

Schrödinger’s Camera: First Steps Towards a Quantum-Based Privacy Preserving Camera

Hannah Kirkland
University of Florida
hkirkland@ufl.edu

Sanjeev J. Koppal
University of Florida
sjkoppal@ece.ufl.edu

Abstract

Privacy-preserving vision must overcome the dual challenge of utility and privacy. Too much anonymity renders the images useless, but too little privacy does not protect sensitive data. We propose a novel design for privacy preservation, where the imagery is stored in quantum states. In the future, this will be enabled by quantum imaging cameras, and, currently, storing very low resolution imagery in quantum states is possible. Quantum state imagery has the advantage of being both private and non-private till the point of measurement. This occurs even when images are manipulated, since every quantum action is fully reversible. We propose a control algorithm, based on double deep Q-learning, to learn how to anonymize the image before measurement. After learning, the RL weights are fixed, and new attack neural networks are trained from scratch to break the system’s privacy. Although all our results are in simulation, we demonstrate, with these first steps, that it is possible to control both privacy and utility in a quantum-based manner.

1. Quantum-Based Privacy for Vision

While machine learning algorithms extract useful knowledge from massive amounts of data, there exist concerns that such techniques can infer sensitive, private information without permission. One of the goals of privacy-preserving computer vision is to allow the machine learning algorithm to access visual information useful for a desired task, but not access additional sensitive information [15].

A major challenge in privacy-preserving computer vision is the privacy-utility trade-off. Blank images are private but regress to data priors for any vision task. Clear images, while useful, do not provide an acceptable level of trust and privacy.

In this paper, we propose a quantum framework where the data, in a quantum sense, is both private and non-private until the point of measurement [23]. Our design is a hy-

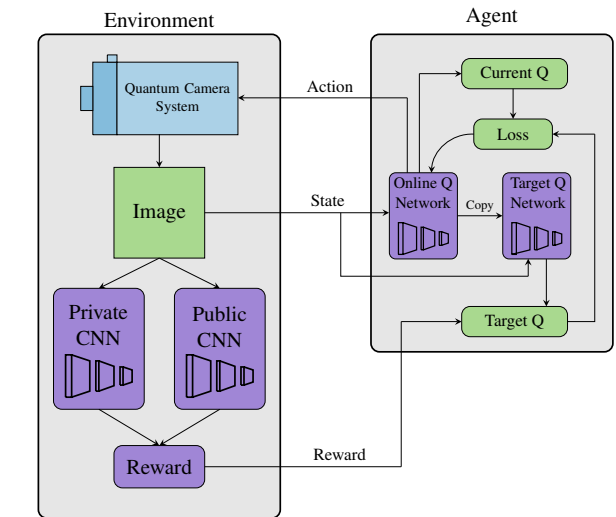


Figure 1. Workflow of our system showing how the quantum camera environment interacts with the DDQN agent. The image data is stored in quantum states, and the agent manipulates these with a set of actions — quantum circuits that are part of our contribution. Adversarial CNNs create the reward and, during testing, new CNNs are trained from scratch against fixed DDQN weights.

brid quantum-silicon system, where conventional machine learning sits in silicon, but the data is stored in quantum states. The algorithm learns what actions to take within the quantum computer *before* measurement, such that the measured image has the desired privacy-utility characteristics.

Our design is unique, and takes advantage of well-known quantum properties that already impact privacy, security and cryptography. One such property is the no-cloning theorem. It states that arbitrary quantum states cannot be perfectly cloned or copied. This is useful for security since it implies an eavesdropper can always be detected [24]. In our design, keeping the data in a quantum state provides this level of trust.

Another property allows lower storage footprints due to quantum computing’s unique property of entangle-

ment. With the Flexible Representation of Quantum Images (FRQI), an M by L image can be stored in $n = \log_2(ML)$ qubits, as opposed to 2^n bits on a traditional computer. This means that while the number of traditional bits required to store an image grows exponentially, the number of qubits needed simply grows linearly [39].

1.1. Contributions

In this paper, we describe our privacy-preserving design with feasibility experiments. Our work presages the impact of quantum technology on imaging and all of our results are in simulation, due to the limits of technology today.

Our key idea is to manipulate imagery in quantum states, measuring them as conventional images only when sensitive, private features are protected. However, there is a challenge to manipulating data stored in quantum states. Even a small number of quantum circuits create an exponentially large space of "actions", i.e. image manipulation steps. Additionally, the action space increases with image size.

We cast the problem of taking the right privacy-preserving actions inside of a quantum framework as a reinforcement learning (RL) task. In this sense, image manipulation is a planning problem, where the goal is a final measured image whose privacy is protected, but many failure states exist along the way. Image measurement at any of these failure states would result in breaching privacy concerns or destroying the public utility.

Given image data stored in a quantum state, we use an RL-based agent to control the quantum privacy preserving camera. The agent's actions provide instructions to the quantum computer, indicating what pixels to redact. The image is then measured and sent to two competing CNNs. The public CNN classifies the image according to the desired non-sensitive classes. The private CNN attempts to classify the image according to the sensitive classes. The reward is generated by rewarding correct public classifications and penalizing correct private classifications, and the agent updates the mask accordingly.

