
This is the **published version** of the bachelor thesis:

Vargas Orellana, Adrian Daniel; Vilariño Freire, Fernando Luis, dir. Future Perspectives in Image Generation:Advancements in GANs and QGANs. 2024. (Enginyeria Informàtica)

This version is available at <https://ddd.uab.cat/record/290067>

under the terms of the  license

Future Perspectives in Image Generation: Advancements in GANs and QGANs

Adrian Vargas Orellana

6 de febrer de 2024

Resum– Este artículo investiga las perspectivas futuras de los modelos de predicción de distribuciones, centrándose en las distribuciones gaussiana, triangular, lognormal y binomial. Compara los modelos de predicción clásicos con sus homólogos cuánticos, explorando los avances realizados en los modelos de predicción cuántica (QPM). Además, explora el potencial de la computación cuántica para mejorar la precisión y eficacia de la predicción de estos modelos. Comenzando con un análisis comparativo básico, el estudio avanza para evaluar modelos de predicción puramente cuánticos antes de proponer un enfoque integrado que combina metodologías clásicas y cuánticas para perfeccionar las técnicas de predicción de la distribución. Esta investigación ofrece valiosas perspectivas sobre el cambiante panorama de la predicción de la distribución y el papel de la computación cuántica en la configuración de su futuro.

Paraules clau– Xarxes Neuronal, Xarxa Neuronal Generativa Adversarial, Computació Quàntica, Models Generatius, Generació d'imatges

Abstract– This paper investigates the future perspectives in distribution prediction models, focusing on Gaussian, triangular, lognormal, and binomial distributions. It compares classical prediction models with their quantum counterparts, exploring the advancements made in Quantum Prediction Models (QPMs). Furthermore, it explores the potential of quantum computing in enhancing the prediction accuracy and efficiency of these models. Beginning with a basic comparative analysis, the study progresses to evaluate purely quantum prediction models before proposing an integrated approach that combines classical and quantum methodologies to refine distribution prediction techniques. This research offers valuable insights into the evolving landscape of distribution prediction and the role of quantum computing in shaping its future.

Keywords– Neural Networks, Generative Adversarial Neural Network, Quantum Computing, Generative Models, Quantum Gain Adversarial Network

1 INTRODUCTION

THE landscape of image generation is undergoing a revolutionary transformation with the infusion of Quantum Neural Networks (QNNs). In the wake of quantum computing's relentless march into uncharted computational territories, the integration of quantum principles into neural networks introduces a fascinating avenue for the evolution of image generation methodologies. This exploration seeks to unravel the vast potential that Quantum Neu-

ral Networks hold in reshaping the dynamics of image synthesis.

Quantum Neural Networks, guided by the principles of quantum mechanics, particularly superposition and entanglement, offer a unique perspective for harnessing computational power beyond the limits of classical computing. This study delves into the intricacies of Quantum Neural Networks, investigating their capacity to redefine the very fabric of image generation. The quantum states, with their inherent ability to exist in multiple configurations simultaneously, present an alluring prospect for pushing the boundaries of creativity and efficiency in the realm of visual synthesis.

This research endeavors to unlock novel methods and unprecedented efficiencies in image generation by tapping into the distinctive features of quantum mechanics. By

- E-mail de contacte: adrian.v.o2002@gmail.com
- Menció realitzada: Enginyeria de Computació
- Treball tutoritzat per: Fernando Vilariño (DCC)
- Curs 2023/24

translating complex image data into quantum states, Quantum Neural Networks promise to revolutionize the speed, precision, and creativity associated with image synthesis. Through this exploration, we aspire to illuminate the path toward a future where quantum-assisted image generation [1] becomes an integral part of the evolving landscape of artificial intelligence.

2 STATE OF ART

The current state of the art of quantum neural networks for image generation is very precarious, taking this into account, the problems solved by this type of models are currently not very developed, even so we have work for the possible future in the field of quantum neural networks, that can be very interesting and useful for the resolution of this type of problems, at the moment we have work that allows us to solve some type of specific problem like reproducing gaussian distributions by means of Quantum Generative Adversarial Network (qGAN's), but for this, first we have to understand the following concepts:

2.1 Generative Adversarial Network

GANs consist of two main neural networks[2]: the generator and the discriminator. These networks are trained simultaneously through a competitive process, making GANs unique in their approach[3].

