THESIS REVIEWER’S REPORT

I. IDENTIFICATION DATA

Thesis title:	**Graph Generative Models for Decoy Targets in Active Directory.**
Author’s name:	**Ondrej Lukas**
Type of thesis :	master
Faculty/Institute:	Faculty of Electrical Engineering (FEE)
Department:	Computer Science
Thesis reviewer:	Fabrizio Biondi
Reviewer’s department:	Avast Software s.r.o.

II. EVALUATION OF INDIVIDUAL CRITERIA

Assignment **challenging**

How demanding was the assigned project?

The assigned project consists of the exploration of Machine Learning (ML) methods for the automated placement of honeypot users into an Active Directory (AD) structure, with the aim to maximize the probability that such users will detect an attacker trying to move laterally. At the same time, the honey users must be sufficiently similar to normal users that the attacked does not find them suspicious.

The project is quite demanding in that it requires an advanced knowledge of both theoretical subjects like ML and attacker modeling and more applied subjects like AD architecture and dataset generation. Apart from mastering multiple computer science subjects, the candidate had to devise a way to create a dataset representative of a real company’s AD to run experiment on and evaluate the quality of such reconstruction, ad detailed in chapters 5 and 6. Finally, the student had to find ways to evaluate the core generative experiments that show the approach proposed by the project.

To conclude, the project required knowledge in various fields both theoretical and practical, understanding of various data analysis and experimental evaluation procedures, and excellent code-writing skills, all of which classify it as a challenging project.

Fulfilment of assignment **fulfilled**

How well does the thesis fulfil the assigned task? Have the primary goals been achieved? Which assigned tasks have been incompletely covered, and which parts of the thesis are overextended? Justify your answer.

The thesis fulfills the assigned task very well. In particular, it covers the task of reconstructing a reasonable AD graph with appropriate detail, which could have been sidestepped by a less conscientious author. Overall, the thesis is thorough in investigating the assigned task and satisfying its requirements. The background is well presented, and helps the reader understand both the concepts required to follow the rest of the thesis and to appreciate the complexity of the task. Similarly, the related work chapter agues the originality of the work, presenting multiple closely related works and presenting the differences with the contribution in the thesis. The thesis presents multiple models for the whole task so that the models can be all tested and evaluated, providing insight on the best architecture to solve the assigned problem. The evaluation sections proceed to execute the experiments and commenting its results, considering the Variational Autoencoder model as the best performing one. Finally, the source code of the thesis has been published, allowing the scientific community to reproduce the results.

Methodology **correct**

Comment on the correctness of the approach and/or the solution methods.

The models proposed do a good job of representing the AD modeling problem, and the encoding devised to model correctly the graphs inside the RNN is likewise quite advanced and aware of the recent state of the art in multi-instance problems. Similarly, the experiments evaluate first the quality of the graph reconstruction itself, and then the performance of the models to add honey users to the graphs, finally delivering strong evidence of the author’s thesis on the possibility of using ML for such tasks.

In tasks such as the one presented by this thesis, where limited related work exists examining the same problem with the same metrics, it is very common for research to use standard metrics that are not necessarily connected with the goals of the project. Hence, I have particularly appreciated the effort in sections 6.1 and 7.1 to present particular metrics and discuss their relevance concerning the specific tasks of graph reconstruction and generative honey user assignment. The selection and discussion of useful metrics is paramount for the experimental results to really be meaningful and lead to the conclusions of the work.

Technical level **A - excellent.**

Is the thesis technically sound? How well did the student employ expertise in the field of his/her field of study? Does the student explain clearly what he/she has done?

The technical level of the thesis is definitely very high. The author has mastered multiple subjects, both concerning practical security concerns and theoretical applications of computer science, and has used multiple concepts with harmony and confidence to devise proper experiments and evaluate the proposed models. The code for the project has been developed using GPU-accelerated Tensorflow 2 technology, representing the state of the art, and has been released for scrutiny and reproducibility. Similarly, both the encoding and the generative models applied represent the state of the art of neural network research, demonstrating the author's mastery of these very advanced concepts.

On the applicability level, it is very positive that the author has applied these advanced concept to a very practical security concern, since this has forced him to consider practical problems and implications that are lost to researchers working mostly on pure ML research.

