Faking and Discriminating the Navigation Data of a Micro Aerial Vehicle Using Quantum Generative Adversarial Networks

Michel Barbeau
Carleton University, Canada
0000-0003-3531-4926

Joaquin Garcia-Alfaro
Telecom SudParis, France
0000-0002-7453-4393

***Abstract*—We show that the Quantum Generative Adversarial Network (QGAN) paradigm can be employed by an adversary to learn generating data that deceives the monitoring of a Cyber-Physical System (CPS) and to perpetrate a covert attack. As a test case, the ideas are elaborated considering the navigation data of a Micro Aerial Vehicle (MAV). A concrete QGAN design is proposed to generate fake MAV navigation data. Initially, the adversary is entirely ignorant about the dynamics of the CPS, the strength of the approach from the point of view of the bad guy. A design is also proposed to discriminate between genuine and fake MAV navigation data. The designs combine classical optimization, qubit quantum computing and photonic quantum computing. Using the PennyLane software simulation, they are evaluated over a classical computing platform. We assess the learning time and accuracy of the navigation data generator and discriminator versus space complexity, i.e., the amount of quantum memory needed to solve the problem.**

I. Introduction

CPSs comprise physical processes monitored and controlled through embedded computing and networked resources. Signals to actuators and feedback from sensors are exchanged with controllers using, e.g., wireless communications. The advantages of such architectures include flexibility and relatively low deployment costs. Nevertheless, the perpetration of cyber-physical attacks must be addressed. The problem is particularly challenging when the CPS consists of disruptive technologies such as MAVs, Unmanned Aerial Vehicles (UAV) and swarming.

Today's cybersecurity solutions, from in-depth defense techniques (e.g., firewalls) to intrusion detection and cryptographic techniques, aim to prevent system breaches from happening. However, several stories of attacks and disruption of CPS exist, e.g., from the Stuxnet worm incident affecting a Iran's *atomic program* [7] to recent incidents in Saudi Arabia affecting Houthi drones [9]. CPS protection solutions must manage and take control over adversarial actions. Protection must be built taking on the adversary mindset, predicting its intentions and adequately mitigating the effects of its actions.

In this paper, we explore the use of the QGAN paradigm to address cyber-physical security issues in the domain of MAVs. A concrete QGAN design is proposed to generate fake MAV navigation data. Initially, the adversary is entirely ignorant about the dynamics of the CPS. From the point of view of the adversary, it is the strength of the approach. A design is also proposed to discriminate between real and fake MAV navigation data. The designs combine classical optimization, qubit quantum computing and photonic quantum computing. We build upon the PennyLane quantum machine learning software platform [3]. In particular, we reuse and adapt ideas from the variational classifier [21] and QGAN [22] examples.

We evaluate our approach using the simulation capabilities of PennyLane. We measure the learning time and accuracy of the navigation data generator and discriminator with respect to the space complexity, i.e., the amount of quantum memory used to solve the problem. At the outset, we acknowledge that the exponentially growing time complexity in the number of qubits of our solution is a barrier to its application on a large scale. In particular, when the calculations are all done in simulation over a classical computing platform. Nevertheless, we show the feasibility of the approach on a small scale and identify hurdles that are likely to be overcome by the upcoming evolution of quantum machine learning.

Section II elaborates further on our problem domain and related work. Section III presents our solution. Section IV provides experimental work. Section V concludes the paper.

II. Problem Domain

The problem domain encompasses CPS controllers, playing the role of defenders, and adversaries. We conceptualize the situation in terms of activities consisting of gathering and hiding knowledge about both defensive and adversarial strategies. We envision the use of new learning theories, in which defenders and adversaries conceal their actions to avoid being profiled for the purpose of thwarting their *cyber-physical battle weapons*. Defenders equipped with Artificial Intelligence (AI) tools, such as machine learning, can identify adversarial actions trying to collect as much knowledge as possible about their targets. The defense starts by learning the weaknesses of the adversaries and offensively mislead their intentions, thwarting their actions in the end. Once the defender knows the adversary, e.g., the behavior performed to identify and disrupt the services, the defender starts offering assets sacrificed to coax the adversary and to manage a potential security breach.

