

Meta-Machine-Learning-Based Quantum Scar Detector

Chen-Di Han,¹ Cheng-Zhen Wang,¹ and Ying-Cheng Lai^{1,2,*}

¹*School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, Arizona 85287, USA*

²*Department of Physics, Arizona State University, Tempe, Arizona 85287, USA*

 (Received 4 November 2022; revised 14 February 2023; accepted 18 May 2023; published 13 June 2023)

A remarkable phenomenon in contemporary physics is quantum scarring in systems whose classical dynamics are chaotic, where certain wave functions tend to concentrate on classical periodic orbits of low periods. Quantum scarring has been studied for more than four decades, but detecting quantum scars still mostly relies on human visualization of the wave-function patterns. The widespread and successful applications of machine learning in many branches of physics suggest the possibility of using artificial neural networks for automated detection of quantum scars. Conventional machine learning often requires substantial training data, but, for quantum scars, this poses a significant challenge: in typical systems the available distinct quantum scarring states are rare. We develop a meta machine-learning approach to accurately detect quantum scars in a fully automated and highly efficient fashion. In particular, taking advantage of some standard large datasets such as Omniglot from the field of image classification, we train a “preliminary” version of the neural network that has the ability to distinguish different classes of noisy images. We then perform few-shot classification to further train the neural network but with a small number of quantum scars. We demonstrate that the meta-learning scheme can find the correct quantum scars from thousands of images of wave functions without any human intervention, regardless of the symmetry of the underlying system. From a general applied point of view, our success opens the door to exploiting meta learning for solving challenging image detection and classification problems in other fields of science and engineering. For example, in microlasing systems, identifying scarring states is critical as these states are desired for directional emission. The task is also important for quantum-dot devices where the scarring states can lead to resonances in the conductance.

DOI: [10.1103/PhysRevApplied.19.064042](https://doi.org/10.1103/PhysRevApplied.19.064042)

I. INTRODUCTION

In the field of quantum chaos that studies the quantum manifestations of classical chaos [1,2], wave-function scarring is a fundamental phenomenon that has been extensively investigated. Here the term “scarring” is referred to as the unusually high concentrations of the wave functions about certain classical periodic orbits. In this regard, if the underlying classical dynamics are integrable with stable periodic orbits, nonuniformly distributed wave functions about these orbits are expected. What is surprising is quantum scarring in systems whose classical dynamics are fully chaotic, where all periodic orbits are unstable, so, intuitively, it does not seem possible for the wave functions to concentrate about them. From another viewpoint, because of the intrinsic ergodicity associated with chaos, a classical trajectory generates a uniform distribution in the phase space (accordingly in the physical space as well), so the intuition would be that the quantum wave functions should also have a uniform distribution in space. It was first

discovered by McDonald and Kaufman [3] when solving the Helmholtz equation in the classically chaotic stadium billiard that there are eigenstates whose wave functions are highly nonuniform and in fact tend to concentrate on some classical unstable periodic orbits. A more detailed study by Heller [4] confirmed the phenomenon, who gave the name “quantum scars” to the nonuniform wave functions. A theory for quantum scars based on the semiclassical Green’s function was developed by Bogomolny [5] and Berry [6].

There are a variety of physical systems in which the phenomenon of quantum scarring can occur. For example, in graphene, the low-energy excitations are governed by the Dirac equation [7–9] and correspond to massless particles, where relativistic quantum scars can arise [10–12]. Solutions of the Dirac equation in classically chaotic billiards also revealed a distinct class of relativistic quantum scars—chiral scars that require two complete cycles for their wave functions to return to their original values [13–16]. Quantum scarring can also occur in open or transport systems such as quantum dots, where the scarred states are referred to as quantum pointer states [17–23]. In microcavity lasing systems, quasibound states

*ying-cheng.lai@asu.edu

in a deformed cavity are scarring states with strong directional emission [24–27]. Quite recently, scarring has been uncovered in quantum many-body systems, where certain special eigenstates were found to concentrate in certain parts of the Hilbert space [28–34], which were named quantum many-body scars and are highly relevant to multiple-qubit systems in quantum information science and technology.

In the study of quantum chaos, identifying quantum scarring states has been a challenging problem. The conventional approach has been “manual,” where one first generates a large number of eigenstates and then visually checks to see if an eigenstate is a scarring state based on information about the classical periodic orbits. This can be an extremely difficult task for two reasons. First, scarring states are rare in typical quantum systems. For example, in a chaotic billiard, only approximately 5% to 10% of the eigenstates are quantum scars. Compounding the difficulty is the different classes of quantum scars, where each class corresponds to a different group of classical periodic orbits. For a given class, quantum scarring states are even more rare. For example, when chiral scars were first discovered [13], more than 10^4 eigenstates were examined visually by human eyes. The second reason is that the wave-function patterns associated with quantum scars can be complicated with a random or noisy component. As a result, labeling the quantum scars is a hard task with significant uncertainties for the human eyes.

In recent years, deep learning has enjoyed great success in visual object recognition, object detection, and many other domains [35]. Deep convolutional neural networks (DCNNs) were introduced for image recognition, the training of which often requires very large datasets [36]. For example, in Ref. [36], the authors considered 1000 classes, which required more than one million images as the training data. In physics, DCNNs have been adopted in different areas to solve problems associated with a large number of images [37], such as jet tagging [38,39], neutrino event classification [40], gravitational lensing [41], and identification of symmetry-breaking states from scanning-electron-microscope images [42]. A convolutionlike neural network structure was also developed to solve problems in quantum many-body physics [43]. Quite recently, a machine-learning-based quantum chaos detector was proposed [44]. In applications where such massive datasets are not available, machine learning based on small datasets was also developed, especially for the classification task [45–47] where only a small number of labeled examples per class (a few shots) are required. Since only small training data are required for classification, such machine learning schemes find broad applications, e.g., learning and detecting rare events [48,49].

The remarkable success of machine learning in image recognition and pattern classification naturally leads to the idea of developing a neural-network-based quantum

scar detector. However, even as the image classification algorithms have become more powerful, detecting and classifying quantum scars remains an open and difficult problem, due to the availability of a small number of quantum scarring patterns and the lack of sufficient distinction between quantum scars and conventional wave functions. Intuitively, since each type of quantum scar corresponds to a unique class of classical periodic orbits, if we select some scarring states as the training data, the neural network so trained should be able to make the correct classification when an image containing a similar pattern is presented as the input to the network. However, this simple approach may not be workable for detecting quantum scars because it is difficult to find them in the first place and obtaining a large number of training images is practically impossible. To develop a machine-learning-based automated quantum scar detector, this difficulty must be overcome. We also note that there were previous works on using convolutional neural networks to distinguish chaotic and regular eigenstates [44,50], and it was mentioned in Ref. [51] that a machine learning tool had been developed for automated detection and classification of a class of scars—perturbation-induced scars. However, a general algorithm applicable to large classes of models is lacking.

In this paper, we exploit meta learning [52] to develop an automated and efficient quantum scar detector. The essence of meta machine learning is to encode “previous experience” into a pretrained neural network so that it can quickly adapt to a new input. Meta learning has found broad applications in regression, classification, and reinforcement learning. Our idea is to exploit some existing data sets to train the neural network, which can be completely unrelated to quantum systems but with images similar to scarring patterns. In this regard, the Omniglot dataset [46] widely used in the field of image classification stands out as an excellent choice. Our strategy is to train a “preliminary” version of the neural network that has the ability to distinguish different classes of noisy images from the Omniglot dataset and then perform few-shot classification to further train the neural network with a small number of quantum scars. A pertinent issue is that standard meta learning algorithms can only be used to classify but not detect patterns. Our solution is to use an ensemble of neural networks. In particular, nonscarring states as input to the neural networks can be used as references with significantly different statistical features than those associated with scarring states, thereby accomplishing the detection task. We demonstrate that the neural network so trained can detect and classify quantum scars with remarkable efficiency and accuracy even with limited training data of actual quantum scars. Our meta-learning-based approach not only solves a long-standing problem in the field of quantum chaos, but can also be generalized to address challenging image detection and classification problems in other fields.

We remark that machine learning has been increasingly used in applied physics. For example, in Ref. [53], a neural network was used to characterize the relationship between the gate voltage and performance in a double quantum-dot device to find an optimal combination of the gate voltages to achieve desired performance. In Ref. [54], a convolutional neural network was used to realize single-exposure absorption imaging of ultracold atoms. Our work is not about using a convolutional neural network on some specific realization of a chaotic billiard. Our meta-learning approach addresses the difficulty of sparse training data and is applicable to a vast range of scarred systems with few modifications.

We also remark that robust and accurate detection of quantum scars has practical applications in physics. For example, it has been known for more than a decade [55,56] that, in microlasing systems, a scarring state is ideal for directional emission. In solid-state devices such as quantum dots, quasiscarring states (pointer states) can lead to Fano resonances in the conductance as a function of the Fermi energy [57]. Identifying scarring states in the dot region is key to understanding experimental characteristics such as the conductance-fluctuation patterns. A quite recent experimental work on the chaotic stadium billiard system [58] revealed a large number of quasibound states in the cavity, including at least three quantum scarring states of relatively high period, which can generate different far-field emission patterns. In this work [58], even in a small interval of size 0.02 in the dimensionless quantity (kR), where k is the wave vector and R is the system size, there are about ten quantum states. In the typical range $kR \sim 100$ for a cavity or quantum-dot system, the number of states can be of the order of thousands. Relying on human visualization is not practical. Our machine-learning-based, automated scarring detector provides an effective solution.

In Sec. II, we describe a prototypical relativistic quantum billiard system and articulate the basic training procedures for neural networks. In particular, we detail the construction of our quantum-scar detector based on meta learning (Sec. II B) and few-shot classification (Sec. II C). In Sec. III, we use two types of chaotic Dirac billiard systems to demonstrate the feasibility and power of meta learning in detecting and classifying relativistic quantum scars. Conclusions and discussions are presented in Sec. IV.

II. METHODS: RELATIVISTIC QUANTUM SCARS AND MACHINE LEARNING ALGORITHMS

A. Relativistic quantum scarring in chaotic billiard systems

We use the standard setting of two-dimensional billiard systems [3,59,60] to generate quantum scars. In the classical limit, a particle moves freely inside the billiard,

experiencing elastic reflections from the hard-wall boundary. The geometric shape of the boundary determines the nature of the classical dynamics. For example, the dynamics in a circular billiard are integrable but those in a stadium, Africa, or heart-shaped billiard are chaotic. To be concrete, we focus on the relativistic quantum regime as described by the massless Dirac equation. For a chaotic billiard, there are standard methods to calculate the relativistic quantum eigenvalues and eigenstates such as those based on the boundary integrals [61,62] and conformal mapping [13,63].

