Exploring exotic configurations with anomalous features with deep learning: Application of classical and quantum-classical hybrid anomaly detection

Kumar J. B. Ghosh [1,*] and Sumit Ghosh [2,3,†]
[1]*E.ON Digital Technology GmbH, 45131 Essen, Germany*
[2]*Institute of Physics, Johannes Gutenberg-University Mainz, 55128 Mainz, Germany*
[3]*Institute of Advance Simulations, Forschungszentrum Jülich GmbH, 52428 Jülich, Germany*

(Received 7 June 2023; revised 23 August 2023; accepted 21 September 2023; published 9 October 2023)

We present the application of classical and quantum-classical hybrid anomaly detection schemes to explore exotic configurations with anomalous features. We consider the Anderson model as a prototype, where we define two types of anomalies—a high conductance in the presence of strong impurity and a low conductance in the presence of weak impurity—as a function of random impurity distribution. Such anomalous outcome constitutes an imperceptible fraction of the data set and is not a part of the training process. These exotic configurations, which can be a source of rich new physics, usually remain elusive to conventional classification or regression methods and can be tracked only with a suitable anomaly detection scheme. We also present a systematic study of the performance of the classical and the quantum-classical hybrid anomaly detection method and show that the inclusion of a quantum circuit significantly enhances the performance of anomaly detection, which we quantify with suitable performance metrics. Our approach is quite generic in nature and can be used for any system that relies on a large number of parameters to find their new configurations, which can hold exotic new features.

DOI: 10.1103/PhysRevB.108.165408

I. INTRODUCTION

In recent years, machine learning has become an integral part of different branches of physics to explore their quantum nature [1–3]. It has shown impeccable performance is dealing with a problem with large degrees of freedom, where extracting an effective model is practically impossible. It has been adopted as a viable alternative for exploring electronic properties [4–7] as well as transport properties [8–11]. Its inherent ability to deal with a high-level nonlinearity makes it quite successful in a highly nontrivial physical problem, such as predicting different phases of matter [12,13] and their topological characterization [14–16]. In addition to playing a crucial role in discovering new materials as well as mapping their quantum features [17], this has been instrumental in designing new experiments to unravel their quantum nature [18]. Although such automatization makes it possible to scan through a huge configuration space, it also has a risk of missing exotic configurations containing significantly new physics. The occurrence of such configurations is statistically insignificant and can be easily overlooked in a learning process. Identifying these rare configurations therefore can hold a key to discovering new physics.

In this article, we present an alternative paradigm, namely, *anomaly detection* [19–22], which is particularly suitable for detecting such special configurations. The main advantage of anomaly detection with respect to conventional classification schemes is that here one does not need the *a priori* knowledge of the data points that are uncharacteristic for a specific data set, namely, *anomaly*. The training is done with normal data. The anomalies are heterogeneous and remain unknown until their occurrence. For example, consider the ECG of a regular heart beat which shows a periodic pattern. An anomaly detection algorithm trained with the normal heart beat can identify the irregularities which have not been observed before and can predict signatures of heart problems [23,24]. Due to the rarity of anomalous events, anomaly-detection data sets are heavily imbalanced. It is, therefore, a highly complex task to formally describe an anomaly [25].

In this work, we demonstrate how anomaly detection can be exploited to reveal subtle features of a condensed matter system which can remain hidden from any conventional regression or classification scheme. We consider the Anderson model, where the distribution of the random impurity constitutes the input parameter space. The output is the conductance of the system which falls down significantly for strong impurity strength. However, for a certain distribution, the system might achieve a significantly larger transmission, which we call *anomaly*. Such occurrence is statistically insignificant and therefore almost impossible to anticipate beforehand. However, a configuration which can provide large conductance in the presence of strong impurity can be quite useful in device design. For example, disorder can enhance dampinglike spin-orbit torque, which is responsible for electrical switching of magnetization [26], or it can enhance the superconducting nature [27] as well. On the other hand, for a weak impurity strength, when the system is expected to have a high transmission, the anomaly is defined as the configuration which suppresses the conductance significantly. Such anomalies can pose a hurdle in quantum optimization, even in the presence

*jb.ghosh@outlook.com
†s.ghosh@fz-juelich.de

of weak impurity [28]. The main objective of the present work is to systematically identify these anomalies with a machine learning algorithm which can be utilized to understand the nature of such unusual configurations.