We are particularly interested in a type of CPS which function is air space surveillance and coastal water monitoring. The application domain of interest includes Micro Aerial

Vehicles (MAVs) and related technologies such as UAV, Unmanned Underwater Vehicles (UUV), formations of MAVs and collaborating MAVs. We focus on scenarios where an adversary targets the components of the CPS and perpetrates covert cyber-physical attacks [26]. The adversary is to operate a stealthy disruption of services. The purpose is disrupting the navigation data of the MAVs and deceive the defender. The role of the defender is to recognize the activities performed by the adversary, i.e., identify the intentions of the adversary and correct the adversarial actions.

A. Covert Attack and Feedback Truthfulness

A covert attack is an aggression on the state of a CPS where the adversary attempts to be invisible [26]. It is assumed that the adversary knows or can learn the dynamics of the CPS. While the attack is being carried out, the perpetrator compensates the impact of the attack over the system by providing fake information to the system operators (e.g., by concealing the effect of the spoofed inputs). Hence, from the point of view of an observer, responsible for detecting the attack, the execution of the CPS looks normal. Assume the scenario shown in Figure 1. It depicts the disruption of the navigation data of a series of MAVs. The manipulation is conducted by a remote adversary via, e.g., GPS jamming and spoofing attacks [1], [12]. The goal of the adversary is to conduct navigation data modifications (e.g., swapping the x, y coordinates of the navigation traces) and hide the disruption to the defender, with additional cyber-physical covert attack models [26].

Assume now the merge of multiple navigation traces of a single MAV over a given period of time. The covert attack conducted by the adversary conceals the alteration of some of the navigation paths. The defender conducts a learning process to guess the adversarial intentions. The defender also prioritizes assets that can get sacrificed as collateral damages (e.g., to offer some tactical victories to the adversary with the aim of reducing the adversarial power in the long term), e.g., by using some game-theoretic ideas. The process allows the defender to get trusted by the adversary, i.e., to make the adversary confident about the success of some perpetrated actions. Practically speaking, the collateral damages allow the defender to reinforce the defensive learning processes, with the aim of handling and correcting the affected system to the original plans before the execution of the adversarial actions.

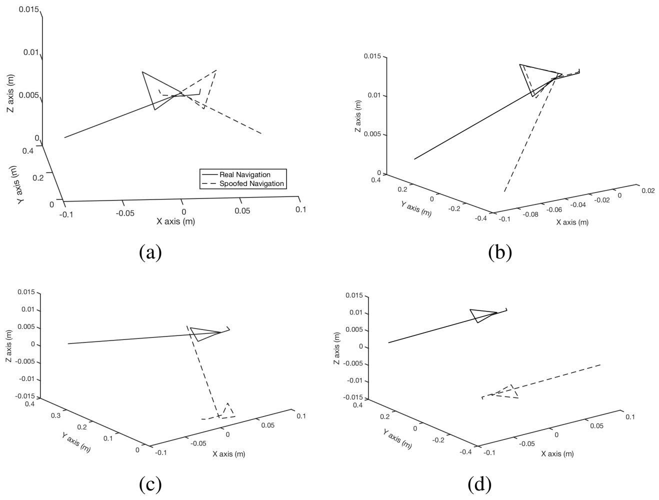

(a)

(b)

(c)

(d)

Fig. 1. MAV navigation data trace disruptions. The adversary perpetrates GPS-like attacks [1], [12], e.g., to swap the navigation coordinates. Solid lines represent genuine navigation data. Dashed lines represent disrupted navigation data. (a) Swapping of the x coordinate. (b) Swapping of the y coordinate. (c) Swapping of the x, z coordinates. (d) Swapping of the x, y, z coordinates.

B. Related Work

Related work include the use of machine learning for cyber-physical protection, management of data quality (w.r.t. feedback-control systems), use of security games and competitive learning. More detailed information follow.

1) CPS Protection Using Machine Learning: The domain of AI, by means of the subfields of search and machine learning, provides a large set of techniques relevant to the resilience of a CPS. Supervised, unsupervised and reinforcement are the three main machine learning paradigms. In supervised machine learning, there are old and new data points. Old data points are labelled, representing classes of data points. Comparing their similarity with old ones, supervised machine leaning assigns labels to new data points. With unsupervised machine learning, the data points are unlabelled. Learning consists of extracting information from data. Data points are grouped together into classes according to similarity. Human experts label the classes.