We consider the situation where the particle has spin 1/2 and a vertical magnetic flux [14] is applied through the billiard region, whose strength α is effectively an externally adjustable parameter: different sets of spinor wave functions can be generated by changing the value of α . Possessing such an experimentally controllable parameter has two advantages. First, from the point of view of machine learning, the neural network can be trained with one set of wave-function patterns and then be tested using datasets from different parameter values, enabling the full power of the learning process to be revealed and exploited. Second, through testing the datasets from different values of α , the statistical properties of the detected quantum scars can be calculated and compared with the semiclassical prediction [14].

Our working examples are the class of billiards with fully chaotic dynamics in the classical limit [13,63], whose boundary is defined by the following conformal mapping of a unit circle in the complex plane $z \equiv x + iy$ to the complex plane $w \equiv u + iv$:

$$w(z) = \frac{z + bz^2 + ce^{i\delta}z^3}{\sqrt{1 + 2b^2 + 3c^2}}. \quad (1)$$

Here b , c , and δ are parameters, and the origin in the z plane maps to the origin in the w plane. For example, for $b = 0.49$, $c = \delta = 0$, the billiard has the shape of a heart, which possesses an inverse symmetry with respect to the x axis. For $b = c = 0.2$ and $\delta = \pi/3$, the billiard has the shape of Africa. For a massless relativistic spin-1/2 particle inside the billiard, the Hamiltonian is [14]

$$\hat{H} = v_F \hat{\sigma} \cdot \left(\hat{\mathbf{p}} - \frac{q}{c} \mathbf{A} \right) + V(u, v) \hat{\sigma}_z, \quad (2)$$

where v_F is the Fermi velocity, $\sigma = (\sigma_x, \sigma_y)$ and σ_z are the Pauli matrices, the hard-wall confining potential $V(u, v)$ is zero inside and infinite outside the billiard region, and \mathbf{A} is the vector potential. If the magnetic flux is Φ , the vector potential is given by

$$\mathbf{A}(u, v) = \frac{\Phi}{2\pi} \left(\frac{\partial F}{\partial v}, -\frac{\partial F}{\partial u} \right), \quad (3)$$

where $\alpha = q\Phi/(hc)$ is the normalized flux strength and the function $F(u, v)$ is determined by the flux profile in the (u, v) plane. For example, for a singular flux through the origin, $F(u, v)$ can be solved from the Poisson-like equation [64]

$$\nabla_{u,v}^2 F = -2\pi\delta(u, v).$$

The eigenvalue problem $H\psi = E\psi$, where ψ is the two-component spinor, can be solved by employing the conformal mapping method [13,14], where the analytically solvable spinor wave functions in the hard-wall circular billiard in the z plane are transformed to those in the w plane through the conformal mapping. The accuracy of this method depends on the size of the wave-function base. In particular, an arbitrarily large number of analytic eigenstates can be written down in the z plane. When transformed into the w plane, a smaller but still large number of eigenstates can be obtained with a preset, desired accuracy. In our work, for each α value in the unit interval, we use 40 000 base eigenstates in the circular billiard in the z plane. For the heart and Africa billiards in the w plane, we use the first 15 000 and the first 10 000 eigenstates, respectively. The eigenstates are labeled by integer n .

We represent each eigenstate $\psi(n)$ by an image, where the probability distribution for a uniformly spaced, rectangular grid of points is evaluated and the probability is set to zero for points outside the billiard region. For the heart-shaped and Africa billiards, the image sizes are 202×232 and 201×148 , respectively. From the point of view of pattern recognition, the eigenstates can be quite distinct. For example, quantum scarring states, the detection of which is the aim of this work, are those whose probability distribution is concentrated about certain classical periodic orbits and are rare. There are also quantum states with certain patterns that do not correspond to any periodic orbits, such as the boundary states. In fact, the majority of the eigenstates do not have any recognizable patterns.

Quantum scarring states are relatively more pronounced in the semiclassical limit $n \gg 1$. Numerically, for both types of chaotic billiards, we find that the first scarring states, which correspond to classical periodic orbits of period two, emerge at $n \approx 200$. As the period increases, the corresponding scarring states occur in a more semiclassical regime. For example, scarring states corresponding to classical periodic orbits of period four or five begin to arise for $n \approx 1000$. Figure 1(a) presents, for the heart billiard, the quantum states with $n = 8480$ and $n = 8523$ for $\alpha = 0.5$, where the $n = 8480$ eigenstate is localized about a squarelike periodic orbit and belongs to a scarring state of period four, while the $n = 8523$ eigenstate is approximately uniform in the billiard with no apparent pattern that can be deemed a nonscarring state. A few distinct classes of quantum scarring states are shown in Fig. 1(b).

It should be noted that, in a classical chaotic system, the number of unstable periodic orbits increases exponentially

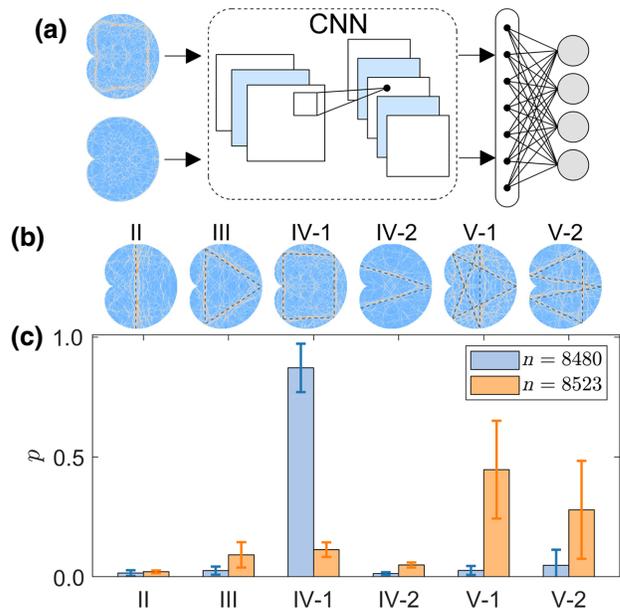


FIG. 1. CNN-based quantum scar detector. (a) Two specific eigenstates (images) from the chaotic heart billiard as the input to the CNN. The upper and lower left states correspond to a scarring and nonscarring state, respectively. The CNN contains convolutional layers and fully connected layers, and the output is an N -dimensional vector whose elements represent the probabilities for the input wave function to belong to the corresponding class. (b) Six types of scarring states in a heart billiard, where the corresponding periods of the classical periodic orbits are denoted by roman numerals. (c) The averages and variances of the output probabilities from a small ensemble of ten neural network realizations.

with the period at the rate of the topological entropy [65]. Within a finite and relatively large period, the number of unstable periodic orbits can be quite large. However, quantum scarring states corresponding to periodic orbits of even moderately large periods tend to be unstable [66]. As a result, the study of quantum scars has mostly been limited to the quantum states that concentrate on classical periodic orbits of low periods.

B. Meta learning

In the physical space (i.e., the billiard domain), the probability distribution of a quantum scarring state focusing on a classical periodic orbit of a low period resembles the image of a simple language character. To develop a machine-learning-based quantum scar detector, we take advantage of this resemblance by exploiting meta learning [52], an image classification algorithm that has attracted a great deal of interest recently. In the field of image classification, an often used and somewhat standard dataset is Omniglot [46], which contains more than one thousand handwritten characters taken from different languages. The basic principle underlying our work is then that a machine

learning algorithm that is able to distinguish the characters in Omniglot can be effectively transferred to detecting quantum scars.

Our specific idea is as follows. Since it is practically infeasible to have a training dataset that contains a large number of quantum scarring states, the problem of identifying quantum scars is similar to the tasks that meta learning algorithms are designed to solve, e.g., image classification for the Omniglot dataset [52]. A meta learning algorithm aims to encode “previous experience” in a pre-trained neural network such that it can quickly adapt to new images. To encode the “previous experience,” training from a large dataset containing many different patterns is necessary. We can thus use the images in Omniglot as the “substituting” training set. Since there is a unique pattern associated with each type of quantum scar, the problem is essentially one of supervised learning. Utilizing certain images from Omniglot as the training data, we can find the images whose patterns correspond to a specific type of quantum scars.

Our meta-learning-based training process consists of three major components. The first component is image processing, as described in Appendix A, where we add noise to each image in the Omniglot dataset to generate a set of images whose patterns resemble those of quantum scars. We then perform image processing for some quantum eigenstates from a chaotic billiard, after which each eigenstate is represented by a matrix of dimension 100×100 with binary elements. The images so created are used as inputs to the CNN.

The second component is to conduct training based on the Omniglot dataset to determine the weights and biases of the CNN. In particular, our CNN consists of four convolutional and two fully connected layers. Each convolutional layer creates a kernel that is convoluted with the layer input to produce a tensor of outputs. The activation function is chosen to be ReLU—the rectified linear unit. After the convolution stage, we apply layer flattening and change the tensor to a large vector. A linear transformation takes this vector to the final output \mathbf{v} , whose dimension is equal to the total number of classes in the system. We use the softmax function to normalize \mathbf{v} :

$$p_i = \frac{\exp(v_i)}{\sum_i \exp(v_i)} \quad (4)$$

where v_i is the i th component of \mathbf{v} and p_i is the normalized probability for the input that belongs to class i . The loss function in the training process is taken to be the cross entropy. The whole process can be implemented in PYTHON with the open-source package Tensorflow [67].

Initially, the weights and biases are randomly chosen from a normal distribution labeled $\theta(w, b)$. In the training process, we first use the modified, noise-contaminated Omniglot dataset to find a pretrained neural network,

denoted $\theta_0(w, b)$, which readily results in different image classes. We use the Reptile algorithm to renew the weights and biases in each step until some criterion is met [68] (see Appendix B for the algorithmic details). To avoid overfitting for the Omniglot dataset, we monitor the test error from each step, as detailed in Appendix D. The training based on the Omniglot dataset is deemed complete when the global minimum is reached.

