

ARTICLE

Open Access

Quantum receiver enhanced by adaptive learning

Chaohan Cui¹, William Horrocks¹, Shuhong Hao², Saikat Guha^{1,3}, Nasser Peyghambarian^{1,2}, Quntao Zhuang^{1,3,4} and Zheshen Zhang^{1,2,3,5}✉

Abstract

Quantum receivers aim to effectively navigate the vast quantum-state space to endow quantum information processing capabilities unmatched by classical receivers. To date, only a handful of quantum receivers have been constructed to tackle the problem of discriminating coherent states. Quantum receivers designed by analytical approaches, however, are incapable of effectively adapting to diverse environmental conditions, resulting in their quickly diminishing performance as the operational complexities increase. Here, we present a general architecture, dubbed the quantum receiver enhanced by adaptive learning, to adapt quantum receiver structures to diverse operational conditions. The adaptively learned quantum receiver is experimentally implemented in a hardware platform with record-high efficiency. Combining the architecture and the experimental advances, the error rate is reduced up to 40% over the standard quantum limit in two coherent-state encoding schemes.

Introduction

Quantum information science (QIS) endows communication^{1–3}, sensing^{4,5}, and computing^{6,7} capabilities unrivaled by their classical counterparts. QIS has also sharpened our understanding of the fundamental limits of information acquisition, transfer, and processing due to the indistinguishability of nonorthogonal quantum states, which in turn place bounds on the rate of optical communications^{8,9} and the precision of sensing¹⁰. Unfortunately, in many scenarios, an appreciable gap separates the performance achievable by routine measurement apparatus and what is allowed by quantum mechanics^{11,12}. A central theme of QIS is thus the quest for protocols that approach the ultimate performance limits.

Quantum receivers are unconventional measurement apparatuses designed to bridge the gap and enable a performance boost in a wide range of information processing tasks modeled as quantum-state discrimination¹³ or parameter estimation¹⁴. Pioneering quantum-receiver

works unveiled that adaptive structures based on quantum circuits configured by feed-forward controls^{15,16} can vastly reduce the error probability in discriminating nonorthogonal quantum states. A landmark development was the Dolinar receiver capable of discriminating two weak coherent states at the fundamental Helstrom bound¹⁷. More recent studies have laid out quantum-receiver structures to benefit a variety of near-term tasks encompassing quantum state tomography^{18–21}, target detection^{22,23}, communication^{24–31}, and computing^{32,33}. Despite these encouraging advances, the design of quantum receivers for general QIP problems remains a formidable challenge, primarily due to the large Hilbert space that quantum states reside in. Indeed, only a few are known to achieve the ultimate performance limit even excluding imperfections. In practical situations, noise and disturbance in a dynamic environment would further bring substantial complexities to the quantum-receiver design³⁴, rendering the traditional design method based on analytic modeling clumsy and impotent.

The rapid development of data science has given rise to efficient tools for addressing complex data-processing problems in a large parameter space, shifting the landscape of data mining³⁵, computer vision³⁶, automated control^{37,38}, and decision-making³⁹. State-of-the-art

Correspondence: Zheshen Zhang (zszh@umich.edu)

¹James C. Wyant College of Optical Sciences, University of Arizona, Tucson, AZ 85721, USA

²Department of Materials Science and Engineering, University of Arizona, Tucson, AZ 85721, USA

Full list of author information is available at the end of the article

data-science methods now help combat the unprecedented challenge of designing quantum protocols and platforms, such as engineering of quantum states^{40,41}, operators⁴², or a combination of both^{43,44}.

Here, we harness reinforcement learning⁴⁵ to design quantum receivers, formulating the quantum receiver enhanced by adaptive learning (QREAL) architecture capable of tackling a series of quantum-state discrimination and data-classification problems. The QREAL architecture enjoys a substantially reduced complexity compared with other proposals based on model-free reinforcement learning^{46,47}, allowing it to be embodied in a photonic platform compatible with telecommunication. We then verify QREAL's capability of tackling different quantum-state discrimination problems with advantages over the standard quantum limit (SQL). By virtue of its capability of adapting to diverse operational conditions, QREAL's advantage over prior quantum receivers is further magnified in the presence of practical imperfections. These features of QREAL render it readily available to enhance long-haul communications, imaging, and sensing systems.