Our contributions are:

1. A privacy-preserving design that uses quantum image processing combined with RL-based learning, where the agent's actions are dynamic in training and testing.
2. Selection of quantum circuits that would be useful in this privacy preserving scenario.
3. Experiments demonstrating that the RL machine learns useful patterns in the large action space, and demonstrating the actual privacy-utility tradeoff by finetuning the CNN classifiers on image data generated by fixed RL policies.

2. Related Work

Privacy Preservation When data elements are defined precisely, then formal guarantees for privacy exist [10].

For imaging, however, such approaches have only found use in image databases, e.g. K-anonymity [29]. Privacy-preserving vision has therefore relied on domain knowledge and features such as pixelation, blurring, face/object replacement, etc.—to degrade sensitive information [1, 3, 7, 9, 12, 40]. However, we consider the case where adversaries can retrain after our privacy-preserving effect. This makes privacy preservation a harder problem, since deep networks can recover identity from severely degraded data (e.g., [32]). These challenges have opened up research questions at the intersection of inference and privacy [13], including methods to avoid capturing sensitive data in the first place [30]. GANs have been utilized to address these challenges. These include using GANs to remove sensitive details from autonomous vehicle data [36]. The GAN approach can have stability issues, however, which the encoder approach helps to alleviate [27]. Another approach beyond GANs is to adaptively sample imagery for robotics. Privacy is maintained by using FPGA or DSP chips to do resolution scaling in hardware, meaning the sensitive data is never stored at a higher resolution. This method is effective but can be difficult to achieve in real-time [18]. In our work, we strike out in a complimentary direction of using RL instead of a GAN encoder. While this solves the stability issue, Double Deep Q-learning convergence becomes the deciding factor for performance.

Learning Inside the Camera Many approaches learn optics and hardware designs to improve post-capture computation accuracy. One example of this strategy can be seen where FGPA and DSP chips are used to scale the resolution of sections of the image pre-capture to achieve privacy [18]. Another example uses programmable physical mask to block out light before it reaches the image sensor [28]. Other optical examples use a series of lenses and masks to make an optical convolution layer [5] or natural chromatic aberrations to improve depth perception [6]. In this paper, the images are in quantum states, and the learned image manipulation happens before measurement.

Quantum Imaging Currently, noise limits the length of quantum circuits and consequently the size of images. Right now a 2×2 pixel image is the largest FRQI image that can be reliably encoded and measured [11]. While research is still being done on how to overcome noise in quantum computers, quantum sensors have developed rapidly and shown promising results. For example, sensors have been developed using quantum dot technology to create extremely sensitive covid-19 tests and image sensors [21, 25]. The quantum dot image sensors show particular promise in the NIR and SWIR wavelengths [22, 25, 26]. Another strategy for realizing quantum cameras involves using photon counting [35]. Smart quantum statistical imaging, using machine learning and photon resolving detection, has been used for superresolving imaging [2].

Quantum Computing and Security Quantum computing allows unique security strengths. For example, qubit transfer between two parties cannot be unknowingly intercepted if the parties use a protocol to detect interference, such as quantum key distribution (QKD). Some security challenges quantum computing does face are methods of imperfect cloning and exploitation of hardware imperfections [34]. The advantages of quantum computers over classical computers for security outweigh the hazards, and we exploit these in our design.

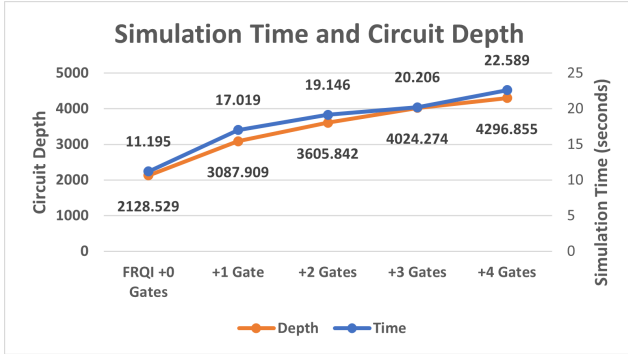


Figure 2. Graph depicting the average simulation time and quantum circuit depth for 16x16 FRQI images. The number of gates on the x-axis refers to the number of additional privacy preserving quantum gates applied after the initial FRQI encoding.

Quantum Image Processing With the development of quantum computing algorithms has come the development of quantum image processing (QImP). In 2003, researchers proposed a method of encoding images in quantum states using a method called Qubit Lattice [33, 37]. As QImP progressed, other methods were proposed with different pros and cons such as Real Ket [19], NEQR [41], MCQI [38], and FRQI [11, 20]. For this project we chose to focus on FRQI due to its low storage footprint and existing image manipulation algorithms. One advantage of QImP over traditional image processing is that certain image manipulations have exponential speedups with quantum computing such as the Fourier and Hadamard transforms [4, 39]. For example, a quantum edge detection algorithm was developed to run with $O(1)$ processing time rather than the $O(2^n)$ processing time required by classical algorithms [4, 39]. Additionally, security methods such as watermarking are already in existence [37].

