Exploring biological neuronal correlations with quantum generative models

Hernandes, Vinicius; Greplova, Eliska

DOI
10.1016/j.xcrp.2025.102682

Publication date
2025

Document Version
Final published version

Published in
Cell Reports Physical Science

Cell Reports
Physical Science

Article

Exploring biological neuronal correlations with quantum generative models

Graphical abstract

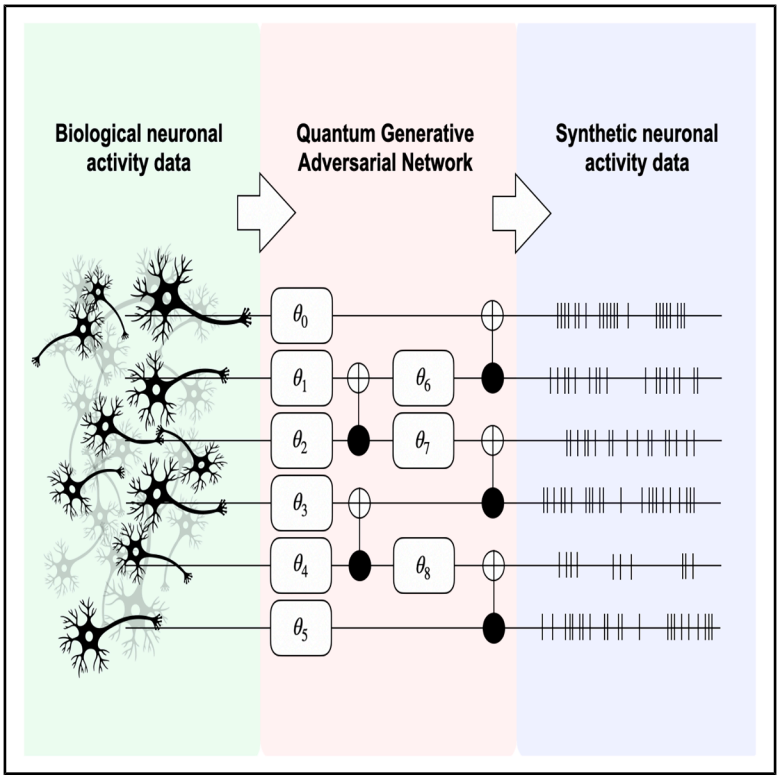

Highlights

- A quantum generative adversarial network (SpiQGAN) generates neuronal activity
- Efficiently models neurons using a reduced number of parameters
- Reproduces spatiotemporal correlations and bursting patterns observed in real data

Authors

Vinicius Hernandes, Eliska Greplova

Correspondence

v.hernandes@tudelft.nl

In brief

Hernandes and Greplova introduce SpiQGAN, an efficient quantum generative adversarial network, to model biological neuronal activity. This quantum approach generates realistic spike trains that capture key spatiotemporal features with significantly fewer parameters than classical deep learning methods, paving the way for more compact and interpretable neuroscience models using quantum computing.

Hernandes & Greplova, 2025, Cell Reports Physical Science 6, 102682
August 20, 2025
https://doi.org/10.1016/j.xcrp.2025.102682

CellPress

Cell Reports
Physical Science

CellPress
OPEN ACCESS

Article

Exploring biological neuronal correlations with quantum generative models

Vinicius Hernandes[1,2,*] and Eliska Greplova[1]
[1]QuTech and Kavli Institute of Nanoscience, Delft University of Technology, Delft, the Netherlands
[2]Lead contact
*Correspondence: v.hernandes@tudelft.nl
https://doi.org/10.1016/j.xcrp.2025.102682

SUMMARY

Understanding how biological neural networks process information is one of the biggest open scientific questions of our time. Advances in machine learning and artificial neural networks have enabled the modeling of neuronal behavior, but classical models often require a large number of parameters and highly task-specific architectures, which can complicate model design and scalability. Quantum computing offers an alternative approach through quantum machine learning, which can achieve efficient training with fewer parameters. In this work, we introduce a quantum generative model framework for generating synthetic data that captures the spatial and temporal correlations of biological neuronal activity. Our model demonstrates the ability to achieve reliable outcomes with fewer trainable parameters compared to classical methods. These findings highlight the potential of quantum generative models to provide new tools for modeling and understanding neuronal behavior, offering a promising avenue for future research in neuroscience.

INTRODUCTION

Exploring information processing within biological neuronal networks remains a core challenge in contemporary science, with direct implications across disciplines like neuroscience, medicine, and deep learning.[1–4] One way to approach this problem is to use computational models that can reproduce the neuronal activity data produced in real systems. Accurate synthetic data can be extremely useful for studying properties such as network connectivity and the response to stimuli under controlled conditions.[5,6]

Several models for neuronal activity have been developed, and many achieve outstanding results in replicating neuronal network correlations. One class of methods that use statistical mechanics tools to model neuronal activity is maximum entropy models, which reliably capture some network correlations by only fitting pairwise interactions.[7,8] Even though numerous adaptations of this technique have been implemented to achieve higher accuracy or to include temporal correlations,[9–14] this approach shows several limitations when addressing larger networks, especially due to the assumption that pairwise correlations are sufficient to encapsulate most of the statistical features of these complex systems.[15,16]

Another effective approach for modeling neuronal activity is to use machine learning (ML) models to produce data that fit the biological network statistics and further investigate the properties of the real system. The ML models do not rely on prior information about the biological system but instead learn to reproduce correlations solely from data. A supervised strategy using convolutional neural networks first showed that a deep learning approach can be successful at generating neural responses from stimuli.[17] However, the model's benefits are hampered by limited accuracy and its dependency on labeled data. With the increasing popularization of generative models, models with superior predictive performance and generalization power were implemented. Models like variational auto-encoders,[18] recurrent neural networks,[19] generative adversarial networks (GANs),[20] and transformers[21] have been used to produce spike trains (binary sequences representing neuronal activity) with high accuracy and good correspondence of spatial and temporal correlations when compared to real data. While each iteration of these models improves in quality, all share the same disadvantage regarding their interpretability. In order to fit the statistics of larger systems, these models need to use a number of trainable parameters that scale unfavorably with the number of simulated neurons. Apart from demanding more computational power, an excessive number of parameters makes the models difficult to analyze or be used as a tool to investigate concrete properties of biological networks.

