

Toward structure-preserving quantum encodings

Arthur J. Parzygnat^{1,2,*} Tai-Danae Bradley^{3,4,†} Andrew Vlasic^{5,‡} and Anh Pham⁶¹*Experimental Study Group, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*²*Deloitte Consulting LLP, Arlington, Virginia 22209, USA*³*SandboxAQ, Palo Alto, California 94301, USA*⁴*Department of Mathematics, The Master's University, Santa Clarita, California 91321, USA*⁵*Deloitte Consulting LLP, Tampa, Florida 33602, USA*⁶*Deloitte Consulting LLP, Atlanta, Georgia 30303, USA*

(Received 11 April 2025; published 8 December 2025)

Harnessing the potential computational advantage of quantum computers for machine learning tasks relies on the uploading of classical data onto quantum computers through what are commonly referred to as quantum encodings. The choice of such encodings may vary substantially from one task to another, and there exist only a few cases where structure has provided insight into their design and implementation, such as symmetry in geometric quantum learning. Here, we propose the perspective that category theory offers a natural mathematical framework for analyzing encodings that respect structure inherent in datasets and learning tasks. We illustrate this with pedagogical examples, which include geometric quantum machine learning, quantum metric learning, topological data analysis, and more. Moreover, such a perspective provides a language in which to ask meaningful and mathematically precise questions for the design of quantum encodings and circuits for quantum machine learning tasks.

DOI: [10.1103/rph8-g15q](https://doi.org/10.1103/rph8-g15q)

I. INTRODUCTION

A critical step in quantum machine learning (QML) for classical data is deciding how to encode the data in a Hilbert space through a quantum circuit [1–6]. Finding the best encoding for a specific task within a large space of options is often a time- and resource-intensive task, and rarely are there any general-purpose guiding principles to choose encodings. Here, we utilize category theory to organize and isolate structure in the dataset and learning task [7]. This restricts the initial unstructured space of all quantum encodings to a structured subspace that is usually significantly smaller, thus enabling the reduction of resources needed to find suitable encodings. By isolating the structure that quantum encodings preserve for a given learning task, such as a symmetry, a metric, or a topology, one identifies the encodings as the space of structure-preserving morphisms in the appropriate category, thereby providing a mathematical model for what constitutes structure preservation.

It is only after encoding the classical data onto a quantum computer that one can try to take advantage of potential speedups due to quantum information processing [8,9]. The specifics of what happens after encoding depends on the

model [10,11]. For example, in variational QML models, embedded data are processed through a parametrized quantum circuit with tunable parameters before obtaining an output through measurement. The measurements can then be used to provide a feedback loop to update the parameters defining the parametrized quantum circuit in order to improve the algorithm for unseen data [12]. Finding optimal parameters is often plagued by barren plateaus, and much recent work is spent on understanding and mitigating the effects of such barren plateaus in variational algorithms [13–17].

Among the different components of a QML model, it is the first step of the procedure, the embedding strategy—sometimes called *state preparation* [18–20] or *data encoding* [11,18,21,22]—that may arguably play the most significant role. More concretely, it involves the choice of a function, often referred to as a (*quantum*) *feature map* [3,19,23–27], that assigns to each data point a quantum state, known as a *feature vector* [24], *data encoding* [3,11,22], or (*quantum*) *embedding* [3,23,27,28], to name a few. In this article, we use *quantum encoding* to refer to the function, but the terminology varies as do the encoding choices, each of which introduces different inductive biases [29]. Notably, there is no general framework for deciding what constitutes a good quantum encoding given a particular dataset or learning task. The notion of “good” is indeed multifaceted and may encompass several (sometimes competing) aspects required for the success of a QML model. For example, previous works have focused on quantum encodings’ expressivity [21,26], expressibility [2], robustness to noise [22], computational cost [18], and effect on model trainability [28].

Another facet, one that we focus on in this paper, is the interplay between quantum encodings and structure in data. One archetypal example is the field of geometric quantum

*Contact author: arthurjp@mit.edu†Contact author: tai.danae@math3ma.com‡Contact author: avlasic@deloitte.com

machine learning, where group theoretic tools aid in constructing quantum encodings that respect a given symmetry [30–32]. This field stems from its successful predecessor of (classical) geometric learning [33,34], whose goal is to characterize and analyze the structure of data from the perspective of symmetry [35]. Generalizing beyond symmetry, how should a quantum encoding be constructed so that a particular structure is preserved? And what is an appropriate definition of “structure” to use in this context?

Similar questions have been asked by Bowles *et al.*, who recently called for a more scientifically rigorous approach to benchmarking in QML [5]:

More studies that focus on questions of structure in data are crucial for the design of meaningful benchmarks: What mathematical properties do real-world applications of relevance have? ... How can we connect them to the mathematical properties of quantum models?

In a similar vein, Larocca *et al.* stated [36]

...models with little to no inductive biases (i.e., with no assumptions about the problem embedded in the model) are likely to have trainability and generalization issues especially for large problem sizes. As such, it is fundamental to develop schemes that encode as much information as available about the problem at hand.

As another example, Thanasilp *et al.* pointed out that many of the difficulties of extracting information from data transferred to quantum states are due to the exponential concentration of quantum kernels. This is caused by various factors, one of which includes the expressivity of quantum encodings [4]:

Our work on embedding-induced concentration suggests that problem-inspired embeddings should be used over problem-agnostic embeddings (which are typically highly expressive and entangling).

...Unstructured data embeddings should generally be avoided and the data structure should be taken into account when designing a data-embedding....

The problem-inspired embeddings that are mentioned here are precisely those types of encodings that preserve a certain structure. Such encodings form the focus of this paper.

Toward addressing questions of structure in data, a quick perusal through the literature shows that a number of mathematical structures have relevancy depending on the dataset and task at hand. Geometric quantum machine learning, as mentioned above, exploits symmetry in data, whereas quantum metric learning prioritizes similarity and distance structure [23]. On the other hand, quantum topological data analysis (TDA) [37], which has been studied in the context of quantum encodings, concerns topological structure in data [38]. Each involves the analysis of a different mathematical structure and the extent to which it is or is not preserved under a passage to a Hilbert space under a non-linear mapping. A methodical study of quantum encodings can therefore benefit from a principled framework in which to think about structured mathematical objects and structure-preserving mappings between them.

Category theory, a relatively modern branch of mathematics, is the study of precisely this [7,39–41] (see also Refs. [42,43] for quantum-focused introductions). In addition to providing a language to analyze mathematical objects and relationships between them, category theory also formalizes the notion of “structure” itself [44], thus lending clarity and rigor to situations that may otherwise not be well understood. Its purview has also extended well beyond mathematics to data analysis where, for instance, it forms the backbone of topological data analysis [45] and the dimensionality reduction technique of Uniform Manifold Approximation and Projection (UMAP) [46], and more recently it has been used to provide guiding principles in the context of (classical) machine learning, such as functional programming [47], probabilistic programming [48], the architecture of software engineering [49], neural network symmetrization [50,51], backpropagation [52], and the general algebraic structure of deep learning [53].

In this article, we provide an overview of various quantum encodings and offer a guided perspective to designing such encodings in a way that respects the structure within a given dataset and learning task. We achieve this through the usage of category theory, which allows one to analyze a variety of mathematical structures simultaneously without the intricate details that are problem specific. Importantly, the reader is not assumed to have prior familiarity with category theory. Instead, this article serves an invitation to those with a background in QML to adopt a categorical perspective when designing encoding schemes. And although we focus primarily on quantum encodings toward the advancement of QML techniques, the ideas presented here are equally applicable to standard (classical) machine learning.

The article is outlined as follows. Section II introduces quantum encodings and walks through several examples while highlighting some of the mathematical structures that are preserved. Section III introduces some basic definitions in category theory and recasts the examples of the previous section in a more formal, categorical context. This will reveal the fact that each example is a special case of a more general concept known as a *forgetful functor*. These observations are finally summarized in Sec. III C in which we reformulate the design setup and design goal of quantum encodings in the language of category theory—the core perspective of the article. Section IV proposes several open questions and avenues for future directions.

II. STRUCTURE-PRESERVING QUANTUM ENCODINGS

Quantum circuits for quantum machine learning tasks can often be decomposed into three essential components: the encoding block, the variational block, and the measurement block, as displayed in Fig. 1 [24]. Each component has a significant influence on the output of a quantum algorithm, though the first step plays a particularly important role. Indeed, just as the quality of data influences the performance of machine learning algorithms, it is desirable to preserve structure inherent in raw data when encoding it onto a quantum computer. And yet, due to the wide variety of quantum algorithms, there are no immediately obvious general guidelines for preserving the latent information within the data in this




FIG. 1. The three essential components of a (variational) quantum machine learning task or algorithm: the encoding block that transfers classical data x onto the quantum computer, the variational block whose parameters θ can be modified so as to optimize some outcome, and the measurement. Our focus here is on the encoding step.

encoding layer. For this reason, we focus our attention on the encoding step, with special attention given to preserving underlying structure within the data.

As we will see below, the meaning of “structure” will depend on the data, the machine learning task, and/or the problem in general. Intuitively, for example, the data could admit geometric or topological structure (such as with data points sampled from a distribution localized near a manifold), algebraic structure (such as symmetries), and/or metric structure (such as when measuring distance based on similarities and dissimilarities between data points). Moreover, many of these structures are not mutually exclusive, which further complicates selecting appropriate quantum encodings.

Several specialized encoding schemes for transferring classical data onto a quantum computer have been constructed for specific cases. For example, a common choice in the context of original quantum algorithms, such as the Deutsch-Jozsa algorithm [54], is bit encoding, where a binary string such as 1011 is assigned to the four-qubit state $|1011\rangle$. Meanwhile, encoding vectors $v = (v_1, \dots, v_d)$ in $\mathbb{R}^d$ into quantum states can be done in many ways. For example, solving systems of linear equations [55], supervised and unsupervised clustering [56], and computing persistent homology in topological data analysis [37] utilize amplitude encoding when working directly with the raw data [57,58]. Other popular choices of encoding schemes in the context of quantum machine learning include angle/rotation encoding, time-evolution encoding, instantaneous quantum polytime encoding, and more [59]. Are these encodings the most ideal ones used for their purposes? What guiding principles should be followed so that one can make better-informed decisions about what types of encodings to use? Can structure that is inherent in the problem be utilized in order to guide those choices? This paper aims to answer that this can indeed be done using appropriate category-theoretic tools [7,39]. Before getting there, however, we first review some examples of encoding techniques and the types of mathematical structures they preserve, saving the categorical explanations for Sec. III A.

A. Generic encodings

En route to surveying examples of quantum encodings, we begin by providing the mathematical definition of quantum encoding used throughout this work. The definition is

meant to capture the idea that classical real-world data can be mapped into a quantum system either as a data-dependent quantum state or a data-dependent quantum circuit. To make this mathematically precise, we first establish some notation and terminology for operators and states in Hilbert space [8,60,61]. Given a Hilbert space $\mathcal{H}$, let $\mathcal{S}(\mathcal{H})$ denote the set of states, i.e., trace-class positive operators with trace 1, and let $\mathcal{U}(\mathcal{H})$ denote the set of unitary operators on $\mathcal{H}$. Although both $\mathcal{S}(\mathcal{H})$ and $\mathcal{U}(\mathcal{H})$ have more structure than that of mere sets [for instance, $\mathcal{S}(\mathcal{H})$ can be viewed as a convex space and $\mathcal{U}(\mathcal{H})$ can be equipped with the structure of a Lie group], our aim is to first focus on the structure present in the data and how that structure can be preserved on $\mathcal{S}(\mathcal{H})$ and $\mathcal{U}(\mathcal{H})$ under quantum encodings. So, we consider them as mere sets for now.

We will often let $\mathcal{X}$ denote a set, called a *data domain* or *feature space*, containing the data of interest, with the notation X typically used for the dataset itself (e.g., a sample), which is usually taken to be finite. The following definition of a quantum encoding is a minimalistic definition, in that it does not refer to any specific task (e.g., classification or regression) or model (e.g., variational quantum circuit, implicit, explicit, data reuploading, etc.) but is general enough to be applicable to most settings.

Definition 1. A *quantum state encoding* from a data domain $\mathcal{X}$ into a Hilbert space $\mathcal{H}$ is a function $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$. A *quantum unitary encoding* is a function $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$. A *quantum encoding* refers to either of these.

The two definitions are related in that every quantum unitary encoding $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$, together with a quantum state $|\psi_0\rangle \in \mathcal{H}$, gives rise to a quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ via

$$\rho(x) = U(x)|\psi_0\rangle\langle\psi_0|U(x)^\dagger. \quad (1)$$

The important point to notice in this general definition is that a quantum encoding is simply a function, which does not necessarily preserve any additional structure. In particular, it is not assumed to be continuous, smooth, distance preserving, or symmetry preserving. An important step in designing quantum encodings, therefore, is to *first* identify appropriate structures on the data domain $\mathcal{X}$ and set of states $\mathcal{S}(\mathcal{H})$ that are relevant to the problem at hand. One *then* seeks to construct a quantum encoding that preserves that structure. Demanding that a certain mathematical structure is preserved places a restriction on the set of all functions from the data domain $\mathcal{X}$ to the set of states $\mathcal{S}(\mathcal{H})$, thus potentially making it easier to search for quantum encodings (cf. Fig. 2).

In the next few sections, we will highlight several familiar illustrations of this, where the relevant structures are symmetry, topologies and smooth structures, distances in the context of topological data analysis, and distances in the context of quantum metric learning.

B. Symmetry

Geometric quantum machine learning provides one example of structure-preserving quantum encodings that aids in the design of quantum circuits [30–32,36,62,63]. In this paradigm, symmetries are described in terms of group actions




FIG. 2. Among all possible set-theoretic functions describing quantum state encodings (dashed arrows) from a data domain $\mathcal{X}$ to a Hilbert space $\mathcal{H}$, demanding that a structure is preserved isolates a subset of quantum encodings (solid arrows), thus potentially simplifying the search for quantum encodings compatible with a given structure.

on both the data domain and the set of quantum states. Quantum encodings that preserve this structure of symmetry are called *equivariant* quantum encodings, and such encodings will be the focus of this section. Since there are fewer quantum encodings that are equivariant with respect to these symmetries as compared to all quantum encodings, one is able to isolate a smaller space of encodings, which in theory reduces the resources needed to find appropriate encodings. More precisely, both the data domain $\mathcal{X}$ and quantum state space $\mathcal{S}(\mathcal{H})$ are equipped with an action of a group G , and a quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ that is compatible with the relevant symmetry is a G -equivariant map. Let us now briefly review how this works by first recalling the formal definition of a group's action on a set.

Definition 2. A G -set is a set $\mathcal{X}$ together with a group homomorphism $\alpha : G \rightarrow \text{Aut}(\mathcal{X})$, where $\text{Aut}(\mathcal{X})$ denotes the group (under function composition) of bijections on $\mathcal{X}$. The map α is referred to as the *group action*. Such a G -set $\mathcal{X}$ is often written as a pair $(\mathcal{X}, \alpha)$.

In this definition, notice the distinction between the underlying set and the additional structure that defines the action. The notation helps to keep this distinction front and center. Indeed, if $\mathcal{X}$ is a G -set, then writing it instead as a pair $(\mathcal{X}, \alpha)$ helps to emphasize the additional structure that $\mathcal{X}$ carries with it, which is useful notation that will reappear later. Unwinding Definition 2 further, “Aut” stands for “automorphisms,” and the bijection (or automorphism) associated with a group element $g \in G$ is often written as $\alpha_g : \mathcal{X} \rightarrow \mathcal{X}$. Its action on an element $x \in \mathcal{X}$ is written as $\alpha_g(x)$, which is another element of $\mathcal{X}$. The group homomorphism property says that the action is compatible with the group operation, in that $\alpha_{gh}(x) = \alpha_g(\alpha_h(x))$ for all $g, h \in G$ and $x \in \mathcal{X}$ as well as $\alpha_{1_G} = \text{id}_{\mathcal{X}}$, where 1_G is the identity in the group G .

As an example, if $\mathcal{H}$ is a Hilbert space, and if $V : G \rightarrow \mathcal{U}(\mathcal{H})$ is a unitary representation of G on $\mathcal{H}$, then $\mathcal{H}$ becomes a G -set because the set of unitary operators on $\mathcal{H}$ can be viewed as a subset of $\text{Aut}(\mathcal{H})$, since unitary operators can be viewed as invertible linear transformations on $\mathcal{H}$. Moreover, this action of G on $\mathcal{H}$ induces an action of G on the space of states $\mathcal{S}(\mathcal{H})$ by sending each group element $g \in G$ to the operator $\text{Ad}_{V_g} \in \text{Aut}(\mathcal{S}(\mathcal{H}))$ defined by

$$\text{Ad}_{V_g}(\sigma) := V_g \sigma V_g^\dagger \quad (2)$$




FIG. 3. This is a variation of Fig. 2 in Ref. [30] based on a binary classification task with discrete symmetries given by $\alpha_{(1,0)}$, a reflection along the line $x_2 = x_1$, $\alpha_{(0,1)}$, an inversion, and $\alpha_{(1,1)}$, a reflection along the line $x_2 = -x_1$. The structure is periodic so that the symmetry is preserved. More details about how to explicitly construct the decision boundaries for this classifier, as well as the associated observable, are provided in Appendix B.

for all $\sigma \in \mathcal{S}(\mathcal{H})$. Here, “Ad” stands for the *adjoint* action. Thus, $(\mathcal{S}(\mathcal{H}), \text{Ad}_V)$ is a G -set as well.

Therefore, if we have a data domain $\mathcal{X}$ with a symmetry, and we want to encode these data onto a quantum system with Hilbert space $\mathcal{H}$, then we might want to encode the data in a way that preserves the symmetry. This leads to our first example of a structure-preserving quantum encoding.

Definition 3. Given a G -set $(\mathcal{X}, \alpha)$, a G -equivariant quantum state encoding of $(\mathcal{X}, \alpha)$ consists of a quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ together with a unitary representation $V : G \rightarrow \mathcal{U}(\mathcal{H})$ of G on the Hilbert space $\mathcal{H}$ such that ρ is G -equivariant, that is,

$$\rho(\alpha_g(x)) = V_g \rho(x) V_g^\dagger \quad (3)$$

for all $x \in \mathcal{X}$ and $g \in G$.

This definition of a G -equivariant quantum state encoding is a special case of the more general notion of a G -equivariant map between two G -sets, defined as follows.

Definition 4. Let $(\mathcal{X}, \alpha)$ and $(\mathcal{Y}, \beta)$ be two G -sets. A G -equivariant map from $(\mathcal{X}, \alpha)$ to $(\mathcal{Y}, \beta)$ is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that

$$f(\alpha_g(x)) = \beta_g(f(x)) \quad (4)$$

for all $x \in \mathcal{X}$ and $g \in G$.

As we will see in Sec. III, Definition 3 amounts to saying that a G -equivariant quantum state encoding is a *morphism* in a particular *category*. In the meantime, observe that Eq. (4), and in particular Eq. (3), imposes a restriction on the set of all functions from one G -set to another since not every function will necessarily satisfy the constraint. Indeed, such constraints have been used in QML to construct equivariant quantum encodings by combining equivariant feature maps

and equivariant gatesets in Ref. [30]. This is illustrated in the following toy example.