In contrast, reinforcement learning rewards or penalizes the learner following the validity of inferred classifications. Learning is from the successes and mistakes. Supervised and reinforcement machine learning are used for system identification and model fitting. Different alternative learning methods exist, based on different considerations on the type of model (e.g., rule-based, support-vector machines, deep learning models) and its properties (e.g., explainable models/decisions, efficiency). The perpetration of control-theoretic attacks [26] may require a system identification phase performed by the adversary. Kernels methods [25], a kind of machine learning, can be used for system identification [19], [20].

2) Data Quality and Source Heterogeneity: Learning methods designed to recognize malicious activities within large complex systems may be sensitive to the consequences of merging data from very heterogeneous sources. The learning process can be impacted by poor data quality, affecting the reliability of the decisions that are made. For example, data sources can be unmalicious but of little confidence. This issue is important for the protection of the learning phase. Adequate risk management calls for the implementation of data quality evaluation. Discriminators must be designed to cope with operational data relatively different and considerably more inaccurate than the training data. Taking into account data quality prevents bias and integrity attacks aiming to deceive a discrimination process [18]. Conversely, discrimination algorithms affected by slight modifications of training data may be too sensitive. They can be exploited more easily by adversaries. There are trade offs that can be taken into account by a security game modeling risk management. An interesting avenue is considering the possibility of a classifier not producing an answer when data quality cannot be guaranteed. This behavior corresponds to the rejection principle. A concept relatively well known by the classification community. However, to the best of our knowledge, its application to security needs more thought [6].

3) Security Games and Competitive Learning: In [2], the authors propose a taxonomy to characterize the different types of attacks targeting AI-based security approaches, for example, machine learning classification and discrimination techniques taking security decisions. This taxonomy focuses on the techniques employed to identify malicious artifacts such as, messages, codes, program inputs and outputs. The work suggests to complement machine learning with game-theoretic security approaches, specially those in which the adversary may alter the training data. These observations have been widely repeated and extended in the information security literature [17].

In this context, the principle of a security game amounts to quantifying the adversarial resources required to attack a system and the defensive capabilities of the latter. Concequently, it is possible to determine optimal configurations to manage the risks of attacks as a quest for defender vs. adversary equilibrium. Along these lines, authors in [18] insists on the importance of using learning techniques with access to real time data. Outdated data increases the risk of making incorrect decisions by a defender [13]. Data must be considered obsolete by a learning process at a certain point in time [4]. The use of utility functions can be provided to formulate more realistic games. Such functions do not necessarily need to be under the classical computing realm. Extended machine learning functions relying on quantum techniques are expected [24].

4) The Quantum Advantage: The time complexity of quantum search techniques are data size independent. Along the same line, quantum machine learning, i.e., the use of quantum computing for machine learning, has great potential because the time complexity of classification is independent of the number of data points. Schuld and Killoran investigated the use of kernel methods [25], that can be used for system identification, for quantum machine learning [24], [23]. Encoding of classical data into a quantum state is needed. A similar approach has been proposed by Havlíček et al. [11].

Schuld and Petruccione [24] discussed in details the application of quantum machine learning classical data generation and quantum data processing. A translation procedure is required to map the classical data, i.e., the data points, to quantum data, enabling quantum data processing, i.e., quantum classification. However, there is a cost associated with translating classical data into the quantum form, which is comparable to the cost of classical machine learning classification. This is right now the main barrier. The approach that will result in real gains is quantum data generation and quantum data processing, there will be no need to translate from classical to quantum data. Quantum data generation requires quantum sensing.