The third component is training based on certain actual quantum scarring states. In particular, after training with Omniglot that yields a neural network $\theta_0(w, b)$, we perform further training by using the few-shot classification algorithm with a small number of actual images of quantum scars as the training dataset (see Sec. II C below). Since $\theta_0(w, b)$ already has the ability to distinguish different classes of noisy images in the Omniglot dataset, the neural network can quickly adapt to quantum scarring states.

When all three components of the training process are completed, the CNN is effectively a mapping from the wave functions to the probability for each class of scars. Figure 1(c) shows the prediction probability for the two images in Fig. 1(a), where the average is taken from an ensemble of ten neural network realizations (see Sec. III A for details) with different combinations of quantum scars as the training data. Since the particular scarring pattern ($n = 8480$) corresponds to a period-four orbit, a high value of the probability for class IV-1 is achieved. For the non-scarring states ($n = 8523$), while there is a relatively large probability to associate it with a period-five scarring state, the variance is large, making the association highly uncertain. The distinct statistical behaviors for scarring and non-scarring states are an initial indication that our meta-learning-based method has the ability to detect quantum scars.

C. Few-shot classification

Physically significant quantum scarring states that have been studied in the past are typically those corresponding to classical periodic orbits of low periods and thus consist of a few classes. From the point of view of image processing, the number of classes to be classified is small and there are only a few images available in each class. The few-shot classification algorithm is designed to deal with such problems [45–47]. Suppose that the target contains N different classes and that each class has at least $K + 1$ images with K being a small integer. We randomly choose $K + 1$ images in each class, train the neural network using the first K images, and then use the trained network to classify the last image. For the Omniglot dataset, a 95% accuracy can be achieved even for $N = 5$ and $K = 1$, and the accuracy can reach 99% if K is increased to $K = 5$. In general, for a fixed N value, increasing the value of K can improve the accuracy. However, for a fixed K value, a larger value of

N leads to a lower accuracy as more classes require more images to train the network.

To test the power of few-shot classification in detecting quantum scars, we take the heart billiard as an example, where six distinct classes of scars are shown in Fig. 1(b). We create an image dataset where each class contains ten images taken from the eigenstates whose energy index $n \in [1000, 15\,000]$. We obtain the pretrained neural network θ_0 by using meta learning and conduct further training using the actual quantum-scar images. To be concrete, we fix $N = 6$ and vary K . For $K = 1$, we randomly choose two images from each class, where the first and second images are used for training and testing, respectively. The training and testing datasets are thus extremely small. To obtain the accuracy, we repeat this process (i.e., training θ_0) 1000 times and calculate the ratio of the number of correctly labeled images to the total number of test images in all the trials. The average ratio and the corresponding variance are obtained by using ten realizations of the neural network θ_0 . Table I lists the accuracy versus the shot number K , where a 90% accuracy is achieved even for $K = 1$. The somewhat lower accuracy than that of classifying the characters in the Omniglot dataset is due to the “noisy” nature of the images of quantum scars, where the scarring patterns are not sharp but have local fluctuations. As the shot number K increases, the scar-detection accuracy improves.

Another quantity characterizing the detection performance is the confusion matrix whose diagonal and off-diagonal elements are the probabilities of correct and incorrect detection of different classes of quantum scars, respectively. Figure 2 shows, for $K = 1$, the confusion matrix for the six classes of quantum scars from the chaotic heart billiard, where the roman letters indicate the periods of the underlying classical periodic orbits and the arabic numbers represent distinct configurations of the orbits. For five out of the six classes (except type V-1), the detection accuracy is close to 100%. The accuracy for type-V-I scars is somewhat lower, where occasionally the machine learning algorithm would erroneously classify such a scar as belonging to class II or III. From Fig. 1(b), it can be seen that this error may be expected for two reasons: (1) the relatively long period of the corresponding classical periodic orbits leads to a relatively low concentration of the spinor wave function around the orbits and (2) the type-V-1 scars have approximately the same edges as those in class II or III.

TABLE I. Few-shot classification accuracy of detecting relativistic quantum scars in the chaotic heart billiard.

Shot number	Accuracy
$K = 1$	$90.16\% \pm 1.41\%$
$K = 2$	$95.30\% \pm 0.68\%$
$K = 5$	$98.58\% \pm 0.32\%$

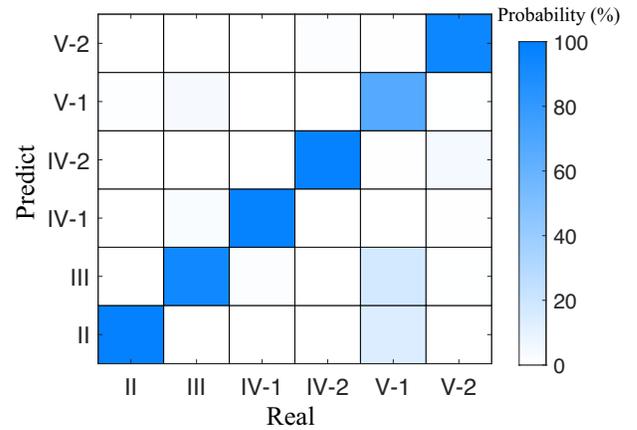


FIG. 2. Confusion matrix associated with detection of quantum scars in the chaotic heart billiard. The horizontal and vertical axes are the labels of the real and detected classes, respectively. The various probabilities are color coded as defined by the color bar on the right, where the diagonal and off-diagonal elements are the probabilities of correct and incorrect detection of different classes of quantum scars, respectively. The detection accuracies (the values of the diagonal elements) are generally close to 100% except for the V-1 scars—see the text for an explanation.

III. RESULTS

We use two types of chaotic Dirac billiard systems to demonstrate the feasibility and power of meta learning in detecting and classifying relativistic quantum scars.

A. Chaotic heart billiard

In Sec. II C, we have described the idea of adapting the method of few-shot classification to detect quantum scars from a chaotic billiard. For a target set of six distinct scars in the chaotic heart billiard, the one-shot accuracy is close to 90%. Here we extend the method to finding desired quantum scars from a large number of wave functions.

There can be a large variety of spatial distributions of the eigenstates of a chaotic Dirac billiard. Visually, most eigenstates have a random distribution and they do not exhibit any discernible pattern. For those eigen-wavefunctions that do, they can be quantum scars associated with periodic orbits of high periods or boundary states. From the point of view of machine-learning-based detection of scars, such a dataset is unbalanced because the number of nonscarring states is much larger than that of the actual scarring states. In computer science, this problem associated with object detection is well known [69], where the number of “real” objects in the training dataset is extraordinarily small. An intuitive solution is to treat the nonscarring states as a new class of objects when implementing the few-shot classification algorithm, but there are limitations due to the unbalanced nature of the dataset.

We develop an alternative approach to addressing this issue of an unbalanced dataset. After the CNN is properly

trained, we input an additional quantum nonscarring pattern to the neural network. Ideally, the prediction should be robust, meaning that the normalized probability should have a high concentration about this pattern. Since this nonscarring image does not belong to any of the classes that the neural network has learned from the training process, when other eigenstates are inputted to the neural network, the overall prediction probability will be lower. Another feature that can be exploited, as can be seen from Fig. 1(c), is that the variances of the output probabilities with nonscarring states as the input are typically large. Our general quantum-scar detector is built on these two features.

For the chaotic heart billiard, the training data are the eigen-wave-functions generated at $\alpha = 0$, where two quantum scars are chosen for each training class: one at low energy ($n \approx 2000$) and another at high energy ($n \approx 10^4$). We perform one-shot classification and train the neural networks. Since each class contains two images, there are 2^6 possible neural networks. The testing data consist of eigen-wave-functions from $\alpha = 0.25$ or $\alpha = 0.5$. For a quantum state n inputted to the i th neural network, the output is denoted as p_j , where j is the scar index. We multiply the results from different neural networks to find the scarring state that maximizes the product:

$$\xi(n) = \max_j \left(\prod_i p_j^{(i)}(n) \right). \quad (5)$$

As an illustrative example, we consider two quantum states: $n = 8480$ (a scarring state) and $n = 8523$ (a non-scarring state) for $\alpha = 0.5$. The average and variance of p for the distinct scar indices are shown in Fig. 1(c). We use an ensemble of ten neural networks for the one-shot ($K = 1$) scheme. For the scarring state $n = 8480$, the predicted p value concentrates on the corresponding type of scar, where the product in Eq. (5) achieves maximum for the IV-1 scar class with $\xi(n = 8480) = 0.23$. For the non-scarring state, the p value does not concentrate at any scar index and the variance from different neural networks is large. We get $\xi(n = 8523) = 4.3 \times 10^{-5}$. There is then a 4-order-of-magnitude difference in the ξ values for the scarring and nonscarring states, rendering accurate and reliable detection of the scarring state.

To detect the scarring states from all available quantum states for $\alpha = 0.5$, we use the ensemble of neural networks and sweep through the states. Figure 3(a) shows $\xi(n)$ versus n for $n \in [4000, 4100]$, where a nearly 8-order-of-magnitude difference in the $\xi(n)$ values emerges among the quantum states. Setting a threshold, e.g., at 5×10^{-3} (the horizontal dashed line), we deem those states whose $\xi(n)$ is above the threshold as quantum scars, as verified by the corresponding distinct scarring patterns above Fig. 3(a). Similar results have been obtained for $n \in [14\,000, 14\,100]$, as shown in Fig. 3(b). The results