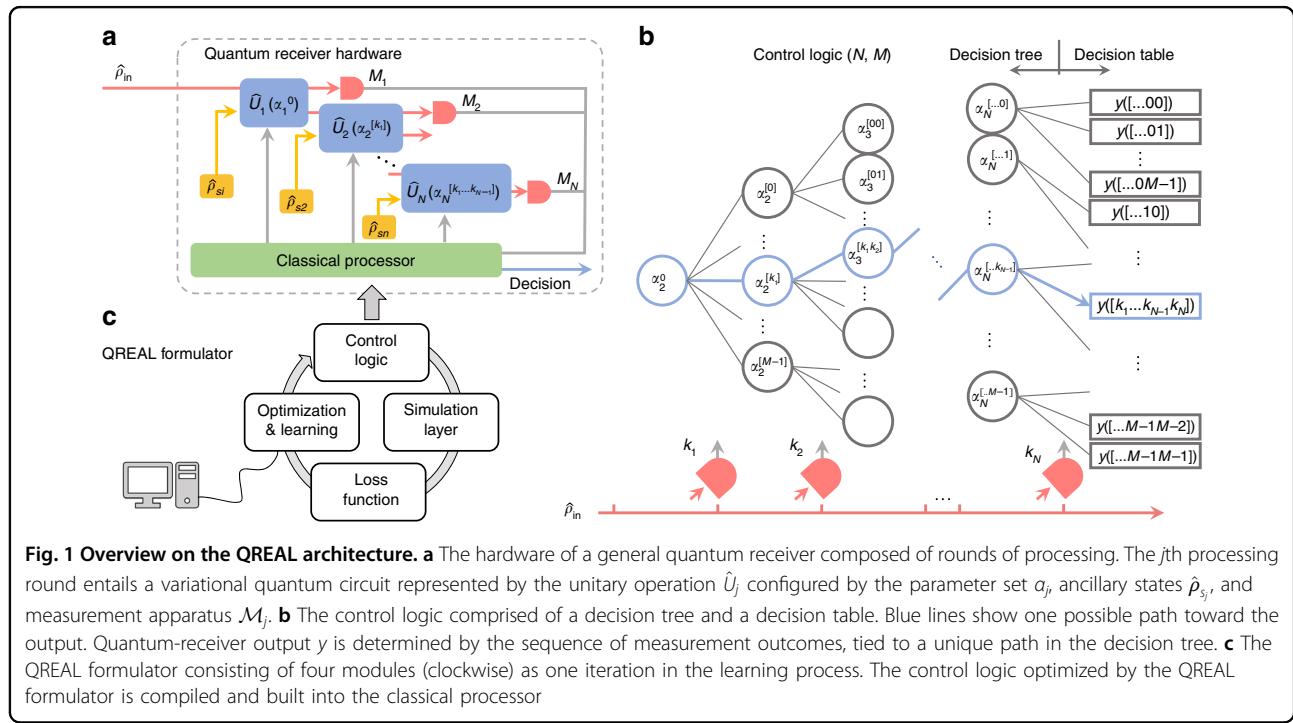
The QREAL Architecture

The general goal of our quantum receiver is to perform a hypothesis-testing task between a set of quantum states, each tagged by a classical entry y . A general quantum-receiver structure depicted in Fig. 1a comprises N rounds of processing and measurements. Following the control logic, the classical processor manages the quantum circuit

and provides a decision based on the measurement history. In the j th round, the quantum state from previous steps is modified by a history-configured unitary operation $\hat{U}_j(a_j^{[k_1 k_2 \dots k_{j-1}]})$ with an ancillary state $\hat{\rho}_{sj}$, then mapped into a result k_j , which leads to the operation in the next round.

Figure 1b illustrates a control logic comprised of an M -ary decision tree with depth N and a decision table, both employed by a quantum receiver specified as QREAL (N, M). Each node, linked from its parent node, contains a variable that is optimized by the QREAL formulator (Fig. 1c). The noise awareness of QREAL is accomplished by simulating a batch of decision trees, each generated with stochastically sampled noise through all rounds. In doing so, the QREAL formulator obtains a collection of probability distributions for the measurement history conditioned on each input state, subject to the noise characteristics. Then, the decision table is derived by Bayesian inference, and the decision tree is adjusted to lower the loss function. After the iterations, an optimized control logic for a specific QIP task is loaded into the classical processor (Detailed description in supplementary section I).

In such an iterative quantum processor with feed-forward, noise is detrimental since its impact may accumulate through operations. Noise can carry different characteristics within the time scale of a round. High-frequency noise uniformly degrades the fidelity of unitary operations in each round, which can be simply treated as



independent and get modeled easily. When noise varies on a time scale much longer than N rounds, it can be regulated by tracking the drift. However, mid-frequency noise affects the system collectively since it varies on a time scale of several rounds. This type of noise obeys stationary statistics across a batch of decision trees. QREAL aims to learn the pattern of mid-frequency noise and accordingly adapt its measurement strategy toward a better overall performance.