As the field of quantum computing rapidly advances, quantum ML (QML) models are rising as an alternative to classical methods, with the possibility of achieving similar results while keeping the parametrized model more compact in terms of trainable parameters.[22–25] Specifically, the field of quantum generative learning has received much attention recently: quantum models have shown better generalization and expressivity for specific tasks when compared to their classical counterparts.[26–28] Since the conception of QML, one class of quantum generative models has been extensively studied: quantum GANs (QGANs).[29] The adversarial approach has proven successful and is being continually improved, producing higher-dimensional data with more stable training routines.[30–34]

CellPress
OPEN ACCESS

Cell Reports
Physical Science
Article

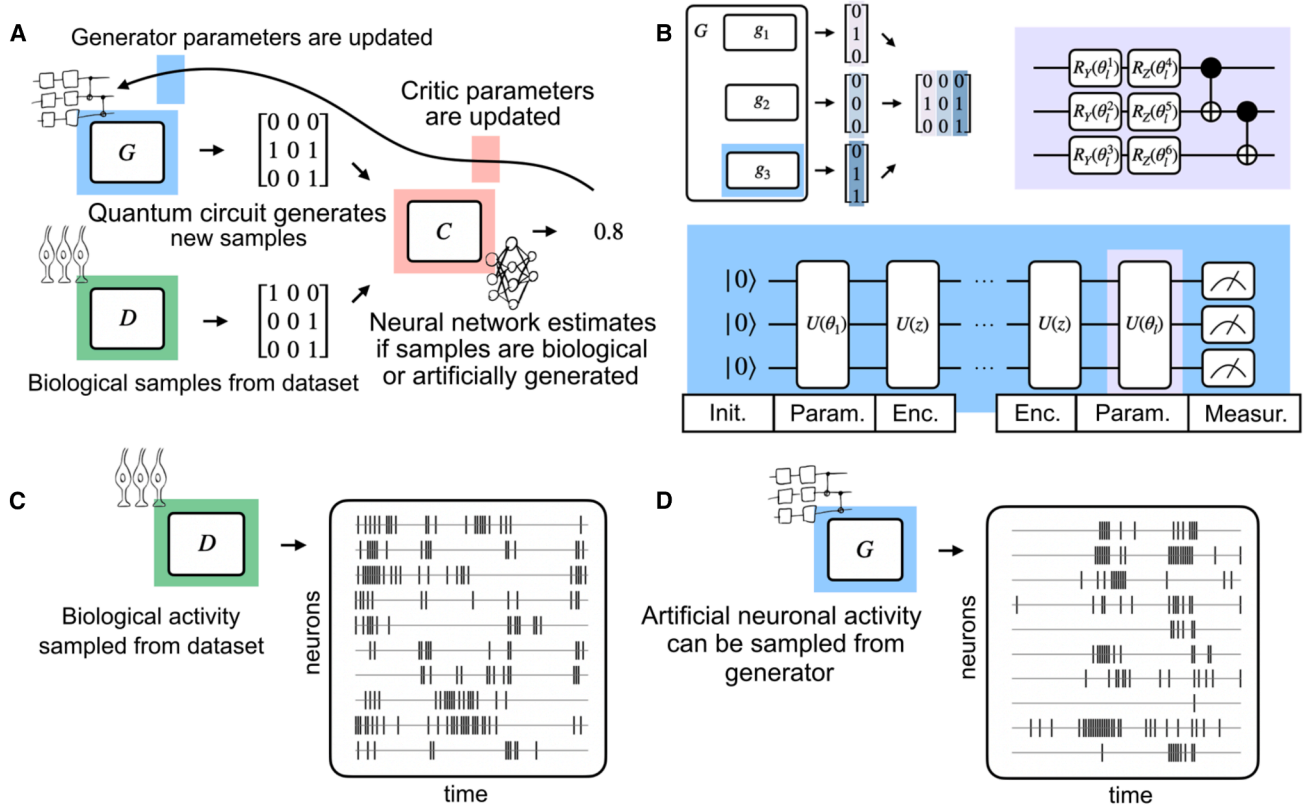

Figure 1. Illustration of the model architecture

(A) Architecture of the model, with generator *G* producing generated samples and dataset *D* producing biological samples, which are both used as input for critic *C*.

(B) Architecture of generator. In the top left corner, the generator composed of several sub-generators is shown. The bottom part shows that each sub-generator is a quantum circuit following a re-uploading scheme. Here, a noise-encoding layer and a parametrized layer are repeated for l layer, with the parametrized layer ansatz of each parametrized layer shown in the top right.

(C and D) After training, the generator can be used to produce samples (D) similar to samples obtained from the biological dataset (C).

This work is inspired by the observation that quantum generative models have shown promise in the replication of discrete distributions.[33,35] Additionally, the salamander retina dataset has been used as a benchmark for distribution learning using quantum Boltzmann machines.[36,37] These observations suggest the possibility for a full reconstruction of both spatial and temporal correlations with the quantum generative model that we present here. We build on our preliminary work[38] and design SpiQGAN, an efficient quantum framework that enables the production of synthetic neuronal data for biological neuronal networks. SpiQGAN generates spike trains of neuronal activity: data that consist of binary activation states of the neurons obtained from recording the response of ganglion cells of the salamander retina to a visual stimulus as a function of time.[39] This dataset represents one of the standard benchmarks in neuronal activation modeling.