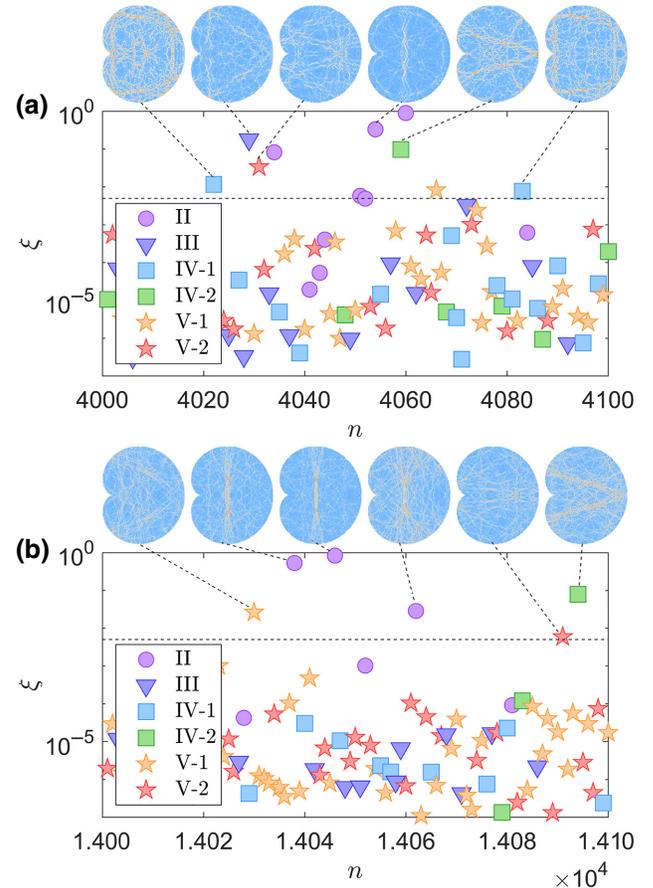


FIG. 3. Emergence of a meta-learning-based quantum scar detector. Shown is $\xi(n)$ versus the eigenstate index n for the chaotic heart billiard for $\alpha = 0.5$. (a) The results are for $n \in [4000, 4100]$ (a relatively low-energy interval), where the scar types are distinguished by the different colored symbols and the horizontal dashed line at 5×10^{-3} indicates the detection threshold. The eigenstates with $\xi(n)$ values above the threshold are deemed to be quantum scars. Some representative scarring states detected are displayed above the panel. The agreement between the detected states and the ground truth is perfect. (b) Same as (a) but for $n \in [14\,000, 14\,100]$ —a higher-energy interval. The empirical threshold is chosen so that approximately 5% of the eigenstates belong to quantum scars.

in Fig. 3 thus indicate the viability of a meta-learning quantum scar detector based on evaluating the values of $\xi(n)$.

To verify if the scarring states found are truly relativistic quantum scars, we exploit a criterion from semiclassical theory. In particular, the semiclassical theory for relativistic quantum billiard systems [13,14] predicts that the recurrent interval in the wave vector must obey the rule

$$\Delta k = \begin{cases} 2\pi(\Delta n - 2W\alpha)/L, & \text{even bounces,} \\ 2\pi(\Delta n - 2W\alpha + \Delta\beta)/L, & \text{odd bounces,} \end{cases} \quad (6)$$

TABLE II. Statistics of quantum scars in the heart billiard system detected by machine learning.

Scar index	Δk	k_0	Number
II ($\alpha = 0.25$)	1.4810	190.9758	271
II ($\alpha = 0.5$)	1.4810	199.8387	303
III ($\alpha = 0.25$)	1.1687	197.3900	61
III ($\alpha = 0.5$)	1.1687	199.4852	51
IV-1 ($\alpha = 0.25$)	1.0843	197.4324	26
IV-1 ($\alpha = 0.5$)	1.0843	192.9035	40

where L is the length of the classical orbit associated with the scarring state and \mathcal{W} is the winding number. For the scarring states corresponding to classical periodic orbits with an odd number of bounces or reflections from the billiard boundary, the quantity $\Delta\beta$ takes on the value $1/2$, so the spacing in the wave vector for such scarring states to occur is π/L .

Table II presents the statistics of different types of quantum scarring states detected. The simulation setting is as follows. The training datasets are the same as those described in Sec. II C, which are associated with zero magnetic flux $\alpha = 0$. Ten neural networks are trained for one-shot classification. The test datasets are associated with $\alpha = 0.25$ and $\alpha = 0.5$ with states whose level index ranges from $n = 4000$ to $n = 15\,000$. The types of scarring states in Table II are those corresponding to classical periodic orbits II, III, and IV-1. The reference wave vector k_0 is chosen from a scarring state with the maximum value of $\xi(n)$ for $n \in [9000, 10\,000]$. For example, for orbits II, there are approximately 300 scarring states. Out of the 11 000 available eigenstates in the chaotic heart billiard, approximately 5% are quantum scars. To verify that the machine-detected scarring states are true quantum scars, we select a few energy intervals and visually identify the scarring states. The percentage of scarring states identified this way is consistent with that of the scarring states found by machine learning.

A convenient semiclassical quantity to characterize the recurrence of a quantum scarring state is the winding number [13]

$$\eta = \frac{|k_n - k_0|}{\delta k} - \left\lfloor \frac{|k_n - k_0|}{\delta k} \right\rfloor, \quad (7)$$

where $\delta k = 2\pi/L$, $[x]$ denotes the largest integer less than or equal to x , and k_0 is the wave vector of a reference scarring state (usually of high energy). Figure 4 shows the value of η for a large number of scarring states of three types: type II, type III, and type IV-1, where type-II scars are the most abundant. As shown in Figs. 4(a) and 4(b) for $\alpha = 0.25$ and $\alpha = 0.5$, respectively, the winding numbers of type-II scars are either zero or one. For this type of scar, because of the even bounce numbers of the underlying classical orbits off the billiard wall, it is not

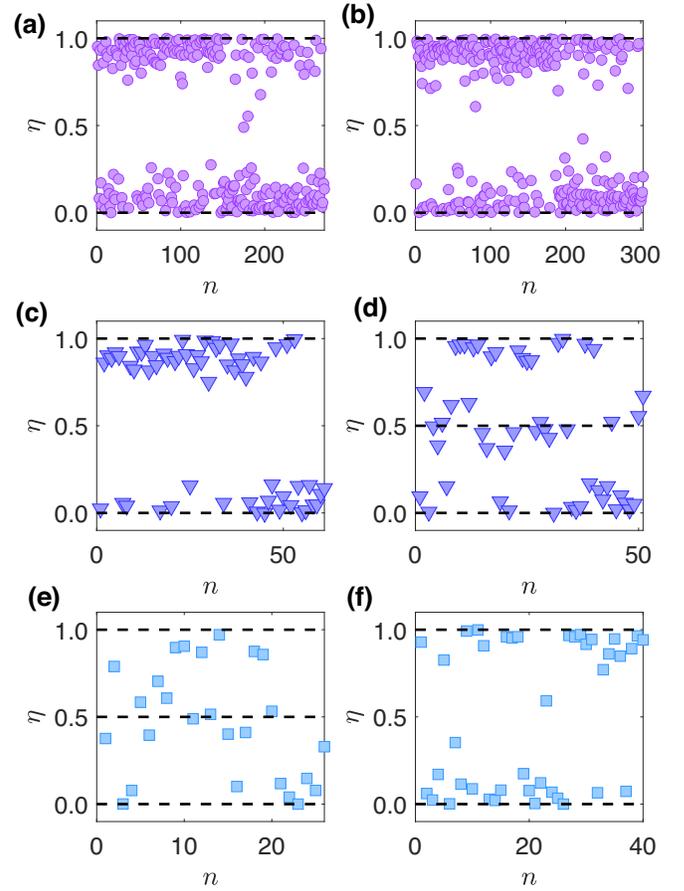


FIG. 4. Semiclassical characterization of the relativistic quantum scarring states detected by the meta-learning algorithm. Shown are the values of the semiclassical winding number η for three types of scars, where the top, middle, and bottom rows correspond to type-II, type-III, and type-IV-1 scars, respectively. A magnetic flux is present: (a),(c),(e) $\alpha = 0.25$ and (b),(d),(f) $\alpha = 0.5$. A geometric-phase-based semiclassical theory [14] stipulates that, for all types of scars, the values of η can be zero or one, and $\eta = 0.5$ is not possible for type-II scars, but can occur for type-III (type-IV-1) scars under magnetic flux $\alpha = 0.5$ ($\alpha = 0.25$). The various quantum scars detected by the machine learning algorithm obey these rules remarkably well, signifying high detection accuracy.

possible for the quantum states to have $\eta = 1/2$. For type-III scars, the winding numbers are one for $\alpha = 0.25$, as shown in Fig. 4(c). However, because of their odd bounces with the billiard wall, scarring states with $\eta = 0.5$ can arise for $\alpha = 0.5$, as shown in Fig. 4(d). For type-IV-1 scars, $\eta = 0.5$ occurs for magnetic flux value $\alpha = 0.25$ but not for $\alpha = 0.5$, as shown in Figs. 4(e) and 4(f), respectively. These features of the relativistic quantum scars have been understood theoretically based on a detailed analysis of the geometric phases [14]. It is remarkable that our meta-learning-based algorithm can reliably detect the different types of scarring states with the correct semiclassical characteristic features.

A few remarks are in order. First, from Fig. 3, it can be seen that most detected quantum scars are associated with classical periodic orbits of low periods, due to the fact that these scarring states are more pronounced than those corresponding to periodic orbits of high periods. This is consistent with the semiclassical theory of quantum scars [5,6].

Second, we use an ensemble of neural networks to combine the accuracy measures, where each is trained based on meta learning with one-shot classification. The training thus requires more than one image per class. For two images per class, we can generate 2^6 combinations, so the training data are sufficient for scar detection. In simulations, we find that choosing two scars with one at relatively low energy and another at high energy can help diversify the neural network outputs to increase the detection accuracy.

Third, Fig. 3 demonstrates the working of the scar detector in the energy interval with the level index ranging from 4000 to 15 000. In a lower-energy interval, e.g., $n \approx 1000$, the degree of quantum scarring in terms of the wavefunction concentration is relatively weak. The difference in the values of $\xi(n)$ between scarring and nonscarring states is especially small, making detecting scars difficult. However, this is expected as quantum scars are more pronounced in the relatively high-energy regime, the so-called semiclassical regime where both quantum and classical behaviors are relevant [5,6].

Fourth, an advantage of the neural-network-based scar detector is that all types of scars can be detected based on a single threshold, as exemplified in Fig. 3. A large threshold means that a small number of states can be found with high accuracy. Decreasing the threshold can lead to more “scarring” states but the accuracy may be compromised. In principle, we can still pick up only those states with the highest ξ values as the scarring states.

B. The Africa billiard

The Africa billiard system has fully chaotic dynamics in the classical limit and exhibits relatively more complicated quantum scarring states than the chaotic heart billiard. Figure 5(a) shows eight types of scars. We extend the few-shot classification algorithm by setting $N = 8$ and selecting 10 images for each class for $\alpha \in [0, 0.1]$ from our quantum-scar dataset. By varying K and repeating the process 1000 times, we obtain the average accuracy, as listed in Table III. Because of the large number of scarring states and the relatively more sophisticated geometric shape of the Africa billiard (than the chaotic heart billiard), the accuracy for one-shot classification is somewhat lower: less than 80%. The accuracy can be improved by increasing the number of images for training.