In this paper, we demonstrate the application of both a classical [25] and quantum-classical hybrid anomaly detection scheme [29–32] for physical problems that manifests anomalous behavior as a complex function of a large number of variables. Taking the Anderson model as a prototype, we systematically show how an anomaly detection scheme can identify the outliers without any prior knowledge of their existence. We consider three methods, namely, isolation forest [33], autoencoder [34–36], and hybrid quantum-classical autoencoder [32], and compare their performances in terms of suitable performance metrics. Our analysis shows that the quantum anomaly detection schemes perform better compared to their classical counterpart due to their inherent ability to deal with the complex feature mapping in the latent dimension. Note that although the present work is focused on the anomalous behavior of conductance due to impurity scattering, the framework is applicable for detecting anomaly in any physical observable as a function of an arbitrarily large number of parameters and therefore would play an instrumental role in discovering new exotic configurations for different physical systems.

II. MODEL AND METHOD

For our study, we consider the Anderson model given by

$$H = \epsilon \sum_i c_i^\dagger c_i + t \sum_{\langle i,j \rangle} c_i^\dagger c_j + \sum_i V_i c_i^\dagger c_i, \tag{1}$$

where $c^\dagger, c$ is the creation/annihilation operator. t is the hopping parameter, which we choose to be -1. ϵ is the on-site energy, which we choose as $-4t$. V_i in the on-site random potential. For this study, we consider a 240×240 scattering region and use it in a two-terminal device configuration (Fig. 1). We choose a total of 80 impurities with the same strength distributed within a 200×200 region in the center [37]. We assign a constant negative value $-V_0$ for all the impurities. The Fermi level is kept at $0.0005t$, which gives a conductance of 1 in the clean limit. The zero bias conductance of the system is given by

$$T = \mathrm{Tr}[\Gamma_1 G^R \Gamma_2 G^A], \tag{2}$$

where $G^{R,A} = [E - H - \Sigma_1^{R,A} - \Sigma_2^{R,A}]^{-1}$ is the retarded/advanced Green's function of the scattering region. $\Gamma_{1,2} = i[\Sigma_{1,2}^R - \Sigma_{1,2}^A]$, where $\Sigma_{1,2}^{R,A}$ is the retarded/advanced self-energy of the left/right electrode. For our calculation, we use the tight-binding code KWANT [38], which uses scattering wave function formalism to obtain these quantities. The local density of states (LDOS) can be directly obtained from the scattering wave function.

First we look at the variation of conductance with the impurity strength. For that, we consider 2000 different impurity configurations (Fig. 1). From Fig. 1(a), one can see that for a strong impurity strength, the conductance of the system is more likely to be close to zero. The exact nature of such a "far tail" can go beyond the log-normal distribution [39] and depends on the material parameters. However, some configurations can also give rise to a high value of conductance due to resonant scattering [40], although the probability of such outcome is quite small. For strong impurity, we label such outcome an *anomaly*. Similar behavior can be observed with a weak impurity. The *anomaly* in this case is a configuration that can completely suppress the current flow, resulting in an insulating behavior. From Fig. 1, one can see that the impurity configurations corresponding to a high (anomaly) and low (normal) value of conductance do not have any characteristic difference. It is therefore impossible to detect such anomalous behavior with any conventional method only from the knowledge of the distribution. In the following, we are going to show how an anomaly detection algorithm can detect such anomalies without any *a priori* knowledge of such outcome.

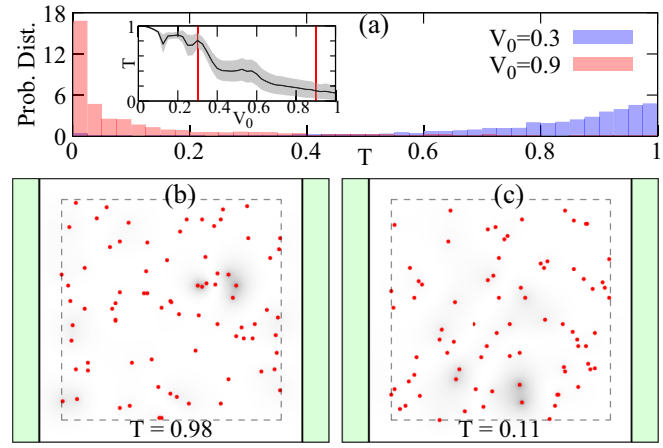

FIG. 1. Distribution of conductance and schematic of the device configuration. (a) Distribution of conductance for $V_0 = 0.9$ (red) and $V_0 = 0.3$ (blue) for 2000 random configurations. Inset: the variation of average conductance (solid black line, gray region shows the rms deviation), with vertical red line denoting $V_0 = 0.3, 0.9$. (b),(c) Two configurations for $V_0 = 0.9$, with LDOS (in gray scale) and conductance (in legend). The green regions show the electrodes. The red dots denote the impurities which are confined within the central region marked by a black dashed line.