III. Faking and Discriminating Navigation Data

Using the Generative Adversarial Network (GAN) framework, we validate that a covert attack can be perpetrated using adversarial learning. A GAN consists of two main entities: a discriminator and a generator [10]. The discriminator is the defender's tool. The generator is the adversary tool. There are genuine (real) data and generated (fake) data. The generator aims at generating data to deceive the discriminator. The discriminator is trained with genuine and generated data. The training process aims to a discriminator able to label genuine or generated data correctly, with high probability of correctness. The adversary wins the game when this probability is at least 50%. To this end, the generator is trained, assuming it can challenge the discriminator with data and access to the verdict. Training is an iterative process. Training iterates until the production of fake data is accepted by the discriminator with high probability.

In a QGAN [5], [15] the data can be quantum. Using a Parrot Mambo MAV, we generate genuine navigation data. The navigation is classical and in continuous domains. Using probability amplitude encoding, the genuine (classical) data is mapped to quantum data and used to train a discriminator, defined as a qubit-quantum circuit. Using a photonic-quantum circuit, we validate that the adversary can learn to generate fake data resembling genuine data, assuming access to nothing else but the verdict of the discriminator.

A. Discriminator Design

We build upon the PennyLane [3] variational classifier [21] and QGAN [22] examples. The elementary circuit design $\mathscr{E}(\omega)$ of Farhi and Neven [8] is used, pictured in Figure 2. Every elementary circuit processes n qubits. In Figure 2, n is three. The circuit formal parameter ω is a n by three matrix of rotation angles. For $i = 0, 1, \ldots, n-1$, the gate $Rot(\omega_{i,0}, \omega_{i,1}, \omega_{i,2})$ applies the x, y and z-axis rotations $\omega_{i,0}$, $\omega_{i,1}$ $\omega_{i,2}$ to qubit $|\psi_i\rangle$. The three rotations can take a qubit from any state to any state. For entanglement purposes, qubit i is connected to qubit $i+1$ modulo n using a CNOT gate.

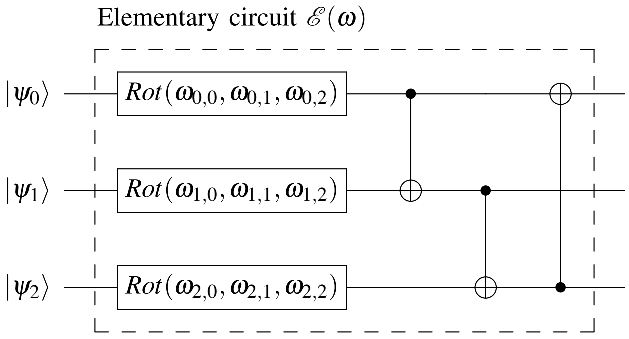
Elementary circuit $\mathscr{E}(\omega)$
$|\psi_0\rangle$
$Rot(\omega_{0,0}, \omega_{0,1}, \omega_{0,2})$
$|\psi_1\rangle$
$Rot(\omega_{1,0}, \omega_{1,1}, \omega_{1,2})$
$|\psi_2\rangle$
$Rot(\omega_{2,0}, \omega_{2,1}, \omega_{2,2})$

Fig. 2. Three-qubit elementary circuit layer.

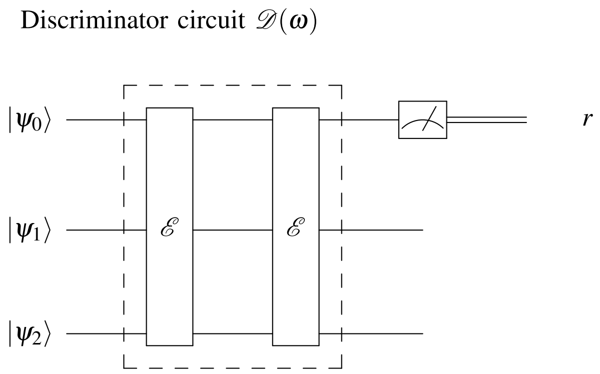

Fig. 3. Discriminator circuit made of two layered elementary circuits.

The discriminator circuit $\mathscr{D}(\omega)$ uses m layers of elementary circuits $\mathscr{E}$. In Figure 3, m is two. Layer 0 accepts the input. Layer i quantum outputs are connected to layer $i+1$ quantum inputs. In this case, the circuit formal parameter ω is a m by n by three matrix of rotation angles. Layer i is actualized with sub-matrix ω_i.