Figure 5(b) shows the confusion matrix under one-shot classification, where the large diagonal values indicate

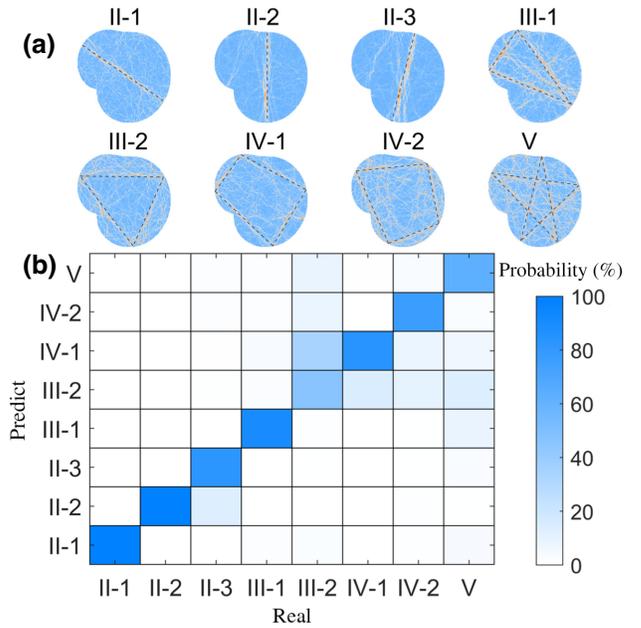


FIG. 5. Types of quantum scarring states in the chaotic Africa billiard and confusion matrix of classification. (a) Eight different types of scars in the Africa billiard and (b) the confusion matrix between the true and predicted labels from one-shot classification. The diagonal elements have large values, indicating high overall accuracies of the predictions of the neural network. Some off-diagonal elements are non-negligible, e.g., those between type-II-2 and type-II-3 scars.

that most predictions are correct. Nonetheless, some off-diagonal values are not negligible. For example, the algorithm appears to be “easily confused” by type-II-2 and type-II-3 scars, because the concentration patterns of their spinor wave functions are similar, as shown in Fig. 5(a). In fact, the different kinds of type-II scars all have a straight line pattern and this leads to the confusion. We also note that the accuracy for type-III-2 scars is the lowest, due to their similarity to the type-V scars.

We build a scar detector based on two-shot classification, due to its reasonable accuracy, as shown in Table III. The number of images per class in the training set is set to be three and the scarring states used for training are associated with the magnetic flux in the interval $\alpha \in [0, 0.1]$. Figure 6 shows $\xi(n)$ for $n \in [6100, 6200]$ and $\alpha = 0.25$. We set the threshold to be 2×10^{-3} and those

TABLE III. Few-shot classification accuracy for scarring states in the chaotic Africa billiard.

Shot number	Accuracy
$K = 1$	$78.14\% \pm 1.64\%$
$K = 2$	$85.44\% \pm 1.15\%$
$K = 5$	$93.00\% \pm 0.96\%$

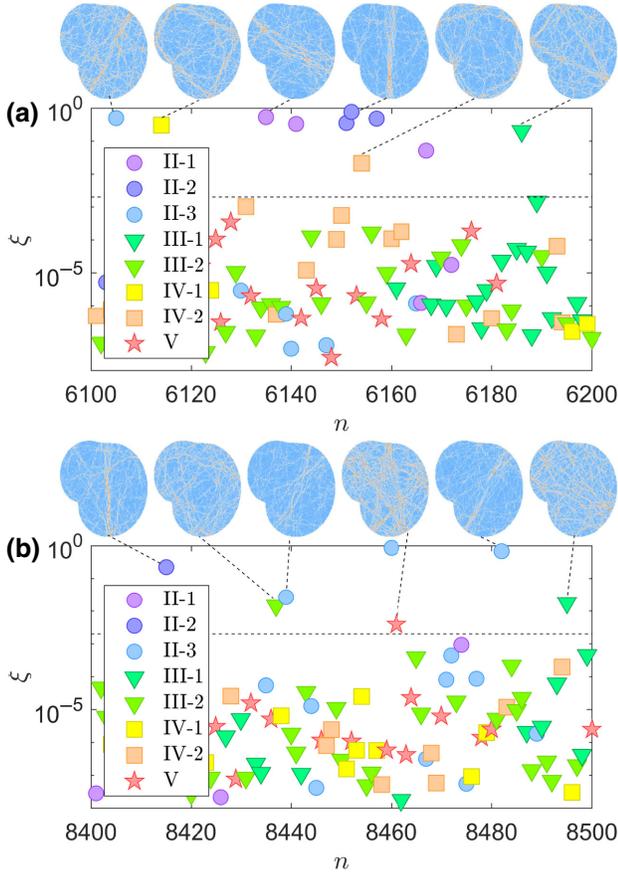


FIG. 6. Machine-learning detection of quantum scars in the chaotic Africa billiard. Shown is ξ versus the mode index for $\alpha = 0.25$ for (a) $n \in [6100, 6200]$ and (b) $n \in [8400, 8500]$, where the horizontal dashed line indicates the detection threshold 2×10^{-3} . Points above the threshold correspond to the detected scarring states with some representatives shown above the panel.

quantum states whose ξ values are larger than the threshold are deemed scarring states. The detection results for $n \in [8400, 8500]$ are shown in Fig. 6(b). As in the case of the chaotic heart billiard, most detected quantum scars correspond to classical periodic orbits of low periods. We then calculate the η values for each type of scar from Eq. (7), where k_0 is chosen to be associated with the eigenstate with the largest $\xi(n)$ value for $n \in [9000, 10000]$. The results for some representative scars detected from the energy-level range $n \in [4000, 10000]$ are shown in Fig. 7 with their basic properties listed in Table IV. For all scars displayed (for $\alpha = 0.25$), the η values are either zero or one, except for the type-IV-2 scars where $\eta = 1/2$ can occur, which agrees well with the semiclassical prediction [14].

IV. DISCUSSION

In spite of the fundamental importance of the phenomenon of quantum scarring, identifying such states from a large number of eigenstates has been a laborious

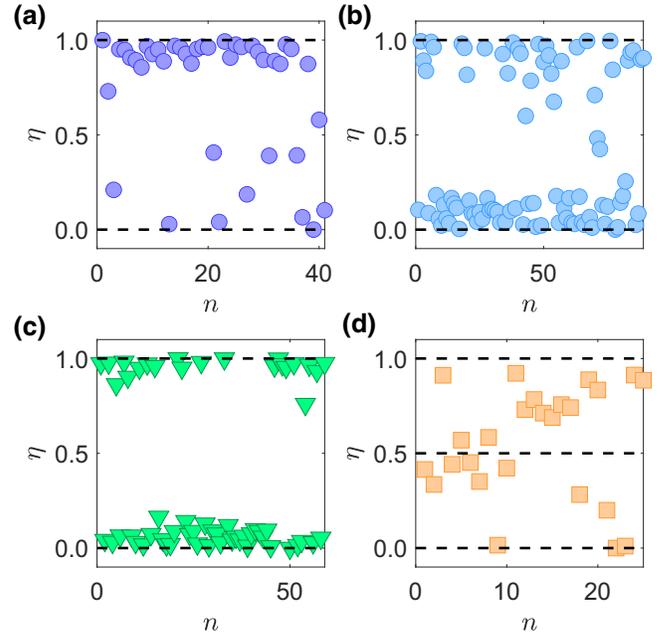


FIG. 7. Semiclassical quantification of the machine-learning-detected relativistic quantum scars in the chaotic Africa billiard. Shown is the semiclassical winding number η for scars of types (a) II-2, (b) II-3, (c) III-1, and (d) IV-2. Because of the presence of a magnetic flux, only IV-2 scars can have $\eta = 0.5$. The statistical behaviors of the detected quantum scars agree well with the semiclassical predictions.

task since the beginning of the field of quantum chaos, relying mostly on going through all the available wavefunction patterns one after another and performing a visual check of each pattern. This task has been well known to practitioners of quantum chaos. To develop a fully automated method to accurately detect quantum scars is thus highly desired, as it will enable a more systematic and comprehensive study of the exotic quantum states. This paper accomplishes this goal by developing a machine-learning-based quantum-scar detector. While it may seem straightforward to adopt some conventional deep neural networks for image recognition to the problem of quantum scar detection, a significant challenge lies in the requirement of extensive training data, as quantum scattering states are rare. We meet this challenge by articulating a meta-learning approach based on few-shot classification of

TABLE IV. Characteristics of representative relativistic quantum scars in the Africa billiard.

Scar index	Δk	k_0	Collected number
II-2	1.6558	198.1712	41
II-3	1.6427	195.5084	89
III-1	1.2313	191.8086	59
IV-2	1.0699	191.5694	25

quantum scarring states, which allows the neural machine to learn the hidden structure of each class and quickly adapt to the new class even with quite a small number of images of the scarring states. We test this approach using two paradigmatic relativistic quantum billiard systems subject to a magnetic flux, which exhibit fully developed chaos in the classical limit. Using a previously developed conformal-mapping method [13,14] to calculate a large number of quantum states in each system, we demonstrate the power of our meta-learning method for accurate detection of quantum scars.

Our fully automated quantum-scar detection framework consists of two major steps. The first step is to use neural networks to classify quantum scars. In particular, we apply few-shot meta learning by designating a number of images per class in the training dataset, e.g., a few different classes of quantum scarring states in a chaotic Dirac billiard. For one-shot classification, the achieved accuracy can already be about 80%, which can be improved by increasing the number of shots. The overall classification accuracy can be characterized by the confusion matrix. In general, the accuracies for quantum scars associated with classical periodic orbits of short periods are higher than those with longer periods. The next step is to exploit the few-shot classification algorithm to detect quantum scars by monitoring the predicted probability difference between scarring and non-scarring states. For quantum scars, due to the high few-shot classification accuracy, the probability distributions are well localized and large with small variances, but the opposite occur for non-scarring states. Utilizing an ensemble of neural networks allows us to define a statistical measure, whereas a quantum eigenstate can be faithfully deemed as a scarring state if this measure exceeds a well-defined threshold. Utilizing a quantity derived from the semiclassical theory, we demonstrate that the quantum scars detected are the correct scarring states. Our meta-learning-based quantum-scar detector is thus accurate, efficient, and fully automated.