One limitation of FRQI is that it stochastically encodes images into quantum states. The accuracy of the pixel grey scale value is directly related to how many times the circuit is run (often referred to as the number of “shots”) [11, 20]. This phenomenon is described by the no-teleportation theorem, which states that it is impossible to exactly reconstruct a quantum state with classical bits. In this paper, we convert this “bug” into a “feature” and use such quantum circuits to

our advantage for privacy applications.

Flexible Representation of Quantum Images The Flexible Representation of Quantum Images (FRQI) is a method of representing images as a quantum state. For a $2^n \times 2^n$ pixel image, $2n + 1$ qubits are needed to encode the image in a quantum state using FRQI. The location of each pixel is stored in the $2n$ qubits while the color information is stored in a single qubit as a scaled value between 0 and $\frac{\pi}{2}$. Using the computational basis $|1\rangle$ and $|0\rangle$, the FRQI image is encoded in a quantum state by:

$$|I(\theta)\rangle = \frac{1}{2^n} \sum_{i=0}^{2^{2n}-1} (\cos \theta_i |0\rangle + \sin \theta_i |1\rangle) \otimes |i\rangle,$$

$$\theta_i \in [0, \frac{\pi}{2}], i = 0, 1, \dots, 2^{2n} - 1$$

Where θ_i is the color information for each pixel location i , and $n = \log_2(\text{image length})$ [11, 20].




Figure 3. 2x2 FRQI image, showing the $|i\rangle$ location information and θ_i color information

3. Experimental Challenges

We faced challenges due to the computational complexity of quantum simulations and circuit size constraint on current NISQ (noisy intermediate-scale quantum era) quantum computers. A 2x2 pixel image is the largest FRQI image that can be reliably recovered from current quantum computers due to noise that arises with longer quantum circuit depth [11]. Using Qiskit’s AerSimulator, the largest reliably simulated FRQI image is 16x16 pixels [11].

Additionally, simulation latency to prepare and run the quantum circuits is a challenge. Fig. 2 shows that the average time to simulate a 16x16 FRQI image with a NVIDIA GeForce GTX 1080 Ti is 11 seconds. The simulation time increases with quantum circuit depth, so that four additional gates after the initial FRQI encoding effectively doubles the simulation time per image. Working within these constraints, rescaled 16x16 EMNIST images were used as the data set with the public utility classes being whether an image represents a ‘letter’ or a ‘number’ and the private, sensitive classes being the exact letter or number (the

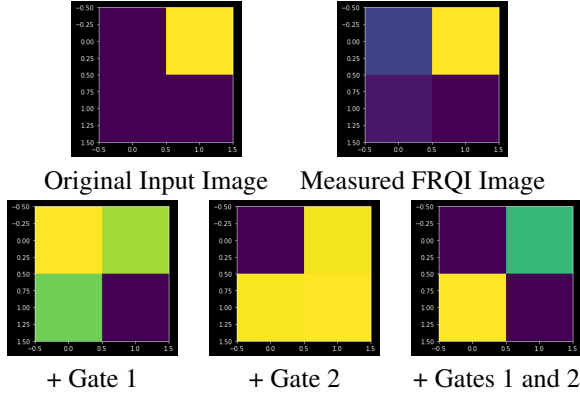


Figure 4. **Real** low-resolution FRQI quantum images run on ‘ibmq_manila’, an IBM Quantum Falcon processor [16]. The first image is the classical input image. The second is the measured FRQI image with no additional gates applied. Gate 1 is the measured FRQI image with an additional gate added (Controlled-RX on p_0, c). Gate 2 is the measured FRQI image with a different additional gate added (Controlled-RX on p_1, c). Gates 1 and 2 is the measured FRQI image with both gates 1 and 2 sequentially applied to the original FRQI image.

uppercase letters and numbers from the balanced EMNIST dataset [8]).

Experimental Proof of Concept Our proposed solution involves the combination of quantum gates applied to an FRQI image to achieve privacy preservation. While 16x16 images are used in our simulation results, providing more options for potential privacy preservation through quantum gates, the same concept can still be seen in 2x2 images. In Fig. 4, a 2x2 FRQI image is run on ‘ibmq_manila,’ which is a 5 qubit IBM Quantum Falcon processor. The input image is compared to the result with no additional gates after FRQI encoding, one controlled-RX gate (applied to either positional qubit), or two controlled-RX gates (applied to both positional qubits). This shows the different outputs that can be achieved with one type of quantum gate, and how they interact together to produce different results in combination.

4. Quantum Privacy Preserving System

In Fig. 1, we show our work-flow, where data from the scene is captured and stored in quantum states. This represents the environment in our reinforcement learning setup, whereas the agent attempts to manipulate the quantum states by controlling internal quantum circuits. We discuss the environment and the agent in detail below.