Example 1. Following the example in Ref. [30, Sec. II A], consider a supervised learning task that distinguishes between two classes of points within a data domain $\mathcal{X} = \mathbb{R}^2$ as shown in Fig. 3. The data domain $\mathcal{X}$ acquires a symmetry determined by the relations

$$y(x_1, x_2) = y(x_2, x_1) = y(-x_1, -x_2), \quad (5)$$

where $y : \mathcal{X} \rightarrow \mathbb{R}$ is a function that determines the class of any input data point via $y(x) > 0$ implies x is of class +1, while $y(x) < 0$ implies x is of class -1. The symmetry depicted in Eq. (5) is modeled by the Klein Four group

$$G = \mathbb{Z}_2 \times \mathbb{Z}_2 = \{(0, 0), (1, 0), (0, 1), (1, 1)\} \quad (6)$$

under addition modulo 2. Suppose this group has the action $\alpha : G \rightarrow \text{Aut}(\mathbb{R}^2)$ on $\mathcal{X} = \mathbb{R}^2$ given by the linear transformations whose associated matrices are

$$\begin{aligned} \alpha_{(0,0)} &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, & \alpha_{(1,0)} &= \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \\ \alpha_{(0,1)} &= \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}, & \alpha_{(1,1)} &= \begin{bmatrix} 0 & -1 \\ -1 & 0 \end{bmatrix}. \end{aligned} \quad (7)$$

Now set $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2$ and set $V : G \rightarrow \mathcal{U}(\mathcal{H})$ to be the representation specified by

$$\begin{aligned} V_{(0,0)} &= \mathbb{1}_2 \otimes \mathbb{1}_2, & V_{(1,0)} &= \text{SWAP}, \\ V_{(0,1)} &= X \otimes X, & V_{(1,1)} &= \text{SWAP}(X \otimes X), \end{aligned} \quad (8)$$

where X is the Pauli matrix $X = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$, $\mathbb{1}_2$ is the 2×2 identity matrix, and SWAP is the unitary operator characterized by $\text{SWAP}(|\psi\rangle \otimes |\phi\rangle) = |\phi\rangle \otimes |\psi\rangle$ for all $|\psi\rangle, |\phi\rangle \in \mathbb{C}^2$.

In this example, a G -equivariant quantum encoding would be a map $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ such that $V_g \rho(x) V_g^\dagger = \rho(\alpha_g(x))$ for all $x \in \mathcal{X}$ and $g \in G$. One example of such an encoding can be obtained from the unitary encoding $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$, which sends a point $x = (x_1, x_2) \in \mathcal{X}$ to

$$U(x_1, x_2) = R_Z(x_1) \otimes R_Z(x_2), \quad (9)$$

where the rotation gate is given by

$$R_Z(\theta) = e^{-i\theta Z/2} = \begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}, \quad (10)$$

where Z is the Pauli matrix $Z = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}$. From this unitary encoding, one can obtain a state encoding by acting on a fiducial state $|\psi_0\rangle \in \mathcal{H}$ of the form

$$|\psi_0\rangle = \sqrt{p}|+, +\rangle - \sqrt{1-p}|-, -\rangle, \quad (11)$$

where $0 < p < 1$. (The value of p chosen to produce Fig. 3 can be found in Appendix B.) Here,

$$|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad \text{and} \quad |-\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \quad (12)$$

are the spin-up and spin-down eigenvectors of X , respectively, and $|\pm, \pm\rangle := |\pm\rangle \otimes |\pm\rangle$. Thus, the resulting state encoding

sends $x = (x_1, x_2)$ to

$$\rho(x) = U(x_1, x_2)|\psi_0\rangle\langle\psi_0|U(x_1, x_2)^\dagger. \quad (13)$$

One can explicitly check that this defines an *equivariant* quantum encoding. More generally, one might ask how such an equivariant encoding can be obtained from the equivariance condition. To see how demanding equivariance reduces the space of possible quantum encodings, it helps to make some assumptions about the form of the encoding. Suppose that the unitary quantum encoding $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$ is of the form

$$U(x) = e^{-i\mathcal{L}(x)}, \quad (14)$$

where $\mathcal{L} : \mathbb{R}^2 \rightarrow \mathbb{M}_4$ is a linear transformation sending each point $x \in \mathcal{X}$ to a Hermitian operator $\mathcal{L}(x) \in \mathbb{M}_4$, where $\mathbb{M}_4$ is the algebra of complex 4×4 matrices. (Encodings of this form are described in more detail in Appendix A.) Moreover, it suffices to assume that this Hermitian operator is traceless so that $-i\mathcal{L}(x) \in \mathfrak{su}(4)$, the Lie algebra of 4×4 special unitary matrices. Such a linear transformation $\mathcal{L}$ is uniquely determined by the image of the two unit vectors

$$L_1 := \mathcal{L}(e_1) \quad \text{and} \quad L_2 := \mathcal{L}(e_2), \quad (15)$$

where $e_1 = (1, 0)$ and $e_2 = (0, 1)$. Therefore, the set of all such linear transformations is isomorphic to the span of all ordered pairs of traceless Hermitian 4×4 matrices, which is a real vector space of dimension $2(4^2 - 1) = 30$.

The equivariance condition singles out a subspace of this vector space as follows. First, notice that the equivariance condition

$$V_g e^{-i\mathcal{L}(x)} V_g^\dagger = e^{-i\mathcal{L}(\alpha_g(x))} \quad (16)$$

at the level of the unitary encoding becomes

$$V_g \mathcal{L}(x) V_g^\dagger = \mathcal{L}(\alpha_g(x)) \quad (17)$$

at the level of the generating elements $\mathcal{L}(x)$ because each V_g is unitary [64]. Let us now see what this says for some of the specific group elements. For $g = (1, 0)$, equivariance implies $\text{SWAP } L_1 \text{ SWAP} = L_2$ (and also $\text{SWAP } L_2 \text{ SWAP} = L_1$, but this is equivalent to the first equation). For $g = (0, 1)$, equivariance implies $(X \otimes X) L_1 (X \otimes X) = -L_1$ and $(X \otimes X) L_2 (X \otimes X) = -L_2$. However, the second of these follows from the previous constraints. Hence, these give us two *operator* constraints

$$\begin{aligned} \text{SWAP } L_1 \text{ SWAP} &= L_2, \\ \{X \otimes X, L_1\} &= 0 \end{aligned} \quad (18)$$

on the space of all quantum encodings parametrized by ordered pairs of Hermitian operators, thereby reducing the space of all possible quantum encodings to structure-preserving quantum encodings. In fact, there is a reduction from the 30-dimensional real vector space of all pairs of Hermitian matrices (L_1, L_2) to an 8-dimensional subspace of pairs $(L_1, \text{SWAP } L_1 \text{ SWAP})$ due to the constraint (18). A basis for this subspace in terms of L_1 consists of

$$\begin{aligned} Z \otimes \mathbb{1}_2, & \quad \mathbb{1}_2 \otimes Z, & Y \otimes \mathbb{1}_2, & \quad \mathbb{1}_2 \otimes Y, \\ X \otimes Z, & \quad Z \otimes X, & X \otimes Y, & \quad Y \otimes X, \end{aligned} \quad (19)$$

which is derived in Appendix B. The example given in Eq. (9) from Ref. [30] corresponds to the choice

$$L_1 = \frac{Z}{2} \otimes \mathbb{1}_2 \quad \text{and} \quad L_2 = \mathbb{1}_2 \otimes \frac{Z}{2} \quad (20)$$

for the generators of the quantum encoding.

More details associated with this example, including how the decision boundaries in Fig. 3 are constructed, are provided in Appendix B.

There are two subtle points to be aware of when designing equivariant encodings. One is that perfectly equivariant encodings (and subsequent equivariant layers) may overfit to perfectly symmetric data. The second point is that such equivariance may also lead to efficient classical replacements of quantum algorithms, i.e., *dequantization* [65,66]. For the first point, due to errors on quantum computers and other factors, the overfitting to perfectly symmetric data might inhibit generalization, and might therefore *not* be ideal for all QML problems [67]. Indeed, it has been observed that having symmetry-breaking layers in data reuploading models may provide more accurate predictions in classification tasks [36,68,69]. (We only briefly mention data reuploading models in Appendix D but otherwise do not make explicit use of them.) However, it is important even in these examples to first isolate the class of equivariant encodings that one can then perturb to get symmetry-breaking gates in a way that is informed by the types of errors that can occur on the quantum computing device. Therefore, isolating the equivariant encodings and gates is an important first step in designing such quantum machine learning models. A discussion about the second point on dequantization is deferred to Sec. IV.

Having briefly presented geometric quantum machine learning as the archetypal example of how structure (namely, symmetry) in a dataset could be used to inform the design of quantum encodings, we now move on to discuss structure preservation in the context of continuity and smoothness.

C. Topologies and smooth structures

Continuity and smoothness (e.g., differentiability) are concepts that can be defined using the tools of topology and manifold theory [70–73], which isolate some of the more flexible structures of Euclidean space dictating how neighborhoods connect together [74]. The underlying topologies and smooth structures on the data domains and quantum state spaces are two examples of mathematical structures that are so natural that they might not always be explicitly stated. Quantum encodings respecting these structures are functions that are continuous and smooth, respectively. In more detail, the data domain $\mathcal{X}$ and the state space $\mathcal{S}(\mathcal{H})$ should each be equipped with a topology and a smooth structure so that the quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ can be chosen to be continuous or smooth (i.e., infinitely differentiable) if we want the encoded quantum states to vary continuously or smoothly based on small variations in the classical data.

We will not need the technicalities of these definitions in what follows, but we do wish to reinforce the notion of adding mathematical structure to a set, which is key concept in this work. To that end, let us briefly recall the definition of a topology and a smooth structure. A *topology* on a set $\mathcal{X}$

is a set $\tau_{\mathcal{X}}$ whose elements are certain subsets of $\mathcal{X}$, called *open* subsets, satisfying a list of axioms [70,73]. When a set $\mathcal{X}$ is equipped with a topology $\tau_{\mathcal{X}}$, the pair $(\mathcal{X}, \tau_{\mathcal{X}})$ is called a *topological space*. Writing it as a pair emphasizes the additional structure that $\mathcal{X}$ carries, although $\tau_{\mathcal{X}}$ is often omitted in the literature for brevity. Meanwhile, a smooth structure can be defined once a set $\mathcal{X}$ is already equipped with a topology. Moreover, this topology should satisfy an additional list of axioms so that $(\mathcal{X}, \tau_{\mathcal{X}})$ defines a *topological manifold* (namely, it should be Hausdorff, second-countable, and locally Euclidean of some fixed dimension) [72]. Then, a *smooth structure* on a topological manifold $(\mathcal{X}, \tau_{\mathcal{X}})$ is a maximal smooth atlas $\mathcal{A}$ on $\mathcal{X}$, and the triple $(\mathcal{X}, \tau_{\mathcal{X}}, \mathcal{A})$ is called a *smooth manifold*. (See Ref. [72] for more details.) Given such structures, only then can one define *continuous* and *smooth* maps between topological spaces and smooth manifolds as functions that preserve the topology and smooth structures [70,72,73]. In the language of category theory, these maps are said to be *morphisms* in particular *categories*, a framing we will revisit in Sec. III.

Technicalities aside, continuity of a quantum encoding is used to ensure that classical data are encoded in such a way that the encoded data do not vary too wildly from the raw classical data, while smoothness is used to ensure even more rigidity and differentiability properties, such as when calculating generators using derivatives [59]. (See also Appendix A.) Most quantum encodings are of this form, and in this section we will list some of the more commonly used examples in the literature. In each of these examples, the data domain $\mathcal{X}$ is a Euclidean space of the form $\mathbb{R}^d$, which is equipped with the standard topology and smooth structure of Euclidean space [72]. Meanwhile, the set of pure states on a Hilbert space $\mathcal{H} = \mathbb{C}^d$ is the complex projective space $\mathbb{CP}^{d-1}$, which also has a standard topology and smooth structure [72,75].

Example 2 (Angle encoding). Angle encoding, also called *rotation encoding*, is a technique explored in Refs. [11,24,76] that utilizes the structure of one-parameter groups and interprets the feature components as parameters that can be used for Hamiltonian evolution. In the case that $\mathcal{X} = \mathbb{R}^d$ is a Euclidean feature space, each component of $x = (x_1, \dots, x_d) \in \mathcal{X}$ is mapped to its own qubit via the map

$$x \mapsto |\Phi(x)\rangle := \left(\bigotimes_{k=1}^d \exp\left(-\frac{ix_k}{2} X_k\right) \right) |0 \cdots 0\rangle, \quad (21)$$

where X_k is the Pauli X gate acting on the k th qubit. The associated quantum state encoding $\rho : \mathbb{R}^d \rightarrow \mathcal{S}(\mathbb{C}^{2^d})$ is then given by

$$\rho(x) := |\Phi(x)\rangle\langle\Phi(x)|. \quad (22)$$

Note that the Pauli X gate may be replaced with a Pauli Y or Pauli Z gate, as in Eq. (9) in the previous section, for example. In fact, one could choose any generalized Pauli operator $\vec{n} \cdot \vec{\sigma}$, where $\vec{n}$ is some unit vector in $\mathbb{R}^3$ and $\vec{\sigma} = (X, Y, Z)$ is the vector of the Pauli gates. All of these quantum state encodings are continuous and smooth. Note that these quantum encodings map a feature space of d dimensions into d qubits, which is a linear scaling (as opposed to the logarithmic scaling of amplitude encoding, which is discussed in Example

3). Moreover, angle encoding has been successfully used in some binary classification tasks [23,27,77]. Note, however, that angle encodings do not make any use of entanglement with respect to the assumed tensor factorization of the qubits because there are no entangling gates in the definition of the encoding. For this reason, it is said to have low expressibility [2], and one can include entangling gates and data reuploading within the quantum circuit to increase expressibility and expressivity [10].

Example 3 (Amplitude encoding). Amplitude encoding is a technique that models entries in the data point array as frequencies (in the statistical sense) of the expected string output [59]. By definition, it takes a vector $x \in \mathbb{R}^{2^d} \setminus \{0\}$ to a vector $\chi(x)$ in $\mathbb{C}^{2^d}$, the Hilbert space of d qubits, and then normalizes it to obtain a genuine quantum state $|\Phi(x)\rangle$. To express this as a succinct formula, if $x = (x_0, x_1, \dots, x_{2^d-1}) \in \mathbb{R}^{2^d}$, then rewrite the indices in binary so that they correspond to a sequence n of d 0s and 1s. For example, if $d = 2$, then n has four possibilities: $n \in \{00, 01, 10, 11\}$. Then, amplitude encoding is defined by the sequence of maps

$$x \mapsto \chi(x) = \sum_{n=0}^{2^d-1} x_n |n\rangle \mapsto |\Phi(x)\rangle = \frac{\chi(x)}{\|\chi(x)\|}, \quad (23)$$

where $\|\chi(x)\|$ denotes the Euclidean norm of $\chi(x)$. Since the data are assumed to be real, one could equivalently normalize the data first onto the sphere S^{2^d-1} inside $\mathbb{R}^{2^d}$ and then map over to $\mathbb{C}^{2^d}$ to get the same resulting quantum state. Although the above description of amplitude encoding provides one mathematical description of the resulting function, there are several proposals for constructing quantum circuits that achieve amplitude encoding, not all of which use the same number of qubits [18,57,58]. We will not discuss these at present, but instead analyze some of the structural properties of the amplitude encoding from Eq. (23). To start, amplitude encoding is both continuous and smooth. Moreover, one of the benefits of amplitude encoding is that a feature space of the form $\mathbb{R}^{2^d}$ is encoded into a Hilbert space of $\log_2(2^d) = d$ qubits. However, in the process of normalization to obtain a genuine quantum state, data points that are far away from each other in the data domain might have their distances so drastically reduced that they become difficult to separate in the encoding (cf. Fig. 4). In fact, the mapping is in general not one to one, which means that two distinct points may be mapped to the same element, thus causing their distance to vanish and therefore become indistinguishable.

There are at least two slight modifications to the amplitude encoding scheme mentioned above that partially address the lack of injectivity. First, one could add an additional entry for the original data and project onto the unit sphere S^{2^d} (instead of S^{2^d-1}) via

$$(x_0, x_1, \dots, x_{N-1}) \mapsto \frac{(x_0, x_1, \dots, x_{N-1}, 1)}{\sqrt{1 + \|x\|^2}}, \quad (24)$$

so that the dimension does not reduce upon encoding [59]. Although the map is now a diffeomorphism onto the open set for the northern hemisphere of S^{2^d} , this does not fully resolve the problem of the difficulty of separation of data mentioned




FIG. 4. The standard form of amplitude encoding is smooth but not dimension preserving. This is because it requires a step that pushes the set of data points in $\mathbb{R}^{2^d}$ onto the unit sphere S^{2^d-1} , thereby reducing the separation between the data points and hence increasing the difficulty in distinguishing between them. For example, although the square $\blacksquare$ is close to the diamond $\blacklozenge$ in the ambient data domain, they are far apart after amplitude encoding. Meanwhile, although $\blacktriangledown$ and $\blacktriangle$ are far apart in the ambient data domain, they are close together after amplitude encoding.

in Fig. 4. This is because data points that are within the unit sphere S^{2^d-1} are mapped into the top part of the northern hemisphere of S^{2^d} within a polar angle of $\frac{\pi}{4}$, while the points outside the unit sphere S^{2^d-1} are all mapped into the northern hemisphere of S^{2^d} with a polar angle between $\frac{\pi}{4}$ and $\frac{\pi}{2}$, thus drastically reducing the distances between points that were originally very far away from each other (cf. Fig. 5).

Another version of amplitude encoding instead first transforms the data into a probability vector by the mapping

$$(x_0, \dots, x_{N-1}) \mapsto \frac{(e^{x_0}, \dots, e^{x_{N-1}}, 1)}{1 + \sum_{j=0}^{N-1} e^{x_j}}, \quad (25)$$

which has image on the probability simplex [38]. The benefit of this over the previous encodings is the direct translation to probabilities, though a similar problem of data separation (due to a lack of distance preservation) remains (cf. Fig. 6). Alternative versions of amplitude encodings are given in Ref. [78], though those also have similar advantages and disadvantages to the examples just given. These problems involving a lack of distance preservation can be circumvented by equipping the target spaces with a modified notion of distance, a topic that we discuss in the next two sections. However, it is not generally possible to transfer a distance function from a data domain to a quantum state space (since encodings are not generally bijective). Instead, it is only possible to go in the opposite direction (cf. Sec. II E). Let us therefore first move on to discuss distances in the context of topological data analysis.




FIG. 5. Although data x_0 in the interval $[-1, 1]$ get mapped into the northern hemisphere of S^1 within polar angle $\frac{\pi}{4}$, all other data get mapped to the region with polar angle between $\frac{\pi}{4}$ and $\frac{\pi}{2}$. This is illustrated by showing how the set of negative integers becomes a sequence with a limit point at $(-1, 0)$ and similarly the image of the positive integers has a limit point at $(1, 0)$. This is shown here for $[-6, 6] \cap \mathbb{Z}$.




FIG. 6. Data x_0 in $\mathbb{R}$ get mapped into the probability 1-simplex. Although values near 0 stay separated, data far away from 0 get clustered near the endpoints. The image of the integers is shown for $[-6, 6] \cap \mathbb{Z}$.

D. Distance for topological data analysis

TDA is a technique to identify geometric features of a data-generating source that are preserved under continuous perturbations [45,79–81]. A quantum encoding transfers the data to a quantum system, which potentially disrupts these topological features. This section will provide guidance on how to minimally disrupt these features by quantifying the distances between data points before and after the encoding.

In more detail, TDA utilizes data embedded in a metric space in order to obtain topological features of a postulated underlying space from which the data are generated. Although a finite set of data points does not inherently acquire a non-trivial topology, the distances between points can be used to construct a simplicial complex [82], called the Vietoris-Rips complex, which is a combinatorial object that contains topological information about this underlying space. We will not need to go through all the components of TDA to get the basic flavor of quantum encodings that can be used for TDA. The important concepts to focus on for our purposes are the structures that are involved in TDA, namely, distance functions, i.e., metrics, whose definition we recall below [83,84]. In this definition, notice how this theme of sets equipped with additional structure makes another appearance.