We discuss a few pertinent issues. First, to build a general machine-learning-based quantum-scar detector, it is necessary to choose some representative quantum scars as the training dataset. The training accuracy can depend on the specific scarring states, e.g., whether they are from the same or different energy regimes. For example, for the chaotic heart billiard system, two scars per class are needed for training. We find that, if one scar is chosen from the low-energy regime and another from the high-energy regime, a high accuracy can be achieved. The intuitive reason is that quantum scarring states from drastically different energy regimes tend to be maximally distinct in their patterns, thereby enhancing the learning capability of the neural network.

The second issue concerns the number of images per class required for the few-shot classification algorithm. Accuracy can be improved by increasing the number

of such images, but the training dataset also becomes larger, thereby increasing the computational complexity. Our empirical experience is that the number of images should be chosen such that the classification accuracy is about 90%. For example, for the chaotic heart billiard with six classes of quantum scars, using one image per class in the training dataset suffices. For the chaotic African billiard, two-shot classification is needed to achieve the desired accuracy. For an ensemble of neural networks, various combinations of the training data are needed. As a result, the number of images per class should be at least one more than the shot number.

The third issue concerns the non-scarring states. An alternative idea is to treat them as another class of “scars.” However, non-scarring quantum states tend to be uniform and do not typically possess any unique feature. Treating them as a class of scarring states can lead to large classification errors.

The fourth issue concerns overfitting in the meta-learning algorithm. In our work, the neural networks are trained based on the Omniglot dataset and transferred to quantum scars. If a neural network performs well in the Omniglot dataset, when executing transfer learning to quantum scarring states, overfitting can arise, reducing the classification accuracy. Our empirical method is to monitor the test error to ensure that it does not exceed a certain (small) threshold.

The fifth issue concerns the necessity of using machine learning for automated detection of quantum scars. Can some filtering techniques be used, e.g., based on counting the probability value of the quantum wave function along a classical periodic orbit? This is in general infeasible because, from an image processing point of view, quantum scars are extremely noisy. Moreover, setting a proper region to estimate the probabilities requires a threshold, which can be difficult as the degree of localization of the scarring wave function depends on the energy. Even worse, the wave function associated with a quantum scarring state may not follow a classical periodic orbit exactly. To our knowledge, exploiting machine learning for quantum-scar detection represents the best strategy at present.

We note that, in Refs. [70,71], the Weyl entropy and Husimi Q value were used to detect quantum scars for a one-dimensional system, where a characterizing quantity was defined based on wave-function localization and a large value indicates a scar. It is difficult to generalize this method to two dimensions for the following reasons. First, for a chaotic cavity, the eigenstates can be found only by numerical means, making it difficult to compute the density matrix. Second, the computation complexity associated with discretizing 10^4 eigenstates in two dimensions can be formidable. Third, localization in two dimensions can have a linear structure, making it difficult to detect scars using a pointwise measure.

Quantum scars can also arise in open systems. For example, in the microlasing system, there can be quasi-bound states that are strongly localized around classical unstable periodic orbits [55,56]. Such states have strong chaos-assisted directional emission. To extend our method to open quantum systems, two modifications are necessary. First, the features should be maximized because, in open quantum systems, the wave function inside the cavity is stronger than that outside. Second, the data distribution between the revised Omniglot dataset and quantum scars should be the same. A possible method is to generate a binary representation of the wave function by setting an appropriate threshold to ensure that the revised Omniglot dataset and quantum scars have the same number of ones.

Taken together, in the field of quantum chaos, finding quantum scars usually relies on human visualization. To develop a machine learning framework typically requires extensive training data, but quantum scarring states are rare, posing a significant challenge. Taking advantage of meta learning, a special class of machine learning for image recognition and classification, we develop a fully automated quantum-scar detector and demonstrate its working with a remarkably small number of scar images for training. Our detection algorithm requires no fine tuning of the parameters of the neural networks and is effective in a broad energy range that contains many thousands of eigenstates. The framework developed in this paper can be readily generalized and extended to solving difficult image processing problems in other disciplines of science and engineering.

The data and code associated with this work are available online [72].

ACKNOWLEDGMENTS

This work is supported by AFOSR under Grants No. FA9550-21-1-0438 and No. FA9550-21-1-0186.

APPENDIX A: DATA PROCESSING

We use the Omniglot dataset to obtain the basic neural network before training with quantum scars, which contains more than one thousand handwritten characters from different languages [46]. While some quantum scars resemble certain characters, their details can be quite different. The left panel in Fig. 8(a) shows the Greek character α in the Omniglot dataset, which has a well-defined structure in that each stroke has a similar width. The left panel of Fig. 8(b) shows a period-four quantum scar, which differs significantly from the α character and is noisy. To use the Omniglot dataset to train the neural networks for detecting quantum scars, image preprocessing is necessary to reduce the difference in details.

An image is represented by a matrix defined on a grid of pixels, so a white pixel can be defined as zero and a black pixel is represented by one. Our data processing

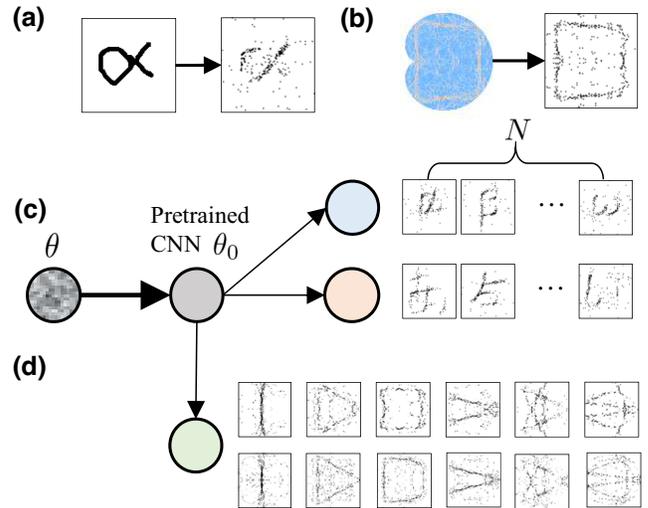


FIG. 8. Schematic illustration of image processing. (a) Left: a representative image in the Omniglot dataset with the Greek character α of size 105×105 pixels. Right: the image after the transform. (b) Left: a quantum scar image of size 202×234 from the chaotic heart billiard. Right: the image after data processing and resizing. With image processing, the detailed patterns of the Omniglot and quantum-scar images become similar. (c),(d) Two steps of the meta-learning algorithm, respectively. The first step (c) is to start from a neural network with random weights and biases θ to find the optimal network denoted as θ_0 that quickly adapts to randomly selected classes in the Omniglot dataset. The second step (d) is to perform training with quantum scars. Because of the similarity between the processed Omniglot character and quantum-scar images, adaptation of the neural network θ_0 to quantum scars can be achieved with only a small number of images.

for each image in the Omniglot dataset consists of the following steps: (i) converting an image to a matrix of dimension 108×108 , (ii) adding a uniform noise between 0 to 1 for each element, (iii) smoothing the image with low-frequency filtering, (iv) setting the maximum 4% of points in the matrix as black and other points as white, and (v) using the PIL package in PYTHON to resize the image to 100×100 . For low-pass filtering, we use fast Fourier transform (FFT) to transform the matrix into a coefficient vector, one component for each frequency, set the coefficients for the top 10% of the frequency range to zero, then perform the inverse FFT. To process the quantum scarring images, we first calculate the density $\psi = \psi_1^2 + \psi_2^2$ of the spinor wave function, assuming that the wave function outside the cavity is zero. We then set the points with the top 4% density values to be one (black) and the other points to zero (white), and use the PIL package in PYTHON to resize the image to 100×100 . The right panels in Figs. 8(a) and 8(b) show the images after processing for the α character and the period-four quantum scars, respectively. It can be seen that the processing has resulted in similar patterns

for the Omniglot character and quantum scar. A heuristic reason for this similarity is that the neural network trained with the Omniglot dataset performs classification by extracting line segments as features. With image processing, both the Omniglot characters and quantum scars contain line segments as features.

APPENDIX B: META-LEARNING ALGORITHM

There are two categories of meta-learning algorithms: metric [47] and optimization [52,68] based. We adopt the optimization-based approach, where θ denotes the training parameter set in the convolutional neural network. We aim to find a pretrained neural network, denoted θ_0 , such that the loss is minimized after operation $U_\tau^k(\theta)$ — k steps of gradient descent based on task τ . Let S be the cross-entropy loss. The goal can then be formulated as $\min_\theta S[U_\tau^k(\theta)]$ for all τ .

Different methods are available to solve the optimization problem. One is model-agnostic meta learning (MAML) [52], which can give higher accuracy. The method requires computing the gradient for the iterative equation $U^k(\theta)$, which is time consuming when $k > 1$. An approximate method is Reptile [68]. For the Omniglot and Mini-imagenet datasets, Reptile is computationally efficient, but the accuracy is slightly compromised in comparison with that of MAML.

Figures 8(c) and 8(d) respectively illustrate the two basic steps involved in our meta-learning procedure. Starting from a neural network with randomly generated weights and biases, denoted θ , we use the images from the Omniglot dataset to train the network to find an optimal network, denoted θ_0 , as shown in Fig. 8(c). In the second step [Fig. 8(d)], we start from θ_0 and perform training with quantum scars. Since the Omniglot dataset is large, θ_0 , which has been trained to adapt to any image class in the dataset, can quickly adapt to quantum scars.

The Reptile method is illustrated in Algorithm 1 below. Initially, all the weights and biases are randomly chosen. A loop is employed to update θ until the desired neural network θ_0 is found. In each epoch, the network θ is updated as follows. The first step is to randomly generate N classes from the processed Omniglot dataset, where each class contains K images, so the labeled dataset has NK images.

```



---


Initialized  $\theta$ 
for iteration = 1, 2, ... do
  Sample task  $\tau_1, \tau_2, \dots, \tau_m$ 
  for  $i = 1, 2, \dots$  do
    compute  $W_i = U_{\tau_i}^k(\theta)$ 
  end for
  Update:  $\theta \leftarrow \theta + \beta \frac{1}{m} \sum_{i=1}^m (W_i - \theta)$ 
end for


---



```

Algorithm 1. Reptile—batched version.