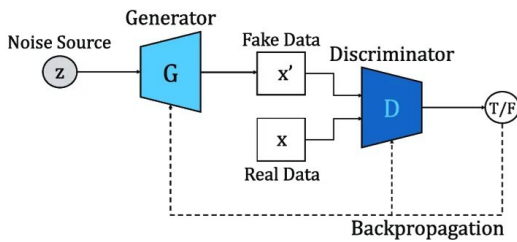


Fig. 1: GAN Schema

2.1.1 Generator

The generator takes random input, commonly referred to as noise, and transforms it into generated data that should be indistinguishable from real data. Mathematically, the generator seeks to learn the underlying distribution of the training data to generate convincing samples.

2.1.2 Discriminator

On the other hand, the discriminator evaluates whether a given sample is real or generated by the generator. Essentially, it acts as a "detective" trying to distinguish between real and generated data. It is trained to improve its ability to make this distinction over time.

2.1.3 Functioning of GANs

The training process of GANs occurs through an adversarial game between the generator and the discriminator. The generator continually strives to enhance its ability to generate more realistic data, while the discriminator improves its ability to differentiate between real and generated data.

This feedback loop continues until the generator can produce data that is indistinguishable from real data.

2.2 Quantum Computing

In the realm of quantum computing, a fundamental tool that plays a pivotal role in our project is Qiskit. Developed by IBM, Qiskit stands out as a comprehensive quantum computing software development framework. It serves as a versatile platform for designing, simulating, and executing quantum circuits on both real quantum devices and classical simulators[4].

2.3 Quantum Machine Learning (QML)

In this section, we delve deeper into the intersection of machine learning and quantum circuits, drawing a parallel between classical machine learning models and their quantum counterparts. Classical machine learning models traditionally consist of layers, inputs, and outputs, where the inputs and outputs share namesakes, and layers are composed of neurons responsible for training the models.

In the quantum realm, a transformative shift occurs when adapting classical machine learning to quantum neural networks. Not only does the paradigm shift involve a change in the type of data from classical to quantum, but the very structure of the models undergoes a metamorphosis. Instead of training through layers of classical neurons, quantum neural networks utilize layers of Anzats for their training process. Anzats, in this context, serve as parameterized quantum circuits, offering a flexible and expressive framework for quantum machine learning tasks[5].

2.4 Generative Adversarial Network

Now, extending this discussion to the realm of generative models, we introduce the concept of Generative Adversarial Networks (GANs). GANs, in classical machine learning, are adept at generating new data instances that closely resemble a given dataset. They consist of a generator network and a discriminator network engaged in an adversarial training process, where the generator strives to create realistic data, and the discriminator aims to distinguish between real and generated data[6].

In the quantum landscape, Quantum Generative Adversarial Networks (QGANs) emerge as a natural extension. Here, the adversarial training process persists, but the underlying mechanisms leverage the principles of quantum computation. QGANs harness the unique capabilities of quantum circuits, offering potential advantages in generating quantum data distributions. The layers of Anzats play a crucial role in the training of QGANs, providing a quantum analogue to the classical GAN framework[6].

This nuanced connection between classical GANs and QGANs exemplifies the ongoing synergy between classical machine learning and quantum computing, opening avenues for exploring generative models in the quantum domain[6].

2.5 Classical-Quantum Generative Diffusion Model (CQGDM)

The proposal of a Classical-Quantum Generative Diffusion Model (CQGDM)[7] marks a significant advancement by combining elements of quantum physics with classical machine learning techniques. In this innovative amalgamation, the diffusion process remains classical, while the denoising phase is executed through quantum dynamics. This hybrid approach represents an evolution in data generation, as the training dataset remains classical, encompassing images, videos, and time series. Formally, the model begins with data sampled from an unknown distribution, subjecting it to a classical stochastic diffusion process that progressively degrades the information to a fully noisy state. The introduction of a computationally tractable quantum prior is a key aspect, and the diffusion process can be implemented with both classical Markov chains [8, 9] and stochastic differential equations. The distinctive innovation lies in leveraging a Quantum Neural Network (QNN) for the denoising phase, harnessing quantum properties to expedite the generation of highly dimensional data, such as images. This fusion of classical and quantum approaches opens new possibilities for efficient data processing and enhanced generative model capabilities.

2.6 Quantum-Classical Generative Diffusion Model (QCQDM)

The Quantum-Classical Generative Diffusion Model (QCQDM)[7] proposes an innovative convergence of quantum and classical elements to revolutionize generative diffusion processes. By leveraging noisy quantum dynamics, quantum noise is harnessed as a valuable resource. This approach is applied to quantum datasets, blending classical information into quantum initial states. Two approaches are introduced for implementing the diffusion process, emphasizing the versatility of quantum noise and the ability to generate non-classical distributions. During the noise removal stage, classical Neural Networks (NNs) play a crucial role, serving as discriminators in the case of quantum distributions. This model holds promising applications, particularly in cybersecurity, enabling quantum attacks and defense.