To achieve the generation of data that maximally resemble the real biological sample, we apply a hybrid QGAN with a quantum generator that produces synthetic activity data and a classical critic that aims to distinguish real data from the dataset[39] from those produced by the quantum generator. The model is trained adversarially, and the outcome is a generator that can reproduce neuronal activity that is, to a high degree, similar to the salamander retina dataset (Figure 1). Compared to classical neural network alternatives, the quantum generator has the advantage of achieving reliable outcomes with a significantly reduced number of trainable parameters, which scale more favorably for increasing systems' sizes: the number of parameters is linear in the number of neurons. In other words, SpiQGAN is able to reproduce the behavior of this complex neuronal dataset in both space and time with significantly fewer trainable parameters than classical ML models, thus forming a stepping stone toward using quantum approaches for more compact and more interpretable models for neuronal behavior.

RESULTS

Distributional similarity between generated and real data

We trained SpiQGAN for $t = \{1, 5, 10, 20, 30\}$ timesteps and $n = \{2, 4, 6, 8, 10\}$ neurons in order to evaluate the quality of the generated data as a function of system size and time trace length.

Please cite this article in press as: Hernandes and Greplova, Exploring biological neuronal correlations with quantum generative models, Cell Reports Physical Science (2025), https://doi.org/10.1016/j.xcrp.2025.102682

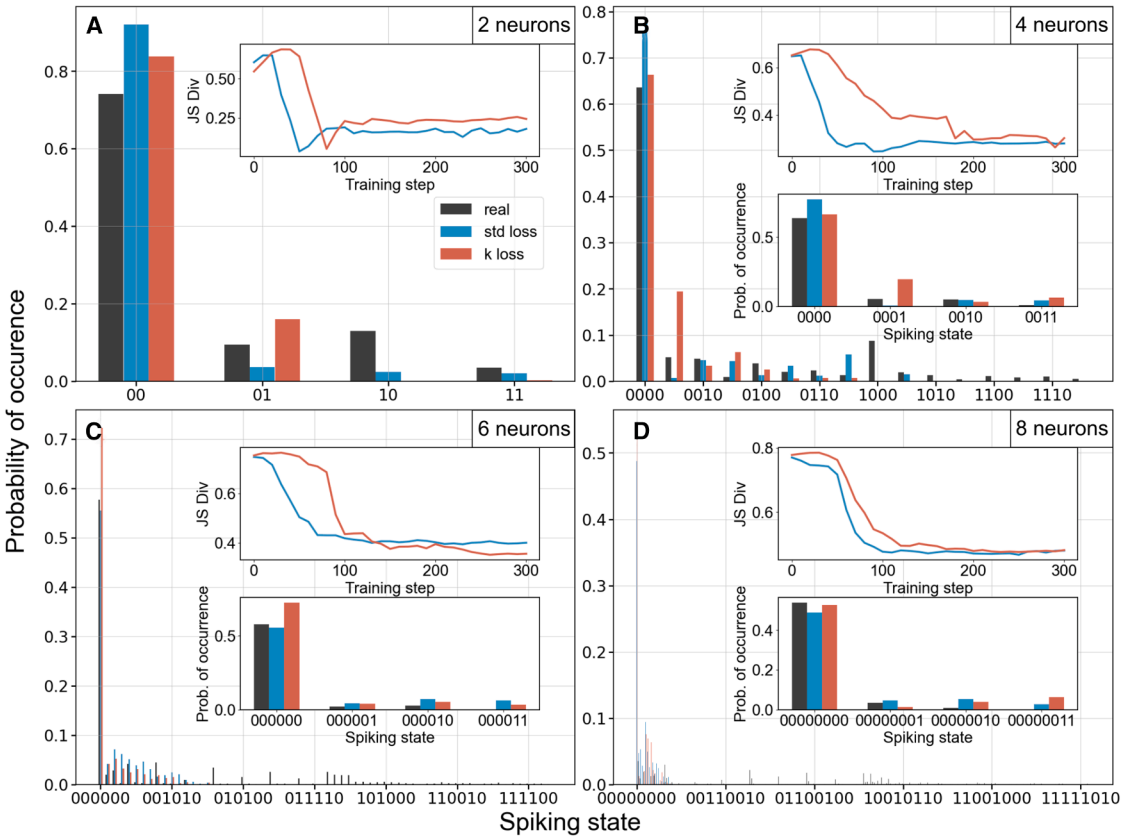

Figure 2. Comparison between distribution of states and JS divergence calculated using generated and real data

Each image show the distribution of spiking states for generated data obtained after training with the *K*-loss (in red) and the standard loss (in blue) and the real distribution of the spiking states (black) for (A) 2, (B) 4, (C) 6, and (D) 8 neurons, all for the case of 1 timestep. The bottom inset shows a zoom of the first four activation states. The top inset shows the JS divergence for all training steps for *K* (red) and standard (blue) loss.

Comparing the distribution of possible states of the simulated data to that of the salamander retina data is a straightforward way to evaluate the quality of our generative model. Reconstruction of these distributions also allows us to calculate distribution distances such as the Jensen-Shannon (JS) divergence. However, direct distribution comparison is particularly challenging: for n neurons and t timesteps, the number of possible (spiking) states is 2^{nt}. We are thus able to directly visually compare distributions only for a small number of neurons.

