Quantum-Enhanced Autonomy: Augmenting Generative AI for Critical Test Scenario Images

Airbus - BMW Group
Quantum Computing Challenge 2024 | Forward Track
version 1.0

1 Introduction

As automated transportation becomes increasingly prevalent in both air and ground mobility, ensuring the reliability and safety of autonomous decisions made by artificial intelligence (AI) has become paramount. Autonomous vehicles require sophisticated scenario navigation abilities for complex road objects and traffic participants, while autonomous aviation systems, particularly in the context of autonomous landing, primarily rely on their runway detection systems under varying conditions. Specifically, one development focus of these AI systems is on their proficiency in handling critical scenarios – rare occurrences that have the potential to cause system failures. As these events are rare in daily life, automatic generation of critical test scenarios [1, 2] is of utmost importance to ensure maximum safety in mobility.

In recent years, the combination of two fields, quantum computing (QC) and AI has given rise to quantum machine learning (QML), which holds great potential for overcoming complex challenges that were previously deemed unsolvable [3, 4, 5]. QML opens up exciting opportunities for autonomous vehicles and aviation systems. Specifically, due to quantum computers' ability to learn and sample from high-dimensional probability distributions [6, 7, 8], quantum generative models have demonstrated either theoretical or empirical advantages over classical algorithms [9, 10, 11] in some cases and are thus particularly promising. The task of critical test scenario generation thus serves as an ideal testing ground for the application of QML in the automotive and aviation industries.

This problem statement focuses on the generation of synthetic images that encapsulate critical scenarios in both autonomous vehicles and aviation systems domains. These critical scenarios encompass a wide range of challenging conditions, such as low visibility due to night-time conditions or adverse weather[12], intricate traffic patterns, and obstructions on runways. Therefore by establishing a repository of images for critical low-visibility scenarios that mimic or replicate, for instance, nighttime conditions, one can advance research in the perception algorithms for both automotive and aviation domains. Thus it will enhance algorithms' capabilities to navigate through dynamic real-world environments. Therefore in this challenge, the focus would be on generating such a repository of critical low-visibility nighttime scenario images by either, performing style transfers on high-visibility daytime images to transform them into low-visibility nighttime images or using other approaches that perform the same task.

2 Classical Generative Modelling Methods

Classical generative modelling methods, primarily rooted in deep learning and probabilistic modeling, have demonstrated significant achievements in generating realistic images [13, 14]. However, these methods exhibit certain limitations, such as slow sampling [15, 16] and security [17, 18] for diffusion models, difficult convergence for GANs [19, 20, 21] and image quality issue of VAEs [22, 23] that motivate the use and exploration of quantum generative modelling approaches. There are various classical generative modelling methods, that have gained relevance for good quality image generation. The current state of the art in generative modelling for images is characterized by a diverse landscape of techniques that have achieved remarkable milestones in generating realistic and high-quality visual content.

Variational Autoencoder (VAE) [24, 25] is one of many popular probabilistic models used for image generation and representation learning. VAEs learn a latent space representation of data and allow for controlled image generation. However, VAEs often struggle with capturing complex data distributions due to their inherent assumption of simple Gaussian latent spaces. This limitation can result in sometimes unrealistic image generation.

Generative Adversarial Network (GAN) [26, 27, 28], also gained popularity in image generation which has demonstrated the capability to produce visually convincing images through adversarial training. GANs consist of a generator and a discriminator that play a two-player min-max game, leading to the generation of high-quality images. Despite their impressive results, GANs are known to be very difficult to train and suffer from issues such as mode collapse, where the generator on many occasions tends to produce limited types of samples, and training instability.

Diffusion Models (DM) [29, 30, 31, 16, 15] have also recently emerged as an effective approach to generate good quality images, offering a unique perspective on image generation by focusing on the process of iteratively transforming a simple distribution into a complex one. This enables them to offer improved stability during training, making them less prone to issues like mode collapse or vanishing gradients. But given the iterative approach of the method, diffusion models also tend to be computationally more expensive and are occasionally prone to generating unrealistic images if the denoising algorithm fails.