Setting up QREAL for decoding weak coherent light

The nonorthogonality of quantum states precludes them from being perfectly discriminated, forbidding the classical receivers operate near the fundamental limits. To bridge the gap, we leverage the QREAL formulator to construct quantum receivers for the problem of discriminating weak nonorthogonal coherent states. To date, experimentally implemented quantum receivers are predominantly devised by static approaches that are incapable of adapting to diverse operational conditions, including various genres of noise and imperfections, rendering their performance far inferior to that of QREAL as we show below.

The quantum-receiver hardware is a fiber-based platform operating at the c-band of optical communication. The components realize suitable functionalities within the general QREAL architecture. Specifically, practical modules of variational quantum circuits and measurements are assembled to manipulate and detect weak coherent states.

We next elaborate on the QREAL architecture with more details on the experimental setup, the modeling of noise patterns, and the learning process. In the quantum-receiver hardware, the ancillaries $\hat{\rho}_{sj}$ are vacuum states.

The unitary operation \hat{U}_j at each round is a reconfigurable quadrature displacement operation with a history-dependent complex variable $\alpha_j^{[k_1 k_2 \dots k_{j-1}]}$ for the displacement. The measurement at each round is a photon-number resolving (PNR) detection.

In the optical part of the hardware, the displacement operation is implemented by interfering the signal and a local oscillator (LO) on a 99:1 beamsplitter. During each round of decoding, the relative phase and amplitude of the LO determine the equivalent displacement applied to the signal. The displaced optical field is then captured by a time-multiplexed superconducting single-photon detector. A 5 kHz clock switches the system status between decoding and phase-locking while synchronizing other supplementary electronic devices.

The classical processor is an FPGA implementing the control logic with a decision tree. The FPGA stores an updatable look-up table in its memory. After acquiring photon arrival times from the detector, the FPGA

searches and determines the displacement for the next round. It then sends voltage signals to change the phase and amplitude of the LO in a proper time window. The FPGA communicates with a desktop computer that runs the QREAL formulator.

In setting up for a QIP task, the QREAL formulator works with the FPGA to estimate on the incoming signal power, dark count rate, and the photon statistics associated with different displacements on each signal. The average interference fringe visibility is obtained at the same time, as a statistical indicator for the fluctuation and noise. In general, the noise model can embody any pattern.

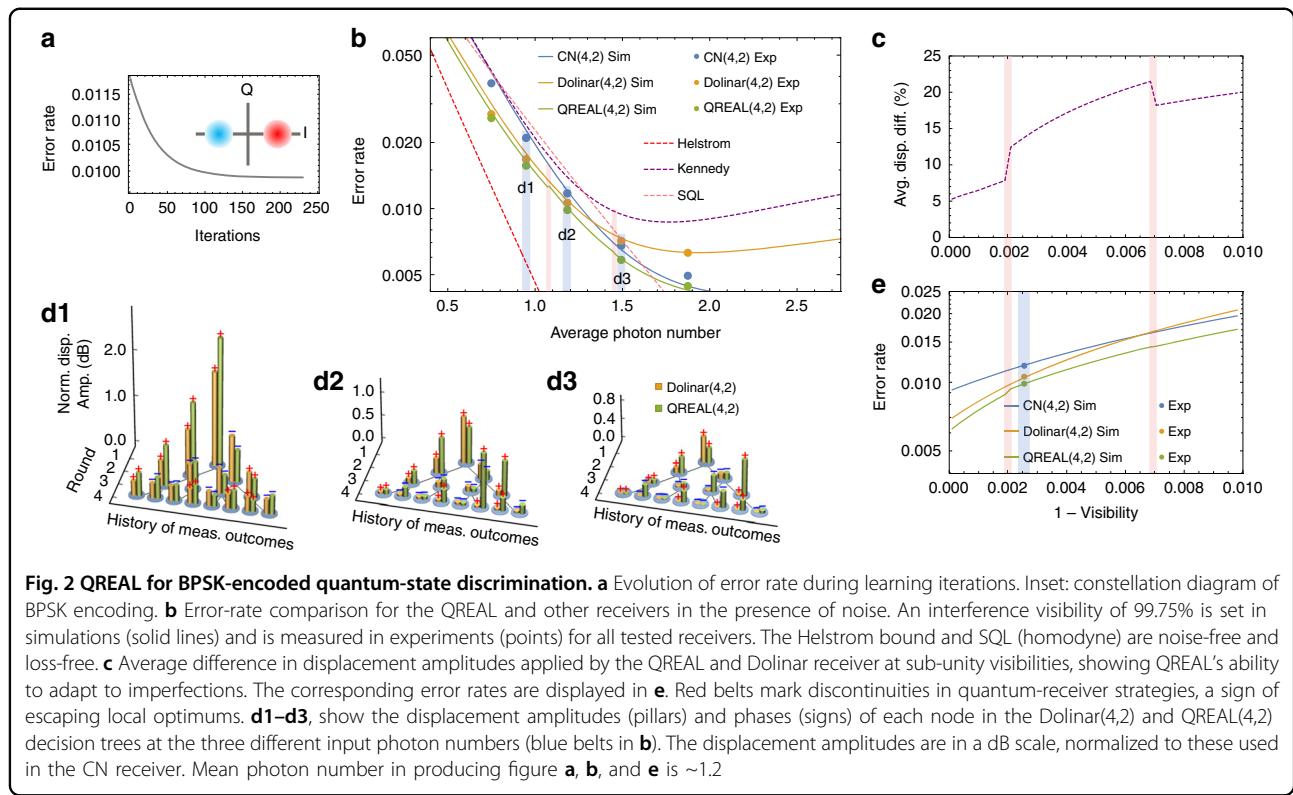
In the current experiment, the preset noise model fits the statistical outcomes by two individual Gaussian distributions on both the phase and amplitude in the displacement operation and one Poissonian distribution for the dark count and ambient light at the detector. The noise is identical for N rounds in each decoding frame but constitutes a stationary stochastic process between frames. The noise model is suitable for our setup since the switch between phase-locking and decoding will eventually add phase fluctuations to each cycle. The intractable loss fluctuation caused by the fiber-stretcher and polarization managements also induces cycle-wise amplitude bias to the LO. Due to such noise, the overall fidelity of displacement operation barely exceeds 99.7%.