2.7 Quantum-Quantum Generative Diffusion Model (QQGDM)

The Quantum-Quantum Generative Diffusion Model (QQGDM)[7] represents an innovative approach that exclusively leverages quantum elements throughout the entire process. From the initial diffusion to the denoising phase, noisy quantum dynamics are employed, utilizing both quantum Markov chains and the Stochastic Schrödinger Equation. This model distinguishes itself by operating entirely within the quantum realm, with Parameterized Quantum Circuits playing a central role in the denoising process. The QQGDM aims not only to preserve but also to amplify quantum advantages, enabling the generation of purely quantum prior distributions and efficiently processing them during the denoising phase. This could potentially lead to exponential advantages in sample complexity and

processing time. With an exclusively quantum focus, the QQGDM represents a significant contribution to the field of quantum generative models.

3 METHODOLOGY

The methodology will aim to discover and outline a detailed roadmap for creating a quantum neural network capable of generating images using quantum algorithms. This approach will meticulously consider crucial aspects, such as quantum encoding of image data, efficient structuring of quantum circuits, integration of reinforcement learning techniques for parameter optimization, and comprehensive performance evaluation compared to classical approaches. To achieve these goals, we will specifically focus on the following aspects:

3.1 Development of Quantum Data Encoding

To train a Quantum Neural Network (QNN) model, the first step involves encoding the data so that it can be used as input[10]. This encoding process transforms classical information into a quantum format understandable by the model, enabling processing within the quantum environment. This encoding phase is crucial to leverage the quantum capabilities of the model during the training process[7].

3.2 Training Classical & Quantum Models

In the realm of machine learning, the training of classical and quantum models stands as a pivotal process shaping the performance and efficacy of predictive algorithms. In classical models, the training phase typically involves the optimization of parameters and the adjustment of weights within neural networks or other algorithmic frameworks. Conversely, the training of quantum models introduces a unique paradigm, necessitating the encoding of classical data into quantum states for processing within a quantum environment. The optimization of quantum circuit parameters becomes a crucial aspect, often involving techniques such as gradient-based optimization or reinforcement learning to refine the model's performance. Understanding the nuances and distinctions in the training methodologies between classical and quantum models is fundamental for harnessing the full potential of these diverse computational approaches.

3.3 Comparative Evaluation of models

Following the training process of classical and quantum models, the subsequent step involves a critical evaluation by comparing the performance of classical and quantum Generative Adversarial Networks (qGANs)[3]. This comparative analysis aims to discern the strengths and weaknesses inherent in each approach, shedding light on the potential advantages offered by quantum algorithms over classical counterparts in the context of generative models. By systematically assessing metrics such as accuracy, computational efficiency, and scalability, researchers can gain valuable insights into the comparative advantages of qGANs, paving the way for informed decision-making and the advancement of quantum-enhanced generative modeling techniques.

4 DEVELOPMENT & PRELIMINARY RESULTS

Given the current state of research in the field and the status of quantum systems, the progress made for this report primarily revolves around the following aspects. The main approach has been to replicate Gaussian distributions and then translate them into a quantum format, with the aim of applying these concepts to image generation. Additionally, both Gaussian and quantum GAN models will be implemented to achieve a normal distribution, intending to conduct a comprehensive comparison between both approaches with pytorch.

4.1 Generative Adversarial Network

The implemented code is a Generative Adversarial Network (GAN)[2] designed to predict a Gaussian distribution. The GAN comprises a generator and a discriminator. The generator is a neural network with three linear layers and leaky ReLU activation functions, while the discriminator has a similar architecture but includes a sigmoid activation function in the final layer. The GAN itself orchestrates the training process, which involves optimization and loss calculation. The goal is to train the generator to produce data indistinguishable from real data sampled from a Gaussian distribution, while the discriminator aims to differentiate between real and generated data. The training loop involves adjusting the generator's weights to deceive the discriminator and updating the discriminator to better distinguish between real and generated samples. The final output includes visualizations of loss and distributions during training. Given the configurations mentioned earlier, the result we obtain is as follows2:

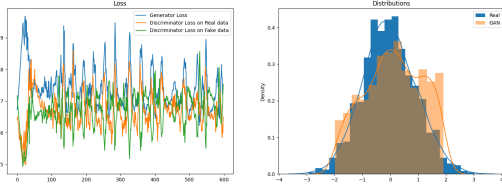


Fig. 2: GAN Results v1

However, it is important to note that, after this initial implementation, a second iteration of the training was carried out using a different model architecture. The results of this second approach showed some improvement over the initial version. It is worth noting that the model was trained in two separate instances: the first consisted of a 1200-step3 training, while the second involved a more extensive 30,000-step4 training.