We use probability amplitude encoding, because we can represent in a given number of qubits an exponential number of data points. Probability amplitude encoding requires normalized data. Let $x_0,\ldots,x_{n-1}$ be the data values, their normal form is

$$v_0 = x_0/\mu, \ldots, v_{n-1} = x_{n-1}/\mu$$

where

$$\mu = \sqrt{x_0^2 + \ldots + x_{n-1}^2}.$$

With probability amplitude encoding, up to 2^n single scalar values can be represented in probability amplitudes in the input circuit quantum state. The input quantum state with probability amplitude encoded data has the following format:

$$|\psi\rangle = \sum_{i=0}^{n-1} v_i \, |i\rangle \tag{1}$$

In Figure 3, layer $m-1$ produces the output expectation r on line 0. The output r ranges in the continuous interval +1 down to -1, respectively corresponding to qubits $|0\rangle$ and $|1\rangle$. Intermediate values represent superpositions of qubits $|0\rangle$ and $|1\rangle$. The output is interpreted as follows. When it is $+1$, the data is accepted as true. When it is -1, the data is rejected and considered fake. The output r is converted to a probability value, in the interval $[0,1]$, using the following conversion:

$$p = \frac{r+1}{2}. \tag{2}$$

When genuine data is submitted on the inputs ($|\psi\rangle$) of the discriminator, the value p in Eq. (2) expresses the *probability of real true* p_R. When fake data submitted, the value p corresponds to the *probability of fake true* p_F.

We aim to a discriminator that maximizes the probability p_R of accepting true data while minimizing the probability p_F of accepting fake data. An optimizer finds a rotation angle matrix ω such that the output of the circuit is approaching $+1$, which corresponds to qubit $|0\rangle$. Using a gradient descent technique, the optimizer iterates with genuine data sets and fake data sets. Gradient descent means that the optimizer tries to minimize the cost represented by the difference $p_F - p_R$.

Definition 1 (Discriminator optimization problem): Given the quantum input state ϕ, probability amplitude encoding fake navigation data, and quantum input state ψ, probability amplitude encoding genuine navigation data, training the discriminator $\mathscr{D}(\omega)$ is the optimization problem that consists of finding the matrix ω ($m \times n \times 3$) that gives the smallest difference $p_F - p_R$.

B. Generator Design

The aim of the generator is to produce fake data that is accepted as true by the discriminator, i.e., the probability of fake true p_F is as close to one as possible. When training the generator, it is assumed that the adversary can submit its fake data to the discriminator and access to the verdict. For the generator, we investigated the three following designs: (1) a generator using a MAV model, (2) a qubit-quantum circuit, and (3) a photonic quantum circuit.

1) MAV model design: A detailed model of the MAV is built and evolved. For example, such a model does exist for the MAV we are using for our experiments [16]. The continuous domain navigation data generated by the MAV model is amplitude-encoded and submitted to the discriminator. According to the output of the discriminator, the MAV model is fine tuned until a high probability of fake data acceptance is reached. The challenge with this approach is that the adversary needs a detailed understanding of the dynamics of the MAV. We aim at a method that does require no knowledge on the part of the adversary about the MAV dynamics. In other words, we aim at an automated learning process. Two alternative designs for such a purpose follow.

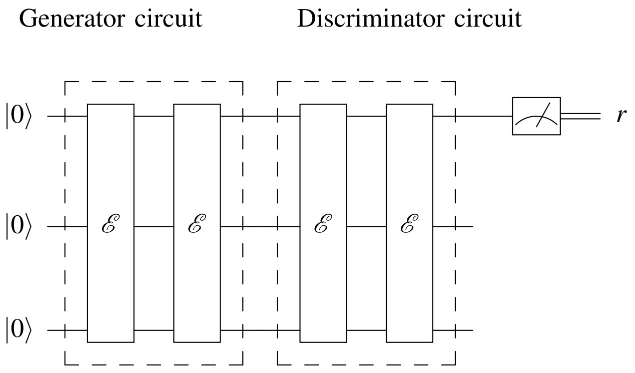

Fig. 4. Generator qubit circuit feeding the discriminator circuit.