Subsequently, based on the pre-estimated noise distributions, the formulator generates a batch of noisy decision trees, each obeying the modeled noise patterns. The formulator then starts to optimize displacements, aimed at reducing the average error rate for the batch of decision trees. (More details are in the supplementary section II, III).

Results

As a first test, we develop and implement the QREAL robust against noise and other imperfections for the binary phase-shift keying (BPSK) format (Fig. 2a, inset). The discrimination of quantum states encoded in the BPSK format has been extensively investigated, with many quantum receivers proposed^{16,17,48} and implemented²⁴ to beat the SQL. The original Dolinar receiver¹⁷ was the first to approach the ultimate Helstrom bound for BPSK quantum-state discrimination in an ideal scenario, given a mean photon number including loss. However, its advantage over SQL quickly fades away under noise, due to the lack of a mechanism to cope with practical non-idealities when the displacement are preset.

In contrast, QREAL(4,2) for BPSK can mitigate the adverse effect caused by noise and other imperfections by adaptively learning their patterns in ~150 iterations (~15 s) (Fig. 2a). The disparity between the learned displacement parameters and those used in the Dolinar(4,2)

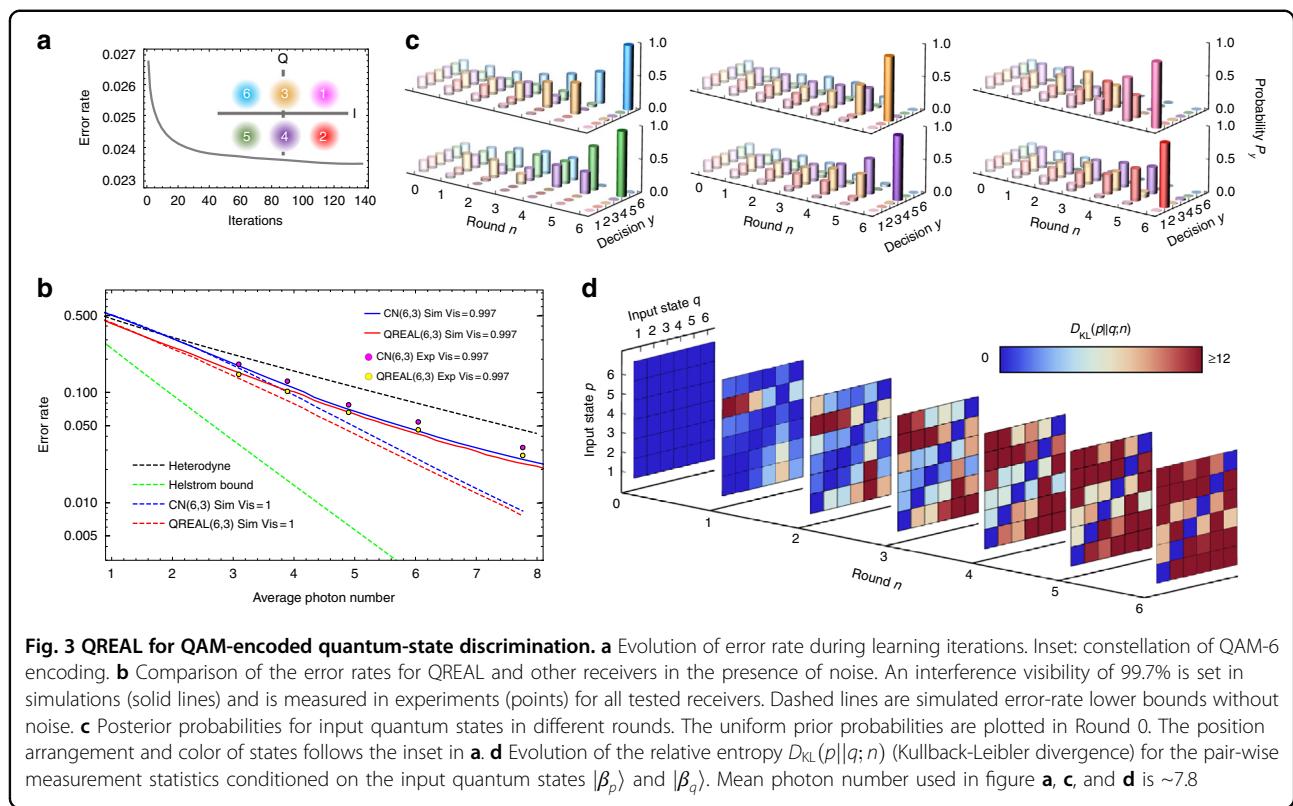


receiver is shown in Fig. 2c, d1–d3. To verify QREAL(4,2)'s robustness against noise, we test QREAL(4,2), the conditional-nulling²⁵ receiver (CN(4,2)), and Dolinar(4,2) subject to an identical noise pattern. The simulation results in Fig. 2b, e show that, while noise quickly degrades the performance of Dolinar(4,2), QREAL(4,2) maintains its advantage over SQL within a broad range of power levels of the input quantum states. In the large photon-number or high-noise regime, QREAL picks a strategy closer to that of the CN receiver, whereas in the small photon-number or low-noise regime, the QREAL learns a strategy