In the first 1200-step phase, we sought to gain an initial insight into the model's ability to learn the Gaussian distribution. This initial training provided a basis on which adjustments were made to the model architecture to address possible limitations identified.

Subsequently, with the experience gained in the initial phase, a second, more extensive training phase was undertaken, carrying out 30,000 steps4. This prolonged approach allowed the model to learn in a deeper and more refined way, improving its ability to generate data that further resembles the desired Gaussian distribution. Together, these

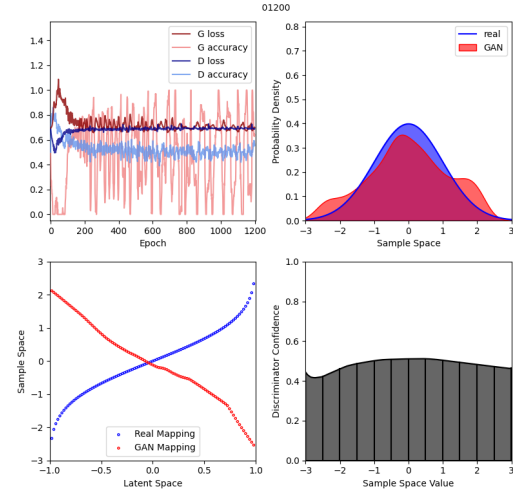


Fig. 3: GAN Results v2 1.2 Steps

two iterations illustrate the iterative and evolutionary process employed to optimise GAN's capabilities in predicting Gaussian distributions. The final results are presented in Figure 3 and 4.

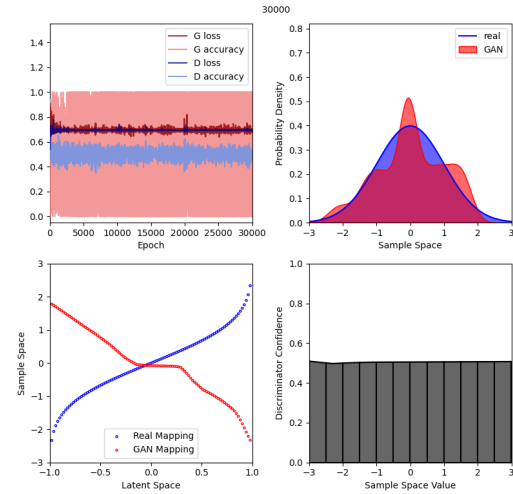


Fig. 4: GAN Results v2 30k Steps

Top Left (Metrics): The generator and discriminator both use binary cross-entropy loss and accuracy. Although generator loss doesn't directly indicate GAN output quality, it's crucial for tracking convergence. If one significantly outperforms the other, the GAN may fail to converge.

Top Right (Probability Density): Compares the observed probability density of the GAN to the real probability density of the standard normal distribution. Ideally, they should align at the end of training, but the GAN shows repetitive over-correction, leading to undulating behavior.

Bottom Left (Learned Mapping): Displays the real mapping used and the mapping learned by the GAN. The GAN tries to mimic a complex mapping with varying success.

Bottom Right (Discriminator Confidence): Shows the discriminator's confidence in identifying real and fake samples. The discriminator effectively recognizes the relative frequency of real and fake samples, with confidence fluctuations at the extremes of the sample

space. Increased confidence in rare real samples prompts undulating behavior in the probability density panel.

It's crucial to recognize that the inherent complexity of Generative Adversarial Networks (GANs) often makes them overkill for modeling relatively straightforward distributions, such as the Gaussian distribution in this context. GANs excel in capturing intricate patterns and generating high-dimensional, realistic data, making them particularly valuable for complex tasks like image synthesis and style transfer. However, for simpler distributions, the intricate interplay between the generator and discriminator might lead to suboptimal convergence and training instability[11].

In response to the challenges faced by traditional GANs in modeling simple distributions, researchers have explored advanced variants, with the Maximum Mean Discrepancy GAN (MMDGAN[12]) being a noteworthy example. MMDGAN[12] introduces a novel approach by incorporating maximum mean discrepancy as a metric to guide the generator towards minimizing the difference between the generated and real data distributions. Such enhancements aim to bolster the performance of GANs in capturing the nuances of simpler data structures, broadening their applicability across a spectrum of machine learning tasks. As the field evolves, ongoing research into novel GAN architectures and training methodologies continues to push the boundaries of what these generative models can achieve, ensuring their adaptability to a diverse range of distributional complexities.