4.1. Quantum States as an Environment

Privacy through Quantum Gates For this project, we explored different methods of redacting or destroying FRQI image information. The privacy for these methods is found after measurement, as quantum computations are inherently

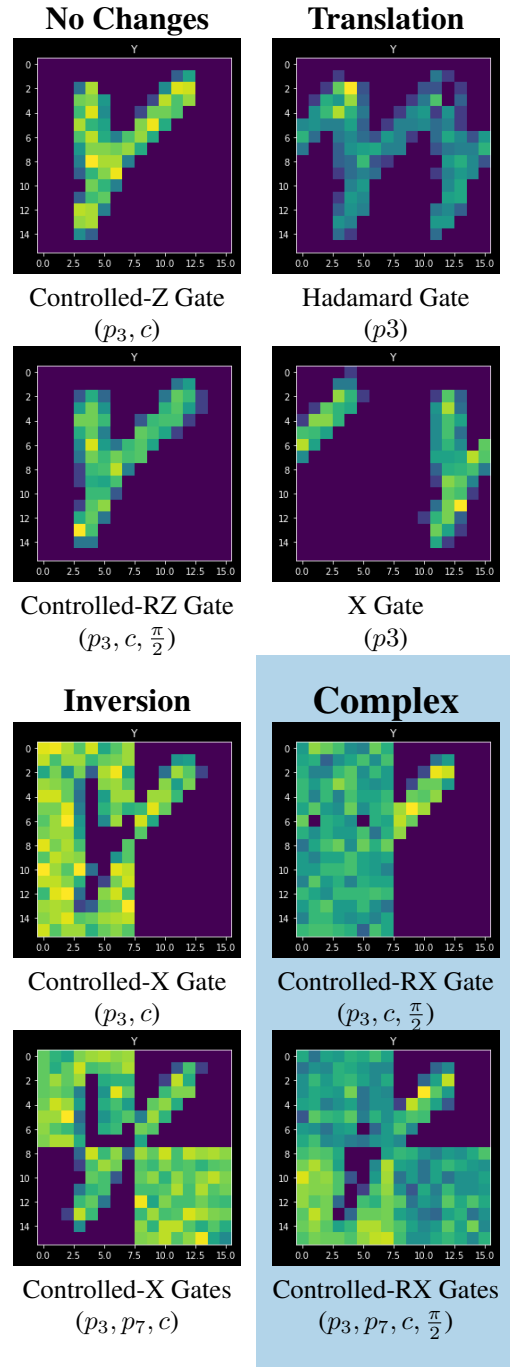


Figure 5. Many quantum gates, such as Controlled-Z gate, do not affect images. Others, such as the Hadamard Gate and the controlled-X gate demonstrate simple transformations. We have selected Controlled-RX gates that can provide sophisticated image manipulation.

reversible. Therefore, we treat the quantum circuit like a black box, using gates and controls to manipulate the unknown image before measuring the result.

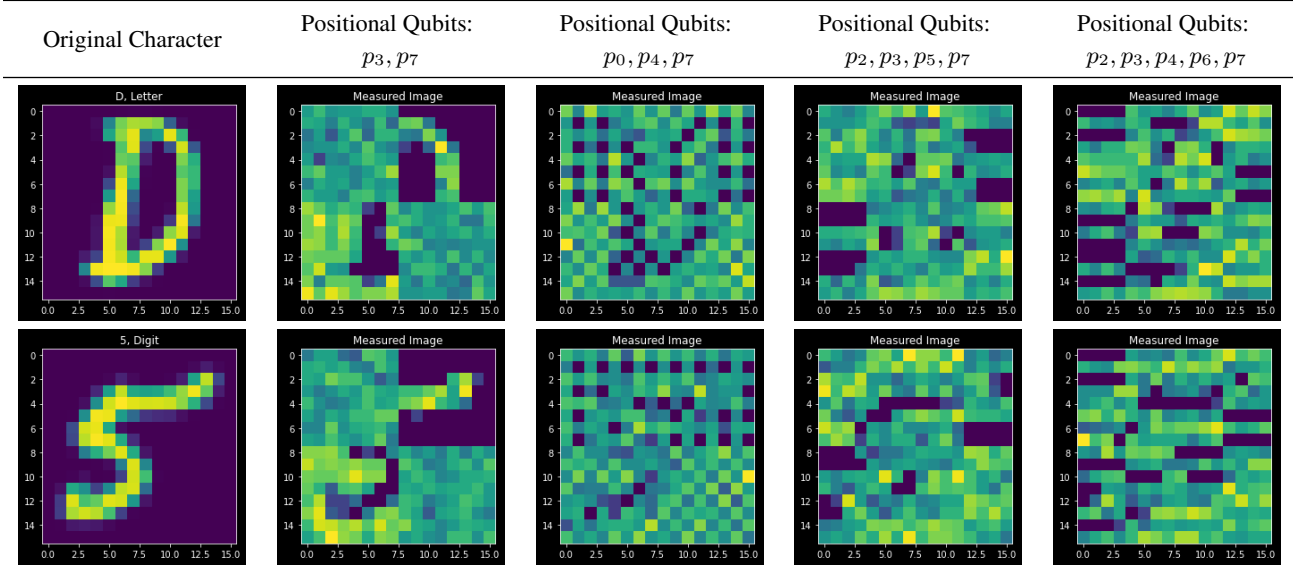
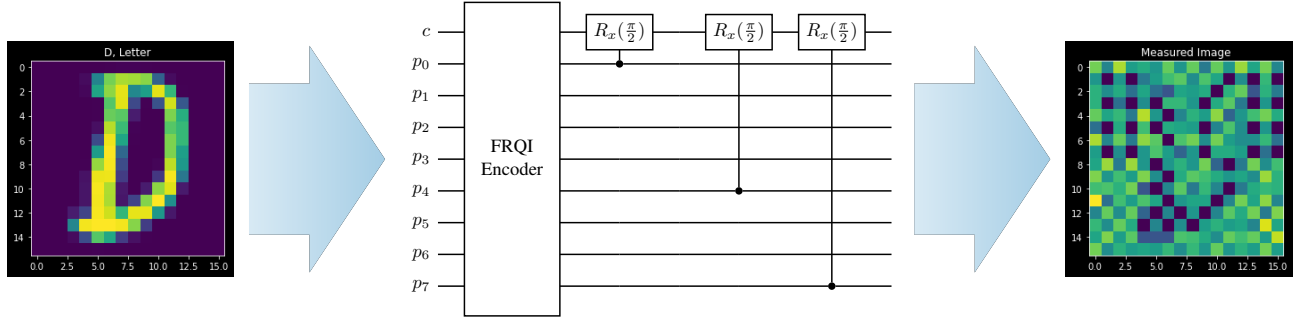