closer to that of the Dolinar receiver. Overall, QREAL outperforms the other two quantum receivers across the entire parameter space, shown by the experimental error rates displayed in colored dots, with error bars smaller than the size of the points. With no additional components between the displacement operation and the detector, the experimental setup achieves two parameters to beat BPSK's loss-less SQL^{49,50}, i.e., an overall efficiency of 85% while maintaining the visibility at 99.75%. Specifically, we achieve a raw bit error rate of <2.5% at a mean photon number of ~0.75, corresponding to an information rate as high as 1.1 bits per received photon. QREAL enjoys an error rate 39% below the SQL at a mean photon number of 0.95 and keeps an advantage over the SQL for mean photon numbers below 1.6. After adapting to the noise, QREAL

reduces up to 14% error rate compared with the best of the CN and Dolinar receivers.

We next develop the QREAL for the quadrature amplitude modulation (QAM) encoding (Fig. 3a, inset) to demonstrate its capability of handling complex tasks in the presence of noise. QAM profits the spectral efficiency by leveraging a larger codeword space but challenges the design of sub-SQL quantum receivers due to its larger parameter space. The QREAL formulator takes ~120 iterations (~10 minutes) to converge the design at a minimal error rate (Fig. 3a). The constructed QREAL surpasses the SQL and beats the error rate of the CN receiver. The learned QREAL(6,3) for QAM-6 decoding is tested back-to-back with CN(6,3). Both the simulation and the experimental results illustrate QREAL's noise mitigation capability, which underpins QREAL's performance gain in the low photon-number regime where noise constitutes a significant portion of error events. In practice, QREAL enables up to 43% (average 32%) error-rate reduction over the SQL (loss-free heterodyne) and up to 19% (average 17%) improvement over the CN receiver, as depicted in Fig. 3b.