4.2 Quantum Generative Adversarial Network

4.2.1 QGAN for a Gaussian Distribution

The project presents a quantum implementation of a Generative Adversarial Network (GAN) using Qiskit. It begins by defining a quantum system to represent and visualize a 1D Gaussian distribution[13]5.

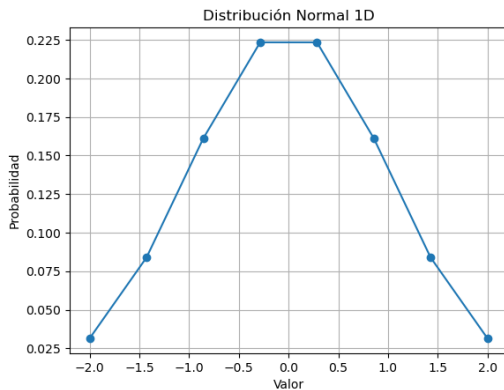


Fig. 5: Gaussian 1D

Then, it constructs a quantum circuit with Hadamard gates and an EfficientSU2 ansatz6.

In this innovative approach, a quantum generator is meticulously crafted and trained in tandem with a discriminator through an adversarial framework, utilizing a quantum sampler to enhance precision. The continual recording of losses

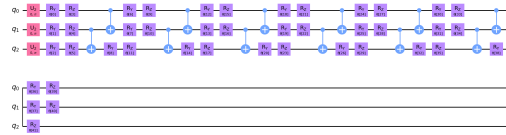


Fig. 6: Ansatz Circuit

for both the generator and discriminator, coupled with visualizing the training progression, vividly underscores the quantum essence of the implementation and its profound influence on the generation of distributions. Ultimately, the unveiling of cumulative distribution functions (CDF) 8 and probability density functions (PDF) 9 not only showcases commendable results but also serves as a conclusive testament to the triumph of the quantum implementation in crafting precisely tailored distributions7.

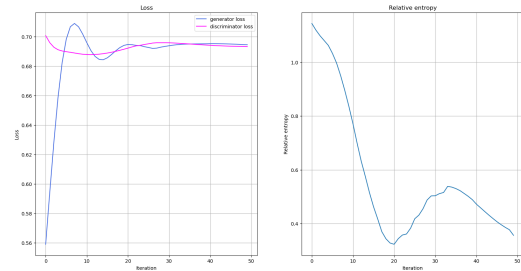


Fig. 7: Loss & Entropy

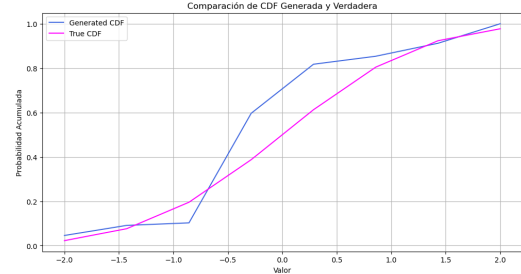


Fig. 8: CDF Real vs GAN

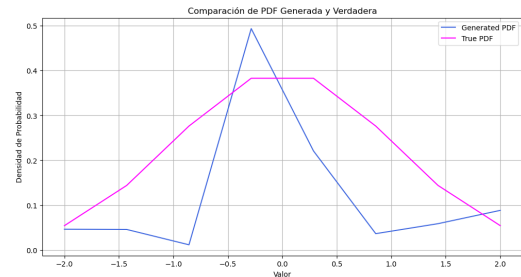


Fig. 9: PDF Real vs GAN

4.2.2 QGAN for a Log-Normal Distribution

The extension of the Quantum Generative Adversarial Network (QGAN) to model a Log-Normal distribution involves adapting the quantum implementation to capture the unique characteristics of this distribution. The Log-Normal distribution is characterized by a probability distribution

of a random variable whose logarithm is normally distributed. To achieve this, the quantum system is configured to represent and visualize a Log-Normal distribution in a manner analogous to the approach taken for the Gaussian distribution[13].

Similar to the Gaussian case, a quantum circuit is constructed using Hadamard gates and an EfficientSU2 ansatz. The quantum generator is then intricately designed and trained alongside a discriminator through the adversarial framework. The quantum sampler is employed to enhance precision, and the training progress is monitored by recording losses for both the generator and discriminator. Visualization of the training progression, coupled with the unveiling of cumulative distribution functions (CDF) and probability density functions (PDF), demonstrates the efficacy of the quantum implementation in generating Log-Normal distributions.