For two neurons, $n = 2$, and one timestep, $t = 1$, the possible states are $\{00, 01, 10, 11\}$. For two neurons and two timesteps, $t = 2$, the possible states are $\{0000, 0001, \ldots, 1110, 1111\}$. In this representation, the first two bits represent neurons 1 and 2 at timestep 1 and the last two the states of these neurons at timestep 2. This means that the distribution of states quickly becomes intractable for an increasing number of neurons or timesteps. Nonetheless, it is very informative to compare distributions directly for low numbers of neurons.

For all cases, regardless of system size, we calculated a series of statistical values useful for evaluating the behavior of the generated and the real data from the salamander retina dataset. Specifically, we calculate the pairwise covariance between the activation state of a pair of neurons; the mean firing rate, which corresponds to how many times a neuron spikes per second; the k-probability, equal to the probability of k neurons being active at the same time; and the autocorrelogram to estimate the correlation between a trace of spikes and itself for delayed timesteps.

First, we consider the case with a unique timestep (one subgenerator quantum circuit) and neurons varying from 2 to 8. In these cases, the distribution is easily numerically tractable. We show the final distribution of generated spiking states compared to the distribution of the real data in Figure 2, with a zoom on the most prominent terms of the distribution in the bottom inset. The probabilities of the spiking states are calculated using the last iteration of the trained circuit. This is coincident with the last value of the JS divergence, which steadily decreases during training, visible in the bottom insets of Figure 2. These results show that for a sufficient number of training steps, the distribution of generated states converges to the distribution of the salamander retina dataset. This distribution convergence is the first indication that the training is working as intended, and the samples produced by the generator match some of the statistics of the real data. Throughout, we compared both standard loss and biologically inspired *K*-loss, which are defined in the methods section. We found that, on average, *K*-loss performed

CellPress
OPEN ACCESS

Cell Reports
Physical Science
Article

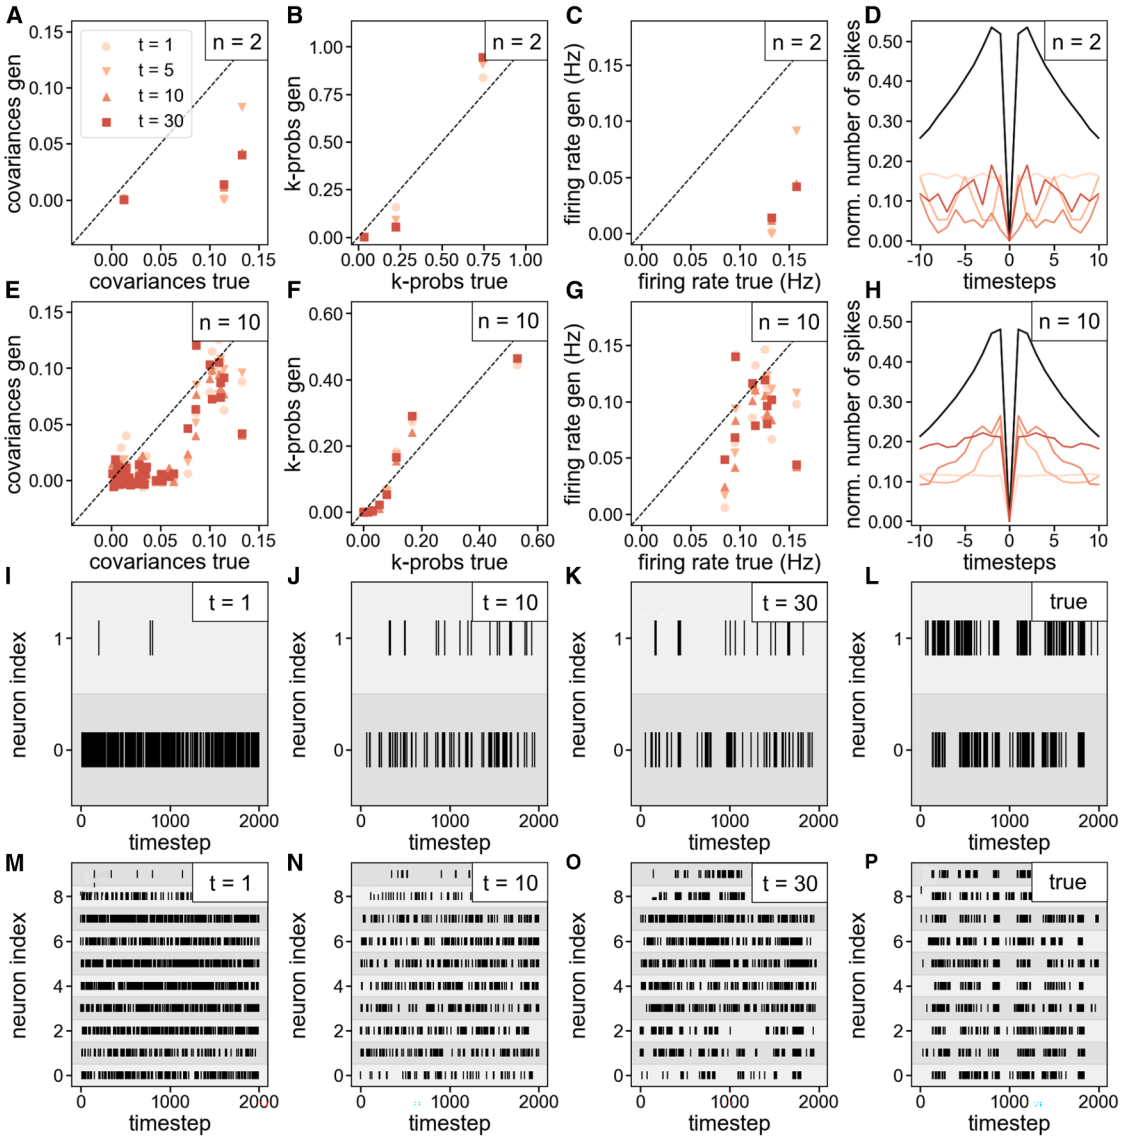

Figure 3. Statistics and generated data for 2 and 10 neurons

(A–H) Statistics for the case of 2 and 10 neurons, with 1, 5, 10, and 30 timesteps represented with different colors in each image. Specifically, (A and E) pairwise covariance, (B and F) k-probability, (C and G) firing rate, and (D and H) autocorrelogram are shown.