A. Classical anomaly detection

Here, we summarize two different classical machine learning methods that we use for anomaly detection. The first method is called isolation forest (IF) [33], which is an unsupervised anomaly detection algorithm that uses a random forest algorithm under the hood to detect outliers in the data set. The algorithm tries to isolate the data points using decision trees such that each observation gets isolated from the others.

The second method is called *autoencoder* (AE) [34–36], which is a deep neural network architecture (Fig. 2). It aims to learn a compressed representation for an input through minimizing its reconstruction error [41,42]. It consists of two parts—an *encoder* (e) and a *decoder* (d). The encoder learns a nonlinear transformation $e : \mathcal{X} \to \mathcal{Z}$, which projects the data from the original high-dimensional input space $\mathcal{X} \equiv \{x\}$ to a lower-dimensional latent space $Z \equiv \{z\}$. For our study, we consider a latent space with four nodes. A decoder learns a nonlinear transformation $d : \mathcal{Z} \to \mathcal{X}$ that projects the latent vectors $z = e(x)$ back into the original high-dimensional input

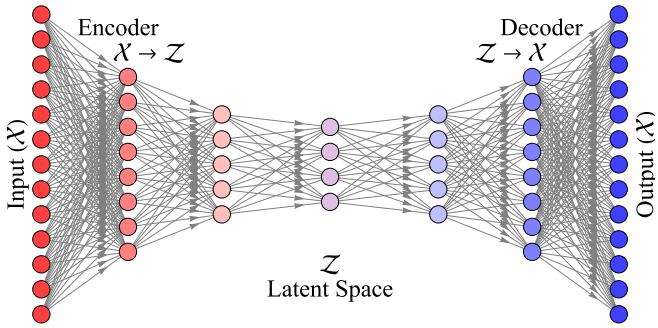

FIG. 2. Schematic representation of an autoencoder. The input data are compressed by the encoder and the decoder expands the compress data to its original size. The intermediate space with compressed dimension is called the latent space.

space $\mathcal{X}$. This transforms the latent vector $z = e(x)$ and reconstructs the original input data as $\hat{x} = d(z) = d[e(x)]$, where $\hat{x}$ is the output corresponding to an input x. One can obtain a more robust decoding of latent vectors with a *variational autoencoder* (VAE) [43], which is a neural network that unifies variational inference approaches with autoencoders. For our study, we focus only on IF and AE.

B. Hybrid quantum-classical autoencoder (HAE)

The quantum-classical hybrid anomaly detection scheme [29–32] is the state-of-the-art approach which utilizes quantum machine learning [44–48] along with its classical counterpart. For our study, we use the hybrid classical-quantum autoencoder (HAE) introduced by Sakhnenko *et al.* in 2022 [32], which significantly enhances the performance metrics of the anomaly detection compared to its fully classical counterpart. The HAE consists of a classical encoder, a parameterized quantum circuit (PQC) [49], and a classical decoder (Fig. 3). The input goes to the PQC via the encoder. The PQC consists of quantum circuits containing different rotation gates. After the blocks of quantum circuits, there are measurements followed by the postmeasurement processing block. After postprocessing, the information is fed into the classical decoder.

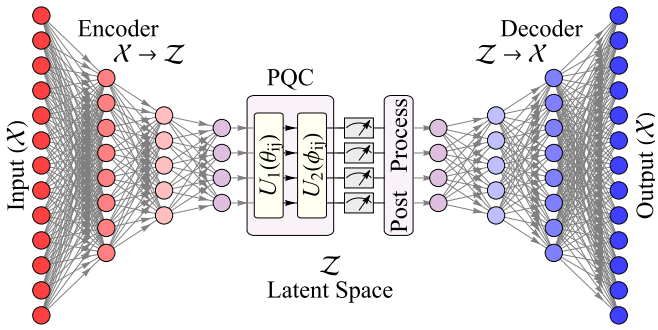

FIG. 3. A schematic diagram of a hybrid classical-quantum autoencoder (HAE) architecture [32]. The data coming from the classical encoder are embedded into the PQC. In PQC, the $U_1(\theta_{ij})$ and $U_2(\phi_{ij})$ are the blocks of quantum circuits containing different rotation gates with rotation parameters θ, ϕ. After the blocks of quantum circuits, there are measurements followed by the postmeasurement processing block (denoted as Post Process). After postprocessing, the information is fed into the classical decoder.