It is noteworthy that the Log-Normal distribution lacks values at or below zero. Consequently, the entropy of this distribution diverges to infinity. This peculiarity, inherent to distributions with non-negative support, underscores the quantum nature's challenge in representing certain types of distributions.

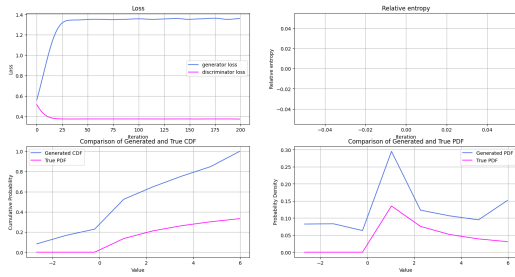


Fig. 10: Log-Normal results

4.2.3 QGAN for a Binomial Distribution

In extending the Quantum Generative Adversarial Network (QGAN) to model a Binomial distribution[13], the project adapts the quantum implementation to accommodate the discrete and bounded nature of this distribution. The Binomial distribution describes the number of successes in a fixed number of independent Bernoulli trials, each with the same probability of success.

The quantum system is configured to represent and visualize a Binomial distribution, incorporating the necessary adjustments to capture the discrete nature of the distribution. The quantum circuit is constructed using Hadamard gates and an appropriate ansatz tailored to the characteristics of the Binomial distribution. The quantum generator and discriminator are trained in a manner similar to previous cases, with a focus on the discrete nature of the distribution.

Monitoring the training progress and visualizing the resulting cumulative distribution functions (CDF) and probability density functions (PDF) provide insights into the quantum implementation's ability to generate Binomial distributions accurately. Unlike continuous distributions, the Binomial distribution has a finite support, and the quantum model's success in capturing this discrete nature is a testament to its versatility.

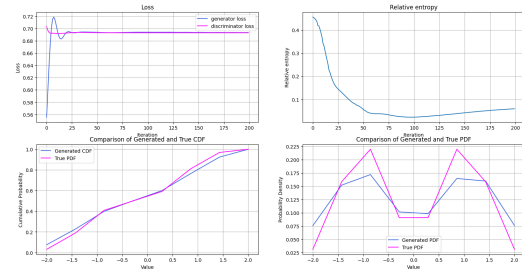


Fig. 11: Binomial results

4.2.4 QGAN for a Triangular Distribution

The Quantum Generative Adversarial Network (QGAN) is further extended to model a Triangular distribution[13], a continuous probability distribution with a triangular shape. The quantum implementation is adapted to accurately represent and visualize the characteristics of the Triangular distribution.

Similar to the previous cases, a quantum circuit is constructed using Hadamard gates and an ansatz designed for the features of the Triangular distribution. The quantum generator and discriminator are trained in an adversarial manner, with the quantum sampler enhancing precision throughout the training process.

It is crucial to note that the Triangular distribution, like the Log-Normal distribution, has values within a finite range. However, distributions with zero values (e.g., at the edges of the support) pose challenges in entropy calculation. In this case, the Triangular distribution does not have a well-defined entropy due to the presence of values with zero probability. This characteristic highlights a limitation in representing distributions with certain features using quantum models and emphasizes the need for careful consideration of distribution-specific properties in quantum generative modeling.

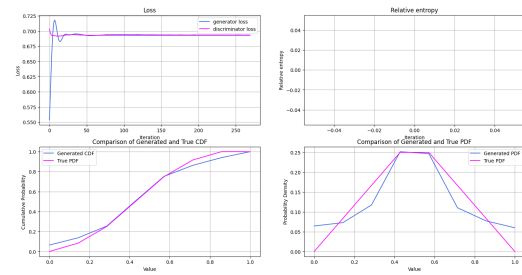


Fig. 12: Triangular results

4.3 Final Tests

Finally, after testing all the models, with different configurations, more epochs, we have achieved the improvement results in the following models:

4.3.1 QGAN Gaussian Distribution

For this case, it has been possible to improve the model, optimising the number of repetitions of the Ansatz, in the tests it has been seen that the optimal anzats is the efficientSU2 for almost all the models, being 3 the optimal number, for a training of 200 epochs we obtain the result of the figure13.