Figure 6. We built a FRQI encoder out of the quantum gates that we selected for privacy-preservation. The 16×16 images are stored in 9 qubits. This corresponds to the efficiency of storing a $2^4 \times 2^4$ pixel image in $2 * 4 + 1$ qubits. Manipulating images in lower dimensional qubits is non intuitive and, further, the size of the action space is 24,157. The action space produces interesting manipulations are shown in the figure for a small sampling of actions (four sets of actions for the letter D and number 5). We learn the best actions for a particular privacy preserving task.

We have explored the space of quantum circuits and selected those that are useful for privacy preservation. A first approach for destroying FRQI image information is by controlling the number of shots, analogous to photon noise in image capture. This comes at the advantage of making the overall computation faster (as the circuit is run fewer times to obtain a noisy image) but at the disadvantage of not having precise control.

The second method is by adding gates to the FRQI image to redact parts of the initial image. Some manipulations can be done with one gate, redacting chunks of the image. Similar methods have been used to translate images or do color manipulations [17, 37].

The gates perform color transformations on select patterns of the image. The patterns are selected by which positional qubits are used as the controls for a controlled X rotation. The rotation angle is selected as $\theta = \frac{\pi}{2}$ to obscure

image information in the selected areas.

Selection of these gates was not trivial. Fig. 5 shows selected examples of manipulations that can be achieved with minimal gates. Only controlled quantum rotations were found to be the most promising for our use case since information is effectively redacted rather than simply translated or inverted. Due to the outputs being similar between controlled-RX and controlled-RY gates, controlled-RX gates were selected to be used in our action set.

Of course, one could follow the approach of many QImp techniques and redact single pixels using the FRQI equation. This comes at the cost of being exponentially more computationally expensive compared to the former two methods, but it allows more granular control. Adding noise can also be achieved this way by randomly picking a number of pixels to redact.

There is an exponentially large number of potentially pri-

vacancy preserving actions that can be taken on an FRQI image. We limited the action space by allowing up to 4 of a total of 28 actions to be selected for a single image. These actions consist of single controlled-RX gates and levels of noise. This limited space already consists of 24,157 unique actions that can be taken on an FRQI agent. If image size, allowed action type, or number of gates is increased this action space size would increase further. Our action space is therefore designed to allow more varied and rich image outputs while limiting the overall size of the action space. Fig. 6 shows a few examples of the possible outputs that can be achieved by combining these gates.

5. Agent Learning

Reinforcement learning was selected in order to fully explore and develop a good policy for selecting privacy preserving quantum actions. A decision on what action to take needs to be made for each image that comes through the simulated quantum camera.

A reward can be constructed to promote actions resulting in public utility and to punish actions resulting in access to sensitive information. The agent can explore the space and fully take advantage of the actions available to maximize the reward and achieve the desired privacy-utility ratio.

Reinforcement Learning with DDQN Double Deep Q Learning (DDQN) is a type of reinforcement Q-learning that results in less overestimation and better performance than the standard Double and Deep Q-Learning (DQN) algorithms [14]. Reinforcement learning’s goal is to learn good policies so an agent can successfully navigate in its environment. It is popular for video games and robotics, where the agent is the game character or robot. The environment in this case is the “world” the agent navigates, and the action space is composed of the things the agent can do like move and jump. For imaging tasks, the environment can be thought of as the imaging system and picture target while the action space can be thought of as the controls the agent can access to manipulate the camera.

To use DDQN in a quantum privacy preserving camera setup, the environment is created to include the “quantum camera” (in our case, represented by an FRQI encoder), privacy gate controls, and CNNs responsible for grading the measured image. The agent includes the target and on-line networks. The action space covers all combinations of the quantum privacy gates allowed in a single step. With this infrastructure, the DDQN model can learn a policy that maximizes the public utility task (in our case, the accuracy of the public classification task) while preserving privacy by minimizing the private task (the accuracy of the private classification task).