To illustrate the internal workflow of QREAL, Fig. 3c plots the evolution of the posterior probability distributions through the processing rounds for each codeword, starting from a uniform prior probability distribution. More information about the incoming quantum state is



acquired as the QREAL executes through the consecutive rounds so that the correct codeword becomes prominent the posterior probability distribution while the incorrect guesses are suppressed. To further understand how the QREAL discriminates between different codewords, we quantitatively compare the distance between the measurement statistics resulted from the 6 different coherent states using the relative entropy $D_{KL}(p||q; n)$, as shown in Fig. 3d. By the end of the first two rounds, the measurement statistics for a subset of the codeword states, e.g., $|\beta_1\rangle$ and $|\beta_5\rangle$, have been adequately disparate so that they are distinguished with confidence. The rest of the codeword states, however, remain unsure due to their close measurement statistics. As the processing proceeds into subsequent rounds, more measurement outcomes lead to distinct measurement statistics for different codeword states, allowing QREAL to pick the correct codeword.

We have developed QREAL to address two quantum-state discrimination problems. First, QREAL for BPSK is shown to outperform both the Dolinar and the CN receivers. Notably, QREALs are proven robust against imperfections so that QREAL's performance advantage over conventional quantum and classical receivers sustains over a large noise region. Such an improvement makes QREAL outperform other quantum receivers in nonideal environments. We also constructed QREAL for QAM to verify its capability of undertaking a complex

QIP task that require optimization in a large parameter space with more than 1000 variables. The QREAL architecture is envisioned to endow new functionalities in the noisy intermediate-scale quantum era by harnessing the hybrid quantum-classical information processing architecture.

Discussion

The quantum-receiver hardware can be augmented with new types of variational quantum circuits and measurements to embrace a wider scope of QIP problems. For example, continuous-variables measurements such as the homodyne detector offers an additional degree of freedom to access a richer set of information²⁰. With a versatile topology, the QREAL formulator can be generalized to most near-term measurement-based quantum applications, including quantum state generation³², tomography¹⁹, variational quantum eigensolvers⁵¹, and quantum-enhanced sensing²³, after proper discretization.

Moving toward practical operational environments, the QREAL formulator can fit the in-situ learning iterations as tests on the hardware also provide the sampled probability distribution. In doing so, the QREAL formulator will be more efficiently in dealing with the exponentially large quantum-state space and, in the meantime, adapt to slowly-varying noise. This feature enables long-time unsupervised stability, which is crucial for real-world

applications. In addition, the QREAL formulator supports artificial neural networks in its decision strategy, for capturing patterns embedded in data⁵², like the variational autoencoder⁵³. In circumstances without prior knowledge or model for the channel and receiver imperfections, model-free reinforcement-learning techniques can be used to design quantum receivers^{46,47}.

A few remarks on the optimization process of QREAL are worth making. A recent paper discovered that the initial parameters, to a large extent, determines the local minimum that a machine-learning algorithm converges to⁵⁴. In our work, the QREAL formulator leverages a greedy algorithm to generate initial parameters that obey a preset criterion, locating near those of the CN receiver. Other sets of initial parameters may exist to further enhance the performance of QREAL, but as a tradeoff they would require more computing resources to identify. In addition, the initialization processes for QREALs tackling more complex quantum information processing problems would call for more systematic studies.

Materials and methods

The classical computing part within the QREAL architecture is achieved by TensorFlow 2.1.0 package on Python 3.7 platform. The communication between classical computing and FPGA is held by TCP/IP protocol via wired local area network (LAN). The fiber-optical setup in front of a superconducting nanowire single photon detector (SNSPD) is compatible with the telecommunication c-band at 1550 nm right. The arrival time of the pulsed output from the SNSPD is recorded by the FPGA analog input port. The other FPGA analog input port is used for synchronization. Two analog output ports of the FPGA are properly amplified and connected to the amplitude and phase modulators. More details on the QREAL architecture and experimental setup are presented in Supplementary Information.

Acknowledgements

We gratefully acknowledge funding support by the National Science Foundation Grants No. CCF-1907918, No. ECCS-1828132, and No. EEC-1941583. Z.Z. acknowledges NSF CAREER Award No. ECCS-2144057. Q.Z. acknowledges support from Defense Advanced Research Projects Agency (DARPA) under Young Faculty Award (YFA) Grant No. N660012014029. C.C. thanks Kaushik Seshadreesan and Masoud Babaeian for their help.

Author details

¹James C. Wyant College of Optical Sciences, University of Arizona, Tucson, AZ 85721, USA. ²Department of Materials Science and Engineering, University of Arizona, Tucson, AZ 85721, USA. ³Department of Electrical and Computer Engineering, University of Arizona, Tucson, AZ 85721, USA. ⁴Department of Electrical and Computer Engineering, University of Southern California, Los Angeles, CA 90089, USA. ⁵Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA

Author contributions

C.C. and Z.Z. conceived the idea, designed the experiment. C.C. programmed the Python code and finished the experiment. W.H. programmed FPGA. C.C., Z.Z., and W.H. prepared the manuscript and supporting materials. S.H., S.G.,

N.P., and Q.Z. contributed to the theoretical development and experimental investigation for the project. Z.Z. supervised the project.