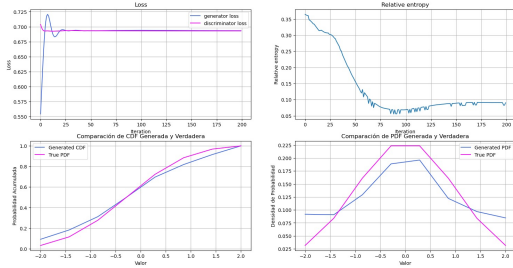


Fig. 13: Gaussian results

As can be seen, there is a difference in the probability distribution between the figure 13 and the figure 13.

4.3.2 QGAN LogNormal Distribution

For this case, we have utilized the same metrics as discussed in 4.3.1, involving 3 repetitions of the efficientSU2 ansatz. However, we have extended the training duration significantly to observe the model's enhancement. These are the outcomes after training the model for nearly 7000 epochs:

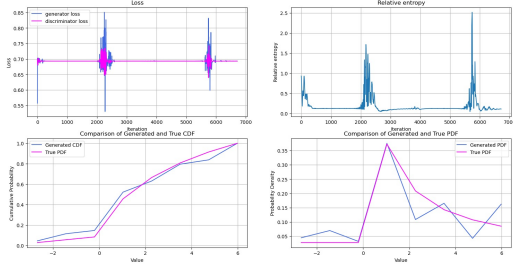


Fig. 14: Q LogNorm results

As can be seen, unlike the 10 figure, the 14 figure has a probability distribution more similar to the real one, even overlapping at several points, whereas the previous one was flawed in this aspect.

4.3.3 QGAN Triangular Distribution

For this case, we have taken the same metrics discussed in 4.3.1, 3 repetitions of the anzats efficientSU2, but we have trained many more epochs, to see the improvement of the model, these are the results of training the model 2000 epochs 15:

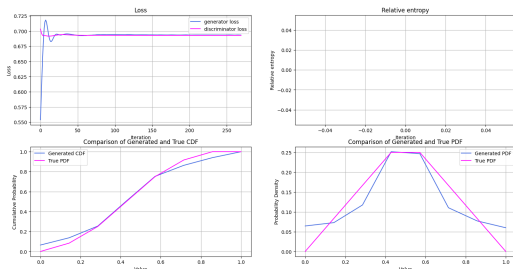


Fig. 15: Q Triang results

As you can see, compared to the figure 12, a change has been made to the distribution to make it look like a Triangular distribution, and in the figure 15 you can see the changes.

4.3.4 QGAN Normal with Zoufal's techniques

Finally, I made a model based on the model presented by Zoufal in his thesis [14] and [13], pre-training the input data, finding the best rotations and then training the model with an anzats, given by him, which is represented in the figure 16.

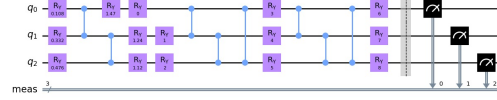


Fig. 16: Zoufal Circuit + Anzats

By pre-training the circuit, 200 epochs, the following result is achieved 17:

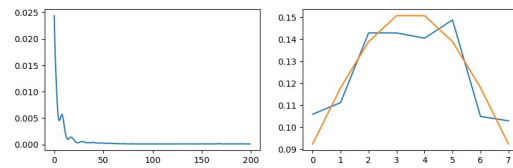


Fig. 17: Zoufal Pretrain Results

Finally, training the model for a Normal Distribution, 2000 epochs, gives the following results 18:

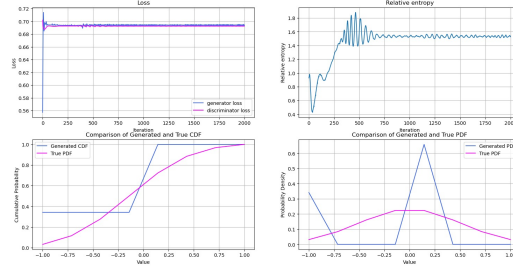


Fig. 18: Zoufal Results

Seeing the figure 18 does not give the expected results, but this is surely due to the bad configuration of the training, or some human error, even so, following this path is very promising given the results exposed by Zoufal, whom I intend to contact, or continue with my tutors to improve these results and achieve those of Zoufal [14].

5 CONCLUSIONS

In this paper, a comparison between classical and quantum GAN models was carried out in order to evaluate their differences and improvements. Both approaches have strengths and weaknesses. For example, the conventional GAN model used in this study does not yield completely accurate results, which is attributed to inherent limitations of GANs themselves. However, after extensive research, it was noted that alternative GAN models, such as MMDGANs, which represent an evolution of classical GANs, are available to improve accuracy. Although this particular problem was not explored further in the context of the comparison between GAN and QGAN, it could be considered as a line of future research.