We demonstrate that DDQN can learn policies for a quantum camera framework by showing model metrics. Fig. 7 shows the target and online Q values converge and

loss decrease, indicating a policy is being learned. The stability of the reward indicates that the model is not overoptimistic.

6. Experimental Design

When designing our privacy preserving quantum framework, we assume temporal consistency. This assumption means thinking of the images like frames in a video, where the letters and numbers contain similar features across time steps. An example of this could be the camera scanning documents containing text. The agent sequentially chooses an action to apply to each character in the document, aiming to preserve the public utility features while redacting the sensitive private features.

Training Methodology The DDQN architecture was used to realize a privacy preserving quantum camera. The agent produces an action that corresponds to which quantum gates are applied to an unknown FRQI image. Once the action gates are applied, the FRQI image is measured and decoded. The resulting damaged image is sent to two CNNs trained on augmented EMNIST data. The data augmentation included downscaling the images for speed and training on common image results from the action gates. One CNN is for identifying the public classes of “letter” and “number,” and one CNN is for identifying the private classes corresponding to each specific letter and number. The reward is determined by CNN accuracy or episode length, depending on the policy. The last two layers of each CNN and the reward are used by the DDQN networks to determine the current and target Q values, as well as the next action to be taken.


Testing Methodology After training a policy, the DDQN agent networks were frozen to assess their performance. This was done by running a test set through the model and saving the generated output image along with the ground truth label at each step. Once the images were generated, the quality of the policy was assessed using the privacy and public utility CNNs to obtain overall testing accuracies. These test accuracies were recorded, then the private and public utility CNNs were finetuned on the generated data, retested, and recorded.

This testing methodology differs slightly from testing methods presented in the original DDQN paper, which relies mainly on reward scores and comparing the learned policy to an empirical best policy [31]. For this use case, we must show that the learned policy is empirically better than chance for it to be considered good. Since we construct and test several different reward functions (as opposed to using a predefined score like in a video game), we must also show that a higher reward corresponds with a better policy outcome. We chose the private and public utility CNN accuracies as a clear, empirical way of comparing the quality of a learned policies and the outcomes of different reward


Q Online and Q Target Values

Mean Loss


Smoothed Reward




(a) Public-Based Reward Policy




(b)




(c)




(a) Public-Private Reward Policy




(b)




(c)




(a) Length-Based Reward Policy




(b)




(c)



(a) Accuracy-Based Reward Policy



(b)



(c)

Figure 7. Here we show the DDQN training graphs for four rewards that we investigate, with one in each row. For each experiment, the convergence of the Q values and the stability of the loss show the machine is learning. The results of these are shown in Table 8.

functions.

6.1. Results

Different RL parameters and rewards were experimented with. Four models were chosen to be tested and the classification CNNs were subsequently finetuned on the image data produced by the tests. The competing classifier CNN models were run on NVIDIA GeForce GTX 1080 Ti's while the DDQN models were run on NVIDIA TITAN RTX's or

NVIDIA TITAN Xp's. The results of these tests can be found in Fig. 8.

The Public-Based Reward Policy gives a set reward whenever the CNN can accurately predict the public utility label. This model achieved the highest performance, as the desired public utility accuracy was above chance while the private task accuracy was below chance. It was most similar to the Gaussian Noise Augmented Baseline for the public utility task, but achieved significantly better perfor-

Privacy and Utility Preservation Performance						
Pure Chance Baseline	50.0%		2.8%			
Data	Public CNN Accuracy		Private CNN Accuracy			
Quantum Augmentation Baseline	88.5%		85.0%			
	Testing	Finetuned	Testing	Finetuned	Training Steps	Training Hours
Quantum Random Action Baseline	49.6%	51.9%	2.7%	2.1%	–	–
Gaussian Blur Augmented EMNIST Baseline	54.1%	83.9%	6.1%	80.7%	–	–
Gaussian Noise Augmented EMNIST Baseline	51.5%	56.5%	3.3%	11.4%	–	–
Public-Based Reward Policy	50.7%	56.5%	4.6%	2.5%	90,397	605
Length-Based Reward Policy	63.9%	51.5%	2.4%	4.3%	28,519	190
Accuracy-Based Reward Policy	64.4%	48.7%	4.2%	2.8%	52,829	348

Figure 8. This table shows the CNN test and finetuned model accuracies for different output image sets. The Pure Chance Baseline shows the odds of pure chance (with the public task having 2 classes and the private task having 36). The Quantum Augmentation Baseline are the CNN models used for finetuning, and shows the test accuracies of the models trained on EMNIST data run through random quantum privacy gates. The Quantum Random Action Baseline was generated by randomly selecting actions from the quantum action space and simulating the output. The Gaussian Blur Test Baseline consists of a dataset with a 4x4 kernel Gaussian blur applied to each image. The Gaussian Noise Test Baseline consists of a dataset with normalized Gaussian noise added to each image. The rows following the baselines represent tests run on the image outputs of selected RL policies. The best performing model according to this metric, the Public-Based Reward Policy, is bolded for clarity.

mance compared to the baseline for the private task. Additionally, the Gaussian Noise Augmented Baseline was able to be improved with finetuning while the Public-Based Reward Policy was not improved with finetuning.