(I–P) Spike traces for 2 and 10 neurons for the case of generated data with 1 (I and M), 10 (J and N), and 30 (K and O) timesteps and for real data (L and P).

slightly better (see Note S1 and Figure S1 for a detailed comparison).

Statistical analysis of generated activity

In Figure 3, we show further statistics used to assess the quality of the generated data for 2 and 10 neurons and varying the number of timesteps, focusing only on the biologically informed *K*-loss from now on. Complete results for all neuron numbers and timesteps, accompanied by a focused comparison between the statistics obtained with two different loss functions, are shown in Note S1 and the figures of the supplemental information (see Figures S1 and S4–S8 for statistical metrics and S9–S13 for activity traces). In Figure 3, we see that the k-probability (Figures 3B and 3F) and the mean firing rate (Figures 3C and 3G) are well fitted by SpiQGAN, while the pairwise covariance (Figures 3A and 3E) and the autocorrelogram (Figures 3D and 3H) show more discrepancy. In Figure 3H, we see that the number of generated spikes as a function of simulated timesteps gets closer to the real distribution as the number of timesteps increases.

Spike train comparison with biological data

A visual comparison of spike trains generated by SpiQGAN and those from the biological dataset is shown in Figure 3 for 2 (Figures 3I–3L) and 10 (Figures 3M–3P) neurons. A visual comparison between the generated spikes and the salamander retina

Please cite this article in press as: Hernandes and Greplova, Exploring biological neuronal correlations with quantum generative models, Cell Reports Physical Science (2025), https://doi.org/10.1016/j.xcrp.2025.102682

samples shows that for an increasing number of timesteps, the QGAN-generated samples start forming bursting clusters, an important feature of the biological dataset.

DISCUSSION

Overall, all SpiQGAN iterations that we implemented achieved a reasonable fit of the key statistical features of the data, especially given the model's simplicity and general purpose design, while maintaining a low number of parameters, which scale favorably (linearly) in the number of neurons. Specifically, for our model with 4 parameterized layers, the total number of trainable parameters is equal to 8 times the number of neurons per timestep. This scaling has the following implication: data generation for a neuronal network with dozens of neurons in our implementation uses hundreds of trainable parameters, compared to thousands or tens of thousands in the case of the traditional ML approaches.[20,21] A detailed comparison benchmarking our quantum generator against classical fully connected generators, which consistently underperformed even with more parameters, can be found in Note S2 and Figures S2 and S3. Moreover, it is clear that models that simulate more neurons presented improved performance, which insinuates that using larger circuits could return even better results. While some metrics, like the mean firing rate and pairwise covariance, show discrepancies, others, such as the k-probability and the similarity between generate and real spike train patterns, including burst-like patterns, indicate that the model captures essential features of the neuronal data.

We have shown that QGANs are able to generate synthetic neuronal activity data that faithfully reproduce both spatial and temporal correlations of the biological dataset. We designed and implemented a resource-efficient SpiQGAN that re-uses the same building block across the model. Additionally, we included a biologically informed loss function to take into account the statistical properties of the generated samples.

This work lays the foundation for the utilization of quantum learning models beyond quantum science, here in neuroscience modeling. In particular, SpiQGAN opens the possibility of running resource-efficient algorithms on quantum computers to beneficially model neuronal activity. With the compact quantum models, the dynamics and interpretation of neuronal activity can be efficiently explored in future work.

METHODS

GANs

GANs, which were introduced by Goodfellow et al.,[40] are a powerful class of generative models that learn to synthesize data samples by framing the learning process as an adversarial game between two neural networks: a generator and a discriminator. The generator network, G, aims to produce data samples that mimic those drawn from the true data distribution P_x. It takes a noise vector z, sampled from a predefined distribution P_z (e.g., a Gaussian or uniform distribution), and transforms it into a synthetic data sample, $G(z)$. Meanwhile, the discriminator network, D, acts as a binary classifier, distinguishing between real samples from the true data distribution and fake samples generated by G.

The training objective is formulated as a minimax game where the generator tries to minimize the probability of the discriminator correctly identifying generated samples while the discriminator simultaneously maximizes its ability to correctly classify the samples:

$$\min_G \max_D E_{x\sim P_x}[\log D(x)] + E_{z\sim P_z}[\log(1 - D(G(z)))]. \quad \text{(Equation 1)}$$

Although GANs have achieved remarkable success in various applications (e.g., image synthesis and text generation), they are often plagued by training instabilities, such as vanishing gradients and mode collapse.[41] These issues arise primarily because the loss function may not provide meaningful gradients when the discriminator is too strong or too weak, leading to poor convergence.

The Wasserstein GAN (WGAN)[41] addresses many of the training challenges associated with standard GANs by leveraging the Wasserstein distance (also known as the Earth-mover distance) to measure the divergence between the true data distribution and the generated data distribution. Unlike the original GAN, the discriminator in WGAN, referred to as a critic, outputs a scalar value instead of a binary classification, quantifying how well the generated samples approximate the real data distribution. The WGAN objective is formulated as

$$\min_G \max_{D\in\mathcal{D}} E_{x\sim P_x}[D(x)] - E_{z\sim P_z}[D(G(z))], \quad \text{(Equation 2)}$$

where $\mathcal{D}$ is the set of all 1-Lipschitz functions, enforced through weight clipping or gradient penalties.[42] By stabilizing the gradients, WGAN significantly improves convergence behavior, allowing the generator to learn a more accurate representation of the target distribution.