It is worth mentioning that a PQC performs much better compared to a classical circuit with equal dimension [32]. Schuld [50,51] showed a connection between the quantum neural networks and kernel methods, where the quantum networks encode and process the data in a high-dimensional Hilbert space through a highly nonlinear feature mapping. This is classically intractable and only can be revealed through the inner products and measurements of the quantum states. In our case, the PQC in HAE also expands the latent space into a higher-dimensional Hilbert space. Therefore, its internal degrees of freedom increase, resulting in a performance boost.

For HAE, we consider the same encoder and decoder as the classical AE which is combined with a four-qubit PQC (Fig. 3). The final quantum state is measured in the Pauli Z basis, and the corresponding expectation value for each qubit constructs the latent space for the anomaly detection. This information is fed to the decoder via the postprocessing module, which expands the compressed data to its original size. The model is implemented using QISKIT [52] for our analysis.

C. Training and testing of IF, AE, and HAE for the anomaly detection

For the isolation forest (IF), we choose a training data set with nominal data points only. After training the model, we apply the trained IF model on the testing data and obtain the predicted labels. For the AE and HAE, we first train the networks with the above training data for 50 epochs with batch size of 16 and learning rate 0.001, and then compute the losses during training. We compute the mean-squared-error loss (MSE), defined as

$$\mathrm{MSE} = \frac{1}{n}\sum_{i=1}^{n}(x_i - \hat{x}_i)^2, \tag{3}$$

where x and $\hat{x}$ are the original and the predicted values of the data points, respectively. By minimizing the above losses for the training data, we search for a suitable threshold with extensive empirical trials. Then we predict the outputs with respect to this threshold from the training data and compute the MSEs with respect to the actual training data. The threshold is then defined by the ratio between the mean and standard deviation of the MSEs. After defining the threshold, we apply the trained model over the testing data and predict the outputs. We compute the MSEs with the predicted data and the actual test data. For each prediction, if the MSE is greater than the threshold, then we label it as an anomaly/outlier; otherwise, it is labeled as a nominal/normal data point. With the above training and testing procedure, we finally compute the precision, recall, and F1 scores by comparing the predicted labels and the actual labels of the testing data.

TABLE I. Performance metrics for anomaly detection IF, AE, and HAE using two data sets ($V_0 = 0.3$ and $V_0 = 0.9$). Each value shows the mean of the corresponding metric with their standard deviations denoted in parentheses.

Set	Metric	IF	AE	HAE
$V_0 = 0.9$	Precision	0.474 (0.013)	0.494 (0.018)	0.484 (0.021)
	Recall	0.675 (0.067)	0.715(0.064)	0.748 (0.071)
	F1 score	0.556 (0.029)	0.583 (0.032)	0.588 (0.037)
$V_0 = 0.3$	Precision	0.431 (0.015)	0.436 (0.009)	0.432 (0.008)
	Recall	0.497 (0.053)	0.595 (0.015)	0.618 (0.026)
	F1 score	0.461 (0.031)	0.503 (0.008)	0.509 (0.014)

D. Performance measures of the model: Precision recall and F1 score

A reliable metric is necessary for measuring the performance of an anomaly-detection model, which should describe the fractions of uncovered anomalies from a mixture of nominal data and outliers. This is usually described by *precision* (fraction of true anomalies of all discovered instances), *recall* (fraction of true anomalies that were discovered), and their harmonic mean *F1 score* [53,54], which are computed based on the counts of true positives (T_P), false positives (F_P), and false negatives (F_N), are defined as follows:

$$precision = \frac{T_P}{T_P + F_P}, \quad recall = \frac{T_P}{T_P + F_N},$$
$$F1\ score = \frac{2 \times precision \times recall}{precision + recall}. \tag{4}$$

An outcome with high *recall* and low *precision* contains more results, but most of them would be wrong (F_P). A low *recall* and high *precision*, on the other hand, correspond to less results, but most of them would be right (T_P). The most desirable outcome is one with both high *recall* and high *precision*, which in turn gives a high *F1 score*.