Data availability

Supporting code and data has been uploaded to the online repository <https://doi.org/10.5281/zenodo.7382847>.

Competing interests

The authors declare no competing interests.

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41377-022-01039-5>.

Received: 13 July 2022 Revised: 8 November 2022 Accepted: 16 November 2022

Published online: 08 December 2022

References

1. Gisin, N. & Thew, R. Quantum communication. *Nat. Photonics* **1**, 165–171 (2007).
2. Scarani, V. et al. The security of practical quantum key distribution. *Rev. Mod. Phys.* **81**, 1301–1350 (2009).
3. Qi, R. Y. et al. Implementation and security analysis of practical quantum secure direct communication. *Light Sci. Appl.* **8**, 22 (2019).
4. Degen, C. L., Reinhard, F. & Cappellaro, P. Quantum sensing. *Rev. Mod. Phys.* **89**, 035002 (2017).
5. Pirandola, S. et al. Advances in photonic quantum sensing. *Nat. Photonics* **12**, 724–733 (2018).
6. Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79 (2018).
7. Arute, F. et al. Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510 (2019).
8. Holevo, A. S. Bounds for the quantity of information transmitted by a quantum communication channel. *Probl. Peredachi Informatsii* **9**, 3–11 (1973).
9. Caves, C. M. & Drummond, P. D. Quantum limits on bosonic communication rates. *Rev. Mod. Phys.* **66**, 481 (1994).
10. Braunstein, S. L. & Caves, C. M. Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.* **72**, 3439–3443 (1994).
11. Nakahira, K., Kato, K. & Usuda, T. S. Optimal discrimination of optical coherent states cannot always be realized by interfering with coherent light, photon counting, and feedback. *Phys. Rev. A* **97**, 022320 (2018).
12. Pereira, J. L., Banchi, L. & Pirandola, S. Bounding the benefit of adaptivity in quantum metrology using the relative fidelity. *Phys. Rev. Lett.* **127**, 150501 (2021).
13. Barnett, S. M. & Croke, S. Quantum state discrimination. *Adv. Opt. Photonics* **1**, 238–278 (2009).
14. Helstrom, C. W. Quantum detection and estimation theory. *J. Stat. Phys.* **1**, 231–252 (1969).
15. Dolinar, S. J. A class of optical receivers using optical feedback. PhD thesis, Massachusetts Institute of Technology, Massachusetts, 1976.
16. Assalini, A., Dalla Pozza, N. & Pierobon, G. Revisiting the Dolinar receiver through multiple-copy state discrimination theory. *Phys. Rev. A* **84**, 022342 (2011).
17. Dolinar, S. J. An Optimum Receiver for the Binary Coherent State Quantum Channel. 115–120 (Research Laboratory of Electronics, MIT, 1973).
18. Andersen, U. L. et al. Hybrid discrete-and continuous-variable quantum information. *Nat. Phys.* **11**, 713–719 (2015).
19. Izumi, S., Neergaard-Nielsen, J. S. & Andersen, U. L. Tomography of a feedback measurement with photon detection. *Phys. Rev. Lett.* **124**, 070502 (2020).
20. Xu, H. C. et al. Experimental quantification of coherence of a tunable quantum detector. *Phys. Rev. Lett.* **125**, 060404 (2020).
21. Burenkov, I. et al. Time-resolving quantum measurement enables energy-efficient, large-alphabet communication. *PRX Quantum* **1**, 010308 (2020).
22. Zhuang, Q. T., Zhang, Z. S. & Shapiro, J. H. Optimum mixed-state discrimination for noisy entanglement-enhanced sensing. *Phys. Rev. Lett.* **118**, 040801 (2017).
23. Zhuang, Q. T. & Zhang, Z. S. Physical-layer supervised learning assisted by an entangled sensor network. *Phys. Rev. X* **9**, 041023 (2019).