The Public-Private Reward Policy add points to the reward when the CNNs can accurately predict the public utility label, but also subtracts points from the reward when the CNNs can accurately predict the private label. The Length-Based Reward Policy awards points based on how long the episode has been running above an accuracy threshold, thereby promoting longer episode times. The Accuracy-Based Reward Policy awards points based on the current episode public utility accuracy, thereby promoting higher public accuracies. These models were able to keep the private task accuracy at or near chance. However, the finetuned public utility accuracy was lower than the accuracy of the Public-Based Reward Policy.

7. Discussion and Limitations

The best performing model was the Public-Based Reward Policy, as it succeeded in keeping the images private and the public task accuracy better than random. This is significant given the state-of-the-art. In [27], which is a conventional, purely silicon-based deep learning method, the difference between their public baseline (blur, 86.9%) and their proposed method (learned encoder, 91.6%) was 4.7%.

Privacy is a hard problem, and we have taken the first step into quantum-based privacy preservation. In light of this, on-par performance with the public task (while doing better on the private task) is a reasonable performance.

We are the first to demonstrate that quantum gates can be used, in concert with a machine learning algorithm, to learn privacy preserving quantum actions for images. However, the sub-field we are investigating here has many challenges that we document, and there are two main limitations:

1. Quantum cameras are currently in the early stages of development, and real quantum images are extremely low-resolution. Therefore simulations have resulted in long training and test times.
2. While our method shows promise, currently we are only slightly better than augmenting images with Gaussian noise. However, due to the sophisticated nature of quantum actions, we believe that, in the future, we will be able to fully exploit the action space to increase privacy protection.

We have selected quantum circuits for privacy, integrated QImP terminology into a machine learning framework and written code for our novel design that we will release so that other researchers may build on our initiative. Quantum computing in vision will have many impacts, and this paper shows that privacy preservation will be one of them.

References

- [1] P. Agrawal and P. Narayanan. Person de-identification in videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 2011. 2
- [2] Narayan Bhusal, Mingyuan Hong, Ashe Miller, Mario A. Quiroz-Juárez, Roberto de J. León-Montiel, Chenglong You, and Omar S. Magaña-Loaiza. Smart quantum statistical imaging beyond the abbe-rayleigh criterion. 8(1):1–8. Number: 1 Publisher: Nature Publishing Group. 2
- [3] Karla Brkic, Ivan Sikiric, Tomislav Hrkac, and Zoran Kalafatic. I know that person: Generative full body and face de-identification of people in images. *CVPRW*, 1(2):4, 2017. 2
- [4] Giacomo Cavalieri and Dario Maio. A quantum edge detection algorithm. Number: arXiv:2012.11036 version: 1. 3
- [5] Julie Chang, Vincent Sitzmann, Xiong Dun, Wolfgang Heidrich, and Gordon Wetzstein. Hybrid optical-electronic convolutional neural networks with optimized diffractive optics for image classification. 8:12324. 2
- [6] Julie Chang and Gordon Wetzstein. Deep optics for monocular depth estimation and 3d object detection. 2
- [7] D. Chen, Y. Chang, R. Yan, and J. Yang. Tools for protecting the privacy of specific individuals in video. *EURASIP Journal on Advanced Signal Processing*, 2006. 2
- [8] Gregory Cohen, Saeed Afshar, Jonathan Tapson, and André van Schaik. EMNIST: an extension of MNIST to handwritten letters. version: 1. 4
- [9] B. Driessen and M. Durmuth. Achieving anonymity against major face recognition algorithms. *Lecture notes in computer science*, 2013. 2
- [10] Cynthia Dwork. Differential privacy: A survey of results. In *International Conference on Theory and Applications of Models of Computation*, pages 1–19. Springer, 2008. 2
- [11] Alexander Geng, Ali Moghiseh, Claudia Redenbach, and Katja Schladitz. Improved FRQI on superconducting processors and its restrictions in the NISQ era. 22(2):104. 2, 3
- [12] R. Gross, L. Sweeney, F. de la Torre, and S. Baker. Semi-supervised learning of multi-factor models for face de-identification. *CVPR*, 2008. 2
- [13] Rakibul Hasan, Eman Hassan, Yifang Li, Kelly Caine, David J Crandall, Roberto Hoyle, and Apu Kapadia. Viewer experience of obscuring scene elements in photos to enhance privacy. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2018. 2
- [14] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. 30(1). Number: 1. 6
- [15] Roberto Hoyle, Luke Stark, Qatrunnada Ismail, David Crandall, Apu Kapadia, and Denise Anthony. Privacy norms and preferences for photos posted online. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 27(4):1–27, 2020. 1
- [16] IBM. IBM quantum. 4
- [17] Abdullah M. Ilyyasu. Towards realising secure and efficient image and video processing applications on quantum computers. 15(8):2874–2974. Number: 8 Publisher: Multidisciplinary Digital Publishing Institute. 5
- [18] Myeung Un Kim, Harim Lee, Hyun Jong Yang, and Michael S. Ryoo. Privacy-preserving robot vision with anonymized faces by extreme low resolution. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 462–467. ISSN: 2153-0866. 2
- [19] Jose I. Latorre. Image compression and entanglement. 3
- [20] Phuc Q. Le, Fangyan Dong, and Kaoru Hirota. A flexible representation of quantum images for polynomial preparation, image compression, and processing operations. 10(1):63–84. 3
- [21] Changhao Li, Rouhollah Soleyman, Mohammad Kohandel, and Paola Cappellaro. SARS-CoV-2 quantum sensor based on nitrogen-vacancy centers in diamond. 22(1):43–49. 2
- [22] Ludong Li, Zheng Lou, and Guozhen Shen. Flexible broadband image sensors with SnS quantum dots/zn2sno4 nanowires hybrid nanostructures. 28(6):1705389. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.201705389>. 2
- [23] Christopher Monroe, DM Meekhof, BE King, and David J Wineland. A “schrödinger cat” superposition state of an atom. *Science*, 272(5265):1131–1136, 1996. 1
- [24] Anirban Pathak. *Elements of Quantum Computation and Quantum Communication*. Taylor & Francis Group. 1
- [25] Vladimir Pejović, Epimitheas Georgitzikis, Jiwon Lee, Itai Lieberman, David Cheyns, Paul Heremans, and Paweł E. Malinowski. Infrared colloidal quantum dot image sensors. 69(6):2840–2850. Conference Name: IEEE Transactions on Electron Devices. 2
- [26] Vladimir Pejović, Epimitheas Georgitzikis, Itai Lieberman, Paweł E. Malinowski, Paul Heremans, and David Cheyns. Photodetectors based on lead sulfide quantum dot and organic absorbers for multispectral sensing in the visible to short-wave infrared range. 32(28):2201424. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/adfm.202201424>. 2
- [27] Francesco Pittaluga, Sanjeev Koppal, and Ayan Chakrabarti. Learning privacy preserving encodings through adversarial training. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 791–799. IEEE. 2, 8
- [28] Francesco Pittaluga and Sanjeev Jagannatha Koppal. Pre-capture privacy for small vision sensors. 39(11):2215–2226. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence. 2
- [29] L. Sweeney. k-anonymity: A model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 2002. 2
- [30] Robert Templeman, Mohammed Korayem, David J Crandall, and Apu Kapadia. Placeavoider: Steering first-person cameras away from sensitive spaces. In *NDSS*, volume 14, pages 23–26, 2014. 2
- [31] Hado van Hasselt, Arthur Guez, and David Silver. Deep reinforcement learning with double q-learning. type: article. 6
- [32] Igor Vasiljevic, Ayan Chakrabarti, and Gregory Shakhnarovich. Examining the impact of blur on recognition by convolutional networks. *arXiv preprint arXiv:1611.05760*, 2016. 2