III. RESULTS

For our study, we consider a two-terminal device configuration with randomly distributed 80 impurities (Fig. 1) distributed over a 200×200 region. We use their coordinates $(x_1, y_1, x_2, y_2, \ldots,$ total of 160 features) as the input and the resulting conductance (T) as the output. We consider two magnitudes of the impurity ($V_0 = 0.3, 0.9$) and calculate the transmission of 5000 different configurations. The nominal (0) and anomalous (1) data are identified with respect to the dominant behavior. Considering the distribution of conductance in these two cases (Fig. 1), for $V_0 = 0.9$, $T > 0.5$ is considered as an anomaly, whereas for $V_0 = 0.3$, $T < 0.5$ is considered as an anomaly. In both cases, the total number of anomalies is less than 10% of the entire data set. From each data set, we prepare four different train-test samples by randomly choosing 900 nominal and 100 anomalous data points. Note that for anomaly detection, the training is done only with nominal data. After training and testing with four different sets, we compute the individual performance metrics (precision, recall, F1 score) and present their respective mean values and standard deviations in Table I.

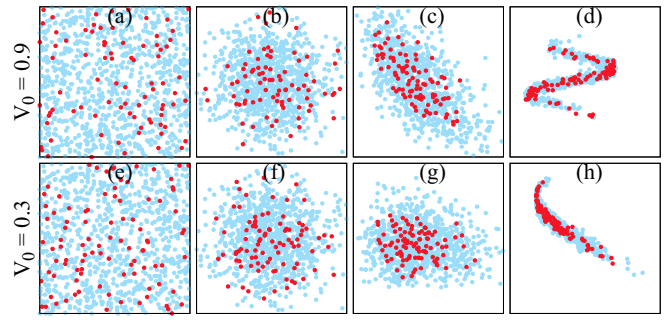

FIG. 4. Visualization of anomalies (red) and nominal data (light blue). (a)–(d) $V_0 = 0.9$ and (e)–(h) $V_0 = 0.3$ data sets with respect to two arbitrary input dimensions. (a) and (e) show the original input data. (b) and (f) show the PCA of the original data. (c) and (g) show the output from the classical encoder, and (d) and (h) show the output from the PQC.

For AE, the input size is 160 (the original data size), which is followed by the encoder containing three layers with 106, 56, and 4 nodes, respectively. The decoder is the mirror image of the encoder, followed by the output layer of size 160 (Fig. 2). The latent size of our AE model is equal to 4 (Fig. 3). The same classical encoder and decoder are used for HAE with a four-qubit PQC. A bigger latent dimension could improve the outcome; however, due to the limitation of computational resources, we cannot consider more than four qubits, which also restricts the latent dimension of classical AE.

Figure 4 shows two data sets with respect to two arbitrary features. Figures 4(a) and 4(e) show the original data where the normal data (light blue) and anomalies (red) are uniformly distributed over the whole space. Figures 4(b) and 4(f) show the output from principal component analysis (PCA) [55], which shows some clustering. This clustering is enhanced with a classical AE [Figs. 4(c) and 4(g)] and even more with a HAE [Figs. 4(d) and 4(h)]. In a higher-dimensional space, such clustering leads to a better isolation of anomalies from the normal data, which is also reflected in their individual performance metric (Table I).

Note that the data set corresponding to $V_0 = 0.9$ shows better performance compared to the data set corresponding to $V_0 = 0.3$. This can be understood from the physical nature of the problem as well. For $V_0 = 0.9$, most of the configurations lead to a localization reducing the conductance. This is reflected in the large occurrence of zero conductance in Fig. 1(a). Consequently, the anomaly is very well defined since the majority of the data has the same characteristics. Compared to that, the distribution due to $V_0 = 0.3$ is relatively flat, which makes the boundary between the nominal and anomalous data more smudged. Nevertheless, both AE and HAE show exceptional performance in both cases. To benchmark our results, we compare them with the results obtained with other data sets. We find that the performance metrics obtained in Table I are comparable to what is observed with standard publicly available data sets, and are even better than the dataset used to detect anomaly for a gas turbine using the same models [32]. From Table I, we see that the HAE is performing better in terms of recall and the F1 scores keeping the precision comparable to the other models. This is expected

due to the inclusion of PQC, as we discussed previously. Due to the limitation of computational resource, we are limited to a four-qubit PQC, where the distinction between the performance of AE and HAE is not very prominent. By increasing the dimension of the PQC and the latent space, one can further enhance the performance of the HAE.

