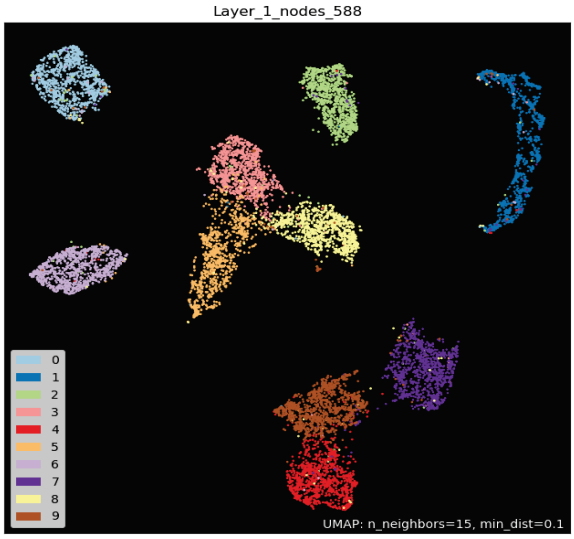

Figure 2. Activation Clusters of 588 Nodes from the First Hidden Layer

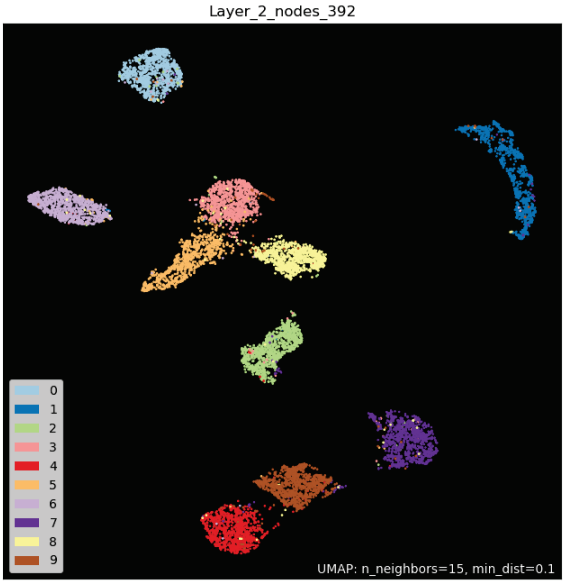

Figure 3. Activation Clusters of 392 Nodes from the Second Hidden Layer

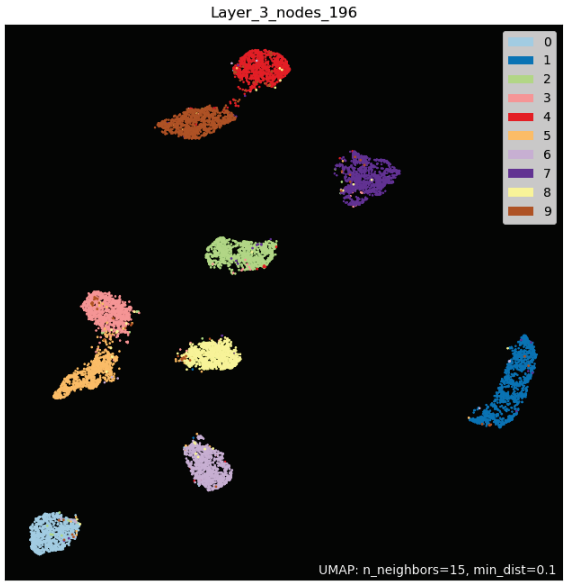

Figure 4. Activation Clusters of 196 Nodes from the Third Hidden Layer

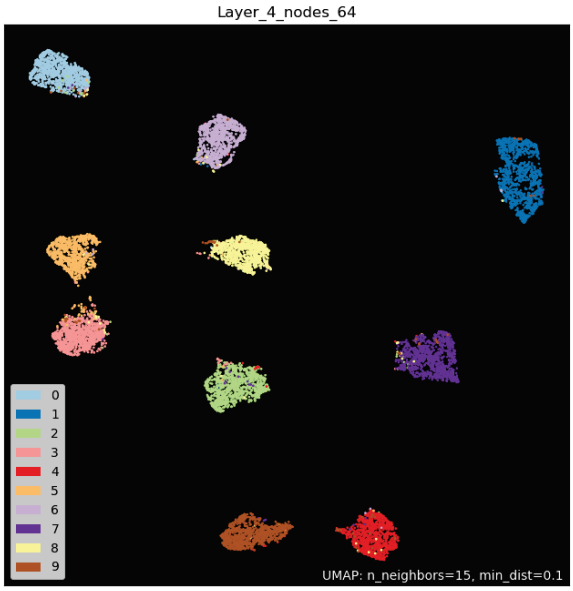

Figure 5. Activation Clusters of 64 Nodes from the Fourth Hidden Layer

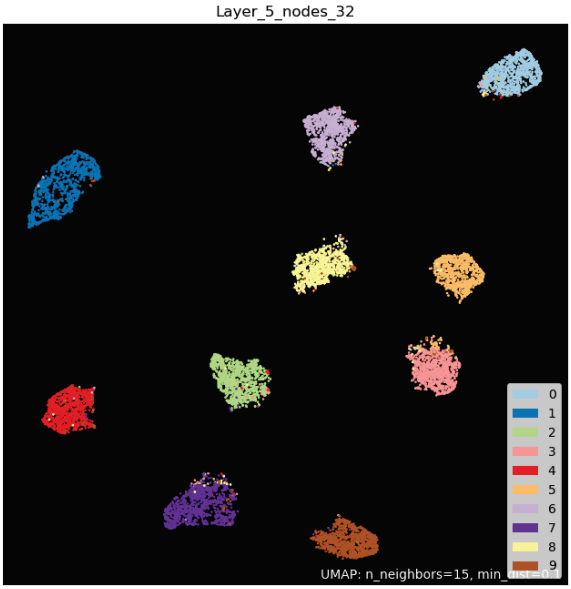

Figure 6. Activation Clusters of 32 Nodes from the Fifth Hidden Layer

The plots in Figures 2 through 6 also show that, with each progressive layer, the neural network improves its ability to recognize the patterns of each digit. The manifolds upon which each class of digit exists becomes more defined and compact, and they also become further separated from each other.

At each layer, the activations of the inputs that were miscategorized were assessed for anomalous characteristics. Initial assessment of the activation values indicate that the nodes may provide insight into where the anomalous characteristics exist, and for each node, the range of values that should provide an alert for further analysis. Tables 1 to 3 show the activation values for multiple nodes in Layer 2.

Table 1. Layer 2 Node 12: Activation Values

Labeled as Digit 9/ Identified as Digit 9	Labeled as Digit 9/ Identified as Digit 8	Labeled as Digit 9/ Identified as Digit 5	Labeled as Digit 9/ Identified as Digit 3