The second step is to generate five tasks, each containing ten images that are randomly chosen and can be repeated. The third step is to perform stochastic gradient descent (SGD) and Adam for each task. Let $U^k(\theta)$ denote the stochastic gradient updating for k steps on cross-entropy loss with the initial parameter set θ , which returns the final parameter set. During this training process, we set $k = 5$. The fourth step is to update θ based on the average. The learning rate β decays linearly with the number of epochs, e.g., $\beta = 10^{-3}$ for the first epoch and $\beta = 0$ for the last epoch. After each epoch, θ is saved and SGD is performed for quantum-scar images. The testing accuracy can then be calculated. The cutoff point can be found based on hyperparameter optimization as described in Appendix D.

APPENDIX C: NEURAL NETWORK ARCHITECTURE

The convolutional neural network used in our study is illustrated in Fig. 9, which contains four convolutional layers. Each layer is followed by batch normalization (BN) to avoid overfitting [73]. We use ReLU as the nonlinear activation function, which is applied after the batch normalization. Each convolutional layer contains 64 kernels of size 3×3 . When applying convolutional kernels, we use stride 2 and zero padding. The whole neural network contains approximately 10^5 training parameters.

APPENDIX D: HYPERPARAMETER OPTIMIZATION

Since the neural network with the parameter set θ is optimized with the processed Omniglot dataset, too many training steps can lead to overfitting for the Omniglot dataset and reduce the training accuracy for quantum scars. Our solution is to monitor the test error at each epoch. Specifically, for each epoch, we save θ , perform a few-shot classification for quantum scars, and calculate the classification error. The error can have large fluctuations due to the small quantum-scar testing set, so we take the average over 200 steps. Figures 10(a) and 10(b) show the classification error versus the number of epochs for the quantum scars in the chaotic heart and African billiards,

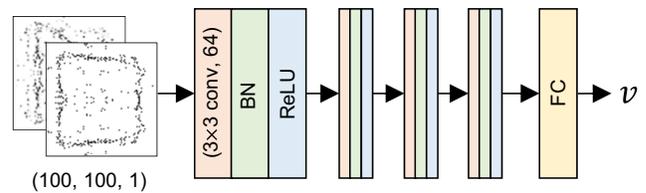


FIG. 9. Detailed structure of the CNN used in our study. The CNN contains four convolutional layers, each layer having a 3×3 convolutional kernel with stride equal to 2 and zero padding, followed by ReLU and batch normalization.

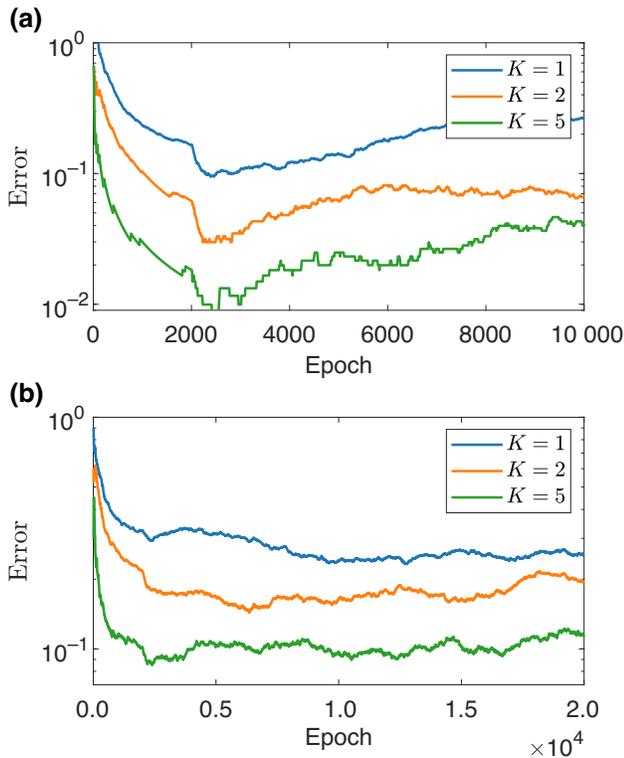


FIG. 10. Few-shot classification errors versus the number of training epochs. The error is the average from a moving window of 200 epochs. (a),(b) The results for the quantum scars from the chaotic heart and Africa billiards, respectively. The cutoff point is determined by the empirical criterion that the testing error reaches the minimum. Further updating the neural network can lead to overfitting.

respectively. For quantum scars from the heart billiard, the number of classes is 6, the total number of iterations is 10^4 , and the error reaches a minimum for about 2000 steps for all K values. The cutoff point is then set to be 2000. For quantum scars from the Africa billiard, the number of classes is 8, the total number of iterations is 2×10^4 , and the cutoff points for $K = 1, 2$, and 5 are 10^4 , 5000, and 2000, respectively. When performing the hyperparameter optimization, the number of available quantum scars is also limited. In our computations, the number of images per class in the quantum scar dataset is fixed to be $K + 1$ to achieve data consistency.

APPENDIX E: META LEARNING VERSUS TRANSFER LEARNING

The dataset of quantum scarring images typically contains a small number of classes, which is the main reason that we choose to exploit meta learning. In machine learning, transfer learning has also been widely used. It is important to differentiate the two types of learning.

Indeed, meta learning and transfer learning share some common features: both requiring a two-step learning process and the trained neural network can then be applied to different datasets. However, there are significant differences between the two methods, which can be understood in a concrete manner by examining the classic problem of image classification.

In transfer learning, the procedure is to train a convolutional neural network on the first dataset to determine the neural network with the maximal validation or testing accuracy. The neural network is then fine-tuned using the target dataset. Note that, even without the target dataset for training, the neural network has already been well trained using the first dataset and has “learned” the rules to extract features. When a different dataset is presented to the neural network, it uses the already learned rules to classify the images. Typically, the datasets in transfer learning are required to be similar in terms of their sizes and the number of classes contained therein.

In meta learning, the task is to classify images in a dataset that has a small number of classes. Training a neural network using such a dataset typically leads to overfitting. The solution is to first identify a *pretrained* neural network according to the criterion that it has the “potential” to be further trained for classifying the images in the target dataset. The chosen network is then trained on the target image dataset. The key difference from transfer learning is that, in meta learning, without the second training step, the pretrained neural network is unable to classify the images in the target dataset, but it can be trained to do so. There is then greater flexibility for the neural network to be applied to datasets of different characteristics.

In our meta-learning framework for detecting and classifying quantum scars, we find it necessary to apply data augmentation by creating a “noisy” version of the Omniglot dataset because, from an imaging point of view, the quantum scarring images are intrinsically noisy.

-
- [1] F. Haake, *Quantum Coherence in Mesoscopic Systems* (Springer, New York, 1991), 583.
 - [2] H.-J. Stöckmann, *Quantum Chaos: An Introduction* (Cambridge University Press, Cambridge, England, 2000).
 - [3] S. W. McDonald and A. N. Kaufman, Spectrum and Eigenfunctions for a Hamiltonian with Stochastic Trajectories, *Phys. Rev. Lett.* **42**, 1189 (1979).
 - [4] E. J. Heller, Bound-State Eigenfunctions of Classically Chaotic Hamiltonian Systems: Scars of Periodic Orbits, *Phys. Rev. Lett.* **53**, 1515 (1984).
 - [5] E. Bogomolny, Smoothed wave functions of chaotic quantum systems, *Phys. D* **31**, 169 (1988).
 - [6] M. V. Berry, Quantum scars of classical closed orbits in phase space, *Proc. R. Soc. A* **423**, 219 (1989).
 - [7] K. S. Novoselov, A. K. Geim, S. V. Morozov, D.-e. Jiang, Y. Zhang, S. V. Dubonos, I. V. Grigorieva, and A. A. Firsov,