IV. CONCLUSION

In this paper, we demonstrate the application of anomaly detection to reveal exotic features of a condensed matter system. We consider an Anderson model which shows two kind of anomalies, i.e., a high transmission at strong impurity strength and a low transmission at weak impurity strength. Here we focus on three different approaches: isolation forest (IF), which is based on the classification scheme random forest; autoencoder (AE), which is based on a classical neural network; and the hybrid classical-quantum auto encoder (HAE), which is a combination of a classical neural network and a parametric quantum circuit. Unlike the classification scheme, here the training is done only on the normal data and the learning algorithm detects the anomalous outcome without any prior knowledge of the anomaly class. The performance of these algorithms is quantified with three different scores, namely, precision, recall, and F1 score. Predicting such high level of nonlinear outcome is only possible via a neural network, which is also reflected in their individual scores. We also demonstrate that the HAE performs better compared to its classical counterpart (AE) due to its inherent ability to deal with highly nonlinear feature mapping [50,51], which cannot be achieved with a classical circuit with the same dimension.

In the context of quantum transport, these anomaly detection schemes can be instrumental in understanding the behavior of Anderson localization, as well as the formation of solitons in a disordered system. The method we present here is quite generic and can be extended to other systems. For example, in the case of optical lattices, this formalism can be exploited to investigate the Anderson localization of light [56]. For an abstract higher-dimensional phase space, where the input parameters are made of different physical observable such as electronic or chemical properties of a system, this approach can reveal new exotic configurations which cannot be explored with any conventional methods, and thus would be instrumental in new physics hidden in remote corners of complex phase space.

ACKNOWLEDGMENTS

S.G. would like to acknowledge funding from Transregional Collaborative Research Center (SFB/TRR) 173, Spin+X.

[1] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Banchi, F. Marquardt, L. Pezzè, and C. Bonato, Learning quantum systems, Nat. Rev. Phys. **5**, 141 (2023).

[2] M. Krenn, J. Landgraf, T. Foesel, and F. Marquardt, Artificial intelligence and machine learning for quantum technologies, Phys. Rev. A **107**, 010101 (2023).

[3] A. Dawid, J. Arnold, B. Requena, A. Gresch, M. Płodzień, K. Donatella, K. A. Nicoli, P. Stornati, R. Koch, M. Büttner, R. Okuła, G. Muñoz-Gil, R. A. Vargas-Hernández, A. Cervera-Lierta, J. Carrasquilla, V. Dunjko, M. Gabrié, P. Huembeli, E. van Nieuwenburg, F. Vicentini *et al.*, Modern applications of machine learning in quantum sciences, arXiv:2204.04198.

[4] A. Chandrasekaran, D. Kamal, R. Batra, C. Kim, L. Chen, and R. Ramprasad, Solving the electronic structure problem with machine learning, npj Comput. Mater. **5**, 22 (2019).

[5] A. Sheverdin, F. Monticone, and C. Valagiannopoulos, Photonic Inverse Design with Neural Networks: The Case of Invisibility in the Visible, Phys. Rev. Appl. **14**, 024054 (2020).

[6] J. Westermayr, M. Gastegger, K. T. Schütt, and R. J. Maurer, Perspective on integrating machine learning into computational chemistry and materials science, J. Chem. Phys. **154**, 230903 (2021).

[7] H. J. Kulik, T. Hammerschmidt, J. Schmidt, S. Botti, M. A. L. Marques, M. Boley, M. Scheffler, M. Todorović, P. Rinke, C. Oses, A. Smolyanyuk, S. Curtarolo, A. Tkatchenko, A. P. Bartók, S. Manzhos, M. Ihara, T. Carrington, J. Behler, O. Isayev, M. Veit *et al.*, Roadmap on machine learning in electronic structure, Electron. Struct. **4**, 023004 (2022).

[8] A. Lopez-Bezanilla and O. A. von Lilienfeld, Modeling electronic quantum transport with machine learning, Phys. Rev. B **89**, 235411 (2014).

[9] K. Li, J. Lu, and F. Zhai, Neural networks for modeling electron transport properties of mesoscopic systems, Phys. Rev. B **102**, 064205 (2020).