0.850808	0.313349	0.427454	0.675887
0.679745	0.755537	0.610466	0.687527
0.849429	0.562109	0.750866	0.470647
0.791211	0.611243	0.467977	0.681889
0.991022	0.237689	0.55892	0.669789
0.720284	0.512944	0.455019	0.604119

Table 2. Layer 2 Node 110: Activation Values

Labeled as Digit 9/ Identified as Digit 9	Labeled as Digit 9/ Identified as Digit 8	Labeled as Digit 9/ Identified as Digit 5	Labeled as Digit 9/ Identified as Digit 3
0	0	0	0
0	0	0	0
0	0	0	0
0	0	0	0
0	0	0	0
0	0	0	0

Table 3. Layer 2 Node 392: Activation Values

Labeled as Digit 9/ Identified as Digit 9	Labeled as Digit 9/ Identified as Digit 8	Labeled as Digit 9/ Identified as Digit 5	Labeled as Digit 9/ Identified as Digit 3
0.189244	0.506983	0.606051	0.401707
0.131373	0.405798	0.537039	0.349792
0.150368	0.211779	0.323686	0.433142
0.176075	0.319907	0.308523	0.374977
0.163472	0.338173	0.169961	0.214903
0.184493	0.500676	0.124963	0.413454

The formal assessment the activation values for each node in each layer using hierarchical clustering techniques such as CLAM/CHAODA is still under development.

4. CONCLUSION

This research explored the potential of using hierarchical clustering techniques to identify anomalous data points within neural network hidden layer activations, and identifying the nodes in which the activations adversely impacted the results.