Definition 5. A metric space is a pair $(\mathcal{X}, d_{\mathcal{X}})$, where $\mathcal{X}$ is a set and $d_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a function satisfying the following conditions:

- (1) $d_{\mathcal{X}}(x_1, x_2) = 0$ if and only if $x_1 = x_2$.
- (2) $d_{\mathcal{X}}(x_1, x_2) = d_{\mathcal{X}}(x_2, x_1)$ for all $x_1, x_2 \in \mathcal{X}$.
- (3) $d_{\mathcal{X}}(x_1, x_2) \leq d_{\mathcal{X}}(x_1, x_3) + d_{\mathcal{X}}(x_2, x_3)$ for all $x_1, x_2, x_3 \in \mathcal{X}$ (the *triangle inequality*).

Such a function $d_{\mathcal{X}}$ is called a *distance function*, or *metric*, on $\mathcal{X}$.

The general pipeline for TDA is as follows (see Fig. 7 for a visual flowchart).

- (1) If $X \subseteq \mathcal{X}$ is a finite subset of $\mathcal{X}$, which is to be interpreted as a finite sample of data points, and $\mathcal{X}$ has a




FIG. 7. The persistent homology pipeline uses data and their distances (X, d_X) to first construct a combinatorial object $K_{\bullet}(X, d_X)$, which is then used to construct an algebraic object $H_{\bullet}(X, d_X)$, which is finally used to construct numerical quantities $\text{dgm}_{\bullet}(X, d_X)$ (see text for more details) [89].

metric $d_{\mathcal{X}}$ on it, then X acquires a metric d_X by the formula $d_X(x_1, x_2) = d_{\mathcal{X}}(x_1, x_2)$ for all $x_1, x_2 \in X$.

- (2) From (X, d_X) , one can construct a combinatorial object called a filtered simplicial complex $K_{\bullet}(X, d_X)$, with one example being the Vietoris-Rips complex [45,85].

- (3) From the filtered simplicial complex, one can construct algebraic objects, which are the associated persistence vector spaces and homologies $H_{\bullet}(X, d_X)$ [45,86,87].

- (4) From the persistence vector spaces, one can construct numerical quantities, the persistence diagrams $\text{dgm}_{\bullet}(X, d_X)$ [88].

The original standard *quantum* TDA protocol [37] jumps in once the filtered simplicial complex $K_{\bullet}(X, d_X)$ has been constructed, and then it focuses on a small component of the persistence diagrams. Namely, rather than computing the full persistence diagrams (which provide a full invariant of the persistent homology up to isomorphism [89]), the quantum TDA algorithm is aimed at computing the persistent Betti numbers, which are the dimensions of the persistent homology groups [90], and finding the eigenvalues and eigenvectors of the combinatorial Laplacian (obtained from the graph associated with the simplicial complex). The quantum TDA algorithm takes as its starting point the simplicial complex, rather than the raw data. Briefly, if S denotes a simplicial complex constructed from a dataset with n elements for some distance threshold ϵ , then each k -simplex $s \in S$ is uniquely determined by its $k+1$ vertices. Supposing that there is a total ordering on the vertices, $v_1, v_2, \dots, v_n$, such a simplex s can therefore be expressed as $s = \{v_{i_1}, \dots, v_{i_{k+1}}\}$, where $i_1 < \dots < i_{k+1}$ and the distinct vertices $v_{i_1}, \dots, v_{i_{k+1}}$ define the simplex s . Therefore, the quantum state $|s\rangle$ associated with the simplex s would be the n -qubit state with the $i_1, \dots, i_{k+1}$ qubits in state $|1\rangle$ and all other qubits in state $|0\rangle$ (this is an example of *bit encoding*, which is discussed in more detail in Appendix E). For example, for $n = 7$ and $k = 3$, the 2-simplex $s = \{v_2, v_4, v_5\}$ would be represented by the state $|s\rangle = |0101100\rangle$.

However, one could ask whether it is possible to begin the quantum algorithm at the beginning of the pipeline by inputting the raw classical data directly onto the quantum computer and applying quantum algorithms to calculate persistence diagrams from the encoded quantum states. Although Ref. [37] suggested that this step could be done using amplitude encoding, Ref. [38] showed that amplitude encoding leads to persistence diagrams that differ quite drastically from the persistence diagrams of the original dataset. The reason for this stems from the fact that amplitude encoding distorts the distances between the data points so much (as discussed in Example 3) that the persistence diagrams obtained from the quantum encoded states no longer resemble the persistence diagrams of the original dataset. The study of how much a persistence diagram is distorted under different mappings is known as *stability* in the (classical) TDA literature, and its study leads to the following notions of structure-preserving map between finite metric spaces.

Definition 6. A *distance nonincreasing* function $f : (\mathcal{X}, d_{\mathcal{X}}) \rightarrow (\mathcal{Y}, d_{\mathcal{Y}})$ from one metric space $(\mathcal{X}, d_{\mathcal{X}})$ to another $(\mathcal{Y}, d_{\mathcal{Y}})$ is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ that satisfies $d_{\mathcal{Y}}(f(x_1), f(x_2)) \leq d_{\mathcal{X}}(x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}$. An *embed-*

ding from $(\mathcal{X}, d_{\mathcal{X}})$ to $(\mathcal{Y}, d_{\mathcal{Y}})$ is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ that satisfies $d_{\mathcal{Y}}(f(x_1), f(x_2)) = d_{\mathcal{X}}(x_1, x_2)$ for all $x_1, x_2 \in \mathcal{X}$; i.e., f is *distance preserving*.

From the categorical perspective that we will introduce soon, quantum encodings that preserve metric structure in either of these ways are said to be *morphisms* in a particular *category*. Additionally, when two mathematical objects are viewed as being equivalent in a way that respects structure via such morphisms that go back and forth between those two objects, then those two objects are said to be *isomorphic*. The following theorem identifies how preserving the metric during the encoding stage guarantees that the constructions in the persistent homology pipeline are isomorphic, whether they are built from the classical data or the quantum encoded data.

Theorem 1. Let $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ be two metric spaces, and let $X \subseteq \mathcal{X}$ be a finite subset equipped with the induced metric d_X from $d_{\mathcal{X}}$. If $f : (X, d_X) \rightarrow (\mathcal{Y}, d_{\mathcal{Y}})$ is an embedding, let $Y := f(X)$ be the image of X under f equipped with the induced metric d_Y from $d_{\mathcal{Y}}$. Then, the filtered simplicial complexes $K_{\bullet}(X, d_X)$ and $K_{\bullet}(Y, d_Y)$ are isomorphic, the persistence vector spaces $H_{\bullet}(X, d_X)$ and $H_{\bullet}(Y, d_Y)$ are isomorphic, and also the persistence diagrams $\text{dgm}_{\bullet}(X, d_X)$ and $\text{dgm}_{\bullet}(Y, d_Y)$ are all isomorphic.

For brevity, we refer the reader to Refs. [45,88,89] for more details on the definitions and proof, the latter of which follows from the fact that each of the arrows in the TDA pipeline in Fig. 7 defines a *functor*, a concept that we will define in Sec. III B. Instead, we will elaborate on the meaning of Theorem 1 in the context of quantum encodings for TDA. The main point is that if the initial quantum encoding preserves the distances exactly with respect to some suitably chosen metric on the set of quantum states, then the associated persistent homology of the quantum states obtained from encoding the classical data must agree with the persistent homology of the original data. Because different types of metrics exist on the space of quantum states (some examples will be given in the next section), one must choose both a suitable metric on the space of quantum states as well as a suitable encoding that preserves the distance from the classical data to the quantum states.

However, if one does not use a distance-preserving quantum encoding, then a more general stability theorem dictates how the persistence homologies and diagrams may change [88,89]. Namely, there is a bound relating the *Gromov-Hausdorff distance* between the original metric space of classical data and the metric space of encoded quantum data to the *bottleneck distance* between the associated persistence diagrams using distance nonincreasing maps. It remains an open problem to explicitly construct a distance nonincreasing quantum encoding that has a persistence diagram close to the one obtained from the original classical data. For brevity, we will not elaborate on the details of this here. Instead, we focus on another setting in which structure could be preserved in the context of quantum metric learning.

E. Distance for metric learning

In the previous section, we saw how topological properties of data can be preserved as long as one uses encodings that are

distance preserving or, more generally, distance nonincreasing. However, this is not always possible, since an encoding might drastically distort distances. For example, faces of an individual from slightly different angles might be considered similar by us, while the vectors of pixels might be widely separated. A similar situation occurs when encoding classical data into a quantum system. The idea of *metric learning* is to encode the notions of similarity, dissimilarity, or relative constraints between data points in order to recover a metric that accurately describes the proximity of data [91–96]. This is most practically done when the metric depends on a reasonable set of parameters that can be optimized, such as in linear metric learning (cf. Appendix C). In most of these situations, recovering a metric employs the mathematical concept of an embedding, that is, a distance-preserving function, between metric spaces, as described in the previous section.

It follows from the definition of a metric space (cf. Definition 5) that an embedding is automatically injective as a function. Moreover, any injective function from an arbitrary set $\mathcal{X}$ into a metric space $(\mathcal{Y}, d_{\mathcal{Y}})$ induces a metric on $\mathcal{X}$, as we will recall in Lemma 1 below. This is an important idea in the context of metric learning because one might not know which metric to use on $\mathcal{X}$, and the space of metrics is overwhelmingly vast and unsuitable for optimization in most cases. Instead, one parametrizes the metric through other means [91]. Namely, one fixes a suitable metric space $(\mathcal{Y}, d_{\mathcal{Y}})$, such as Euclidean space in the case of classical machine learning or the state space of a quantum system in the case of quantum machine learning (where the choice of metric in the quantum case depends on many factors) [23,95]. What varies, then, is the encoding map $f : \mathcal{X} \rightarrow \mathcal{Y}$, which is assumed to be an injective function. One chooses f according to some class of models, and the following mathematical fact allows us to define a metric $d_{\mathcal{X}}$ on $\mathcal{X}$ from the metric $d_{\mathcal{Y}}$ on $\mathcal{Y}$ and the encoding f .

Lemma 1. Let $\mathcal{X}$ be a set, $(\mathcal{Y}, d_{\mathcal{Y}})$ a metric space, $\mathcal{X} \xrightarrow{f} \mathcal{Y}$ a function, and $d_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty)$ the function

$$d_{\mathcal{X}}(x_1, x_2) := d_{\mathcal{Y}}(f(x_1), f(x_2)) \quad (26)$$

defined for all $x_1, x_2 \in \mathcal{X}$. Then, $d_{\mathcal{X}}$ is a metric on $\mathcal{X}$ if and only if f is one to one. In such a case, $(\mathcal{X}, d_{\mathcal{X}}) \xrightarrow{f} (\mathcal{Y}, d_{\mathcal{Y}})$ is an embedding.

The method of obtaining a metric on $\mathcal{X}$ in this manner via some injective function f is called *global* (possibly non-linear) metric learning [94,95,97]. In general, the parameters describing such a metric can come from two sources: (1) a family of metrics on the codomain $\mathcal{Y}$ (such as the Minkowski metrics induced by the L_p norm [91]) and (2) a family of injective maps $f : \mathcal{X} \rightarrow \mathcal{Y}$ (such as a space of injective linear transformations in linear metric learning, as described in Appendix C). The proof of Lemma 1 is given in Appendix D. The lemma itself motivates the following definition.

Definition 7. Given a set $\mathcal{X}$, a metric space $(\mathcal{Y}, d_{\mathcal{Y}})$, and a one-to-one function $\mathcal{X} \xrightarrow{f} \mathcal{Y}$, the metric $d_{\mathcal{X}}$ on $\mathcal{X}$ constructed in Lemma 1 is called the *pullback metric* or the *embedding metric*.

We can now use this idea to make rigorous sense of quantum metric learning [23]. In this setting, the codomain $\mathcal{Y}$ is

taken to be $\mathcal{S}(\mathcal{H})$, the convex space of states on a Hilbert space $\mathcal{H}$. Utilizing Lemma 1, one uses a quantum encoding map $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ together with a metric on $\mathcal{S}(\mathcal{H})$ to define the distances between data points in $\mathcal{X}$ via the embedding metric $d_{\mathcal{X}}$. Intuitively, one hopes to arrive at a metric $d_{\mathcal{X}}$ such that $(\mathcal{X}, d_{\mathcal{X}})$ is close to $(\mathcal{X}, d_{\mathcal{X}_{\text{true}}})$, where $d_{\mathcal{X}_{\text{true}}}$ is some idealized but unknown metric and where closeness can be defined with respect to the Gromov-Hausdorff distance, for example [83]. There are many options for the models used for the quantum encoding, including those from Appendix A. Moreover, there are several options for metrics on the space of states $\mathcal{S}(\mathcal{H})$. The following definition provides some possibilities of metrics on quantum states [98–103].

Definition 8. Fix a Hilbert space $\mathcal{H}$ and let ρ and σ be two density matrices on $\mathcal{H}$. The *trace/ ℓ_1 distance* between ρ and σ is

$$d_{\text{Tr}}(\rho, \sigma) = \|\rho - \sigma\|_1 \equiv \text{Tr} \left[\sqrt{(\rho - \sigma)^\dagger (\rho - \sigma)} \right]. \quad (27)$$

The *Hilbert-Schmidt/Frobenius/ ℓ_2 distance* is

$$d_{\text{HS}}(\rho, \sigma) = \sqrt{\text{Tr}[(\rho - \sigma)^\dagger (\rho - \sigma)]}. \quad (28)$$

The *Bures/Helstrom/infidelity distance* is

$$d_B(\rho, \sigma) = \sqrt{2(1 - \sqrt{F(\rho, \sigma)})}, \quad (29)$$

where

$$F(\rho, \sigma) = \left(\text{Tr} \left[\sqrt{\sqrt{\rho} \sigma \sqrt{\rho}} \right] \right)^2 \quad (30)$$

denotes the *fidelity* between the states $\rho, \sigma \in \mathcal{S}(\mathcal{H})$.

These distance functions (as well as scalar multiples of them) provide several examples of metrics on the space $\mathcal{S}(\mathcal{H})$ of quantum states. There are several other interesting possibilities, which can be found in Refs. [104–107] and the references therein. Let $d_{\mathcal{H}}$ be the notation for any one of these metrics. The goal of (classical) global metric learning is then to find an injective function $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ from which we can pull back the metric $d_{\mathcal{H}}$ to the data domain $\mathcal{X}$. One may additionally add on a parameter space Θ to define a function of the form $\rho : \mathcal{X} \times \Theta \rightarrow \mathcal{S}(\mathcal{H})$, from which a specific parameter $\theta \in \Theta$ is chosen to optimize some machine learning task. For example, this is essentially what is done in Ref. [23].

However, there are a few important remarks to be made about quantum metric learning as used in Ref. [23]. First, it is in fact a special case of global nonlinear metric learning where the space for the encoded data is the state space of a quantum system (as opposed to a standard Euclidean space $\mathbb{R}^d$, for example) [91,95]. Second, many of the encodings used in Ref. [23] are technically not embeddings in the sense defined here. This is because the quantum encodings are not always *injective* functions. We will soon illustrate this in an example, which shows that the distances between *different* points can vanish. This is important because then Lemma 1 fails, and we need a modification. This is obtained by the notion of a *semimetric space* (also called a *fuzzy metric space* and *pseudometric space*) [83,91,108],

whose importance was also recognized in the development of UMAP [46].

Definition 9. A *semimetric space* is a pair $(\mathcal{X}, d_{\mathcal{X}})$ as in the definition of a metric space, but $d_{\mathcal{X}}$ satisfies the same properties in Definition 5 except that $d_{\mathcal{X}}(x_1, x_2) = 0$ no longer implies $x_1 = x_2$; i.e., there may exist distinct points whose distance is zero.

The importance of this definition is to allow for the flexibility of quantum encodings that are not injective. Although the sampled *data* should perhaps be embedded in a one-to-one fashion, it may be that the larger data domain might not strictly embed under the encoding. This is illustrated in Fig. 8, where we have reproduced a version of Fig. 4 of Ref. [23]. The next lemma shows that a semimetric space structure on the domain of an encoding can be obtained even when the encoding is not necessarily injective.

Lemma 2. Let $\mathcal{X}$ be a set, $(\mathcal{Y}, d_{\mathcal{Y}})$ a metric space, $\mathcal{X} \xrightarrow{f} \mathcal{Y}$ a function, and $d_{\mathcal{X}} : \mathcal{X} \times \mathcal{X} \rightarrow [0, \infty)$ the function from Eq. (26). Then, $d_{\mathcal{X}}$ is a semimetric on $\mathcal{X}$.

Proof. The same argument in the proof of Lemma 1 applies with the exception of the last paragraph.

We leave the details behind Fig. 8 to Appendix D so that we may first summarize our examples and then provide the general categorical framework that encompasses them.

F. Summarizing our examples

As we conclude this brief survey of examples, observe how each paradigm in Secs. II B–II E involved the choice of a mathematical structure that is paired with the data domain $\mathcal{X}$ and quantum state space $\mathcal{S}(\mathcal{H})$. Finding a structure-preserving quantum encoding is thus a multistep process. One first selects a structure of interest, then one equips the data domain and set of states with that structure, and finally one looks for a function $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ with the property that it preserves the structure.

Although relatively straightforward, the perspective put forth in this article is that category theory provides a natural language in which to recast these design steps, potentially lending additional mathematical clarity and insight. Said explicitly, a categorical perspective reveals that designing a structure-preserving quantum encoding amounts to looking for mathematical objects $\mathcal{X}'$ and $\mathcal{S}(\mathcal{H})'$ in a particular *category* that are, in fact, $\mathcal{X}$ and $\mathcal{S}(\mathcal{H})$ equipped with extra structure, an idea that can be formalized by invoking the categorical concept of a *functor*. One then requires that the desired encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$, which is *a priori* blind to any structure, corresponds to a *morphism* $\rho' : \mathcal{X}' \rightarrow \mathcal{S}(\mathcal{H})'$ in that category. This mapping ρ' can be thought of as an improved version of ρ that not only *sees* the mathematical structure but also *respects* it. This, too, can be made precise in the language of category theory. Our goal in the next section is to provide the reader with the categorical tools needed to grasp this perspective.

III. CATEGORY THEORY CAPTURES STRUCTURE

Informally speaking, a category consists of objects and relationships, called morphisms, between them. For example,




FIG. 8. (a) Two classes of data on $\mathcal{X} = [-2, 2]$ are presented. One class has orange data points labeled $\circ$ clustered near 0 and the other class has blue data points labeled $\times$ clustered near 2 and -2. (b) The image of the quantum encoding $\rho : [-2, 2] \rightarrow \mathcal{S}(\mathbb{C}^2)$, with $\mathcal{S}(\mathbb{C}^2)$ visualized as the Bloch ball and the encoding ρ defined in Eq. (D10) using the optimized parameters $\theta_1 = 0.31$, $\theta_2 = 1.48$, $\theta_3 = 0$, and $\theta_4 = 0$. (c) A visualization of the pullback metric from the quantum encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathbb{C}^2)$ back onto the domain $\mathcal{X}$. The metric has been uniformly rescaled so that the maximum distance between points is 1. (d) The quantum circuit architecture from Ref. [23] that we used to achieve the quantum encoding shown in panel (b).

there is a category of all topological spaces and continuous functions between them as well as a category of all G -sets and G -equivariant maps between them for any group G . Definitions and examples will be presented more formally below, but as mentioned above, categories can be used to describe data domains and quantum state spaces together with any additional structure that may be imposed on them. As already hinted, examples of such structures include symmetries, topologies, smooth structures for differentiability, metrics, linear algebraic structure such as convex combinations, and more. Meanwhile, a structure-preserving quantum encoding is a morphism in one (or more) of these categories.