Parametrized quantum circuits and QGANs

As quantum computing has advanced, QML has emerged as a promising frontier. One key concept is parametrized quantum circuits (PQCs). PQCs consist of a sequence of quantum gates with parameters that are classically optimized. PQCs can encode complex quantum states and can be used to approximate complex distributions.

Building on this foundation, QGANs extend the GAN framework into the quantum domain by incorporating quantum components such as quantum generators, quantum discriminators, or both. In QGANs, the generator may be implemented as a PQC, which is trained to generate samples that match the desired distribution.

Implementation of the quantum generator and the classical critic

SpiQGAN uses a quantum generator to model the spike activity patterns of retinal ganglion cells. Specifically, we employ a patch WQGAN approach, where the quantum generator is divided into several sub-generators, each corresponding to a different timestep. Each sub-generator shares the same PQC architecture but has independent trainable parameters, allowing for flexibility in capturing the temporal dynamics of neuronal activity.

The generator begins with a random initial quantum state $|z\rangle$, which is mapped to the final state $|g\rangle$ using a data re-uploading scheme.[43] The quantum circuit consists of five layers, where each layer applies a sequence of parametrized unitaries $U(\theta_i)$ and noise-encoding unitaries $U(z)$. The parametrized unitary $U(\theta_i)$ is implemented using rotation gates around the y and z axes (R_Y and R_Z) and entangling operations (CNOT gates) between adjacent qubits, while the encoding block applies R_X rotations to each qubit to encode a sampled noise vector. The generator outputs a sequence of activity states for multiple neurons over several timesteps by concatenating the outputs from all sub-generators. All quantum circuit results in this work are obtained from classical numerical simulations of ideal quantum hardware. This allows us to benchmark the behavior of the model without the influence of noise. This comes with the cost of longer training times, given the inefficiency of calculating the quantum gradient. The training times range from approximately 2 h for the smallest system to 5 days for the larger systems. We also tested generators with a higher number of variational layers per sub-generator to evaluate whether deeper circuits would improve performance. However, the improvements were not significant, and training became significantly more expensive due to the cost of quantum gradient estimation. For this reason, we fixed the number of layers to five in all reported experiments, balancing expressivity and trainability.

The critic in our QGAN framework is a fully connected classical neural network. The network consists of the following:

(1) An input layer matching the size of the generated samples.
(2) A hidden layer with 64 neurons using ReLU activation.
(3) An output layer without an activation function, which directly provides a scalar value representing the divergence between the real and generated distributions.

Training procedure

SpiQGAN is trained by optimizing two separate loss functions for the generator and the critic. The critic's loss function aims to maximize the difference between its outputs for real samples x and generated samples $G(z)$:

$$L_C = \frac{1}{2B}\sum_j \left(C(G(z_j)) - C(x_j)\right), \qquad \text{(Equation 3)}$$

whereas the generator's objective is to minimize the critic's evaluation of the generated samples:

$$L_G = -\frac{1}{B}\sum_j C(G(\boldsymbol{z_j})) - K\left(\sum_i G(z_j)^i - x_j^i\right), \qquad \text{(Equation 4)}$$

with B being the batch size and the term $K\left(\sum_i G(z_j)^i - x_j^i\right)$ added to the standard Wasserstein's generator loss function, inspired by maximum entropy models,[11] corresponding to the difference between the number of spikes in a fake sample and those in a real sample. This loss was named K-loss, and by setting $K = 0$, the standard loss is retrieved. Training alternates between two updates of the critic and one update of the generator, ensuring stable convergence. The Adam optimizer is employed with learning rates of 0.05 for the generator and 0.002 for the critic.

Dataset and evaluation metrics

The dataset comprises neuronal spike activity recorded from retinal ganglion cells in a salamander retina.[39] It contains 297 repetitions of a 19 s natural movie, recorded as binary spike events, where 1 indicates a spike and 0 indicates no spike. The goal is to generate synthetic data that replicate these binary spike patterns while maintaining important statistical properties. To evaluate the performance of the QGAN, we used the following statistical metrics.

(1) Pairwise covariance: measures the extent to which two neurons fire together. High covariance suggests that the neurons are more likely to spike simultaneously.
(2) Mean firing rate: the average rate at which a neuron fires spikes over time. This metric helps ensure that the generated data match the overall activity level of the real data.
(3) k-probability: the probability distribution over the number of spikes (k) in a given time window. Matching this distribution ensures that the generated data capture the variability in spike counts.
(4) Autocorrelogram: a measure of the temporal structure of the spike train, representing the correlation of a neuron's spike times with itself over different time lags. This metric is crucial for capturing the temporal dynamics of neuronal activity.

To comprehensively assess the model's performance, we conducted experiments varying the number of neurons $n = \{2, 4, 6, 8, 10\}$ and timesteps $t = \{1, 2, 5, 10, 20, 30\}$. For small-scale systems, alongside the metrics listed above, we computed the exact probabilities of all possible spiking states, which were used to compare the generated and real data distributions using distance measures such as JS divergence, which, for the two probability distributions P and Q, is defined as

$$\mathrm{JS}(P\|Q) = \frac{1}{2}\mathrm{KL}(P\|M) + \frac{1}{2}\mathrm{KL}(Q\|M), \qquad \text{(Equation 5)}$$

where $M = \frac{1}{2}(P+Q)$, and KL denotes the Kullback-Leibler divergence:

$$\mathrm{KL}(P\|Q) = \sum_i P(i)\log\frac{P(i)}{Q(i)}. \qquad \text{(Equation 6)}$$

RESOURCE AVAILABILITY

Lead contact

Requests for further information and resources should be directed to and will be fulfilled by the lead contact, Vinicius Hernandes (v.hernandes@tudelft.nl).