Investigating a methodology to identify adversarial inputs is important because automated image classification and object detection can provide additional support to the war fighter when conducting UAV missions.

Conducting this research enabled the author to become proficient in using Python to implement a basic neural network, creating visualizations of the output, and using anomaly detection algorithms to analyze the output. One of the challenges encountered when implementing the code to support this research was that the CHAODA code was difficult to read and implement. CHAODA requires the guidance and periodic support from the original creators in order to implement it according to design.

Recommended future work should include:

1. Preprocessing the data by applying principal component analysis (PCA) before inputting it into the neural network. PCA reduces the dimensionality of the data, which in turn would improve the neural network runtime and classification accuracy.
2. Export or recreate the CHAODA code into a stable code base. This minimizes the reliance on open-source code created, maintained, and periodically altered by graduate students.
3. Create adversarial examples and explore how the neural network interprets the data prior to classification.
4. Observe the accuracy of the neural network and CHAODA algorithm in identifying the adversarial examples as anomalous data.

REFERENCES

[1] D. Gettinger, "Drones in the Defense Budget: Navigating the Fiscal Year 2018 Budget Request," Center for the Study of the Drone at Bard College, Annandale-on-Hudson, NY, October 2017. Accessed on May 2021. Available: https://dronecenter.bard.edu/files/2018/01/Drones-Defense-Budget-2018-Web.pdf

[2] U.S. Department of Transportation, "UAS Sightings Report." Accessed on May 2021. Available: https://www.faa.gov/uas/resources/public_records/uas_sightings_report

[3] U.S. Department of Transportation, "FAA Aerospace Forecast Fiscal Years 2021–2041." Accessed on May 2021. Available: https://www.faa.gov/Data_research/Aviation/Aerospace_forecasts/Media/Unmanned_aircraft_systems.Pdf

[4] S. Gorman, Y.J. Dreazen, and A. Cole, "Insurgents Hack U.S. Drones," *Wall Street Journal*, 17 December 2009. Accessed on May 2021. Available: https://www.wsj.com/articles/SB126102247889095011

[5] K. McCaney, "US Drone Fleet Infected by Virus," *Government Computer News (GCN)*, 11 October 2011. Accessed on May 2021. Available: https://gcn.com/cybersecurity/2011/10/us-drone-fleet-infected-by-virus/282462/

[6] M. McNulty, "New Video: Drone crash lands in Manhattan," *New York Post*, 3 October 2013. Accessed on May 2021. Available: https://nypost.com/2013/10/03/video-captures-drones-flight-above-manhattan/

[7] C. Thompson, "Drug traffickers are hacking US surveillance drones to get past border patrol," Business Insider, 30 December 2015. Accessed on May 2021. Available: https://www.businessinsider.com/drug-traffickers-are-hacking-us-border-drones-2015-12

[8] I.J. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and Harnessing Adversarial Examples," arXiv:1412.6572, 20 March 2015. DOI: https://doi.org/10.48550/arXiv.1412.6572

BIBLIOGRAPHY

Brems, M., "A One-Stop Shop for Principal Component Analysis," *Towards Data Science*, 17 April 2017. Accessed on March 2022. Available: https://towardsdatascience.com/a-one-stop-shop-for-principal-component-analysis-5582fb7e0a9c

Coenen, A. and A. Pearce, "*Understanding UMAP*," People + AI Research (PAIR), Google Research, 2020. Accessed on March 2022. Available: https://pair-code.github.io/understanding-umap/

Ishaq, N., T. J. Howard and N. M. Daniels, "Clustered Hierarchical Anomaly and Outlier Detection Algorithms," *2021 IEEE International Conference on Big Data (Big Data)*, Orlando, FL, USA, 2021, pp. 5163-5174, DOI: 10.1109/BigData52589.2021.9671566.

Office of the Secretary of Defense, "Unmanned Aerial Vehicles Roadmap (2002-2027)", Department of Defense, Washington, DC, December 2002.

INITIAL DISTRIBUTION LIST

Internal
Codes:
1033 Corporate Research and Information Center (CRIC)

External
Defense Technical Information Center (DTIC)

Total: 2