On the other hand, a notable advantage of classical GANs is that, although they do not generate completely accurate models, they produce results that more closely resemble reality, almost continuously. In contrast, in the tests performed with QGAN, discrete models were obtained. In our case, we discretised at 8 points to minimise the computational cost, since the higher the number of discretisation points, i.e. the more continuous the function, the higher the computational cost. Although this approach has an apparent limitation, it demonstrates an effective and more efficient performance than classical GANs.

6 NEXT STEPS

To conclude, I believe that, as I mentioned earlier, a more effective comparison could be achieved by comparing more advanced models. In the case of classical GANs, it would be interesting to contrast them with MMDGANs. As for QGANs, different ansatz could be explored, such as the one mentioned by Zoufal in his thesis [14], or in his academic paper [13], as well as applying data pre-treatment to improve training. However, as evidenced by the final tests, the implementation of the latter strategy has proved unsuccessful. Nevertheless, I am convinced that this could be a promising avenue for future research. It is my intention to pursue the path proposed by Zoufal and to further explore these ideas once my dissertation is completed. And this is just the beginning, because if these models can achieve good results, who is to say that in the near future we will not be able to generate images with QGAN, just as with classical GAN? Although perhaps these expectations are too ambitious, great ideas often come from ambition.

7 CODE AVIABILITY

The full code used during all these first steps of the project can be found on our Github public repository.

8 ACKNOWLEDGMENT

Thank you very much to everyone involved in this project for giving me support and motivating me to continue, even knowing the hard work I've had behind the scenes. Thanks to my family for their support, especially to my sister Alicia, who always motivates me. Also, to my mentors Fernando and Matías. Finally and most importantly, to friends like Yeray who motivate me to reach the next level, and to Laia who inspires me and helps me to do my best, even though we haven't spoken for a long time.

Muchas gracias a todos los involucrados en este proyecto por brindarme apoyo y motivarme a seguir adelante, incluso sabiendo el arduo trabajo que he tenido detrás de escena. Gracias a mi familia por su apoyo, especialmente a mi hermana Alicia, quien siempre me motiva. También a mis tutores Fernando y Matías. Finalmente y mas importante, a amigos como Yeray que me motivan a alcanzar el siguiente nivel, y a Laia que me inspira y ayuda a dar lo mejor de mí, aunque hace tiempo no hablamos.

Repositorio de GitHub: <https://github.com/adriend1102/QML>

REFERÈNCIES

- [1] Y. Chen, N. Akhtar, N. Al, H. Haldar, and A. Mian, "On quantifying and improving realism of images generated with diffusion."
- [2] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, Eds., vol. 27. Curran Associates, Inc., 2014. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf
- [3] M. Zaheer, C.-L. Li, B. Póczos, and R. Salakhutdinov, "Gan connoisseur: Can gans learn simple 1d parametric distributions?"
- [4] R. P. D. Silva, "Quantum circuits with qiskit," 2023. [Online]. Available: <https://ssrn.com/abstract=4359664>
- [5] L. Lamata, "Quantum reinforcement learning with quantum photonics," pp. 1–8, 2 2021.
- [6] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial networks," 2014.
- [7] M. Parigi, S. Martina, and F. Caruso, "Quantum-noise-driven generative diffusion models," 8 2023. [Online]. Available: <http://arxiv.org/abs/2308.12013>
- [8] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, Eds., vol. 33. Curran Associates, Inc., 2020, pp. 6840–6851. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2020/file/4c5bcfec8584af0d967f1ab10179ca4b-Paper.pdf
- [9] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *Proceedings of the 32nd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, F. Bach and D. Blei, Eds., vol. 37. Lille, France: PMLR, 07–09 Jul 2015, pp. 2256–2265. [Online]. Available: <https://proceedings.mlr.press/v37/sohl-dickstein15.html>
- [10] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, "Quantum embeddings for machine learning," 2020.
- [11] M. Zaheer, C.-L. Li, B. Póczos, and R. Salakhutdinov, "Gan connoisseur: Can gans learn simple 1d parametric distributions?"
- [12] C.-L. Li, W.-C. Chang, Y. Cheng, Y. Yang, and B. Póczos, "Mmd gan: Towards deeper understanding of moment matching network," 2017.

- [13] C. Zoufal, A. Lucchi, and S. Woerner, “Quantum generative adversarial networks for learning and loading random distributions,” 3 2019. [Online]. Available: <http://arxiv.org/abs/1904.00043><http://dx.doi.org/10.1038/s41534-019-0223-2>
- [14] C. Zoufal, “Generative quantum machine learning,” 2021.