In Sec. III A, we will present the formal definition of a category with examples that appear in classical and quantum machine learning, focusing particularly on the examples listed throughout Sec. II. As hinted previously in Sec. II F, the main perspective of structured quantum encodings requires the concept of a functor (specifically, a *forgetful functor*), which is a particular kind of passage between categories. As such, functors will be described in Sec. III B, and many relevant examples will be provided. With this background in hand, Sec. III C will then reveal the main perspective proposed in this paper, namely, that reducing the search space from general quantum encodings to structure-preserving quantum encodings is achieved by lifting the problem from a structure-less category to one containing more structure in the presence of some forgetful functor.

A. Categories

Before presenting the formal definition of a category [7], which may appear abstract at first glance, it is helpful to know that a classic example of a category involves sets and functions between them—including quantum encodings, one of the most basic ingredients in QML applications. Let us

walk through this example slowly, making some simple observations about functions along the way. To start, observe that for any two sets $\mathcal{X}$ and $\mathcal{Y}$, there is a set of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ from $\mathcal{X}$ to $\mathcal{Y}$. Functions, though quite minimalistic, nevertheless exhibit some structure when viewed as a whole. Namely, for any two functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathcal{Z}$, their composite is also a function $g \circ f : \mathcal{X} \rightarrow \mathcal{Z}$. Moreover, this composition is associative in the sense that given three functions $f : \mathcal{X} \rightarrow \mathcal{Y}$, $g : \mathcal{Y} \rightarrow \mathcal{Z}$, and $h : \mathcal{Z} \rightarrow \mathcal{W}$, the composites $h \circ (g \circ f)$ and $(h \circ g) \circ f$ are equal. Finally, for each set $\mathcal{X}$, there is a particular function $\text{id}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}$ defined by $\text{id}_{\mathcal{X}}(x) = x$ for all $x \in \mathcal{X}$ that acts as a unit/identity for this composition since $\text{id}_{\mathcal{Y}} \circ f = f \circ \text{id}_{\mathcal{X}} = f$ for all functions $f : \mathcal{X} \rightarrow \mathcal{Y}$. These observations show us that we have some mathematical “objects” (sets) and “arrows” or “morphisms” (functions) between them that interact together in a reasonable way via the composition rule. This is a classic example of a category and is denoted by **Set**. In this way, a data domain $\mathcal{X}$ and the set $\mathcal{S}(\mathcal{H})$ of states of a quantum system are objects in the category **Set**. Furthermore, a quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ as presented in Definition 1 is a morphism in **Set**.

The formal definition of a category abstracts the example of sets and functions in the following way.

Definition 10. A category **C** consists of

- (1) a collection of *objects*, denoted as $\mathcal{X}, \mathcal{Y}, \dots$
- (2) for every pair of objects $\mathcal{X}, \mathcal{Y}$, a set of *morphisms* from $\mathcal{X}$ to $\mathcal{Y}$, depicted as arrows: $f : \mathcal{X} \rightarrow \mathcal{Y}$
- (3) a *composition rule*: for any morphisms $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathcal{Z}$, a specified morphism $g \circ f : \mathcal{X} \rightarrow \mathcal{Z}$.

Moreover, these items must satisfy the following axioms:

- (1) Composition is *associative*; that is, for any three composable morphisms $f : \mathcal{X} \rightarrow \mathcal{Y}$, $g : \mathcal{Y} \rightarrow \mathcal{Z}$, and $h : \mathcal{Z} \rightarrow \mathcal{W}$, the composites $h \circ (g \circ f)$ and $(h \circ g) \circ f$ are equal.

(2) There exist *identity morphisms*; that is, for every object $\mathcal{X}$ there exists a morphism $\text{id}_{\mathcal{X}} : \mathcal{X} \rightarrow \mathcal{X}$ such that $f \circ \text{id}_{\mathcal{X}} = \text{id}_{\mathcal{Y}} \circ f = f$ for all morphisms $f : \mathcal{X} \rightarrow \mathcal{Y}$.

Below are three additional examples of categories that have already appeared in our analysis of quantum encodings.

Example 4 (G-sets and G-equivariant functions). For a fixed group G , there is a category $G\text{-Set}$ whose objects are G -sets, where a morphism between G -sets is a G -equivariant function. The following observations show that this is indeed a category:

- (1) The composite of two G -equivariant functions is again G -equivariant.
- (2) Composition is associative due to the associativity of function composition.
- (3) For each G -set $(\mathcal{X}, \alpha)$, the identity function $\text{id}_{\mathcal{X}}$ on $\mathcal{X}$ is G -equivariant.

We have already seen a class of examples coming from geometric quantum machine learning [30,31,36], namely, G -equivariant encodings $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$, where G acts on the data domain $\mathcal{X}$ by some action α and G acts on $\mathcal{S}(\mathcal{H})$ by the adjoint action of some unitary representation of G on $\mathcal{H}$ (recall Sec. II B). Thus, such a G -equivariant encoding ρ is a morphism in the category $G\text{-Set}$. Moreover, since the composite of equivariant functions is still an equivariant function, this is what enables one to build equivariant layers in quantum circuits and still obtain an overall equivariant function.

Example 5 (Topological spaces and continuous functions). There is a category $\mathbf{Top}$ whose objects are topological spaces, where a morphism between spaces is a continuous function. The following observations show that this is indeed a category:

- (1) The composite of two continuous functions is again continuous.
- (2) Composition is associative due to the associativity of function composition.
- (3) For each topological space $(\mathcal{X}, \tau_{\mathcal{X}})$, the identity function $\text{id}_{\mathcal{X}}$ on $\mathcal{X}$ is continuous.

In Sec. II C, we saw that angle encodings and amplitude encodings are two examples of morphisms in the category $\mathbf{Top}$ when the data domain $\mathcal{X}$ (being Euclidean space) and the set of states $\mathcal{S}(\mathcal{H})$ are given their standard topologies [70,72,109]. These two examples of encodings are also morphisms in the category consisting of smooth manifolds and smooth maps between them.

Example 6 (Metric spaces and embeddings). There is a category $\mathbf{Met}^{\text{emb}}$ whose objects are metric spaces, where a morphism between metric spaces is an embedding. The following observations show that this is indeed a category:

- (1) The composite of two embeddings is an embedding.
- (2) Composition is associative due to the associativity of function composition.
- (3) For each metric space $(\mathcal{X}, d_{\mathcal{X}})$, the identity function $\text{id}_{\mathcal{X}}$ on $\mathcal{X}$ is an embedding.

The category $\mathbf{Met}^{\text{emb}}$ appeared in Sec. II D for TDA and Sec. II E for quantum metric learning, where one equips the

set of states $\mathcal{S}(\mathcal{H})$ with a metric and uses a quantum state encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ to define a metric on the data domain $\mathcal{X}$.

Our discussion surrounding TDA in Sec. II D likewise involved metric spaces but additionally included more general morphisms between them. Namely, we considered distance nonincreasing functions rather than embeddings. More generally, one might wish to adjust the extent to which a function preserves distances, which can lead to different categories having the same objects but different morphisms. The following example summarizes this observation.

Example 7 (Metric spaces and distance nonincreasing functions). There is a category $\mathbf{Met}^{\text{nif}}$ whose objects are metric spaces $(\mathcal{X}, d_{\mathcal{X}})$. A morphism in $\mathbf{Met}^{\text{nif}}$ between metric spaces is a distance nonincreasing function. As before, this is a category since the following criteria are satisfied:

- (1) The composite of two distance nonincreasing functions is a distance nonincreasing function.
- (2) Composition is associative due to the associativity of function composition.
- (3) For each metric space $(\mathcal{X}, d_{\mathcal{X}})$, the identity function $\text{id}_{\mathcal{X}}$ on $\mathcal{X}$ is distance nonincreasing.

Notice that embeddings between metric spaces are also distance nonincreasing. For this reason, $\mathbf{Met}^{\text{emb}}$ is called a *subcategory* of $\mathbf{Met}^{\text{nif}}$, meaning that all objects and morphisms of $\mathbf{Met}^{\text{emb}}$ are also objects and morphisms of $\mathbf{Met}^{\text{nif}}$, and the identities and composition rules agree on these common objects and morphisms.

Each of the four examples presented here describes categories whose objects are sets with structure (a group action, a topology, and a metric) and whose morphisms are functions that preserve (part of) that structure. We have also seen that some quantum encodings, when chosen appropriately, can be viewed as morphisms in these categories. *A priori*, however, a quantum encoding is merely a function between sets, as was introduced in Definition 1. In QML, therefore, one sometimes works within the category $\mathbf{Set}$, while at other times one works in a different category, such as those listed in Examples 4–7. So, to take account of when a passage is being made from one category to another, we must understand mappings between categories. This leads to the concept of a functor.

B. Functors

While categories are useful for organizing mathematical objects and relationships (morphisms) between them, much of mathematics concerns relationships between *categories* themselves. This involves assigning the objects and morphisms of one category to another in a manner that respects morphism composition and identities.

To situate this in the context of QML, think back to G -equivariant quantum state encodings for some group G . Given a unitary representation $V : G \rightarrow \mathcal{U}(\mathcal{H})$ of G , one obtains a G -equivariant function $\rho : (\mathcal{X}, \alpha) \rightarrow (\mathcal{S}(\mathcal{H}), \text{Ad}_V)$ which, as we have seen, is a morphism in the category $G\text{-Set}$. Here, there is an obvious way to pass from $G\text{-Set}$ to the category $\mathbf{Set}$: simply discard, or forget, the group actions associated with each set (cf. Fig. 9), and assign each G -equivariant function to the function itself. Although straightforward, let us denote this assignment by $F : G\text{-Set} \rightarrow \mathbf{Set}$, so that for an

$$\begin{array}{ccc}
G\text{-Set} & \xrightarrow{\text{forget action}} & \mathbf{Set} \\
\\
\text{data domain} & (\mathcal{X}, \alpha) \xrightarrow{\text{forget } \alpha} & \mathcal{X} \\
\\
\text{state space} & (\mathcal{S}(\mathcal{H}), \text{Ad}_V) \xrightarrow{\text{forget } \text{Ad}_V} & \mathcal{S}(\mathcal{H}) \\
\\
\text{arbitrary } G\text{-set} & (\mathcal{Y}, \beta) \xrightarrow{\text{forget } \beta} & \mathcal{Y}
\end{array}$$

FIG. 9. Forgetting the action on a G -set and just remembering the underlying set defines a functor from $G\text{-Set}$ to $\mathbf{Set}$. This functor can be applied to every G -set, including data domains and state spaces, regardless of the G -actions they come equipped with.

arbitrary G -set $(\mathcal{X}, \alpha)$, we have $F(\mathcal{X}, \alpha) = \mathcal{X}$ and for any G -equivariant function f , we have $F(f) = f$. Then, F behaves well with respect to the two categories in the following ways.

First, for any two composable G -equivariant functions f and g , we have, somewhat trivially,

$$F(g \circ f) = g \circ f = F(g) \circ F(f). \quad (31)$$

In words, this means that composing two G -equivariant functions and *then* discarding the group actions on their domains and codomains results in the same function obtained by *first* discarding the group actions and then composing the underlying functions. In more concise terms, F preserves the composition rules of the two categories.

Second, for any G -set $(\mathcal{X}, \alpha)$, we have that

$$F(\text{id}_{(\mathcal{X}, \alpha)}) = \text{id}_{(\mathcal{X}, \alpha)} = \text{id}_{\mathcal{X}} = \text{id}_{F(\mathcal{X}, \alpha)}. \quad (32)$$

The first and third equalities follow from the definition of F on morphisms and objects, respectively. The second equality follows from the fact that the identity functions on $\mathcal{X}$, when $\mathcal{X}$ is viewed as a G -set or as a (regular) set, are both defined by $x \mapsto x$ for all $x \in \mathcal{X}$. So, forgetting the group action and *then* considering the identity function on the underlying set results in the same function obtained by *first* considering the G -equivariant identity function and then forgetting about the group action. Succinctly put, F preserves identity morphisms.

In summary then, $F : G\text{-Set} \rightarrow \mathbf{Set}$ is an assignment on the objects and morphisms of one category to another that preserves the categorical structure. This is an example of a functor.

Definition 11. A functor $F : \mathbf{C} \rightarrow \mathbf{D}$ between categories consists of

- (1) an object $F(\mathcal{X})$ in $\mathbf{D}$ for every object $\mathcal{X}$ in $\mathbf{C}$
- (2) a morphism $F(f) : F(\mathcal{X}) \rightarrow F(\mathcal{Y})$ in $\mathbf{D}$ for every morphism $f : \mathcal{X} \rightarrow \mathcal{Y}$ in $\mathbf{C}$.

Moreover, these items must satisfy the following axioms:

- (1) Composition is preserved; that is, $F(g \circ f) = F(g) \circ F(f)$ for all morphisms $f : \mathcal{X} \rightarrow \mathcal{Y}$ and $g : \mathcal{Y} \rightarrow \mathcal{Z}$ in $\mathbf{C}$.

- (2) Identities are preserved; that is, $F(\text{id}_{\mathcal{X}}) = \text{id}_{F(\mathcal{X})}$ for every object $\mathcal{X}$ in $\mathbf{C}$.

A few more examples of functors will appear below, but to help connect these upcoming examples to the goals of this article, let us first establish some terminology relating to

properties that functors may satisfy. Let $\mathbf{C}(\mathcal{X}, \mathcal{Y})$ denote the set of all morphisms from object $\mathcal{X}$ to object $\mathcal{Y}$ in a category $\mathbf{C}$. Given another category $\mathbf{D}$, if a functor $F : \mathbf{C} \rightarrow \mathbf{D}$ has the property that for all objects $\mathcal{X}$ and $\mathcal{Y}$ in $\mathbf{C}$, the function $\mathbf{C}(\mathcal{X}, \mathcal{Y}) \rightarrow \mathbf{D}(F(\mathcal{X}), F(\mathcal{Y}))$ defined by $f \mapsto F(f)$ is injective, then F is said to be *faithful*. If the function $f \mapsto F(f)$ is surjective, then F is said to be *full*. If it is a bijection, then F is said to be *fully faithful*. As the next example illustrates, we have already seen an example of a functor that is faithful but not, however, full.

Example 8. The functor $F : G\text{-Set} \rightarrow \mathbf{Set}$ described prior to Definition 11 (which “forgets” the group action of G -sets and G -equivariant functions to obtain the underlying sets and functions between them) is faithful since two G -equivariant functions are equal whenever they are equal as functions on the underlying sets. However, F is not full since there exist functions between the underlying sets of G -sets that are not equivariant. The lack of fullness is what allows us to narrow down the large space of quantum encodings to a significantly smaller subset of structure-preserving (in this case, equivariant) encodings. Hence, the smaller space of equivariant encodings generally simplifies the search for optimal encodings.

One of the major themes throughout this article has been that many mathematical objects consist of sets equipped with extra structure, so the ability to forget or omit that structure certainly extends far beyond group theory. Below are additional examples of functors relating to our discussions of quantum encodings, which are analogous to the one mentioned in Example 8.

Example 9 (From metric spaces to sets). There are functors $\mathbf{Met}^{\text{emb}} \rightarrow \mathbf{Set}$ and $\mathbf{Met}^{\text{nif}} \rightarrow \mathbf{Set}$ that both assign to every metric space $(\mathcal{X}, d_{\mathcal{X}})$ its underlying set of points $\mathcal{X}$, thus forgetting the distance function $d_{\mathcal{X}}$. Moreover, the functors also both assign every embedding and distance nonincreasing function to itself. This functor is faithful since two embeddings or two distance nonincreasing functions are equal whenever they are equal as functions on the underlying sets. However, neither of these functors is full because not every function is distance preserving nor distance nonincreasing. The lack of fullness in this case restricts the set of all functions to those that do preserve distances or are distance nonincreasing.

Example 10 (From topological spaces to sets). There is a functor $\mathbf{Top} \rightarrow \mathbf{Set}$ that assigns to every topological space $(\mathcal{X}, \tau_{\mathcal{X}})$ its underlying set $\mathcal{X}$ of points, thus forgetting the topology $\tau_{\mathcal{X}}$. Moreover, the functor assigns every continuous function to itself. As in the previous examples, this functor is faithful but not full. Again, the lack of fullness is what causes the space of morphisms between objects to become smaller when adding structure in the sense that the set of *continuous* functions between topological spaces is in general smaller than the set of *all* functions between the underlying sets.

Perhaps not surprisingly, each of the functors listed in Examples 8–10 are commonly referred to as “forgetful functors.” Loosely speaking, a functor $\mathbf{C} \rightarrow \mathbf{D}$ is called *forgetful* if it drops, or omits, some or all of the structure of the objects in $\mathbf{C}$.

Though this is not a completely rigorous definition, the notion of a forgetful functor as used in this article will always be of this kind, so a formal definition of this notion is not needed here. The takeaway, though, is that each of the functors in the previous examples was also *faithful*, and faithfulness can be used to formalize the notion of “structure” itself. One may say that objects in a category $\mathbf{C}$ are objects in a category $\mathbf{D}$ *equipped with extra structure* if there exists a faithful functor $\mathbf{C} \rightarrow \mathbf{D}$ [110]. In the special case when $\mathbf{D} = \mathbf{Set}$ is the category of sets, $\mathbf{C}$ is further said to be a *concrete category* [39]. All of the examples discussed so far—**Set**, **G-Set**, **Met^{emb}**, **Met^{nif}**, and **Top**—are concrete categories, as their objects are always sets $\mathcal{X}$ equipped with extra structure (a group action α , a metric $d_{\mathcal{X}}$, or a topology $\tau_{\mathcal{X}}$), and their morphisms are functions that preserve that structure. Although not all categories are concrete [111], we are primarily focusing on those that are.

At last, the reward of introducing this high level of abstraction is that we now have precise mathematical language to describe the design of structure-preserving quantum encodings.

C. Quantum encoding from a categorical perspective

Let us finally reformulate the QML encoding scheme for states from a categorical perspective. (A similar reformulation exists for unitary quantum encodings.) This idea constitutes the main perspective proposed in this paper.

Design setup: For a given learning task, start with a data domain $\mathcal{X}$ and choose a Hilbert space $\mathcal{H}$. Identify a concrete category $\mathbf{C}$ and equip $\mathcal{X}$ and the set of states $\mathcal{S}(\mathcal{H})$ with extra structure to obtain objects $\mathcal{X}'$ and $\mathcal{S}(\mathcal{H})'$ in $\mathbf{C}$ such that $F(\mathcal{X}') = \mathcal{X}$ and $F(\mathcal{S}(\mathcal{H})') = \mathcal{S}(\mathcal{H})$ are in the image of a faithful forgetful functor $F : \mathbf{C} \rightarrow \mathbf{Set}$.

Design goal: Find a function $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ with the property that there exists a morphism $\rho' : \mathcal{X}' \rightarrow \mathcal{S}(\mathcal{H})'$ in $\mathbf{C}$ such that $F(\rho') = \rho$. The morphism ρ' is the desired, structure-preserving quantum encoding.

The process described in the design goal can be illustrated diagrammatically. Among all of the set-theoretic functions $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ in **Set**, there might exist a few that can be “lifted” to a morphism $\rho' : \mathcal{X}' \rightarrow \mathcal{S}(\mathcal{H})'$ in $\mathbf{C}$, the category of structure-preserving morphisms, as in the following diagram.

$$\begin{array}{ccc} \mathbf{C} & & \mathcal{X}' \text{ --- } \rho' \text{ --- } \mathcal{S}(\mathcal{H})' \\ \downarrow F & & \downarrow \quad \quad \downarrow \\ \mathbf{Set} & & \mathcal{X} \xrightarrow{\rho} \mathcal{S}(\mathcal{H}) \end{array}$$

Since the functor F is faithful, and rarely full, the space of all encodings ρ in **Set** is reduced to a smaller space of structure-preserving encodings ρ' in $\mathbf{C}$. In this way, searching for those encodings ρ that lift to such a ρ' might provide a method by which to simplify the search space when designing quantum encodings that are adapted to the problem being considered.