2) Qubit-quantum Circuit: The fake data can be generated with a qubit-quantum circuit, with an architecture as the one pictured in Figure 4. The generator circuit is similar to the discriminator circuit, pictured in Figure 3. For the generator circuit, the inputs are all at $|0\rangle$. The optimization is done on the rotation angles, using the verdict of the discriminator r. The learning process is automatic. The generator outputs the navigation data with entropy. The outputs of the generator are directly connected to the inputs of the discriminator. The generated navigation data is encoded in the probability amplitudes of the quantum state produced by the generator. Although it works, the navigation data must be transformed to a quantum format. Hence, the data is unusable for practically perpetrating an attack. Indeed, qubit-circuit outputs, obtained through measurements, collapse to zeros and ones. To be usable in an attack scenario, the data needs to be transformed from classical continuous domains. An alternative design for such a purpose follows next.

3) Photonic quantum circuit: The generator combines photonic quantum computing [14] and qubit-quantum computing. Photonic devices are trained to generate photon numbers corresponding to navigation data accepted by the discriminator. A photonic quantum circuit is shown in Figure 5. It has a single line, called a *qumode*. The input of the circuit, $|0\rangle$, is the zero energy level. There are two Gaussian devices. There is a displacement gate D, with parameter α, and a rotation gate R, with parameter ϕ. They change the circuit energy level and expected numbers of output photons. The measurement gate determines the average number of photons at the output of the circuit.

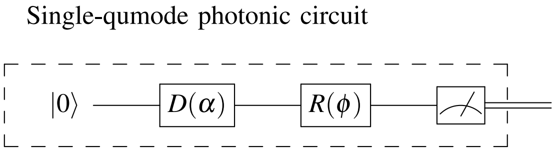

Fig. 5. Photonic-cricuit model.

We use photonic devices to generate fake navigation data. The output is amplitude-encoded and submitted to the discriminator. The photonic-quantum circuit is optimized on the parameters α and ϕ such that the probability of acceptance of the fake data by the discriminator is high.

The architecture pictured in Figure 6 shows a generator feeding a discriminator circuit through a probability amplitude encoder $\mathscr{A}$, including normalization. The MAV navigation data set is amplitude encoded according to Eq. (1). Since n qubits can amplitude-encode 2^n datum, a n-qubit discriminator is fed by a generator with 2^n qumodes. In Figure 6, n is two.

The generator is initialized with arbitrary displacements and rotation angles (α and ϕ). A gradient descent optimizer is used to minimize the cost represented by the term $-p_F$. The outcome of the optimization of the generator is two column vectors of displacements and rotation angles, 2^n rows each, actualizing the generator circuit such that the probability that fake data is recognized as true is high.

Definition 2 (Generator optimization problem): Given the quantum input state ψ, probability amplitude encoding fake navigation data, the discriminator $\mathscr{D}(\omega)$, actualized with rotation angle matrix ω, training the generator $\mathscr{G}(\alpha,\phi)$ is the optimization problem that consists of finding the column vectors of the rotation angles α and ϕ (2^n rows each) that gives the smallest difference $-p_F$.

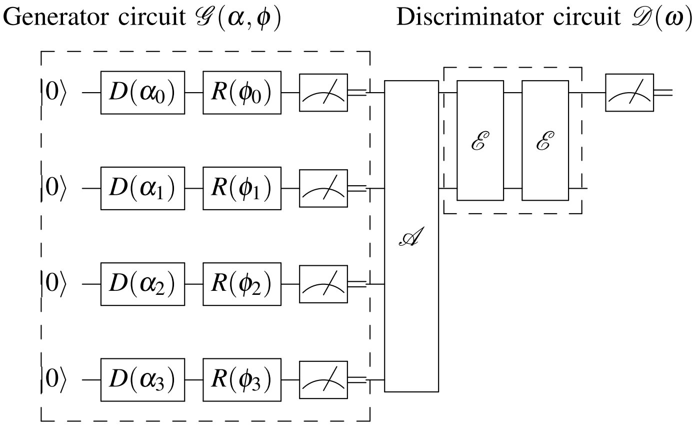

Fig. 6. Generator qubit circuit feeding the discriminator circuit (n is two) .