Materials availability

This study did not generate new unique reagents.

Data and code availability

- The data used in this work are deposited at Zenodo[44]: https://doi.org/10.5281/zenodo.13388723.
- Accompanying code is available at GitLab[45]: https://gitlab.com/QMAI/papers/spiqgan.
- Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

ACKNOWLEDGMENTS

We acknowledge useful discussions with Amira Abbas, Antón Rodriguez-Otero, Thomas Spriggs, Dimphna Meijer, and Geeske van Woerden. We acknowledge the Kavli Institute of Nanoscience Delft Synergy Grant. This work is part of the project Engineered Topological Quantum Networks (project no. VI.Veni.212.278) of the research program NWO Talent Programme Veni Science Domain 2021, which is financed by the Dutch Research Council (NWO).

AUTHOR CONTRIBUTIONS

E.G. designed the project with input from V.H. V.H. wrote the code, ran the simulations, analyzed the data, and created the images with input from E.G. E.G. and V.H. co-wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at https://doi.org/10.1016/j.xcrp.2025.102682.

Received: October 14, 2024
Revised: April 18, 2025
Accepted: June 10, 2025

REFERENCES

1. Rieke, F., Warland, D., Ruyter van Steveninck, R. de, and Bialek, W. (1999). Spikes: Exploring the Neural Code (MIT Press).
2. White, A., Williams, P.A., Hellier, J.L., Clark, S., Dudek, F.E., and Staley, K. J. (2010). EEG spike activity precedes epilepsy after kainate-induced status epilepticus. Epilepsia *51*, 371–383.
3. Hassabis, D., Kumaran, D., Summerfield, C., and Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. Neuron *95*, 245–258.
4. Saxe, A., Nelli, S., and Summerfield, C. (2021). If deep learning is the answer, what is the question? Nat. Rev. Neurosci. *22*, 55–67.
5. Wen, J., Peitz, M., and Brüstle, O. (2022). A defined human-specific platform for modeling neuronal network stimulation in vitro and in silico. J. Neurosci. Methods *373*, 109562.
6. Senk, J., Kriener, B., Djurfeldt, M., Voges, N., Jiang, H.-J., Schüttler, L., Gramelsberger, G., Diesmann, M., Plesser, H.E., and van Albada, S.J. (2022). Connectivity concepts in neuronal network modeling. PLoS Comput. Biol. *18*, e1010086.
7. Tkacik, G., Schneidman, E., Berry II, M.J., and Bialek, W. (2006). Ising models for networks of real neurons. Preprint at arXiv. https://doi.org/10.48550/arXiv.q-bio/0611072.
8. Schneidman, E., Berry, M.J., Segev, R., and Bialek, W. (2006). Weak pairwise correlations imply strongly correlated network states in a neural population. Nature *440*, 1007–1012. https://doi.org/10.1038/nature04701.
9. Tang, A., Jackson, D., Hobbs, J., Chen, W., Smith, J.L., Patel, H., Prieto, A., Petrusca, D., Grivich, M.I., Sher, A., et al. (2008). A Maximum Entropy Model Applied to Spatial and Temporal Correlations from Cortical Networks In Vitro. J. Neurosci. *28*, 505–518. https://doi.org/10.1523/JNEUROSCI.3359-07.2008.
10. Marre, O., El Boustani, S., Frégnac, Y., and Destexhe, A. (2009). Prediction of spatiotemporal patterns of neural activity from pairwise correlations. Phys. Rev. Lett. *102*, 138101.
11. Tkačik, G., Marre, O., Mora, T., Amodei, D., Berry II, M.J., and Bialek, W. (2013). The simplest maximum entropy model for collective behavior in a neural network. J. Stat. Mech. *2013*, P03011.
12. Granot-Atedgi, E., Tkačik, G., Segev, R., and Schneidman, E. (2013). Stimulus-dependent maximum entropy models of neural population codes. PLoS Comput. Biol. *9*, e1002922.
13. Tkačik, G., Marre, O., Amodei, D., Schneidman, E., Bialek, W., and Berry, M.J. (2014). Searching for collective behavior in a large network of sensory neurons. PLoS Comput. Biol. *10*, e1003408.
14. Delamare, G., and Ferrari, U. (2022). Time-dependent maximum entropy model for populations of retinal ganglion cells. Physical Sciences Forum *5*, 31. MDPI.
15. Roudi, Y., Nirenberg, S., and Latham, P.E. (2009). Pairwise maximum entropy models for studying large biological systems: when they can work and when they can't. PLoS Comput. Biol. *5*, e1000380.
16. Köster, U., Sohl-Dickstein, J., Gray, C.M., and Olshausen, B.A. (2014). Modeling higher-order correlations within cortical microcolumns. PLoS Comput. Biol. *10*, e1003684.
17. McIntosh, L.T., Maheswaranathan, N., Nayebi, A., Ganguli, S., and Baccus, S.A. (2016). Deep Learning Models of the Retinal Response to Natural Scenes. In Proceedings of the 30th International Conference on Neural Information Processing Systems. NIPS'16 (Curran Associates Inc.), pp. 1369–1377.
18. Pandarinath, C., O'Shea, D.J., Collins, J., Jozefowicz, R., Stavisky, S.D., Kao, J.C., Trautmann, E.M., Kaufman, M.T., Ryu, S.I., Hochberg, L.R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. Nat. Methods *15*, 805–815.
19. Bellec, G., Wang, S., Modirshanechi, A., Brea, J., and Gerstner, W. (2021). Fitting summary statistics of neural data with a differentiable spiking network simulator. In Advances in Neural Information Processing Systems, A. Beygelzimer, Y. Dauphin, P. Liang, and J.W. Vaughan, eds. (Curran Associates Inc).
20. Molano-Mazon, M., Onken, A., Piasini, E., and Panzeri, S. (2018). Synthesizing realistic neural population activity patterns using Generative Adversarial Networks. In International Conference on Learning Representations.
21. Le, T., and Shlizerman, E. (2022). STNDT: Modeling Neural Population Activity with Spatiotemporal Transformers. In Advances in Neural Information Processing Systems, A.H. Oh, A. Agarwal, D. Belgrave, and K. Cho, eds. (Curran Associates Inc.).
22. Schuld, M., and Petruccione, F. (2018). Learning with Quantum Models. In Supervised Learning with Quantum Computers (Springer International Publishing), pp. 247–272. https://doi.org/10.1007/978-3-319-96424-9_8.
23. Schuld, M., Bocharov, A., Svore, K.M., and Wiebe, N. (2020). Circuit-centric quantum classifiers. Phys. Rev. *101*, 032308.
24. Dunjko, V., and Wittek, P. (2020). A non-review of quantum machine learning: trends and explorations. Quantum Views *4*, 32.
25. Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., and Woerner, S. (2021). The power of quantum neural networks. Nat. Comput. Sci. *1*, 403–409.
26. Benedetti, M., Garcia-Pintos, D., Perdomo, O., Leyton-Ortega, V., Nam, Y., and Perdomo-Ortiz, A. (2019). A generative modeling approach for benchmarking and training shallow quantum circuits. npj Quantum Inf. *5*, 45. https://doi.org/10.1038/s41534-019-0157-8.
27. Du, Y., Hsieh, M.-H., Liu, T., and Tao, D. (2020). Expressive power of parametrized quantum circuits. Phys. Rev. Res. *2*, 033125.
28. Du, Y., Tu, Z., Wu, B., Yuan, X., and Tao, D. (2022). Power of quantum generative learning. Preprint at arXiv. https://doi.org/10.48550/arXiv.2205.04730.