- [33] Salvador E. Venegas-Andraca and Sougato Bose. Storing, processing, and retrieving an image using quantum mechanics. In *Quantum Information and Computation*, volume 5105, pages 137–147. SPIE. 3
- [34] Henning Weier, Harald Krauss, Markus Rau, Martin Fürst, Sebastian Nauerth, and Harald Weinfurter. Quantum eavesdropping without interception: an attack exploiting the dead time of single-photon detectors. 13(7):073024. Publisher: IOP Publishing. 3
- [35] Osian Wolley, Thomas Gregory, Sebastian Beer, Takafumi Higuchi, and Miles Padgett. Quantum imaging with a photon counting camera. 12(1):8286. Number: 1 Publisher: Nature Publishing Group. 2
- [36] Zuobin Xiong, Wei Li, Qilong Han, and Zhipeng Cai. Privacy-preserving auto-driving: A GAN-based approach to protect vehicular camera data. In *2019 IEEE International Conference on Data Mining (ICDM)*, pages 668–677. IEEE. 2
- [37] Fei Yan, Abdullah M. Ilyasu, and Zhengang Jiang. Quantum computation-based image representation, processing operations and their applications. 16(10):5290–5338. Number: 10 Publisher: Multidisciplinary Digital Publishing Institute. 3, 5
- [38] Fei Yan, Abdullah M. Ilyasu, and Salvador E. Venegas-Andraca. A survey of quantum image representations. 15(1):1–35. 3
- [39] Xi-Wei Yao, Hengyan Wang, Zeyang Liao, Ming-Cheng Chen, Jian Pan, Jun Li, Kechao Zhang, Xingcheng Lin, Zhehui Wang, Zhihuang Luo, Wenqiang Zheng, Jianzhong Li, Meisheng Zhao, Xinhua Peng, and Dieter Suter. Quantum image processing and its application to edge detection: Theory and experiment. 7(3):031041. 2, 3
- [40] X. Yu, K. Chinomi, T. Koshimizu, N. Nitta, Y. Ito, and N. Babaguchi. Privacy protecting visual processing for secure video surveillance. *IEEE ICIP*, 2008. 2
- [41] Yi Zhang, Kai Lu, Yinghui Gao, and Mo Wang. NEQR: a novel enhanced quantum representation of digital images. 12(8):2833–2860. 3