The learning process is automatic. The output of the photonic quantum circuit is classical and in the continuous domain. It is directly usable by the adversary to generate fake navigation data during a covert attack. The circuit complexity is although in $\mathscr{O}(2^n)$.

IV. Performance

The performance of the photonic-circuit design described in Section III has been validated through simulation on a classical computing platform. Simulations were conducted using an Intel Xeon 32-core 2.70 GHz server, with 256 GB of memory. We generated genuine navigation data for a Parrot Mambo

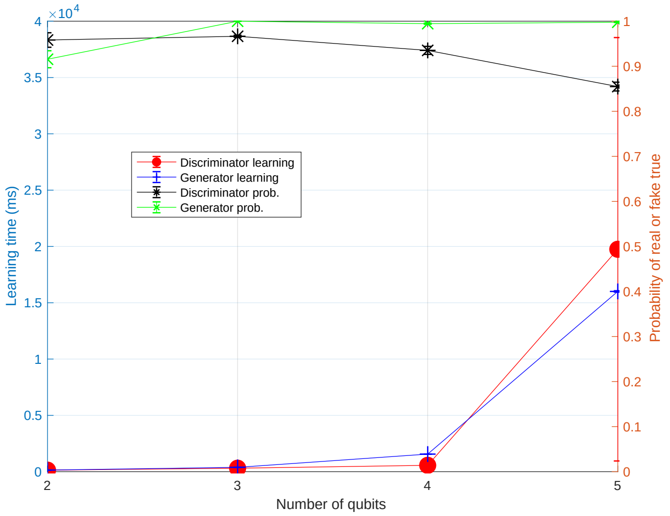

Fig. 7. Learning time (ms) versus the number of qubits for the discriminator and generator.

MAV. In the scenario, the MAV takes off one meter. Does two circles on the horizontal plane, then lands. The navigation data consists of x, y and z velocity triples. The whole scenario generates less than 64 real number values.

Figure 7 plots the discriminator and generator learning time (ms) versus the number of qubits available. The x axis represents the number of qubits. The left y axis refers to the learning time (ms). The right y axis shows the corresponding probability of real true, for the discriminator, and probability of fake true, for the generator. Hundred optimization iterations were done for each case. Negligible error margins are included, but not visible since they are very tiny. The discriminator is trained with six different genuine navigation data sets. A navigation data set is picked at random at every optimization iteration. The discriminator optimization time grows exponentially. Due to the exponential complexity of the generator circuit (in $\mathcal{O}(2^n)$), the optimization time also grows exponentially. On our simulation platform, it becomes unpractical from six qubits. The learning time becomes in the order of days. Amplitude encoding has also $\mathcal{O}(2^n)$ time complexity, but it is only executed once at the start of the optimization process.

V. Conclusion

We have investigated the use of QGAN designs to generate fake MAV navigation data. We assume the same approach to discriminate between genuine and fake MAV navigation data. The goal pursued by the adversary is to generate fake data that is accepted as true by a trained discriminator. On the other hand, the discriminator must accept with high probabilities true navigation data and reject fake one. The elaborated quantum circuits have been evaluated running on a a classical computing platform. As demonstrated in Figure 7, the exponentially growing time complexity in the number of qubits is a barrier to scalability. We identified hurdles that must be overcome by the upcoming evolution of quantum machine learning. The main hurdle for the adversary is the generation of navigation data in classical continuous domains, i.e., real numbers, and the cost of the transformation into the quantum format at every optimization iteration. Further research is needed to improve and find alternatives to the design depicted in Figure 6. The source code for the examples presented in this paper is available at https://github.com/jgalfaro/mirrored-QGANMAV.

Acknowledgments — Work partially supported by the Natural Sciences and Engineering Research Council of Canada and the European Commission under grant agreement 830892 (H2020 SPARTA project).

References

[1] Michel Barbeau, Joaquin Garcia-Alfaro, and Evangelos Kranakis. Geocaching-inspired resilient path planning for drone swarms. In *IEEE MiSARN 2019, co-located with IEEE INFOCOM 2019 – IEEE Conference on Computer Communications*, France, 2019.