In conclusion, though classical generative modelling methods have made remarkable strides in image generation, their limitations, however, motivate the exploration of quantum generative modelling.

Leveraging the unique properties of quantum physics, complexity of Hilbert space and better data representation, quantum generative modelling might have the potential to overcome some of the above mentioned shortcomings of the classical methods and to accelerate the generation of better, high-quality, diverse, and coherent images.

3 Quantum Generative Modelling Methods

In recent years, the field of quantum generative modeling has witnessed a surge in innovation, with various algorithms and techniques making their way into the spotlight. Notable examples include the Quantum Circuit Born Machine (QCBM)[32, 33], the Quantum Boltzmann Machine [34, 35], and Quantum Generative Adversarial Networks (QGANs) [36, 37, 38]. These algorithms harness the unique properties of quantum computing to generate data and probability distributions, offering a fresh perspective on generative modeling in the quantum realm.

Generative adversarial networks employ an adversarial learning protocol mimicing the dynamics of a two-player game where the players correspond to a parameterized generator and a parameterized discriminator. QGANs harness the inherent principles of superposition and entanglement, potentially leading to enhanced efficiency in representing and sampling complex probability distributions. This unique aspect of quantum information may also provide greater flexibility to both the generator and discriminator, potentially allowing them to converge more efficiently to the Nash equilibrium.

Notably, Loyds and Weedbrook [36] have demonstrated that in scenarios where the data is classical and high-dimensional, and both the generator and discriminator are quantum, QGANs can potentially achieve an exponential advantage over their classical counterparts. This result hints at the transformative potential of quantum generative modeling in scenarios where classical methods struggle to cope with the complexities of high-dimensional classical data.

Despite the exciting advancements, the field of quantum generative modelling is still in the early stages. While these quantum generative modeling techniques hold a lot of potential, it is necessary to conduct more real-world tests and to better understand their capabilities and limitations. In this context, it is of paramount importance to test different types of related quantum algorithms for specific industrially relevant situations.

4 Case Study

For both vehicles and aviation systems, motion planning is more difficult during the night compared to during the day. In this task, we consider the images of critical test scenarios to be images of road and runways during the night. We wish to generate the images by performing image-to-image translation, meaning to translate images of a source domain (day) to the target domain (night).

The task should be solved under the unsupervised learning setting, meaning that images are not provided in pairs, since obtaining paired data is typically difficult and expensive. If the participants find difficulty in constructing an unsupervised learning solution, they may alternatively perform supervised learning with paired data constructed or found with extra resources, however at a cost of potential missing score in the final evaluation.

We have two datasets, one for autonomous vehicles and the other for autonomous aviation systems. For autonomous vehicles, we have an open dataset, the Berkeley DeepDrive dataset (BDD100K). [1] This dataset is widely used for image classification, detection, and segmentation tasks. It consists of

[1] https://paperswithcode.com/dataset/bdd100k, https://arxiv.org/pdf/1805.04687.pdf

$\sim 100,000$ images at 720p resolution, split into $\sim 70,000$ training, $\sim 10,000$ validation, and $\sim 20,000$ test sets. The images are taken under different weather conditions, scenes and times of day, as shown in tables 1,2,3.

clear	overcast	undefined	snowy	rainy	partly cloudy	foggy	total
42690	10009	9276	6318	5808	5619	143	79863

Table 1: training and validation images divided by weather

city street	highway	residential	parking lot	undefined	tunnel	gas stations	total
49628	19878	9327	426	414	156	34	79863

Table 2: training and validation images divided by scene

daytime	night	dawn/dusk	undefined	total
41986	31900	5805	172	79863

Table 3: training and validation images divided by time of day

For autonomous aviation (autonomous landing) system, we have an open dataset, Landing Approach Runway Detection (LARD) [2] [3]. It consists of high-quality aerial images for the task of runway detection during approach (final preparatory step before landing) and landing phases. Most of the dataset is composed of synthetic images produced with a virtual globe tool (Google Earth Studio), but also contains manually labelled images from real landing footages. The training set includes synthetic images of several specific runways regrouped by airports indicated in the filenames. The test set includes both synthetic images as well as a specific archive for real images. In total there are around 15k real and synthetic images of different sizes like, 2448×2648, 3840×2160 or 1920×1080 from multiple runways and airports.