[10] F. Kotthoff, F. Pollmann, and G. De Tomasi, Distinguishing an Anderson insulator from a many-body localized phase through space-time snapshots with neural networks, Phys. Rev. B **104**, 224307 (2021).

[11] K. J. B. Ghosh and S. Ghosh, Classical and quantum machine learning applications in spintronics, Digit. Discov. **2**, 512 (2023).

[12] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, Nat. Phys. **13**, 431 (2017).

[13] S. Tibaldi, G. Magnifico, D. Vodola, and E. Ercolessi, Unsupervised and supervised learning of interacting topological phases from single-particle correlation functions, SciPost Phys. **14**, 005 (2023).

[14] N. Claussen, B. A. Bernevig, and N. Regnault, Detection of topological materials with machine learning, Phys. Rev. B **101**, 245117 (2020).

[15] N. Käming, A. Dawid, K. Kottmann, M. Lewenstein, K. Sengstock, A. Dauphin, and C. Weitenberg, Unsupervised machine learning of topological phase transitions from experimental data, Mach. Learn.: Sci. Technol. **2**, 035037 (2021).

[16] E. Koridon, J. Fraxanet, A. Dauphin, L. Visscher, T. E. O'Brien, and S. Polla, A hybrid quantum algorithm to detect conical intersections, arXiv:2304.06070.

[17] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, Nat. Phys. **14**, 447 (2018).

[18] M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, and A. Zeilinger, Automated Search for new Quantum Experiments, Phys. Rev. Lett. **116**, 090405 (2016).

[19] S. Omar, A. Ngadi, and H. H. Jebur, Machine learning techniques for anomaly detection: An overview, Intl. J. Comput. Applic. **79**, 33 (2013).

[20] A. Valdes, R. Macwan, and M. Backes, Anomaly detection in electrical substation circuits via unsupervised machine learning, in *2016 IEEE 17th International Conference on Information Reuse and Integration (IRI)* (IEEE, Pittsburgh, PA, USA, 2016), pp. 500–505.

[21] Y. Cui, P. Bangalore, and L. B. Tjernberg, An anomaly detection approach based on machine learning and SCADA data for condition monitoring of wind turbines, in *2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)* (IEEE, Boise, ID, USA, 2018), pp. 1–6.

[22] K. Kottmann, F. Metz, J. Fraxanet, and N. Baldelli, Variational quantum anomaly detection: Unsupervised mapping of phase diagrams on a physical quantum computer, Phys. Rev. Res. **3**, 043184 (2021).

[23] J. L. P. Lima, D. Macedo, and C. Zanchettin, Heartbeat anomaly detection using adversarial oversampling, in *2019 International Joint Conference on Neural Networks (IJCNN)* (IEEE, Budapest, Hungary, 2019), pp. 1–7.

[24] R.-G. Colt, C.-H. Várady, R. Volpi, and L. Malagò, Automatic feature extraction for heartbeat anomaly detection, arXiv:2102.12289.

[25] G. Pang, C. Shen, L. Cao, and A. van den Hengel, Deep learning for anomaly detection: A review, ACM Computing Surveys **54**, 1 (2021).

[26] S. Ghosh and A. Manchon, Spin-orbit torque in two-dimensional antiferromagnetic topological insulators, Phys. Rev. B **95**, 035422 (2017).

[27] V. D. Neverov, A. E. Lukyanov, A. V. Krasavin, A. Vagov, and M. D. Croitoru, Correlated disorder as a way towards robust superconductivity, Commun. Phys. **5**, 177 (2022).

[28] B. Altshuler, H. Krovi, and J. Roland, Anderson localization makes adiabatic quantum optimization fail, Proc. Natl. Acad. Sci. **107**, 12446 (2010).

[29] J. Romero, J. P. Olson, and A. Aspuru-Guzik, Quantum autoencoders for efficient compression of quantum data, Quantum Sci. Technol. **2**, 045001 (2017).

[30] D. Herr, B. Obert, and M. Rosenkranz, Anomaly detection with variational quantum generative adversarial networks, Quantum Sci. Technol. **6**, 045004 (2021).

[31] T. Schlegl, P. Seeböck, S. M. Waldstein, U. Schmidt-Erfurth, and G. Langs, Unsupervised anomaly detection with generative adversarial networks to guide marker discovery, arXiv:1703.05921.