24. Tsujino, K. et al. Quantum receiver beyond the standard quantum limit of coherent optical communication. *Phys. Rev. Lett.* **106**, 250503 (2011).
25. Chen, J. et al. Optical codeword demodulation with error rates below the standard quantum limit using a conditional nulling receiver. *Nat. Photonics* **6**, 374–379 (2012).
26. Becerra, F. E. et al. Experimental demonstration of a receiver beating the standard quantum limit for multiple nonorthogonal state discrimination. *Nat. Photonics* **7**, 147–152 (2013).
27. Becerra, F. E., Fan, J. & Miggdall, A. Photon number resolution enables quantum receiver for realistic coherent optical communications. *Nat. Photonics* **9**, 48–53 (2015).
28. Ferdinand, A. R., DiMario, M. T. & Becerra, F. E. Multi-state discrimination below the quantum noise limit at the single-photon level. *npj Quantum Inf.* **3**, 43 (2017).
29. Burenkov, I. A., Tikhonova, O. V. & Polyakov, S. V. Quantum receiver for large alphabet communication. *Optica* **5**, 227–232 (2018).
30. Izumi, S. et al. Experimental demonstration of a quantum receiver beating the standard quantum limit at telecom wavelength. *Phys. Rev. Appl.* **13**, 054015 (2020).
31. Rengaswamy, N. et al. Belief propagation with quantum messages for quantum-enhanced classical communications. *npj Quantum Inf.* **7**, 97 (2021).
32. Su, D. Q., Myers, C. R. & Sabapathy, K. K. Conversion of Gaussian states to non-Gaussian states using photon-number-resolving detectors. *Phys. Rev. A* **100**, 052301 (2019).
33. Bourassa, J. E. et al. Blueprint for a scalable photonic fault-tolerant quantum computer. *Quantum* **5**, 392 (2021).
34. DiMario, M. T. & Becerra, F. E. Phase tracking for sub-shot-noise-limited receivers. *Phys. Rev. Res.* **2**, 023384 (2020).
35. Han, J. W., Kamber, M. & Pei, J. *Data Mining: Concepts and Techniques*. 3rd edn. (London: Elsevier, 2012).
36. Forsyth, D. A. & Ponce, J. *Computer Vision: A Modern Approach*. (Boston: Prentice Hall, 2012).
37. Abbeel, P. et al. An application of reinforcement learning to aerobatic helicopter flight. Proceedings of the 19th International Conference on Neural Information Processing Systems. Vancouver, British Columbia, Canada: MIT Press, 2007.
38. Murphy, M. P. et al. The littleDog robot. *Int. J. Robot. Res.* **30**, 145–149 (2011).
39. Gibney, E. Google AI algorithm masters ancient game of Go. *Nature* **529**, 445–446 (2016).
40. Knott, P. A. A search algorithm for quantum state engineering and metrology. *N. J. Phys.* **18**, 073033 (2016).
41. Cui, C. H. et al. Wave-function engineering for spectrally uncorrelated biphotons in the telecommunication band based on a machine-learning framework. *Phys. Rev. Appl.* **12**, 034059 (2019).
42. Moleky, S. et al. Inverse design in nanophotonics. *Nat. Photonics* **12**, 659–670 (2018).
43. Arrazola, J. M. et al. Machine learning method for state preparation and gate synthesis on photonic quantum computers. *Quantum Sci. Technol.* **4**, 024004 (2019).
44. Krenn, M., Erhard, M. & Zeilinger, A. Computer-inspired quantum experiments. *Nat. Rev. Phys.* **2**, 649–661 (2020).
45. Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction*. 2nd edn. (Cambridge: MIT Press, 2018).
46. Bilgis, M. et al. Real-time calibration of coherent-state receivers: learning by trial and error. *Phys. Rev. Res.* **2**, 033295 (2020).
47. Vargas, E. M. et al. Quantum sequential hypothesis testing. *Phys. Rev. Lett.* **126**, 180502 (2021).
48. Kennedy, R. S. A Near-Optimum Receiver for the Binary Coherent State Quantum Channel. 219–225 (Research Laboratory of Electronics, MIT, 1973).
49. Cook, R. L., Martin, P. J. & Geremia, J. M. Optical coherent state discrimination using a closed-loop quantum measurement. *Nature* **446**, 774–777 (2007).
50. DiMario, M. T. & Becerra, F. E. Robust measurement for the discrimination of binary coherent states. *Phys. Rev. Lett.* **121**, 023603 (2018).
51. Cerezo, M. et al. Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).
52. Sentís, G. et al. Unsupervised classification of quantum data. *Phys. Rev. X* **9**, 041029 (2019).
53. Pu, Y. C. et al. Variational autoencoder for deep learning of images, labels and captions. Proceedings of the 30th International Conference on Neural Information Processing Systems. Barcelona, Spain: Curran Associates Inc, 2016, 2352–2360.
54. Bittel, L. & Kliesch, M. Training variational quantum algorithms is NP-hard. *Phys. Rev. Lett.* **127**, 120502 (2021).