The participants are allowed to perform appropriate pre-processing of both datasets on their own. The final submitted trained models should take images of the source domain, and/or random seeds and/or additional tunable parameters to output images of the target domain. It is possible to add parameters if deemed appropriate which are not part to the dataset, e.g., random numbers, additional variables assigned to the output images.

5 Submission Guidelines and Key Performance Metrics

In this section, we offer participants precise directives that augment the general guidelines within the framework of Quantum Generative Modelling. We also emphasize to follow the general submission guidelines provided on the challenge website. The evaluation of the submitted approaches will be based on the following Key Performance Metrics.

- The designed algorithm is expected to appropriately utilise quantum components at least on a partial level. This may include quantum-hardware-dependent but also quantum inspired solutions.
- The motivation and (potential) advantage of the quantum solution shall be demonstrated in fair comparison with classical state-of-the-art solutions.

[2] https://arxiv.org/pdf/2304.09938.pdf, https://github.com/deel-ai/LARD
[3] https://share.deel.ai/s/H4iLKRmLkdBWqSt?path=%2Flard%2F1.0.0

- The feasibility of the algorithm needs to be proven, so it needs to be formulated in a commonly used language. Scripts / notebooks for testing on at least one of the datasets should be provided. A classical simulation of the quantum components may be used. Emulation or usage of quantum hardware may provide a more substantial proof feasibility but is not required.
- We are interested not only in the performance of the methods on the datasets, but also studies on the (speed) computational complexity of the algorithm, together with an analysis of the requirements on the quantum resources and scalability.
- The quality of generated images should be evaluated by the submitter, based on the recommendations provided by domain experts from Airbus and BMW. By default, Fréchet Inception Distance (FID) will be used. However, the participants may use alternative metrics if they favor the training and evaluation of the algorithm (Need to justify the choice, and provide the code implementation of the chosen metrics.)
- When developing industrial applications around generative modelling, usually the generated images need to be evaluated by human experts. Thus, domain experts at Airbus and BMW will contribute to the evaluation of the generated images.
- Provide an assessment of how many training images are used and how long the training process may take.
- Assess how well the algorithm generalizes to the data from both aerospace and automotive domains.

References

[1] Moritz Klischat and Matthias Althoff. Generating critical test scenarios for automated vehicles with evolutionary algorithms. In *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 2352–2358, 2019.

[2] Ian Rhys Jenkins, Ludvig Oliver Gee, Alessia Knauss, Hang Yin, and Jan Schroeder. Accident scenario generation with recurrent neural networks. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 3340–3345, 2018.

[3] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195–202, September 2017.

[4] Vedran Dunjko and Hans J Briegel. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics*, 81(7):074001, July 2018.

[5] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, April 2015.

[6] Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, October 2019.

[7] Hui Wang, Jian Qin, Xing Ding, Ming-Cheng Chen, Si Chen, Xiang You, Yu-Ming He, Xiao Jiang, L. You, Z. Wang, C. Schneider, Jelmer J. Renema, Sven Höfling, Chao-Yang Lu, and Jian-Wei Pan. Boson sampling with 20 input photons and a 60-mode interferometer in a 10^{14}-dimensional hilbert space. *Phys. Rev. Lett.*, 123:250503, Dec 2019.

[8] Niklas Pirnay, Ryan Sweke, Jens Eisert, and Jean-Pierre Seifert. Superpolynomial quantum-classical separation for density modeling. *Physical Review A*, 107(4):042416, April 2023.