[32] A. Sakhnenko, C. O'Meara, K. J. B. Ghosh, C. B. Mendl, G. Cortiana, and J. Bernabé-Moreno, Hybrid classical-quantum autoencoder for anomaly detection, Quantum Mach. Intell. **4**, 27 (2022).

[33] F. T. Liu, K. M. Ting, and Z.-H. Zhou, Isolation forest, in *2008 Eighth IEEE International Conference on Data Mining* (IEEE, Pisa, Italy, 2008), pp. 413–422.

[34] P. Baldi, Autoencoders, unsupervised learning, and deep architectures, in *Proceedings of the 2011 International Conference on Unsupervised and Transfer Learning Workshop - Volume 27* (JMLR.org, Washington, USA, 2011), pp. 37–50.

[35] A. Makhzani, J. Shlens, N. Jaitly, I. Goodfellow, and B. Frey, Adversarial autoencoders, arXiv:1511.05644.

[36] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, and P.-A. Manzagol, Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion, J. Mach. Learn. Res. **11**, 3371 (2010).

[37] This corresponds to more than 10^{368} configurations.

[38] C. W. Groth, M. Wimmer, A. R. Akhmerov, and X. Waintal, Kwant: A software package for quantum transport, New J. Phys. **16**, 063065 (2014).

[39] B. K. Nikolić and V. Z. Cerovski, Structure of quantum disordered wave functions: Weak localization, far tails, and mesoscopic transport, Eur. Phys. J. B **30**, 227 (2002).

[40] T. O. Wehling, S. Yuan, A. I. Lichtenstein, A. K. Geim, and M. I. Katsnelson, Resonant Scattering by Realistic Impurities in Graphene, Phys. Rev. Lett. **105**, 056802 (2010).

[41] A. Asperti and M. Trentin, Balancing reconstruction error and Kullback-Leibler divergence in variational autoencoders, IEEE Access **8**, 199440 (2020).

[42] M. Sabokrou, M. Fathy, and M. Hoseini, Video anomaly detection and localisation based on the sparsity and reconstruction error of auto-encoder, Electron. Lett. **52**, 1122 (2016).

[43] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, arXiv:1312.6114.

[44] M. Schuld and F. Petruccione, *Supervised Learning with Quantum Computers* (Springer International Publishing, 2018).

[45] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, Phys. Rev. A **98**, 032309 (2018).

[46] A. Mari, T. R. Bromley, J. Izaac, M. Schuld, and N. Killoran, Transfer learning in hybrid classical-quantum neural networks, Quantum **4**, 340 (2020).

[47] L. Funcke, T. Hartung, K. Jansen, S. Kühn, and P. Stornati, Dimensional expressivity analysis of parametric quantum circuits, Quantum **5**, 422 (2021).

[48] Z. Tabi, B. Bako, D. T. R. Nagy, P. Vaderna, Z. Kallus, P. Haga, and Z. Zimboras, Hybrid quantum-classical autoencoders for end-to-end radio communication, in *2022 IEEE/ACM 7th Symposium on Edge Computing (SEC)* (IEEE, Seattle, WA, USA, 2022), pp. 468–473.

[49] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, Quantum Sci. Technol. **4**, 043001 (2019).

[50] M. Schuld and N. Killoran, Quantum Machine Learning in Feature Hilbert Spaces, Phys. Rev. Lett. **122**, 040504 (2019).

[51] M. Schuld, Supervised quantum machine learning models are kernel methods, arXiv:2101.11020.

[52] G. Aleksandrowicz, T. Alexander P. Barkoutsos, L. Bello, Y, Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, S. De La Puente González, E. De La Torre, D. Ding, E. Dumitrescu *et al.*, *Qiskit: An Open-source Framework for Quantum Computing* (Zenodo, 2019), https://zenodo.org/record/2562110.

[53] N. Chinchor and B. Sundheim, Muc-5 evaluation metrics, in *Fifth Message Understanding Conference (MUC-5): Proceedings of a Conference*, Baltimore, Maryland (1993), https://aclanthology.org/M93-1007.

[54] C. J. Van Rijsbergen, Foundation of evaluation, J. Doc. **30**, 365 (1974).

[55] Principal Component Analysis for Special Types of Data, *Principal Component Analysis* (Springer, New York, 2002), pp. 338–372.

[56] M. Segev, Y. Silberberg, and D. N. Christodoulides, Anderson localization of light, Nat. Photon. **7**, 197 (2013).