This seemingly abstract design goal is made to be compatible with our earlier examples.

Example 11. Let $\mathcal{X}' := (\mathcal{X}, \alpha)$ be a data domain that is a G -set, i.e., an object of **G-Set**, which means that the data domain is equipped with a symmetry coming from a group action. Also, if $G \ni g \mapsto V_g \in \mathcal{U}(\mathcal{H})$ is a unitary representation of a group G on the Hilbert space $\mathcal{H}$, with the induced action of G on $\mathcal{S}(\mathcal{H})$ via $G \ni g \mapsto \text{Ad}_{V_g}$, then $\mathcal{S}(\mathcal{H})' := (\mathcal{S}(\mathcal{H}), \text{Ad}_V)$ is an object of **G-Set**. If $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ is a quantum state encoding that is G -equivariant, which means that the symmetry is reflected in the quantum encoding, then the quantum encoding “lifts” to a morphism $\rho' : \mathcal{X}' \rightarrow \mathcal{S}(\mathcal{H})'$ in **G-Set**. In this case, the general diagram we drew earlier becomes

$$\begin{array}{ccc} \mathbf{G-Set} & & (\mathcal{X}, \alpha) \text{ --- } \rho' \text{ --- } (\mathcal{S}(\mathcal{H}), \text{Ad}_V) \\ \downarrow F & & \downarrow \quad \quad \downarrow \\ \mathbf{Set} & & \mathcal{X} \xrightarrow{\rho} \mathcal{S}(\mathcal{H}) \end{array}$$

where $F : \mathbf{G-Set} \rightarrow \mathbf{Set}$ is the forgetful functor from Example 8. Note that the map ρ in **G-Set** is the same ρ as in **Set** when viewed as a function (hence the same notation), which is because F is faithful. Thus, by restricting to the subset of those ρ that lift from **Set** to **G-Set**, we search within a smaller space of suitable embeddings of G -equivariant morphisms in **G-Set**.

This categorical proposal extends far beyond the setting of geometric quantum machine learning. We have already described other examples involving topological structure, smooth structure, and metric structure. Additionally, there is a category **LieOPS**, defined in Appendix A, that identifies the structure preserved under quantum unitary encodings $U : \mathbb{R}^d \rightarrow \mathcal{U}(\mathcal{H})$ of the form

$$U(x) = e^{-i\mathcal{L}(x)}, \quad (33)$$

where $\mathcal{L} : \mathbb{R}^d \rightarrow \mathcal{B}(\mathcal{H})$ is a linear transformation sending each point $x \in \mathbb{R}^d$ to a Hermitian operator $\mathcal{L}(x)$ acting on a Hilbert space $\mathcal{H}$ [here, $\mathcal{B}(\mathcal{H})$ denotes the set of bounded operators on $\mathcal{H}$]. Such encodings are common in the quantum machine learning literature, and we have already seen several examples in the context of angle encoding and in geometric quantum machine learning. As proven in Appendix A, such morphisms are obtained by first equipping sets with a structure related to one-parameter subgroups (OPSs). In essence, the morphisms in **LieOPS** preserve these one-parameter subgroups, and Theorem A2 proves that all unitary encodings $U : \mathbb{R}^d \rightarrow \mathcal{U}(\mathcal{H})$ that are morphisms in **LieOPS** (i.e., that preserve one-parameter subgroups) are of the form in Eq. (33). This structure reduces the space of quantum encodings to a *finite-dimensional* space. Combining this with other structures allows one to further reduce the space of encodings relevant in a quantum machine learning task. Indeed, we have already seen an example worked out in Sec. IIB in the setting of geometric quantum machine learning, where the preservation of one-parameter subgroups together with equivariance led to a simple description of Lie algebra generators allowed for equivariant encodings.

IV. DISCUSSION, OPEN QUESTIONS, AND OUTLOOKS

Our primary focus in this perspective paper was to illustrate how category theory can be used to organize structure in order to guide the encoding of data onto quantum systems. Different datasets and tasks admit structures that are relevant to specific problems, and identifying the appropriate category of such structures enables one to isolate a subset of quantum encodings preserving those structures. Thus, identifying structure is the first step toward finding encodings that optimize desired properties. We supported our claim with the example of geometric quantum machine learning [30–32,36]. Here, the relevant structure isolated from the category of group actions on sets is equivariance. We illustrated this perspective through several other examples to highlight that geometric quantum machine learning forms a single instance of this categorical formalism. Since finding a mathematical framework to isolate what we mean by structure in quantum machine learning tasks has been a long-standing problem [5,36,112], we anticipate that such a perspective will offer insight toward extending the successes of geometric quantum machine learning to situations where other structures are preserved.

A. Extending the successes of category theory from classical to quantum settings

There is a lot of supporting evidence to suggest that a categorical perspective is worth pursuing. First of all, there is already an active discussion regarding the effectiveness of categorical techniques in *classical* machine learning and computer science more broadly. In the body of this work, we mentioned two quintessential success stories where category theory has greatly enhanced data analysis techniques, namely, UMAP and TDA [45,46]. Others include functional programming [47], probabilistic programming [48], the architecture of software engineering [49], neural network symmetrization [50,51], backpropagation [52], and understanding the general algebraic structure of deep learning [53]. The plethora of examples illustrates that category theory has the extraordinary flexibility to describe a variety of tasks across many classical machine learning domains.

The effectiveness of category theory bringing clarity and advancements to *classical* machine learning domains suggests that it could also be capitalized in areas of *quantum* machine learning that have been otherwise challenging to address. Two recent examples where algorithms have been improved due to isolating relevant structure is in variational circuit design by utilizing the Riemannian geometry of the special unitary group [113] and isolating algebraic structure for block encodings [114].

A concrete example of how category theory could potentially benefit quantum machine learning by isolating the appropriate structure that goes outside the context of geometric quantum machine learning is TDA performed on quantum computers, a topic we discussed earlier. In Sec. IID, we argued that if one encodes the raw classical data directly onto a quantum computer, in order to preserve the topological features of the original data, a categorical perspective suggests that the encoding should preserve the pointwise distances

between the data point as much as possible. This leads to a concrete mathematical problem, which is to isolate an encoding $\rho : X \rightarrow \mathcal{S}(\mathcal{H})$ from the space of data points X to the state space of a quantum system $\mathcal{S}(\mathcal{H})$ that best preserves the distances, i.e., comes close to being an embedding. To formalize this more, one chooses metrics on X and $\mathcal{S}(\mathcal{H})$. Since often the encoding is done using pure quantum states, one may view the space of pure quantum states as a complex projective space with the Fubini-Study metric. The mathematical problem then becomes to find an encoding ρ that minimizes the distortion between the original distances in X and the quantum encoded distances in complex projective space. The analog of this problem for *Euclidean space* is known as classical multidimensional scaling [115], but the present example requires extending this approach to complex projective space [116]. An optimal solution ρ_{opt} in the context of quantum encodings for the Fubini-Study metric would provide a *lower bound* on the distortion. In other words, *every* quantum encoding would have distortion at least as great as the optimal encoding. This approach then informs us how to search for an encoding that minimizes the distortion. By taking a proposed encoding and calculating its distortion, one could infer how far the encoding is from the optimal one. Namely, an encoding that has distortion that is close to the optimal one would therefore approximately preserve the topological features of the data by the stability theorems of TDA. Such investigations are part of ongoing work, and the results will be reported elsewhere.

B. Pushing beyond open challenges with a categorical perspective

There are many open questions and avenues of research that a categorical perspective brings to surface. The first set of questions involves the construction of a full quantum algorithm that not only preserves the structure of the data, but also leverages that structure. In this perspective paper, we have only provided the basic ingredients for utilizing category theory in the context of QML at the level of the state encoding. A next step is to incorporate learning tasks, such as classification or regression, which could itself help isolate the relevant structure. Another step is to carefully select measurements that respect the structure. Putting these components together would help design the entire circuit with the structure in mind. In this process of developing a full quantum model, one could account for the structure by identifying the appropriate category that preserves the structure.

Another important set of questions involves quantum advantage. The current viewpoint is that structure in a dataset or task is important in determining possible quantum advantage [3–5,36,117]. For instance, Bowles *et al.*, through a number of experiments, argued that linearly separable data do not lend themselves to quantum advantage [5]. As another example, the recently developed method of dequantization by Tang shows that certain classes of quantum algorithms perform just as well as some classical algorithms due to the additional structure of the dataset or problem [20,65,118–120]. Thus, a fruitful direction for the QML community is to isolate structure that would render QML algorithms difficult to simulate classically. For example, an understanding of mathematical structure, as given by equivariance, has led to deeper insights

into the barren plateau problem and classical simulability [17,121]. In particular, it seems natural to find the trade-offs between structure preservation, expressivity, generalization, performance, efficiency, and classical simulability, in order to determine a potential quantum advantage [17,66], especially for noisy intermediate scale quantum devices [122]. Alternatively, should we be asking a different question besides quantum advantage and instead aim at achieving different goals [112]? We hope that the categorical perspective offered in this paper may help materialize such questions more concretely and hence push the field forward.

In summary, the perspective offered in this paper highlights that thinking categorically about quantum machine learning models provides a guide for isolating mathematical structure within the pipeline of quantum algorithm design. In particular, it can help practitioners make informed choices when benchmarking models [5]. Should a quantum circuit allow data to be distorted in terms of their original geometry or symmetries? What structures should be preserved for a given task or problem? What is the category of such structure-preserving encodings, circuits, and measurements? If the structure is well studied, such as equivariance or metric embeddings, then the QML user needs only to identify that structure and capitalize on the mathematics that has presumably already been heavily developed for that structure. In some instances, it could be the case that the particular combinations of structures involved might lead to areas of mathematics that deserve more attention. This would therefore not only be fruitful for quantum machine learning, but it would also potentially allow those structures to be further developed and applied beyond these settings and encourage more cross-disciplinary collaborations.

ACKNOWLEDGMENTS

A.J.P. thanks Cheyne Glass, Seth Lloyd, and Alexander Schmidhuber for discussions. We thank Zachary Bradshaw for discussions and comments on our first version of the manuscript. We also thank the anonymous reviewers for their helpful feedback.

A.J.P. has received financial support from Deloitte in his involvement with this project. He carried out this project as a consultant to Deloitte and not as part of his MIT responsibilities.

DATA AVAILABILITY

No data were created or analyzed in this study.

APPENDIX A: ONE-PARAMETER GROUPS FOR QUANTUM ENCODINGS

Many of the unitary encodings $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$ considered in quantum machine learning are of the form

$$U(x) = e^{-i\mathcal{L}(x)}, \quad (\text{A1})$$

where $\mathcal{L} : \mathcal{X} \rightarrow \mathcal{B}(\mathcal{H})$ is a linear transformation sending each point $x \in \mathcal{X}$ (from a real vector space $\mathcal{X}$) to a Hermitian operator $\mathcal{L}(x)$ acting on a Hilbert space $\mathcal{H}$, where $\mathcal{B}(\mathcal{H})$ denotes the set of bounded operators on $\mathcal{H}$ [21,30,59]. One then obtains a quantum state encoding by fixing some fiducial state

$|\psi_0\rangle \in \mathcal{H}$ via the formula $\rho(x) = U(x)|\psi_0\rangle\langle\psi_0|U(x)^\dagger$. Examples include the one from Sec. II B on geometric quantum machine learning as well as angle encoding from Sec. II C. These types of examples of quantum encodings can also be viewed as preserving a certain structure. The type of structure preserved is a bit more technical than the types of structures we have considered in the body of the paper, so we include the discussion here in this Appendix for completeness and because of the fact that unitary encodings of the form in Eq. (A1) are ubiquitous in the quantum machine learning community. We begin with a reminder of one-parameter subgroups of Lie groups (all Lie groups here will be matrix Lie groups for simplicity) [31,64,123–125].

Definition A1. Let G be a Lie group. A continuous one-parameter subgroup of G is a continuous group homomorphism $\gamma : \mathbb{R} \rightarrow G$, where $\mathbb{R}$ is viewed as a group under addition, i.e.,

- (1) γ is a continuous function,
- (2) $\gamma(0) = 1$, and
- (3) $\gamma(s+t) = \gamma(s)\gamma(t)$ for all $s, t \in \mathbb{R}$.

Recall that every one-parameter subgroup $\gamma : \mathbb{R} \rightarrow G$ of a Lie group G is generated by a unique element $M \in \mathfrak{g}$ in the Lie algebra $\mathfrak{g}$ by the exponential map $\exp : \mathfrak{g} \rightarrow G$. Namely,

$$\gamma(t) = \exp(tM) \quad (\text{A2})$$

for all $t \in \mathbb{R}$. In particular, this implies γ is not only continuous but smooth as well. Note that by differentiating the curve γ at $t = 0$ (cf. Sec. 1.3 in Ref. [71]), one obtains $M = \frac{d}{dt}\gamma|_{t=0}$. This example is precisely what is used when one has a data domain $\mathcal{X}$ of the form $\mathcal{X} = \mathbb{R}$, since in this case $x \mapsto e^{-ixL}$, with $L \in \mathfrak{su}(2^n)$ for a quantum unitary encoding mapping one-dimensional classical data onto n qubits. It is also worth noting that the exponential map $\exp : \mathbb{R}^d \rightarrow \mathbb{R}^d$ from the Lie algebra of the additive Lie group $\mathbb{R}^d$ to the additive Lie group $\mathbb{R}^d$ is the identity map, $\exp(x) = x$ for all $x \in \mathbb{R}^d$.

However, for classical data that are provided in more than one dimension, we need to be careful about the ordering of operators obtained from such exponentials. To expound on this and to make more rigorous comparisons later, we recall the following result relating Lie group homomorphisms and Lie algebra homomorphisms. (For this theorem, recall that $[V, W] := VW - WV$ denotes the commutator.)

Theorem A1. Let G be any connected Lie group, let H be any Lie group, and let $f : G \rightarrow H$ be a Lie group homomorphism. Then, there exists a unique real-linear map $\mathcal{M} : \mathfrak{g} \rightarrow \mathfrak{h}$ such that

$$f(\exp(W)) = \exp(\mathcal{M}(W)) \quad (\text{A3})$$

for all $W \in \mathfrak{g}$. Moreover, $\mathcal{M}$ is a Lie algebra homomorphism, meaning that

$$[\mathcal{M}(W_1), \mathcal{M}(W_2)] = \mathcal{M}([W_1, W_2]) \quad (\text{A4})$$

for all $W_1, W_2 \in \mathfrak{g}$.

This theorem is well known (see Theorem 3.28 in Ref. [64], for example), so we will not give a proof here because we will need a generalization later whose proof will be different. In contrast to the setting of Theorem A1, the types of morphisms

between Lie groups that describe quantum encodings of the form in Eq. (A1) are *not* Lie group homomorphisms in general, primarily due to the lack of commutativity of arbitrary unitary gates in quantum circuits. A simple example illustrating this can be seen with a two-dimensional data domain $\mathcal{X} = \mathbb{R}^2$. Suppose that $U : \mathcal{X} \rightarrow \mathcal{U}(\mathbb{C}^2)$ is a quantum unitary encoding of the form

$$U((s, t)) = e^{-isX - itY} \quad (\text{A5})$$

for all $(s, t) \in \mathbb{R}^2$. Then, this defines a perfectly reasonable quantum unitary encoding map that additionally satisfies Eq. (A3), as will be explained momentarily. However, note that Eq. (A5) does not define a group homomorphism because

$$\begin{aligned} U((s, 0))U((0, t)) &= e^{-isX} e^{-itY} \\ &\neq e^{-i(sX + tY)} \\ &= U((s, 0) + (0, t)) \end{aligned} \quad (\text{A6})$$

in general.

In more detail, we still have the following properties associated with Eq. (A5). First, let $\mathcal{M} : \mathbb{R}^2 \rightarrow \mathfrak{su}(2)$ be the linear map uniquely determined by $\mathcal{M}((1, 0)) = -iX$ and $\mathcal{M}((0, 1)) = -iY$, so that $\mathcal{M}((s, t)) = -isX - itY$ for all $(s, t) \in \mathbb{R}^2$. Since the exponential map $\exp : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ going from the Lie algebra of the Lie group $\mathbb{R}^2$ (with group structure given by addition) to the Lie group $\mathbb{R}^2$ coincides with the identity map, Eq. (A3) still holds and it agrees with Eq. (A5). Therefore, the assumption that $f : G \rightarrow H$ is a group homomorphism in Theorem A1 is not necessary for Eq. (A3) to hold. Second, for every one-dimensional additive subgroup $\mathbb{R}_\alpha \subset \mathbb{R}^2$ generated by some nonzero vector α , i.e.,

$$\mathbb{R}_\alpha = \{r\alpha : \alpha \in \mathbb{R}^2 \setminus \{0\}, r \in \mathbb{R}\}, \quad (\text{A7})$$

the restriction of U to this subgroup does indeed define a Lie group homomorphism $U|_{\mathbb{R}_\alpha} : \mathbb{R}_\alpha \rightarrow \mathcal{U}(\mathbb{C}^2)$ since

$$\begin{aligned} U(r\alpha)U(u\alpha) &= e^{-irW} e^{-iuW} \\ &= e^{-i(r+u)W} \\ &= U(r\alpha + u\alpha), \end{aligned} \quad (\text{A8})$$

where $W := \alpha \cdot (X, Y)$ is the dot product of α with the vector of operators (X, Y) . The second equality in Eq. (A8) holds because every matrix W commutes with itself. In other words, the one-parameter subgroup $\mathbb{R} \xrightarrow{\alpha} \mathbb{R}^2$ of $\mathbb{R}^2$ sending $r \in \mathbb{R}$ to $\alpha(r) := r\alpha$ is a one-parameter subgroup of $\mathbb{R}^2$ that gets pushed forward to a one-parameter subgroup

$$\mathbb{R} \xrightarrow{\alpha} \mathbb{R}^2 \xrightarrow{U} \mathcal{U}(\mathbb{C}^2) \quad (\text{A9})$$

of $\mathcal{U}(\mathbb{C}^2)$ that sends $r \in \mathbb{R}$ to $U(r\alpha)$ [126].

The previous discussion hints that the quantum unitary encodings (A1) are more closely related to one-parameter subgroups when restricted to one-dimensional Lie subalgebras. We make this precise by introducing the category **LieOPS** of Lie groups and one-parameter subgroup homomorphisms, which makes use of the notion of a smooth map between smooth manifolds [71,72].

Definition A2. Let **LieOPS** be the category (cf. Sec. III A) whose objects are Lie groups and where a morphism $f : G \rightarrow H$ from a Lie group G to a Lie group H is a *one-parameter*

subgroup homomorphism; i.e., f is a smooth map and for every one-parameter subgroup $\mathbb{R} \xrightarrow{\gamma} G$ in G , the composite $\mathbb{R} \xrightarrow{\gamma} G \xrightarrow{f} H$ is a one-parameter subgroup in H .

One can show that **LieOPS** is indeed a category. It is important to note that although every Lie group homomorphism is an OPS homomorphism, the converse is not true. Namely, not every OPS homomorphism is a Lie group homomorphism, with Eq. (A5) providing an explicit example. The following theorem is a generalization of Theorem A1 and provides a characterization for how OPS homomorphisms can always be written in a form analogous to quantum unitary encodings as in Eq. (A1).

Theorem A2. Let $G \xrightarrow{f} H$ be an OPS homomorphism of Lie groups. Then, there exists a unique linear map $\mathcal{M} : \mathfrak{g} \rightarrow \mathfrak{h}$ such that

$$f(\exp(x)) = \exp(\mathcal{M}(x)) \quad (\text{A10})$$

for all $x \in \mathfrak{g}$.

Our proof of Theorem A2 will be different than the proof of Theorem A1 given in Ref. [64] because we cannot assume that $f : G \rightarrow H$ is a group homomorphism. We assume some tools from differential geometry in the following proof [71,72,127].