Formal and language level, scope of thesis **B - very good.**

Are formalisms and notations used properly? Is the thesis organized in a logical way? Is the thesis sufficiently extensive? Is the thesis well-presented? Is the language clear and understandable? Is the English satisfactory?

The thesis is well presented and all formalisms are conveyed using the appropriate notation and descriptions. I found the Background section sufficient to understand the thesis' contribution, making the thesis self-contained. I had no problem understanding and following the academic English of the thesis, and understanding the choices of the author and the proposed methodology and conclusions.

On the scope of the thesis, I would have appreciated a more thorough discussion of the practical applications of the results. More specifically: the thesis takes a very practical problem and proposes an ML-based solution, concluding by evaluating the different performance of different ML models to solve the problem. However, it does not discuss in much detail the reverse transformation from the theoretical solution proposed by the ML models to implementation in a real-life protection environment. For instance, would it be sufficient to generate the honey users once? How would this change with the evolution of the AD structure? What is a reasonable probability estimation of catching an attacker based on the application of this technique? While some of these points are quickly touched in the Future Work section, I feel they should have been given more thought for completeness.

Selection of sources, citation correctness **B - very good.**

Does the thesis make adequate reference to earlier work on the topic? Was the selection of sources adequate? Is the student's original work clearly distinguished from earlier work in the field? Do the bibliographic citations meet the standards?

The thesis presents related work mostly on the problem definition and techniques used, since the only comparable work is the recent paper by Siniosoglou et al. hence there is not much more to compare with on the specific task. This shows the originality of the task and its separation from other solutions in the field. The bibliography, while not extensive, is appropriate for the thesis and properly formatted and referenced in the text.

Additional commentary and evaluation (optional)

Comment on the overall quality of the thesis, its novelty and its impact on the field, its strengths and weaknesses, the utility of the solution that is presented, the theoretical/formal level, the student's skillfulness, etc.

No additional comment.

III. OVERALL EVALUATION, QUESTIONS FOR THE PRESENTATION AND DEFENSE OF THE THESIS, SUGGESTED GRADE

Summarize your opinion on the thesis and explain your final grading. Pose questions that should be answered during the presentation and defense of the student's work.

To summarize, the thesis is a meticulous work to encode a practical security problem into a formal model and to compare different ML models to find a solution, and succeeds in the task of convincing the reader of the performance of such ML models to solve the problem. The author demonstrates a high level of mastery of both security and ML subjects, together with coding skills in state-of-the-art GPU-accelerated ML implementation languages. The author shows judgment and awareness in choosing and discussing evaluation metrics for the experiments that really represent the performance of the models in the practical problem space, and correctly interprets the experimental results to conclude the superiority of the Variational Autoencoder model for the task. While I would have appreciated more thought on how to apply the results in the practical case, this can be addressed in the discussion or in future work.

During the thesis defense, I would propose the following questions:

1) "The core problem of the thesis concerns placing honey users into an Active Directory structure to maximize the probability of detecting an attacker trying to perform lateral movement. Can you comment on how the placement of such honey user in the directory affects the probability that an attacker is detected? What are good and bad practices for the task of creating honey users and placing them in the AD structure?"
2) "In Section 4.3.4 you introduce a loss function for the training of the ML models. Please detail the reasoning behind the choice of this loss function, what alternative function shapes could you have used, and what would have been the impact on the experimental results."
3) "Please detail the limitations that you have encountered on creating a dataset for the thesis without access to a real-size AD structure, how you overcame them, and what limitations remain in the approach that you have chosen."
4) "In both the graph reconstruction experiments and the generative experiments, you have proposed quality metrics for the evaluation of the three proposed models and rated the models according to these metrics. Please comment on the metrics you have chosen, how they really represent the relative performance of the models, and what decisions you have taken into consideration while determining which metrics to use"
5) "The thesis mostly concludes by declaring the superior performance of the Variational Autoencoder model into solving the model of the problem. Can you comment on what concerns and challenges would raise by applying the results of the thesis to the AD of a real medium-sized company, for instance considering the evolution of the AD structure, technical concerns, et cetera? Would the solution have to be generated only once or regularly? Would it affect other security and day-to-day operations in the company? Would it be expensive or complex to deploy or maintain?"

The grade that I award for the thesis is **A - excellent.**

Date: **2.9.2020** Signature: Fabrizio Biondi