29. Dallaire-Demers, P.-L., and Killoran, N. (2018). Quantum generative adversarial networks. Phys. Rev. A *98*, 012324. https://doi.org/10.1103/PhysRevA.98.012324.

30. Zoufal, C., Lucchi, A., and Woerner, S. (2019). Quantum Generative Adversarial Networks for learning and loading random distributions. npj Quantum Inf. *5*, 103. https://doi.org/10.1038/s41534-019-0223-2.

31. Situ, H., He, Z., Wang, Y., Li, L., and Zheng, S. (2020). Quantum generative adversarial network for generating discrete distribution. Inf. Sci. *538*, 193–208. https://doi.org/10.1016/j.ins.2020.05.127.

32. Huang, H.-L., Du, Y., Gong, M., Zhao, Y., Wu, Y., Wang, C., Li, S., Liang, F., Lin, J., Xu, Y., et al. (2021). Experimental quantum generative adversarial networks for image generation. Phys. Rev. Appl. *16*, 024051.

33. Tsang, S.L., West, M.T., Erfani, S.M., and Usman, M. (2023). Hybrid Quantum-Classical Generative Adversarial Network for High Resolution Image Generation. IEEE Trans. Quantum Eng. *4*, 1–19. https://doi.org/10.1109/TQE.2023.3319319.

34. Zhou, N.-R., Zhang, T.-F., Xie, X.-W., and Wu, J.-Y. (2023). Hybrid quantum–classical generative adversarial networks for image generation via learning discrete distribution. Signal Process. Image Commun. *110*, 116891.

35. Chaudhary, S., Huembeli, P., MacCormack, I., Patti, T.L., Kossaifi, J., and Galda, A. (2023). Towards a scalable discrete quantum generative adversarial neural network. Quantum Sci. Technol. *8*, 035002.

36. Kappen, H.J. (2020). Learning quantum models from quantum or classical data. J. Phys. A: Math. Theor. *53*, 214001.

37. Huijgen, O., Coopmans, L., Najafi, P., Benedetti, M., and Kappen, H.J. (2024). Training quantum Boltzmann machines with the β-variational quantum eigensolver. Mach. Learn, Sci. Technol. *5*, 025017.

38. Hernandes, V., and Greplova, E. (2023). Modeling Neuronal Activity with Quantum Generative Adversarial Networks. In 2023 IEEE International Conference on Quantum Computing and Engineering (QCE) (IEEE), pp. 330–331.

39. Marre, O., Tkacik, G., Amodei, D., Schneidman, E., Bialek, W., and Berry, M. (2017). Multi-electrode Array Recording from Salamander Retinal Ganglion Cells (Institute of Science and Technology Austria). https://doi.org/10.15479/AT:ISTA:61.

40. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. Preprint at arXiv. https://doi.org/10.48550/arXiv.1406.2661.

41. Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein generative adversarial networks. In International conference on machine learning (PMLR), pp. 214–223.

42. Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A.C. (2017). Improved training of wasserstein gans. Preprint at arXiv. https://doi.org/10.48550/arXiv.1704.00028.

43. Pérez-Salinas, A., Cervera-Lierta, A., Gil-Fuster, E., and Latorre, J.I. (2020). Data re-uploading for a universal quantum classifier. Quantum *4*, 226.

44. Hernandes, V. and Greplova, E. (2024). Data and scripts used in: "Exploring Biological Neuronal Correlations with Quantum Generative Models". Zenodo. https://doi.org/10.5281/zenodo.13388723.

45. SpiQGAN GitLab (2025). Available at https://gitlab.com/QMAI/papers/spiqgan.