Proof. First, let $\mathcal{M} : \mathfrak{g} \rightarrow \mathfrak{h}$ be the *differential* (sometimes called the *pushforward*) of f at the identity $1_G \in G$, i.e., $\mathcal{M} := Df|_{1_G}$. To be somewhat self-contained, let us recall how $\mathcal{M}$ is defined (the following definition is valid whenever f is a smooth map between smooth manifolds, not necessarily Lie groups). For any smooth curve $\gamma : \mathbb{R} \rightarrow G$ with $\gamma(0) = 1_G$ and $\frac{d}{dt}\gamma|_{t=0} = x \in \mathfrak{g}$, we have $\frac{d}{dt}(f \circ \gamma)|_{t=0} = \mathcal{M}(x)$. The fact that $\mathcal{M}$ defines a linear map follows from Exercise 17 in Sec. 1.3 of Ref. [71] or Proposition 3.6 of Ref. [72].

Now suppose that γ is a one-parameter subgroup of G with $\frac{d}{dt}\gamma|_{t=0} = x$. By Eq. (A2) for γ , this means that

$$\gamma(t) = \exp(tx) \quad (\text{A11})$$

for all $t \in \mathbb{R}$. Since f is an OPS homomorphism, the composite $\mathbb{R} \xrightarrow{\gamma} G \xrightarrow{f} H$ is a one-parameter subgroup of H . Hence, by Eq. (A3) for $f \circ \gamma$, there exists a unique element $y \in \mathfrak{h}$ such that

$$(f \circ \gamma)(t) = \exp(ty) \quad (\text{A12})$$

for all $t \in \mathbb{R}$. Putting these results together, we have

$$\begin{aligned} y &= \frac{d}{dt} \exp(ty)|_{t=0} = \frac{d}{dt} (f \circ \gamma)|_{t=0} \quad \text{by Eq. (A12)} \\ &= Df|_{1_G} \left(\frac{d}{dt} \gamma|_{t=0} \right) \quad \text{by the chain rule} \\ &= Df|_{1_G} \left(\frac{d}{dt} \exp(tx)|_{t=0} \right) \quad \text{by Eq. (A11)} \\ &= \mathcal{M}(x) \quad \text{by definition of } \mathcal{M}. \end{aligned} \quad (\text{A13})$$

Since x was arbitrary, this together with applying f to Eq. (A11) proves that

$$f(\exp(tx)) = \exp(t\mathcal{M}(x)) \quad (\text{A14})$$

for all $t \in \mathbb{R}$ and $x \in \mathfrak{g}$.

We mention an immediate corollary relevant for quantum unitary encodings by taking $G = \mathbb{R}^d$ to be the additive group and $\mathfrak{g} = \mathbb{R}^d$ to be its associated Lie algebra. Note that in this case, $\exp : \mathfrak{g} \rightarrow G$ is the identity map as a function.

Corollary A1. Let $\mathbb{R}^d \xrightarrow{f} G$ be an OPS homomorphism of Lie groups. Then, there exists a unique linear map $\mathcal{M} : \mathbb{R}^d \rightarrow \mathfrak{g}$ such that

$$f(x) = \exp(\mathcal{M}(x)) \quad (\text{A15})$$

for all $x \in \mathbb{R}^d$. Conversely, given any linear map $\mathcal{M} : \mathbb{R}^d \rightarrow \mathfrak{g}$, the function $f : \mathbb{R}^d \rightarrow G$ specified by Eq. (A15) defines an OPS homomorphism $\mathbb{R}^d \xrightarrow{f} G$.

Corollary A1 says that any quantum unitary encoding $U : \mathbb{R}^d \rightarrow \mathcal{U}(\mathbb{C}^n)$ of the form

$$U(x) = e^{-i\mathcal{L}(x)}, \quad (\text{A16})$$

where $\mathcal{L} : \mathbb{R}^d \rightarrow \mathcal{B}(\mathbb{C}^n)$ is a linear map such that $\mathcal{L}(x)$ is self-adjoint for each $x \in \mathbb{R}^d$, defines a morphism $\mathbb{R}^d \xrightarrow{U} \mathcal{U}(\mathbb{C}^n)$ in **LieOPS**. It is important to note that the unitary quantum encoding U does *not* define a Lie group homomorphism in general. This is because $U(x+y) = e^{-i\mathcal{L}(x+y)}$ is *not* in general equal to $U(x)U(y) = e^{-i\mathcal{L}(x)}e^{-i\mathcal{L}(y)}$. Because the group structure as a whole is not preserved, this type of structure falls outside the context of geometric quantum machine learning. Instead, what is preserved is the group structure *when restricted* to any one-dimensional subspace in $\mathbb{R}^d$.

The benefit of Theorem A2 in the context of quantum machine learning is that it offers a *finite-dimensional* model of quantum encodings and shows what structures are preserved in the process. In fact, one can combine this structure together with symmetry to further reduce the space of quantum encodings. This is done in an illustrative example in Sec. IIB, with more details in Appendix B.

APPENDIX B: GEOMETRIC QUANTUM MACHINE LEARNING

In this Appendix, we go into more details for Example 1 from Sec. IIB. We begin by describing how Fig. 3 is generated. We then prove exactly how the space of *equivariant* unitary quantum encodings is reduced to an eight-dimensional space.

The example involves a learning task that distinguishes between two classes, labeled as -1 and $+1$, of points within a dataset $X \subseteq \mathbb{R}^2$. The labeling is determined by a binary *classifier*, which is a function $c : \mathcal{X} \rightarrow \{-1, 0, +1\}$, where the 0 element is included to allow for an undecided class. The set of all possible data $\mathcal{X}$ has a symmetry determined by the relations

$$c(x_1, x_2) = c(x_2, x_1) = c(-x_1, -x_2). \quad (\text{B1})$$

Such a binary classifier factors through the quantum encoding $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ if and only if there exists an observable $O \in \mathcal{B}(\mathcal{H})$ such that the spectrum of O is $\{-1, +1\}$ and

$$c(x) = \begin{cases} -1, & \text{if } y(x) < 0, \\ 0, & \text{if } y(x) = 0, \\ +1, & \text{if } y(x) > 0, \end{cases} \quad (\text{B2})$$

where $y : \mathcal{X} \rightarrow \mathbb{R}$ is the function defined by

$$y(x) = \text{Tr}[\rho(x)O], \quad (\text{B3})$$

which is the expectation value of O in the state $\rho(x)$. The functional y is called the *quantum classifier* and the set $y^{-1}(0) \subseteq \mathcal{X}$ is called the *decision boundary* of the classifier [128]. The quantum classifier is *invariant* if and only if

$$y(x_1, x_2) = y(x_2, x_1) = y(-x_1, -x_2) \quad (\text{B4})$$

for all $x_1, x_2 \in \mathbb{R}$. Note that invariance of the quantum classifier (B3) implies invariance of the classifier (B2), i.e., Eq. (B1) holds.

Now let $|\psi_0\rangle \in \mathcal{H} = \mathbb{C}^4$ be the initial state

$$|\psi_0\rangle = \sqrt{p}|+, +\rangle - \sqrt{1-p}|-, -\rangle \quad (\text{B5})$$

with $p = 0.99$, let O be the observable

$$O = X \otimes X, \quad (\text{B6})$$

and let $U : \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$ be the quantum unitary encoding as in Eq. (9). In this case, the quantum classifier can be computed explicitly as

$$y(x_1, x_2) = \cos(x_1) \sin(x_2) + 2\sqrt{p(1-p)} \sin(x_1) \sin(x_2). \quad (\text{B7})$$

The orange region in Fig. 3 consists of the points (x_1, x_2) such that $y(x_1, x_2) < 0$, while the blue region consists of the points (x_1, x_2) such that $y(x_1, x_2) > 0$. The boundary between the two regions consists of the points (x_1, x_2) such that $y(x_1, x_2) = 0$, which is the decision boundary.

Having shown how Fig. 3 is obtained, we next analyze how equivariance reduces the space of possible quantum encodings. In Example 1, we only looked at *one* such encoding U based on Eq. (9). Meanwhile, a large set of unitary quantum encodings $U : \mathbb{R}^2 \rightarrow \mathcal{U}(\mathbb{C}^4)$ are of the form $U(x) = e^{-i\mathcal{L}(x)}$ for some linear transformation $\mathcal{L} : \mathbb{R}^2 \rightarrow \mathfrak{su}(4)$ as in Eq. (14). Such a linear transformation $\mathcal{L}$ is uniquely determined by the value of $\mathcal{L}$ on the basis elements e_1, e_2 of $\mathbb{R}^2$. Namely, $L_1 := \mathcal{L}(e_1)$ and $L_2 := \mathcal{L}(e_2)$. The fact that L_1 is a traceless Hermitian matrix says that L_1 is of the form

$$L_1 = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ \overline{a_{12}} & a_{22} & a_{23} & a_{24} \\ \overline{a_{13}} & \overline{a_{23}} & a_{33} & a_{34} \\ \overline{a_{14}} & \overline{a_{24}} & \overline{a_{34}} & a_{44} \end{bmatrix}, \quad (\text{B8})$$

where $a_{11}, a_{22}, a_{33}, a_{44} \in \mathbb{R}$, $a_{12}, a_{13}, a_{14}, a_{23}, a_{24}, a_{34} \in \mathbb{C}$, and $a_{11} + a_{22} + a_{33} + a_{44} = 0$, and similarly for the matrix L_2 . This parametrization leads to a $2(4 + 2(6) - 1) = 30$ -dimensional vector space of unitary quantum encodings $U : \mathbb{R}^2 \rightarrow \mathcal{U}(\mathbb{C}^4)$ that are of the form $U(x) = e^{-i\mathcal{L}(x)}$ for some linear transformation $\mathcal{L} : \mathbb{R}^2 \rightarrow \mathfrak{su}(4)$.

The equivariance constraint $L_2 = \text{SWAP } L_1 \text{ SWAP}$ from Eq. (18) says that L_2 is uniquely determined by L_1 , thereby reducing the dimension of the space of unitary quantum encodings to 15. Combining this with the equivariance constraint

$\{X \otimes X, L_1\} = 0$ from Eq. (18) yields

$$\begin{bmatrix} \overline{a_{14}} & \overline{a_{24}} & \overline{a_{34}} & a_{44} \\ \overline{a_{13}} & \overline{a_{23}} & a_{33} & a_{34} \\ \overline{a_{12}} & a_{22} & a_{23} & a_{24} \\ a_{11} & a_{12} & a_{13} & a_{14} \end{bmatrix} = - \begin{bmatrix} a_{14} & a_{13} & a_{12} & a_{11} \\ a_{24} & a_{23} & a_{22} & \overline{a_{12}} \\ a_{34} & a_{33} & \overline{a_{23}} & \overline{a_{13}} \\ a_{44} & \overline{a_{34}} & \overline{a_{24}} & \overline{a_{14}} \end{bmatrix}. \quad (\text{B9})$$

Therefore, L_1 is of the form

$$L_1 = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ \overline{a_{12}} & a_{22} & a_{23} & -\overline{a_{13}} \\ \overline{a_{13}} & -a_{23} & -a_{22} & -\overline{a_{12}} \\ -a_{14} & -a_{13} & -a_{12} & -a_{11} \end{bmatrix}, \quad (\text{B10})$$

where

$$a_{12}, a_{13} \in \mathbb{C}, \quad a_{14}, a_{23} \in i\mathbb{R}, \quad a_{11}, a_{22} \in \mathbb{R}. \quad (\text{B11})$$

The set of such L_1 satisfying such conditions is an eight-dimensional real vector space. One can explicitly check that the matrices

$$\begin{array}{cccc} Z \otimes \mathbb{1}_2, & \mathbb{1}_2 \otimes Z, & X \otimes Z, & Z \otimes X, \\ Y \otimes \mathbb{1}_2, & \mathbb{1}_2 \otimes Y, & X \otimes Y, & Y \otimes X, \end{array} \quad (\text{B12})$$

are of this form. Since these matrices are linearly independent, this set of matrices forms a basis of this eight-dimensional subspace.

APPENDIX C: CLASSICAL LINEAR METRIC LEARNING AND THE MAHALANOBIS METRIC

Earlier, we saw that quantum metric learning can be described within a category of metric spaces and distance-preserving functions (embeddings). We can focus the discussion on metric spaces that have underlying sets $\mathbb{R}^n$ and the functions between them are linear. This leads to metrics known as *Mahalanobis distance metrics*.

In more detail, let $\mathcal{X} = \mathbb{R}^n$ and equip it with a Mahalanobis distance metric, where $(\mathcal{Y} = \mathbb{R}^m, d_{\mathcal{Y}})$ is a Euclidean space (so that $d_{\mathcal{Y}}$ is the Euclidean distance) and $\mathbb{R}^n \xrightarrow{V} \mathbb{R}^m$ is an injective linear map, described by some $m \times n$ matrix V with linearly independent columns [95,96], in which case

$$d_{\mathcal{X}}(x_1, x_2) = \sqrt{(x_1 - x_2)^T (V^T V) (x_1 - x_2)}, \quad (\text{C1})$$

where V^T denotes the transpose of V . Although this example is less relevant for *quantum* metric learning, it is worth spending a few moments to illustrate the categorical structure associated with this in the context of *classical* metric linear learning.

Let **Mahal** be the category whose objects are pairs $(\mathbb{R}^m, A)$, with $m \in \mathbb{N}$ and A a positive-definite $m \times m$ matrix. A morphism in **Mahal** from $(\mathbb{R}^m, A)$ to $(\mathbb{R}^n, B)$ is an $n \times m$ matrix V with trivial kernel such that $V^T B V = A$. The trivial kernel condition means that the nullspace of A is 0 [129]; i.e., the matrix V defines an injective linear transformation from $\mathbb{R}^m$ to $\mathbb{R}^n$. The category **Mahal** is *equivalent* [130] to the subcategory of **Met**^{emb} whose objects are Euclidean spaces equipped with the Mahalanobis distance metric and whose morphisms are required to be linear. To see this, first define

the Mahalanobis metric on $\mathbb{R}^m$ associated with A by

$$d_A(x_1, x_2) := \sqrt{(x_1 - x_2)^T A (x_1 - x_2)}. \quad (\text{C2})$$

Then, the assumption $V^T B V = A$ shows that

$$\begin{aligned} d_A(x_1, x_2) &= \sqrt{(x_1 - x_2)^T V^T B V (x_1 - x_2)} \\ &= \sqrt{(Vx_1 - Vx_2)^T B (Vx_1 - Vx_2)} \\ &= d_B(Vx_1, Vx_2), \end{aligned} \quad (\text{C3})$$

which proves that V is distance preserving. In this setting, if one is given data in $\mathbb{R}^m$ and the goal is to learn a Mahalanobis metric on $\mathbb{R}^m$ that captures the similarity in the data, then the task can be described as finding a linear map $\mathbb{R}^m \xrightarrow{V} \mathbb{R}^n$, since pulling back the Euclidean metric from $\mathbb{R}^n$ will give a Mahalanobis metric on $\mathbb{R}^m$.

APPENDIX D: METRIC LEARNING AND SEMIMETRIC SPACES

In this Appendix, we prove Lemma 1 and then provide a more detailed description of the quantum metric learning example shown in Fig. 8.

Proof. Without assuming anything about f ,

$$\begin{aligned} d(x_2, x_1) &= d_{\mathcal{Y}}(f(x_2), f(x_1)) \\ &= d_{\mathcal{Y}}(f(x_1), f(x_2)) \\ &= d_{\mathcal{X}}(x_1, x_2) \end{aligned} \quad (\text{D1})$$

shows that $d_{\mathcal{X}}$ is symmetric. Furthermore,

$$\begin{aligned} d_{\mathcal{X}}(x_1, x_2) &= d_{\mathcal{Y}}(f(x_1), f(x_2)) \\ &\geq d_{\mathcal{Y}}(f(x_1), y) + d_{\mathcal{Y}}(y, f(x_2)) \end{aligned} \quad (\text{D2})$$

holds for all $y \in \mathcal{Y}$. In particular, for any $x_3 \in \mathcal{X}$,

$$\begin{aligned} d_{\mathcal{X}}(x_1, x_2) &\geq d_{\mathcal{Y}}(f(x_1), f(x_3)) + d_{\mathcal{Y}}(f(x_3), f(x_2)) \\ &= d_{\mathcal{X}}(x_1, x_3) + d_{\mathcal{X}}(x_3, x_2), \end{aligned} \quad (\text{D3})$$

which shows that $d_{\mathcal{X}}$ satisfies the triangle inequality. It is also true that $d_{\mathcal{X}}(x, x) = 0$ for all $x \in \mathcal{X}$.

Now, if $x_1, x_2 \in \mathcal{X}$ satisfy

$$0 = d_{\mathcal{X}}(x_1, x_2) = d_{\mathcal{Y}}(f(x_1), f(x_2)), \quad (\text{D4})$$

then $f(x_1) = f(x_2)$. Hence, f is one to one if and only if $d_{\mathcal{X}}(x_1, x_2) = 0$ implies $x_1 = x_2$ for all x_1, x_2 that satisfy Eq. (D4).

The goal, first outlined in Ref. [23], is to classify a certain collection of data points on the interval $\mathcal{X} = [-2, 2]$ into two classes. Although the given data $X \subset \mathcal{X}$ shown in Fig. 8 are provided as a subset of the Euclidean space $\mathbb{R}$, the data cannot have the induced metric from its embedding in Euclidean space because they are not linearly separable. Therefore, in order to arrive at a metric that more accurately reflects the similarity between points from one class and dissimilarity between points from different classes, we may apply a variation of techniques used in nonlinear metric learning [91]. Reference [23] takes the perspective to map these data points into the space of quantum states equipped with the Hilbert-Schmidt metric with the hope of more accurately reflecting the appropriate classification.

Let $\{A, B\}$ be a partition of X , where the elements of A are labeled as class $+1$ and the elements of B are labeled as class -1 . The associated classifier $c : \mathcal{X} \rightarrow \{-1, 0, 1\}$ satisfies

$$c(x) = \begin{cases} +1, & \text{if } x \in A, \\ -1, & \text{if } x \in B. \end{cases} \quad (\text{D5})$$

Now, since the space of all states on $\mathcal{H}$ is a convex space, set

$$\rho_A := \frac{1}{\#A} \sum_{a \in A} \rho(a) \quad \text{and} \quad \rho_B := \frac{1}{\#B} \sum_{b \in B} \rho(b), \quad (\text{D6})$$

where $\#A$ and $\#B$ denote the cardinalities of A and B , respectively. These density matrices are taken to be the empirical centroid density matrices for our classification task. In this case, we wish to *maximize* the Hilbert-Schmidt distance between these two density matrices (cf. Definition 8). Equivalently, we wish to *minimize* the cost function

$$C(\rho_A, \rho_B) = 1 - \frac{1}{2} d_{\text{HS}}(\rho_A, \rho_B). \quad (\text{D7})$$

Note that finding an embedding that maximizes this distance is not a trivial task because the space of encodings $\mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ is infinite dimensional (in the sense that it defines an infinite-dimensional smooth space of paths [70,131–134]). To make this maximization procedure more approachable, we look at a *parametrized* family of encodings $\rho : \Theta \times \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$, which are referred to as *data reuploading models* [10,21,135]. (We will not discuss data reuploading models in more detail in this work.) More specifically, we take $\Theta = [0, 2\pi]^4$ and write $\theta = (\theta_1, \theta_2, \theta_3, \theta_4) \in \Theta$ and $\rho_\theta(x) := \rho(\theta, x)$. Moreover, we factor ρ through a parametrized unitary quantum encoding $U : \Theta \times \mathcal{X} \rightarrow \mathcal{U}(\mathcal{H})$, where

$$U_\theta(x) = \left(\prod_{n=1}^4 R_X(x) R_Y(\theta_n) \right) R_X(x), \quad (\text{D8})$$

which is depicted in Fig. 8 and where our ordering convention on products of operators is

$$\prod_{n=1}^N A_n = A_N \cdots A_1. \quad (\text{D9})$$

Combining the unitary encoding with the ground state gives the quantum state encoding

$$\rho_\theta(x) = |x\rangle\langle x|, \quad (\text{D10})$$

where

$$|x\rangle = U_\theta(x)|0\rangle. \quad (\text{D11})$$

Set

$$O := \rho_A - \rho_B, \quad (\text{D12})$$

which is an observable that depends on θ (since ρ_A and ρ_B depend on θ), and define the quantum classifier $y : \mathcal{X} \rightarrow \mathbb{R}$ to be

$$y(x) = \text{Tr}[\rho(x)O]. \quad (\text{D13})$$

Lemma D1. Let ρ_A and ρ_B be two distinct qubit density matrices. Then, there exists a $\lambda \in (0, 1]$ such that the eigenvalues of $\rho_A - \rho_B$ are $\{-\lambda, +\lambda\}$.