[9] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert. On the quantum versus classical learnability of discrete distributions. *Quantum*, 5:417, mar 2021.

[10] Brian Coyle, Daniel Mills, Vincent Danos, and Elham Kashefi. The born supremacy: quantum advantage and training of an ising born machine. *npj Quantum Information*, 6(1), jul 2020.

[11] Manuel S. Rudolph, Ntwali Bashige Toussaint, Amara Katabarwa, Sonika Johri, Borja Peropadre, and Alejandro Perdomo-Ortiz. Generation of high-resolution handwritten digits with an ion-trap quantum computer, 2022.

[12] Zhaoyang Wang, Dan Zhao, and Yunfeng Cao. Visual navigation algorithm for night landing of fixed-wing unmanned aerial vehicle. *Aerospace*, 9, 10 2022.

[13] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans, 2018.

[14] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis, 2021.

[15] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):10850–10869, 2023.

[16] Hanqun Cao, Cheng Tan, Zhangyang Gao, Yilun Xu, Guangyong Chen, Pheng-Ann Heng, and Stan Z. Li. A survey on generative diffusion model, 2023.

[17] Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwag, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting training data from diffusion models, 2023.

[18] Yihan Cao, Siyu Li, Yixin Liu, Zhiling Yan, Yutong Dai, Philip S. Yu, and Lichao Sun. A comprehensive survey of ai-generated content (aigc): A history of generative ai from gan to chatgpt, 2023.

[19] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans, 2016.

[20] Naveen Kodali, Jacob Abernethy, James Hays, and Zsolt Kira. On convergence and stability of gans, 2017.

[21] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for GANs do actually converge? In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3481–3490. PMLR, 10–15 Jul 2018.

[22] Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Adversarial variational bayes: Unifying variational autoencoders and generative adversarial networks. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2391–2400. PMLR, 06–11 Aug 2017.

[23] Ruoqi Wei, Cesar Garcia, Ahmed El-Sayed, Viyaleta Peterson, and Ausif Mahmood. Variations in variational autoencoders - a comparative evaluation. *IEEE Access*, 8:153651–153670, 2020.

[24] Diederik P Kingma and Max Welling. Auto-encoding variational bayes, 2022.

[25] Carl Doersch. Tutorial on variational autoencoders, 2021.

[26] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014.

[27] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks, 2020.

[28] Chaoyue Wang, Chang Xu, Chaohui Wang, and Dacheng Tao. Perceptual adversarial networks for image-to-image transformation. *IEEE Transactions on Image Processing*, 27(8):4066–4079, aug 2018.

[29] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pages 2256–2265. PMLR, 2015.

[30] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

[31] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32, 2019.

[32] Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6), dec 2018.

[33] Marcello Benedetti, Delfina Garcia-Pintos, Oscar Perdomo, Vicente Leyton-Ortega, Yunseong Nam, and Alejandro Perdomo-Ortiz. A generative modeling approach for benchmarking and training shallow quantum circuits. *npj Quantum Information*, 5(1), may 2019.

[34] Mohammad H. Amin, Evgeny Andriyash, Jason Rolfe, Bohdan Kulchytskyy, and Roger Melko. Quantum boltzmann machine. *Physical Review X*, 8(2), may 2018.

[35] Feng Hu, Ban-Nan Wang, Ning Wang, and Chao Wang. Quantum machine learning with D-wave quantum computer. *Quantum Engineering*, 1(2), June 2019.

[36] Seth Lloyd and Christian Weedbrook. Quantum generative adversarial learning. *Physical Review Letters*, 121(4), jul 2018.

[37] Pierre-Luc Dallaire-Demers and Nathan Killoran. Quantum generative adversarial networks. *Physical Review A*, 98(1), jul 2018.

[38] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner. Quantum generative adversarial networks for learning and loading random distributions. *npj Quantum Information*, 5(1), nov 2019.