- Electric field effect in atomically thin carbon films, *Science* **306**, 666 (2004).
- [8] K. S. Novoselov, A. K. Geim, S. V. Morozov, D. Jiang, M. I. Katsnelson, I. Grigorieva, S. Dubonos, and A. A. Firsov, Two-dimensional gas of massless Dirac fermions in graphene, *Nature* **438**, 197 (2005).
- [9] A. C. Neto, F. Guinea, N. M. Peres, K. S. Novoselov, and A. K. Geim, The electronic properties of graphene, *Rev. Mod. Phys.* **81**, 109 (2009).
- [10] L. Huang, Y.-C. Lai, D. K. Ferry, S. M. Goodnick, and R. Akis, Relativistic Quantum Scars, *Phys. Rev. Lett.* **103**, 054101 (2009).
- [11] Y.-C. Lai, L. Huang, H.-Y. Xu, and C. Grebogi, Relativistic quantum chaos - an emergent interdisciplinary field, *Chaos* **28**, 052101 (2018).
- [12] L. Huang, H.-Y. Xu, C. Grebogi, and Y.-C. Lai, Relativistic quantum chaos, *Phys. Rep.* **753**, 1 (2018).
- [13] H. Xu, L. Huang, Y.-C. Lai, and C. Grebogi, Chiral Scars in Chaotic Dirac Fermion Systems, *Phys. Rev. Lett.* **110**, 064102 (2013).
- [14] C.-Z. Wang, L. Huang, and K. Chang, Scars in Dirac fermion systems: The influence of an Aharonov–Bohm flux, *New J. Phys.* **19**, 013018 (2017).
- [15] M.-Y. Song, Z.-Y. Li, H.-Y. Xu, L. Huang, and Y.-C. Lai, Quantization of massive Dirac billiards and unification of nonrelativistic and relativistic chiral quantum scars, *Phys. Rev. Res.* **1**, 033008 (2019).
- [16] Z.-Y. Li and L. Huang, Quantization and interference of a quantum billiard with fourfold rotational symmetry, *Phys. Rev. E* **101**, 062201 (2020).
- [17] R. Akis, D. Ferry, and J. Bird, Wave Function Scarring Effects in Open Stadium Shaped Quantum Dots, *Phys. Rev. Lett.* **79**, 123 (1997).
- [18] P. Blanchard and R. Olkiewicz, Decoherence-Induced Continuous Pointer States, *Phys. Rev. Lett.* **90**, 010403 (2003).
- [19] W. H. Zurek, Decoherence, einselection, and the quantum origins of the classical, *Rev. Mod. Phys.* **75**, 715 (2003).
- [20] D. Ferry, R. Akis, and J. Bird, Einselection in Action: Decoherence and Pointer States in Open Quantum Dots, *Phys. Rev. Lett.* **93**, 026803 (2004).
- [21] R. Brunner, R. Akis, D. Ferry, F. Kuchar, and R. Meisels, Coupling-Induced Bipartite Pointer States in Arrays of Electron Billiards: Quantum Darwinism in Action?, *Phys. Rev. Lett.* **101**, 024102 (2008).
- [22] D. Ferry, L. Huang, R. Yang, Y.-C. Lai, and R. Akis, Open quantum dots in graphene: Scaling relativistic pointer states, *J. Phys. Conf. Ser.* **220**, 012015 (2010).
- [23] M. Zwolak and W. H. Zurek, Complementarity of quantum discord and classically accessible information, *Sci. Rep.* **3**, 1 (2013).
- [24] J. U. Nöckel and A. D. Stone, Ray and wave chaos in asymmetric resonant optical cavities, *Nature* **385**, 45 (1997).
- [25] S. Shinohara, T. Harayama, T. Fukushima, M. Hentschel, T. Sasaki, and E. E. Narimanov, Chaos-Assisted Directional Light Emission from Microcavity Lasers, *Phys. Rev. Lett.* **104**, 163902 (2010).
- [26] H. Cao and J. Wiersig, Dielectric microcavities: Model systems for wave chaos and non-Hermitian physics, *Rev. Mod. Phys.* **87**, 61 (2015).
- [27] Q. Song, L. Ge, B. Redding, and H. Cao, Channeling Chaotic Rays into Waveguides for Efficient Collection of Microcavity Emission, *Phys. Rev. Lett.* **108**, 243902 (2012).
- [28] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, and M. Greiner, *et al.*, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [29] C. J. Turner, A. A. Michailidis, D. A. Abanin, M. Serbyn, and Z. Papić, Weak ergodicity breaking from quantum many-body scars, *Nat. Phys.* **14**, 745 (2018).
- [30] M. Schechter and T. Iadecola, Weak Ergodicity Breaking and Quantum Many-Body Scars in Spin-1 *XY* Magnets, *Phys. Rev. Lett.* **123**, 147201 (2019).
- [31] S. Choi, C. J. Turner, H. Pichler, W. W. Ho, A. A. Michailidis, Z. Papić, M. Serbyn, M. D. Lukin, and D. A. Abanin, Emergent SU(2) Dynamics and Perfect Quantum Many-Body Scars, *Phys. Rev. Lett.* **122**, 220603 (2019).
- [32] C.-J. Lin, V. Calvera, and T. H. Hsieh, Quantum many-body scar states in two-dimensional Rydberg atom arrays, *Phys. Rev. B* **101**, 220304 (2020).
- [33] S. Dooley, Robust Quantum Sensing in Strongly Interacting Systems with Many-Body Scars, *PRX Quantum* **2**, 020330 (2021).
- [34] P.-F. Zhang, H. Dong, Y. Gao, L.-T. Zhao, J. Hao, J. Y. Desautels, Q.-J. Guo, J.-C. Chen, J.-F. Deng, B.-B. Liu, *et al.*, Many-body Hilbert space scarring on a superconducting processor, *Nat. Phys.* **19**, 120 (2023).
- [35] Y. LeCun, Y. Bengio, and G. E. Hinton, Deep learning, *Nature* **521**, 436 (2015).
- [36] A. Krizhevsky, I. Sutskever, and G. E. Hinton, Imagenet classification with deep convolutional neural networks, *Adv. Neu. Info. Proc. Sys.* **25**, 1097 (2012).
- [37] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [38] P. Baldi, K. Bauer, C. Eng, P. Sadowski, and D. Whiteson, Jet substructure classification in high-energy physics with deep neural networks, *Phys. Rev. D* **93**, 094034 (2016).
- [39] L. de Oliveira, M. Kagan, L. Mackey, B. Nachman, and A. Schwartzman, Jet-images - deep learning edition, *J. High Energy Phys.* **2016**, 69 (2016).
- [40] A. Aurisano, A. Radovic, D. Rocco, A. Himmel, M. Messier, E. Niner, G. Pawloski, F. Psihas, A. Sousa, and P. Vahle, A convolutional neural network neutrino event classifier, *J. Instrum.* **11**, 09001 (2016).
- [41] F. Lanusse, Q. Ma, N. Li, T. E. Collett, C.-L. Li, S. Ravanbakhsh, R. Mandelbaum, and B. Póczos, CMU DeepLens: Deep learning for automatic image-based galaxy–galaxy strong lens finding, *Mon. Not. R. Astron. Soc.* **473**, 3895 (2018).
- [42] Y. Zhang, A. Mesaros, K. Fujita, S. Edkins, M. Hamidian, K. Ch’ng, H. Eisaki, S. Uchida, J. S. Davis, and E. Khatami, *et al.*, Machine learning in electronic-quantum-matter imaging experiments, *Nature* **570**, 484 (2019).
- [43] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).

- [44] Y. A. Kharkov, V. E. Sotskov, A. A. Karazeev, E. O. Kiktenko, and A. K. Fedorov, Revealing quantum chaos with machine learning, *Phys. Rev. B* **101**, 064406 (2020).
- [45] L. Fei-Fei, R. Fergus, and P. Perona, One-shot learning of object categories, *IEEE Trans. Pat. Ana. Mach. Intel.* **28**, 594 (2006).
- [46] B. M. Lake, R. Salakhutdinov, and J. B. Tenenbaum, Human-level concept learning through probabilistic program induction, *Science* **350**, 1332 (2015).
- [47] O. Vinyals, C. Blundell, T. Lillicrap, and D. Wierstra, *et al.*, Matching networks for one shot learning, *Adv. Neural Inf. Process. Syst.* **29**, 3630 (2016).
- [48] L. Kaiser, O. Nachum, A. Roy, and S. Bengio, Learning to remember rare events, [ArXiv:1703.03129](https://arxiv.org/abs/1703.03129) (2017).
- [49] Y.-X. Wang, D. Ramanan, and M. Hebert, in *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2019), p. 9925.
- [50] D. Huber, O. V. Marchukov, H.-W. Hammer, and A. G. Volosniev, Morphology of three-body quantum states from machine learning, *New J. Phys.* **23**, 065009 (2021).
- [51] J. Keski-Rahkonen, A. Ruhanen, E. J. Heller, and E. Räsänen, Quantum Lissajous Scars, *Phys. Rev. Lett.* **123**, 214101 (2019).
- [52] C. Finn, P. Abbeel, and S. Levine, in *International Conference on Machine Learning* (PMLR, 2017), p. 1126.
- [53] J. P. Zwolak, T. McJunkin, S. S. Kalantre, J. Dodson, E. MacQuarrie, D. Savage, M. Lagally, S. Coppersmith, M. A. Eriksson, and J. M. Taylor, Autotuning of Double-Dot Devices In Situ with Machine Learning, *Phys. Rev. Appl.* **13**, 034075 (2020).
- [54] G. Ness, A. Vainbaum, C. Shkedrov, Y. Florshaim, and Y. Sagi, Single-Exposure Absorption Imaging of Ultracold Atoms Using Deep Learning, *Phys. Rev. Appl.* **14**, 014011 (2020).
- [55] S. Shinohara, T. Harayama, T. Fukushima, M. Hentschel, T. Sasaki, and E. E. Narimanov, Chaos-Assisted Directional Light Emission from Microcavity Lasers, *Phys. Rev. Lett.* **104**, 163902 (2010).
- [56] H. Cao and J. Wiersig, Dielectric microcavities: Model systems for wave chaos and non-Hermitian physics, *Rev. Mod. Phys.* **87**, 61 (2015).
- [57] A. P. S. de Moura, Y.-C. Lai, R. Akis, J. P. Bird, and D. K. Ferry, Tunneling and Nonhyperbolicity in Quantum Dots, *Phys. Rev. Lett.* **88**, 236804 (2002).
- [58] F. Li, Y. Wu, P. Chen, L. Feng, Z. Wang, Y. Ren, and Y. Wang, High-quality wave-chaotic microlasers from deformed halide perovskite cavities, *ACS Photon.* **9**, 2431 (2022).
- [59] M. V. Berry, in *A Half-Century of Physical Asymptotics and Other Diversions: Selected Works by Michael Berry* (World Scientific, 1987), p. 307.
- [60] M. C. Gutzwiller, *Chaos in Classical and Quantum Mechanics* Vol. 1 (Springer, New York, 2013).
- [61] M. V. Berry and M. Wilkinson, Diabolical points in the spectra of triangles, *Proc. R. Soc. A* **392**, 15 (1984).
- [62] M. V. Berry and R. Mondragon, Neutrino billiards: Time-reversal symmetry-breaking without magnetic fields, *Proc. R. Soc. A* **412**, 53 (1987).
- [63] M. Robnik, Quantising a generic family of billiards with analytic boundaries, *J. Phys. A* **17**, 1049 (1984).
- [64] M. V. Berry and M. Robnik, Statistics of energy levels without time-reversal symmetry: Aharonov-Bohm chaotic billiards, *J. Phys. A* **19**, 649 (1986).
- [65] E. Ott, *Chaos in Dynamical Systems* (Cambridge University Press, Cambridge, UK, 2002), 2nd ed.
- [66] T. M. Antonsen, E. Ott, Q. Chen, and R. N. Oerter, Statistics of wave-function scars, *Phys. Rev. E* **51**, 111 (1995).
- [67] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, and M. Isard, *et al.*, in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (2016), p. 265.
- [68] A. Nichol, J. Achiam, and J. Schulman, On first-order meta-learning algorithms, [ArXiv:1803.02999](https://arxiv.org/abs/1803.02999) (2018).
- [69] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, in *Proceedings of the IEEE International Conference on Computer Vision* (2017), p. 2980.
- [70] D. Villaseñor, S. Pilatowsky-Cameo, M. A. Bastarrachea-Magnani, S. Lerma-Hernández, and J. G. Hirsch, Quantum localization measures in phase space, [ArXiv:2103.07480](https://arxiv.org/abs/2103.07480) (2021).
- [71] S. Pilatowsky-Cameo, D. Villaseñor, M. A. Bastarrachea-Magnani, S. Lerma-Hernández, L. F. Santos, and J. G. Hirsch, Identification of quantum scars via phase-space localization measures, *Quantum* **6**, 644 (2022).
- [72] https://github.com/hanchendi/ML_quantum_scar.
- [73] S. Ioffe and C. Szegedy, in *International Conference on Machine Learning* (PMLR, 2015), p. 448.