Proof. Since ρ_A and ρ_B are 2×2 self-adjoint matrices, let λ_1 and λ_2 denote the two eigenvalues of $\rho_A - \rho_B$. Since the eigenvalues of ρ_A and ρ_B are between 0 and 1, the eigenvalues of $\rho_A - \rho_B$ are between -1 and 1 . Since $\rho_A \neq \rho_B$, at least one of the two eigenvalues cannot be 0 because $\rho_A - \rho_B$ is self-adjoint. Since

$$\text{Tr}[\rho_A - \rho_B] = \text{Tr}[\rho_A] - \text{Tr}[\rho_B] = 1 - 1 = 0, \quad (\text{D14})$$

the sum of the two eigenvalues of $\rho_A - \rho_B$ must vanish, i.e., $\lambda_1 + \lambda_2 = 0$. In other words, $\lambda_2 = -\lambda_1$. Setting $\lambda = |\lambda_1|$ proves the claim.

Proposition D1. Let $X \subset \mathcal{X}$ be a finite training dataset together with a partition $A \cup B = X$, with elements of A and B labeled as class 1 and -1 , respectively. Let $\rho : \mathcal{X} \rightarrow \mathcal{S}(\mathcal{H})$ be any quantum encoding such that ρ_A and ρ_B , as given by Eq. (D6), are distinct density matrices. Set $O = \rho_A - \rho_B$ and let $\{-\lambda, +\lambda\}$ denote the set of eigenvalues of O , with $\lambda > 0$. Finally, let c be the classifier in Eq. (D5) and let $y : \mathcal{X} \rightarrow \mathbb{R}$ denote the associated quantum classifier as in Eq. (D13). Then,

$$\text{sgn}(y(x)) = c(x) \quad (\text{D15})$$

for all $x \in X$, where $\text{sgn} : (-\infty, 0) \cup (0, \infty) \rightarrow \{-1, +1\}$ denotes the sign function $\text{sgn}(x) = \frac{x}{|x|}$.

Proposition D1 shows that the quantum classifier agrees with the *fidelity classifier* used in Ref. [23] upon a rescaling. Meanwhile, the *Helstrom classifier* of Ref. [23] uses the observable $\Pi_+ - \Pi_-$, where Π_+ and Π_- are the projections onto the positive and negative eigenspaces of $\rho_A - \rho_B$, respectively.

Now, upon minimizing the cost function in Eq. (D7), we obtain specific values of $\theta_1, \theta_2, \theta_3$, and θ_4 for the parametrized encoding $\rho : \Theta \times \mathcal{X} \rightarrow \mathcal{S}(\mathbb{C}^2)$. The associated encoding is not actually an injective map into the space of quantum states, as shown in Fig. 8. This means that pulling back the metric gives a semimetric, rather than a metric, on $\mathcal{X}$ (cf. Definition 9 and Lemma 2). This semimetric can be visualized in the same way as a metric, namely, by illustrating the distance between points on a discretized grid from $\mathcal{X} \times \mathcal{X}$ [cf. Fig. 8(c)]. We note that the need for semimetrics is not a phenomenon unique to quantum systems, as they already appear in linear metric learning, such as in dimensional reduction [91].

APPENDIX E: ADDITIONAL EXAMPLES OF FUNCTORS AND NATURAL TRANSFORMATIONS

This section contains an additional excursion into category theory that is relevant to some of the ideas touched upon in the main text. The mathematical details below lie somewhat outside of the summarized perspective in Sec. III C, but we include the discussion here for the interested reader. And key to this discussion is the concept of a *natural transformation*, which is essentially a mapping between functors. Historically, natural transformations played a pivotal role in the birth of category theory, which originally arose in the context of homology theory in algebraic topology [7]. Although we did not make explicit use of natural transformations in the main body of this work, it is interesting to point out that bit encoding

can be described as a natural transformation that is closely related to a certain forgetful functor from **Vect** to **Set**, which will be described momentarily. Moreover, the notion of G -equivariance can also be described as a natural transformation. This Appendix is intended to go into details on these points in order to provide additional illustrations of categorical concepts through examples coming from quantum information theory and machine learning. In order to proceed, we will need to define two concepts: natural transformations, of course, and that of composing functors. We present the latter first.

Definition E1. Let $F : \mathbf{C} \rightarrow \mathbf{D}$ and $G : \mathbf{D} \rightarrow \mathbf{E}$ be two functors. The *composite* $G \circ F : \mathbf{C} \rightarrow \mathbf{E}$ is the functor that sends each object $\mathcal{X}$ in $\mathbf{C}$ to $G(F(\mathcal{X}))$ and sends each morphism $f : \mathcal{X} \rightarrow \mathcal{Y}$ in $\mathbf{C}$ to the morphism $G(F(f))$.

So, just as functions can be composed—and more generally just as morphisms in a category can be composed—functors can also be composed. This notion will appear in our discussion on bit encodings later. For now, we first give the definition of a natural transformation.

Definition E2. Let $\mathbf{C}$ and $\mathbf{D}$ be categories, and let $F, G : \mathbf{C} \rightarrow \mathbf{D}$ be two functors. A *natural transformation* η from F to G , written $\eta : F \Rightarrow G$, associates with each object $\mathcal{X}$ in $\mathbf{C}$ a morphism $\eta_{\mathcal{X}} : F(\mathcal{X}) \rightarrow G(\mathcal{X})$ in $\mathbf{D}$ such that the diagram

$$\begin{array}{ccc} F(\mathcal{X}) & \xrightarrow{F(f)} & F(\mathcal{Y}) \\ \eta_{\mathcal{X}} \downarrow & & \downarrow \eta_{\mathcal{Y}} \\ G(\mathcal{X}) & \xrightarrow{G(f)} & G(\mathcal{Y}) \end{array} \quad (\text{E1})$$

in $\mathbf{D}$ commutes, i.e., $G(f) \circ \eta_{\mathcal{X}} = \eta_{\mathcal{Y}} \circ F(f)$, for every morphism $\mathcal{X} \xrightarrow{f} \mathcal{Y}$ in $\mathbf{C}$. The commutativity of diagram (E1) is often referred to as *naturality*.

The level of abstraction in this definition warrants several examples, and we focus on some relevant examples from quantum machine learning and quantum algorithms. Our first example arises from bit encoding, which is typically the first and most familiar encoding learned when studying quantum algorithms [8,9,54]. This example will involve both the concepts of functor composition and a natural transformation.

Example E1 (Bit encoding). There is a category **Vect** whose objects are complex vector spaces and whose morphisms are $\mathbb{C}$ -linear transformations. Moreover, there is a familiar passage from **Set** to **Vect** that firstly associates with a set $\mathcal{X}$ the vector space $(\mathbb{C}_{\text{fs}}^{\mathcal{X}}, 0, +, \cdot)$ generated by it, which will be defined in the next paragraph. Secondly, this familiar passage also associates with a set-theoretic function $f : \mathcal{X} \rightarrow \mathcal{Y}$ a linear transformation $\mathbb{C}_{\text{fs}}^{\mathcal{X}} \rightarrow \mathbb{C}_{\text{fs}}^{\mathcal{Y}}$. This passage defines a functor $G : \mathbf{Set} \rightarrow \mathbf{Vect}$, and we will first describe its two ingredients one at a time. Afterward, we will construct the natural transformation of bit encoding, which involves both G and the forgetful functor $F : \mathbf{Vect} \rightarrow \mathbf{Set}$.

Mathematically, the vector space $(\mathbb{C}_{\text{fs}}^{\mathcal{X}}, 0, +, \cdot)$ is the vector space of *finitely supported* complex-valued functions on $\mathcal{X}$. Namely, each vector v in $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ is by definition a function $v : \mathcal{X} \rightarrow \mathbb{C}$ such that $v(x) = 0$ for all but finitely many $x \in \mathcal{X}$.

(If the set $\mathcal{X}$ is finite, then every function will automatically be finitely supported.) The zero vector 0 is the function that assigns 0 to every $x \in \mathcal{X}$. The sum of two functions $v, w \in \mathbb{C}_{\text{fs}}^{\mathcal{X}}$ is the function $v + w$ whose value on $x \in \mathcal{X}$ is $(v + w)(x) := v(x) + w(x)$. Meanwhile, if v is such a function and $\lambda \in \mathbb{C}$ is a scalar, then $\lambda \cdot v$, written as λv for short, is the function whose value on $x \in \mathcal{X}$ is $\lambda v(x)$.

The elements $x \in \mathcal{X}$ define a basis δ_x for $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$, where the function δ_x acts by $\delta_x(x') = \delta_{xx'}$ that is 1 if $x' = x$ and 0 if $x' \neq x$. If we were to use Dirac notation to express this vector, we could write $|x\rangle$ so that $\langle x'|x\rangle = \delta_{xx'}$, but that might be considered slightly abusive since we have not given $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ the structure of a Hilbert space (although one can be defined in terms of our basis when $\mathcal{X}$ is finite [136]). Moreover, we occasionally use the notation $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ to refer to the vector space $(\mathbb{C}_{\text{fs}}^{\mathcal{X}}, 0, +, \cdot)$, even though we should remember that the structure of a zero vector, vector addition, and scalar multiplication are implicitly assumed when referring to $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ as a vector space.

It is useful to view this from another perspective more familiar to the quantum information theorist. Let $\mathcal{X} = \mathbb{Z}_2^n = \{0, 1\}^n$ be the set of all arrays $x = (x_0, x_1, \dots, x_{n-1})$ whose entries are either 0 or 1 , so that we can think of x as a binary representation of a number $\{0, 1, 2, \dots, 2^n - 1\}$. Then, the vector space $\mathbb{C}^{\mathcal{X}}$ generated by $\mathcal{X}$ is precisely (i.e., naturally isomorphic to)

$$\overbrace{\mathbb{C}^2 \otimes \dots \otimes \mathbb{C}^2}^{n \text{ times}} \cong \mathbb{C}^{2^n}, \quad (\text{E2})$$

which is the complex vector space of n qubits. (There is no need to include the subscript fs in $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ in this case because $\mathcal{X}$ is a finite set with 2^n elements.) From this perspective, each element $x = (x_0, x_1, \dots, x_{n-1})$ is associated with the vector

$$|x_0 x_1 \dots x_{n-1}\rangle = |x_0\rangle \otimes |x_1\rangle \otimes \dots \otimes |x_{n-1}\rangle. \quad (\text{E3})$$

This is what *bit encoding* achieves (technically, bit encoding is more accurately described as a natural transformation, and we will make this more precise soon).

Having described how the passage from **Set** to **Vect** associates with each set a vector space, let us next describe how we also obtain linear transformations from set-theoretic functions. Let $f : \mathcal{X} \rightarrow \mathcal{Y}$ be a function between sets. From this function, define the linear transformation $L_f : \mathbb{C}_{\text{fs}}^{\mathcal{X}} \rightarrow \mathbb{C}_{\text{fs}}^{\mathcal{Y}}$ as follows. By writing the basis elements of $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$ and $\mathbb{C}_{\text{fs}}^{\mathcal{Y}}$ as $|x\rangle$ and $|y\rangle$, respectively, the linear transformation L_f sends each basis vector $|x\rangle$ to $|f(x)\rangle$, which is another basis vector in $\mathbb{C}_{\text{fs}}^{\mathcal{Y}}$. Since a linear transformation is uniquely determined by its action on a basis [129], this defines the linear transformation L_f on all of $\mathbb{C}_{\text{fs}}^{\mathcal{X}}$. This technique of transforming $f : \mathcal{X} \rightarrow \mathcal{Y}$ to $L_f : \mathbb{C}_{\text{fs}}^{\mathcal{X}} \rightarrow \mathbb{C}_{\text{fs}}^{\mathcal{Y}}$ is often used in constructing quantum algorithms such as the Deutsch-Jozsa algorithm, which tests whether a function f is constant or balanced [8,9,54].

So far, we have described a functor $G : \mathbf{Set} \rightarrow \mathbf{Vect}$ that takes a set $\mathcal{X}$ and constructs the vector space $(\mathbb{C}_{\text{fs}}^{\mathcal{X}}, 0, +, \cdot)$ of complex-valued functions on $\mathbb{C}$. Now, there is *also* a functor going the other way that takes a vector space $(V, 0, +, \cdot)$ and forgets the vector space structure, leaving only the underlying set V . In more detail, there is a functor $\mathbf{Vect} \rightarrow \mathbf{Set}$ that assigns to each vector space V [technically, $(V, 0, +, \cdot)$, where 0 is the zero vector, $+$ is addition, and $\cdot$ is scalar multiplication]

its underlying set V of vectors, thus forgetting the linear structure (namely, the 0 vector, vector addition $+$, and scalar multiplication $\cdot$). Moreover, this functor assigns every linear transformation to itself, since every linear transformation is a function.

From these two functors, we can compose them (cf. Definition E1) in either order. One of these two ways of composing the functors will be used to exhibit bit encoding. Namely, taking a set $\mathcal{X}$, constructing the associated vector space $F(\mathcal{X}) = (\mathbb{C}_{\text{fs}}^{\mathcal{X}}, 0, +, \cdot)$, and then forgetting the vector space structure gives a set $G(F(\mathcal{X})) = \mathbb{C}_{\text{fs}}^{\mathcal{X}}$.

We now have *two* functors from **Set** to **Set**. On the one hand, we have $G \circ F$, which we just described. On the other hand, we also have the functor $\text{id}_{\text{Set}} : \text{Set} \rightarrow \text{Set}$ that acts as the identity on all objects and morphisms. We can now realize bit encoding as the natural transformation $\eta : \text{id}_{\text{Set}} \Rightarrow G \circ F$ defined by sending a set $\mathcal{X}$ to a particular function $\eta_{\mathcal{X}} : \text{id}_{\text{Set}}(\mathcal{X}) \rightarrow G(F(\mathcal{X}))$, which is a function of the form $\eta_{\mathcal{X}} : \mathcal{X} \rightarrow \mathbb{C}_{\text{fs}}^{\mathcal{X}}$. The definition of this function is precisely the standard bit encoding, which sends an element $x \in \mathcal{X}$ to the vector $\eta_{\mathcal{X}}(x) = |x\rangle \in \mathbb{C}_{\text{fs}}^{\mathcal{X}}$. So what does the naturality property (E1) tell us about bit encoding? If $f : \mathcal{X} \rightarrow \mathcal{Y}$ is any (classical) function between sets, let $L_f : \mathbb{C}_{\text{fs}}^{\mathcal{X}} \rightarrow \mathbb{C}_{\text{fs}}^{\mathcal{Y}}$ denote associated linear transformation determined uniquely by how it acts on the basis $\{|x\rangle\}$. Naturality of η then says that $\eta_{\mathcal{Y}}(f(x)) = L_f(\eta_{\mathcal{X}}(x))$, which reads $|f(x)\rangle = L_f|x\rangle$, a completely *natural* condition used in most quantum algorithms, and one which we in fact already used when defining the linear transformation L_f by how it acts on a basis.

As our next example, one might have suspected that the very definition of an equivariant map as in Eq. (4) looks a lot like naturality. This is indeed the case after one views G -sets as functors and is analogous to how representations can be viewed as functors [137].

Example E2 (G-equivariant maps revisited). Let G be a group and let $\mathbb{B}G$ be the category consisting of only a single object, denoted by $\bullet$. The set of morphisms from $\bullet$ to itself in the category $\mathbb{B}G$ is defined to be the set G . The composition in $\mathbb{B}G$ is then taken to be the multiplication operation in G . One can check that $\mathbb{B}G$ is a category with these definitions. In this way, a group can be viewed as a category. A G -set can then be viewed as a functor $\alpha : \mathbb{B}G \rightarrow \text{Set}$. Indeed, the unique object $\bullet$ in $\mathbb{B}G$ gets sent to some set, call it $\mathcal{X}$. Moreover, since α is a functor, it sends a morphism in $\mathbb{B}G$, which is an element g of the group G , to some function $\alpha_g : \mathcal{X} \rightarrow \mathcal{X}$. Now, since the element g is invertible, there exists a g^{-1} such that $gg^{-1} = 1_G = g^{-1}g$, where 1_G is the unit element of G . Hence, by functoriality of α , we have that $\text{id}_{\mathcal{X}} = \alpha_e = \alpha_{gg^{-1}} = \alpha_g \circ \alpha_{g^{-1}}$, and similarly $\text{id}_{\mathcal{X}} = \alpha_{g^{-1}} \circ \alpha_g$, which proves that α_g is a bijection with inverse $\alpha_{g^{-1}}$.

Now, consider *two* such G -sets, viewed as functors $\alpha, \beta : \mathbb{B}G \rightarrow \text{Set}$, where $\mathcal{X} := \alpha(\bullet)$ and $\mathcal{Y} := \beta(\bullet)$. A natural transformation $f : \alpha \Rightarrow \beta$ assigns to the only object $\bullet$ of $\mathbb{B}G$ a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ (abusively denoted by the same letter). Naturality of f as a natural transformation then gives exactly the G -equivariant condition (4).