[2] Marco Barreno, Blaine Nelson, Anthony D Joseph, and J Doug Tygar. The security of machine learning. *Machine Learning*, 81(2):121–148, 2010.

[3] Ville Bergholm, Josh Izaac, Maria Schuld, Christian Gogolin, Carsten Blank, Keri McKiernan, and Nathan Killoran. Pennylane: Automatic differentiation of hybrid quantum-classical computations, 2018.

[4] Samuel Rota Bulò, Battista Biggio, Ignazio Pillai, Marcello Pelillo, and Fabio Roli. Randomized prediction games for adversarial machine learning. *IEEE transactions on neural networks and learning systems*, 28(11):2466–2478, 2016.

[5] Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *Phys. Rev. A*, 98:012324, Jul 2018.

[6] Claudio De Stefano, Carlo Sansone, and Mario Vento. To reject or not to reject: that is the question-an answer in case of neural classifiers. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 30(1):84–94, 2000.

[7] Nicolas Falliere, Liam Murchu, and Eric Chien. *W32.Stuxnet dossier*, symantec security response. http://j.mp/2jaM6uM, 2011.

[8] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors, 2018.

[9] Alex Gatopoulos. Houthi drone attacks in saudi *show new level of sophistication*. http://j.mp/2LMqR3H, May 2019.

[10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.

[11] Vojtěch Havlíček, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.

[12] Ahmad Y Javaid, Farha Jahan, and Weiqing Sun. Analysis of global positioning system-based attacks and a novel global positioning system spoofing detection/mitigation algorithm for unmanned aerial vehicle simulation. *Simulation*, 93(5):427–441, 2017.

[13] Murat Kantarcıoǧlu, Bowei Xi, and Chris Clifton. Classifier evaluation and attribute selection against active adversaries. *Data Mining and Knowledge Discovery*, 22(1-2):291–335, 2011.

[14] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry Fields: A Software Platform for Photonic Quantum Computing. *Quantum*, 3:129, March 2019.

[15] Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Phys. Rev. Lett.*, 121:040502, Jul 2018.

[16] MathWorks. Quadcopter Project. https://www.mathworks.com/help/aeroblks/quadcopter-project.html. Accessed: 2019-06-20.

[17] Nicolas Papernot, Patrick McDaniel, Arunesh Sinha, and Michael Wellman. Towards the science of security and privacy in machine learning. *arXiv preprint arXiv:1611.03814*, 2016.

[18] Nicolas Papernot, Patrick McDaniel, Xi Wu, Somesh Jha, and Ananthram Swami. Distillation as a defense to adversarial perturbations against deep neural networks. In *2016 IEEE Symposium on Security and Privacy (SP)*, pages 582–597. IEEE, 2016.

[19] Gianluigi Pillonetto, Francesco Dinuzzo, Tianshi Chen, Giuseppe De Nicolao, and Lennart Ljung. Kernel methods in system identification, machine learning and function estimation: A survey. *Automatica*, 50(3):657–682, 2014.

[20] Gianluigi Pillonetto and Giuseppe De Nicolao. A new kernel-based approach for linear system identification. *Automatica*, 46(1):81–93, 2010.

[21] Maria Schuld. Example Q3 - Variational classifier. https://github.com/XanaduAI/pennylane/blob/master/examples/Q3_variational-classifier.ipynb. Accessed: 2019-06-11.

[22] Maria Schuld. Example Q4 - Quantum Generative Adversarial Network. https://github.com/XanaduAI/pennylane/blob/master/examples/Q4_quantum-GAN.ipynb. Accessed: 2019-06-11.

[23] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Phys. Rev. Lett.*, 122:040504, Feb 2019.

[24] Maria Schuld and Francesco Petruccione. *Supervised Learning with Quantum Computers*. Quantum science and technology. Springer, 2018.

[25] John Shawe-Taylor and Nello Cristianini. *Kernel Methods for Pattern Analysis*. Kernel Methods for Pattern Analysis. Cambridge University Press, 2004.

[26] Roy Smith. Covert Misappropriation of Networked Control Systems: Presenting a Feedback Structure. *IEEE Control Systems*, 35(1):82–92, Feb 2015.