-
- [1] S. Aaronson, Read the fine print, *Nat. Phys.* **11**, 291 (2015).
 - [2] S. Sim, P. D. Johnson, A. Aspuru-Guzik, Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, *Adv. Quantum Technol.* **2**, 1900070 (2019).
 - [3] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, Power of data in quantum machine learning, *Nat. Commun.* **12**, 1 (2021).
 - [4] S. Thanasilp, S. Wang, M. Cerezo, and Z. Holmes, Exponential concentration in quantum Kernel methods, *Nat. Commun.* **15**, 5200 (2024).
 - [5] J. Bowles, S. Ahmed, and M. Schuld, Better than classical? The subtle art of benchmarking quantum machine learning models, *arXiv:2403.07059*.
 - [6] M. A. Khan, M. N. Aman, B. Sikdar, Beyond bits: A review of quantum embedding techniques for efficient information processing, *IEEE Access* **12**, 46118 (2024).
 - [7] S. M. Lane, *Categories for the Working Mathematician* (Springer Science & Business Media, New York, NY, 2013), Vol. 5.
 - [8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary Edition (Cambridge University Press, Cambridge, UK, 2011).
 - [9] G. Benenti, G. Casati, D. Rossini, and G. Strini, *Principles of Quantum Computation and Information: A Comprehensive Textbook* (World Scientific, New Jersey, US, 2019).
 - [10] S. Jerbi, L. J. Fiderer, H. Poulsen Nautrup, J. M. Kübler, H. J. Briegel, and V. Dunjko, Quantum machine learning beyond kernel methods, *Nat. Commun.* **14**, 517 (2023).
 - [11] M. Schuld, Supervised quantum machine learning models are kernel methods, *arXiv:2101.11020*.
 - [12] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
 - [13] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
 - [14] M. Larocca, S. Thanasilp, S. Wang, K. Sharma, J. Biamonte, P. J. Coles, L. Cincio, J. R. McClean, Z. Holmes, and M. Cerezo, Barren plateaus in variational quantum computing, *Nat. Rev. Phys.* **7**, 174 (2025).
 - [15] M. Ragone, B. N. Bakalov, F. Sauvage, A. F. Kemper, C. Ortiz Marrero, M. Larocca, and M. Cerezo, A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits, *Nat. Commun.* **15**, 7172 (2024).
 - [16] E. Fontana, D. Herman, S. Chakrabarti, N. Kumar, R. Yalovetzky, J. Heredge, S. H. Sureshbabu, and M. Pistoia, Characterizing barren plateaus in quantum ansätze with the adjoint representation, *Nat. Commun.* **15**, 7171 (2024).
 - [17] M. Cerezo, M. Larocca, D. García-Martín, N. L. Diaz, P. Braccia, E. Fontana, M. S. Rudolph, P. Bermejo, A. Ijaz, S. Thanasilp, E. R. Anschuetz, and Z. Holmes, Does provable

- absence of barren plateaus imply classical simulability? *Nat. Commun.* **16**, 7907 (2025).
- [18] I. F. Araujo, D. K. Park, F. Petruccione, and A. J. da Silva, A divide-and-conquer algorithm for quantum state preparation, *Sci. Rep.* **11**, 6329 (2021).
- [19] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, Circuit-centric quantum classifiers, *Phys. Rev. A* **101**, 032308 (2020).
- [20] E. Tang, Quantum principal component analysis only achieves an exponential speedup because of its state preparation assumptions, *Phys. Rev. Lett.* **127**, 060503 (2021).
- [21] M. Schuld, R. Sweke, and J. J. Meyer, Effect of data encoding on the expressive power of variational quantum-machine-learning models, *Phys. Rev. A* **103**, 032430 (2021).
- [22] R. LaRose and B. Coyle, Robust data encodings for quantum classifiers, *Phys. Rev. A* **102**, 032420 (2020).
- [23] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, Quantum embeddings for machine learning, [arXiv:2001.03622](https://arxiv.org/abs/2001.03622).
- [24] M. Schuld and N. Killoran, Quantum machine learning in feature Hilbert spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [25] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [26] T. Goto, Q. H. Tran, K. Nakajima, Universal approximation property of quantum machine learning models in quantum-enhanced feature spaces, *Phys. Rev. Lett.* **127**, 090506 (2021).
- [27] J. Kim and S. Bekiranov, Generalization performance of quantum metric learning classifiers, *Biomolecules* **12**, 1576 (2022).
- [28] S. Thanasilp, S. Wang, N. A. Nghiem, P. Coles, and M. Cerezo, Subtleties in the trainability of quantum machine learning models, *Quantum Mach. Intell.* **5**, 21 (2023).
- [29] M. Cerezo, G. Verdon, H.-Y. Huang, L. Cincio, and P. J. Coles, Challenges and opportunities in quantum machine learning, *Nat. Comput. Sci.* **2**, 567 (2022).
- [30] J. J. Meyer, M. Mularski, E. Gil-Fuster, A. A. Mele, F. Arzani, A. Wilms, and J. Eisert, Exploiting symmetry in variational quantum machine learning, *PRX Quantum* **4**, 010328 (2023).
- [31] M. Ragone, P. Braccia, Q. T. Nguyen, L. Schatzki, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Representation theory for geometric quantum machine learning, [arXiv:2210.07980](https://arxiv.org/abs/2210.07980).
- [32] Q. T. Nguyen, L. Schatzki, P. Braccia, M. Ragone, P. J. Coles, F. Sauvage, M. Larocca, and M. Cerezo, Theory for equivariant quantum neural networks, *PRX Quantum* **5**, 020328 (2024).
- [33] K. Fukushima, Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position, *Biol. Cybern.* **36**, 193 (1980).
- [34] M. M. Bronstein, J. Bruna, T. Cohen, and P. Veličković, Geometric deep learning: Grids, groups, graphs, geodesics, and gauges, [arXiv:2104.13478](https://arxiv.org/abs/2104.13478).
- [35] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, Geometric deep learning: Going beyond euclidean data, *IEEE Signal Process. Mag.* **34**, 18 (2017).
- [36] M. Larocca, F. Sauvage, F. M. Sbahi, G. Verdon, P. J. Coles, and M. Cerezo, Group-invariant quantum machine learning, *PRX Quantum* **3**, 030341 (2022).
- [37] S. Lloyd, S. Garnerone, P. Zanardi, Quantum algorithms for topological and geometric analysis of data, *Nat. Commun.* **7**, 10138 (2016).
- [38] A. Vlasic and A. Pham, Understanding the mapping of encode data through an implementation of quantum topological analysis, *Quantum Inf. Comput.* **23**, 1091 (2023).
- [39] E. Riehl, *Category Theory in Context* (Dover Publications, Garden City, NY, 2016).
- [40] P. Perrone, Notes on category theory with examples from basic mathematics, [arXiv:1912.10642](https://arxiv.org/abs/1912.10642).
- [41] T. Leinster, *Basic Category Theory*, Vol. 143 (Cambridge University Press, Cambridge, UK, 2014).
- [42] C. Heunen and J. Vicary, *Categories for Quantum Theory: An Introduction* (Oxford University Press, Oxford, UK, 2019).
- [43] B. Coecke and A. Kissinger, *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning* (Cambridge University Press, Cambridge, UK, 2017).
- [44] L. Corry, *Modern Algebra and the Rise of Mathematical Structures*, 2nd ed. (Birkhäuser, Basel, 2004).
- [45] G. Carlsson, Topology and data, *Bull. Am. Math. Soc.* **46**, 255 (2009).
- [46] L. McInnes, J. Healy, J. Melville, UMAP: Uniform manifold approximation and projection for dimension reduction, [arXiv:1802.03426](https://arxiv.org/abs/1802.03426).
- [47] T. Hagino, A categorical programming language, Ph.D. thesis, University of Edinburgh, 1987.
- [48] C. Heunen, O. Kammar, S. Staton, and H. Yang, A convenient category for higher-order probability theory, in *32nd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)* (IEEE, New York, NY, 2017), pp. 1–12.
- [49] C. F. Strnadl, The mathematical syntax of architectures, [arXiv:2004.03719](https://arxiv.org/abs/2004.03719).
- [50] R. Cornish, Neural network symmetrisation in concrete settings, [arXiv:2412.09469](https://arxiv.org/abs/2412.09469).
- [51] R. Cornish, Stochastic neural network symmetrisation in markov categories, [arXiv:2406.11814](https://arxiv.org/abs/2406.11814).
- [52] B. Fong, D. I. Spivak, R. Tuyéras, Backprop as functor: A compositional perspective on supervised learning, in *Proceedings of the 34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '19* (IEEE Press, New York, NY, 2021).
- [53] B. Gavranović, P. Lessard, A. Dudzik, T. von Glehn, J. G. Araújo, and P. Veličković, Position: Categorical deep learning is an algebraic theory of all architectures, in *Proceedings of the 41st International Conference on Machine Learning* (PMLR, 2024), Vol. 235, pp. 15209–15241.
- [54] D. Deutsch and R. Jozsa, Rapid solution of problems by quantum computation, *Proc. R. Soc. London A* **439**, 553 (1992).
- [55] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Phys. Rev. Lett.* **103**, 150502 (2009).
- [56] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum algorithms for supervised and unsupervised machine learning, [arXiv:1307.0411](https://arxiv.org/abs/1307.0411).
- [57] L. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, [arXiv:quant-ph/0208112](https://arxiv.org/abs/quant-ph/0208112).
- [58] P. Kaye and M. Mosca, Quantum networks for generating arbitrary quantum states, in *Optical Fiber Communication Conference and International Conference on Quantum Information* (Optica Publishing Group, Washington, DC, 2001), p. PB28.

- [59] M. Schuld and F. Petruccione, *Machine Learning with Quantum Computers*, 2nd ed. (Springer, Cham, 2021), Vol. 676.
- [60] J. von Neumann, *Mathematical Foundations of Quantum Mechanics: New Edition* (Princeton University Press, New York, NY, 2018).
- [61] B. C. Hall, *Quantum Theory for Mathematicians*, Graduate Texts in Mathematics (Springer, New York, NY, 2013), Vol. 267.
- [62] J. R. Glick, T. P. Gujarati, A. D. Córcoles, Y. Kim, A. Kandala, J. M. Gambetta, and K. Temme, Covariant quantum kernels for data with group structure, *Nat. Phys.* **20**, 479 (2024).
- [63] Z. P. Bradshaw, E. N. Evans, M. Cook, and M. L. LaBorde, Learning equivariant maps with variational quantum circuits, *Phys. Rev. Appl.* **23**, 044007 (2025).
- [64] B. C. Hall, *Lie Groups, Lie Algebras, and Representations*, 2nd ed. (Springer, Cham, 2015).
- [65] E. Tang, Dequantizing algorithms to understand quantum advantage in machine learning, *Nat. Rev. Phys.* **4**, 692 (2022).
- [66] E. R. Anschuetz, A. Bauer, B. T. Kiani, and S. Lloyd, Efficient classical algorithms for simulating symmetric quantum systems, *Quantum* **7**, 1189 (2023).
- [67] C. Tüysüz, S. Y. Chang, M. Demidik, K. Jansen, S. Vallecorsa, and M. Grossi, Symmetry breaking in geometric quantum machine learning in the presence of noise, *PRX Quantum* **5**, 030314 (2024).
- [68] I. N. M. Le, O. Kiss, J. Schuhmacher, I. Tavernelli, and F. Tacchino, Symmetry-invariant quantum machine learning force fields, *New J. Phys.* **27**, 023015 (2025).
- [69] M. F. Langer, S. N. Pozdnyakov, and M. Ceriotti, Probing the effects of broken symmetries in machine learning, *Mach. Learn.: Sci. Technol.* **5**, 04LT01 (2024).
- [70] J. R. Munkres, *Topology: A First Course*, 2nd ed. (Prentice-Hall, Inc., Saddle River, NJ, 2000).
- [71] J. Baez and J. P. Muniain, *Gauge Fields, Knots and Gravity*, Series on Knots and Everything (World Scientific, Singapore, 1994), Vol. 4.
- [72] J. M. Lee, *Introduction to Smooth Manifolds*, Graduate Texts in Mathematics, 2nd ed. (Springer, New York, 2013), Vol. 218.
- [73] T.-D. Bradley, T. Bryson, J. Terilla, *Topology: A Categorical Approach* (MIT Press, Cambridge, MA, 2020).
- [74] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed. (McGraw-Hill, New York, 1976).
- [75] Defining smoothness on the full state space $\mathcal{S}(\mathcal{H})$ is a bit subtle because $\mathcal{S}(\mathcal{H})$ is not technically a smooth manifold, though it is a convex space [109]. Despite the lack of smoothness in the usual sense, for most of the examples we consider in this paper, the quantum state encodings will factor through the subspace of pure states, which is a smooth manifold (it is a complex projective space). For this reason, we will not concern ourselves too much with going into the more technical details regarding the smooth structure on $\mathcal{S}(\mathcal{H})$.
- [76] A. Skolik, J. R. McClean, M. Mohseni, P. van der Smagt, and M. Leib, Layerwise learning for quantum neural networks, *Quantum Mach. Intell.* **3**, 5 (2021).
- [77] G. De Luca, A. Vlasic, M. Vitz, and A. Pham, Empirical power of quantum encoding methods for binary classification, *Quantum Mach. Intell.* **7**, 72 (2025).
- [78] M. Arnott, D. Papaioannou, K. McDowall, P. Lolur, and B. Baldé, Reverse map projections as equivariant quantum embeddings, [arXiv:2407.19906](https://arxiv.org/abs/2407.19906).
- [79] G. Carlsson and M. Vejdemo-Johansson, *Topological Data Analysis with Applications* (Cambridge University Press, Cambridge, UK, 2021).
- [80] R. Ghrist, Barcodes: The persistent topology of data, *Bull. Am. Math. Soc.* **45**, 61 (2008).
- [81] C. Glass and E. Vidaurre, Topological data analysis via undergraduate linear algebra, [arXiv:2406.17045](https://arxiv.org/abs/2406.17045).
- [82] J. R. Munkres, *Elements of Algebraic Topology* (CRC Press, 1984).
- [83] D. Burago, Y. Burago, and S. Ivanov, *A Course in Metric Geometry*, Graduate Studies in Mathematics (American Mathematical Society, Providence, RI, 2001), Vol. 33.
- [84] V. Bryant, *Metric Spaces: Iteration and Application* (Cambridge University Press, Cambridge, UK, 1985).
- [85] J.-C. Hausmann, On the Vietoris-Rips complexes and a cohomology theory for metric spaces, in *Prospects in Topology (AM-138): Proceedings of a Conference in Honor of William Browder* (Princeton University Press, Princeton, 1996), Vol. 138, pp. 175–188.
- [86] A. Zomorodian and G. Carlsson, Computing persistent homology, *Discrete Comput. Geom.* **33**, 249 (2005).
- [87] G. Carlsson, Topological pattern recognition for point cloud data, *Acta Numer.* **23**, 289 (2014).
- [88] D. Cohen-Steiner, H. Edelsbrunner, and J. Harer, Stability of persistence diagrams, *Discrete Comput. Geom.* **37**, 103 (2007).
- [89] F. Mémoli and K. Singhal, A primer on persistent homology of finite metric spaces, *Bull. Math. Biol.* **81**, 2074 (2019).
- [90] H. Edelsbrunner, D. Letscher, and A. Zomorodian, Topological persistence and simplification, *Discrete Comput. Geom.* **28**, 511 (2002).
- [91] A. Bellet, A. Habrard, and M. Sebban, *Metric Learning*, Synthesis Lectures on Artificial Intelligence and Machine Learning (Springer, Cham, 2015).
- [92] E. Xing, M. Jordan, S. J. Russell, and A. Ng, Distance metric learning with application to clustering with side-information, in *Advances in Neural Information Processing Systems*, edited by S. Becker, S. Thrun, and K. Obermayer (MIT Press, Cambridge, MA, 2002), Vol. 15, pp. 521–528.
- [93] K. Q. Weinberger and L. K. Saul, Distance metric learning for large margin nearest neighbor classification, *J. Mach. Learn. Res.* **10**, 207 (2009).
- [94] S. Chopra, R. Hadsell, Y. LeCun, Learning a similarity metric discriminatively, with application to face verification, in *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, Vol. 1 (IEEE, New York, NY, 2005), pp. 539–546.
- [95] B. Kulis, Metric learning: A survey, *FNT Mach. Learn.* **5**, 287 (2012).
- [96] M. Kaya and H. Ş. Bilge, Deep metric learning: A survey, *Symmetry* **11**, 1066 (2019).
- [97] D. Kedem, S. Tyree, F. Sha, G. Lanckriet, and K. Q. Weinberger, Non-linear metric learning, in *Advances in Neural Information Processing Systems*, edited by F. Pereira, C. J. Burges, L. Bottou, and K. Q. Weinberger (Curran Associates, Inc., Red Hook, NY, 2012), Vol. 25.
- [98] M. Hillery, Nonclassical distance in quantum optics, *Phys. Rev. A* **35**, 725 (1987).

- [99] D. Dieks and P. Veltkamp, Distance between quantum states, statistical inference and the projection postulate, *Phys. Lett. A* **97**, 24 (1983).
- [100] S. L. Braunstein and C. M. Caves, Statistical distance and the geometry of quantum states, *Phys. Rev. Lett.* **72**, 3439 (1994).
- [101] D. Bures, An extension of Kakutani's theorem on infinite product measures to the tensor product of semifinite w^* -algebras, *Trans. Am. Math. Soc.* **135**, 199 (1969).
- [102] A. Uhlmann, The "transition probability" in the state space of a $*$ -algebra, *Rep. Math. Phys.* **9**, 273 (1976).
- [103] M. Hübner, Explicit computation of the Bures distance for density matrices, *Phys. Lett. A* **163**, 239 (1992).
- [104] A. R. Kuzmak, Measuring distance between quantum states on a quantum computer, *Quantum Inf. Process.* **20**, 269 (2021).
- [105] J. Watrous, *The Theory of Quantum Information* (Cambridge University Press, Cambridge, UK, 2018).
- [106] K. Życzkowski and W. Słomczyński, The Monge distance between quantum states, *J. Phys. A: Math. Gen.* **31**, 9095 (1998).
- [107] V. V. Dodonov, O. V. Man'ko, V. I. Man'ko, and A. Wünsche, Energy-sensitive and "classical-like" distances between quantum states, *Phys. Scr.* **59**, 81 (1999).
- [108] D. Spivak, Metric realization of fuzzy simplicial sets (2009), http://www.dspivak.net/metric_realization090922.pdf.
- [109] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement*, 2nd ed. (Cambridge University Press, Cambridge, UK, 2017).
- [110] nLab authors, Structure (2024), <https://ncatlab.org/nlab/revision/structure>.
- [111] As an example, the real line $\mathbb{R}$ can be viewed as a category whose objects are elements $x, y, \dots \in \mathbb{R}$ and where there is a morphism $x \rightarrow y$ whenever $x \leq y$. This is a category due to the transitivity and reflexivity of the partial order, but it is not concrete since morphisms are not structure-preserving functions.
- [112] M. Schuld and N. Killoran, Is quantum advantage the right goal for quantum machine learning? *PRX Quantum* **3**, 030101 (2022).
- [113] R. Wiersema and N. Killoran, Optimizing quantum circuits with Riemannian gradient flow, *Phys. Rev. A* **107**, 062421 (2023).
- [114] P. Kuklinski, B. Rempfer, J. Elenewski, and K. Obenland, Efficient block-encodings require structure, [arXiv:2509.19667](https://arxiv.org/abs/2509.19667).
- [115] K. V. Mardia, Some properties of classical multi-dimensional scaling, *Commun. Stat. Theory Methods* **7**, 1233 (1978).
- [116] J. A. Perea, Multiscale projective coordinates via persistent cohomology of sparse filtrations, *Discrete Comput. Geom.* **59**, 175 (2018).
- [117] S. Aaronson and A. Ambainis, The need for structure in quantum speedups, *Theory Comput.* **10**, 133 (2014).
- [118] E. Tang, A quantum-inspired classical algorithm for recommendation systems, in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC '19* (ACM, New York, NY, 2019).
- [119] J. Cotler, H.-Y. Huang, and J. R. McClean, Revisiting dequantization and quantum advantage in learning tasks, [arXiv:2112.00811](https://arxiv.org/abs/2112.00811).
- [120] S. Shin, Y. S. Teo, and H. Jeong, Dequantizing quantum machine learning models using tensor networks, *Phys. Rev. Res.* **6**, 023218 (2024).
- [121] L. J. Henderson, K. Beer, S. Karuvade, R. Gupta, A. White, and S. Shrapnel, Quantum advantage without exponential concentration: Trainable kernels for symmetry-structured data, [arXiv:2509.14337](https://arxiv.org/abs/2509.14337).
- [122] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [123] R. Gilmore, *Lie Groups, Lie Algebras, and Some of Their Applications* (Dover Publications, Garden City, NJ, 2006).
- [124] A. M. Bincer, *Lie Groups and Lie Algebras. A Physicist's Perspective* (Oxford University Press, Oxford, 2013).
- [125] R. W. Carter, I. G. MacDonald, and G. B. Segal, *Lectures on Lie Groups and Lie Algebras*, London Mathematical Society Student Texts (Cambridge University Press, Cambridge, UK, 1995).
- [126] There is a slight abuse of notation here that is commonly done in category theory. The vector $\alpha \in \mathbb{R}^2$ is identified with the additive map $\alpha : \mathbb{R} \rightarrow \mathbb{R}^2$ sending $1 \in \mathbb{R}$ to α .
- [127] J. W. Milnor, *Topology from the Differentiable Viewpoint*, Princeton Landmarks in Mathematics and Physics (Princeton University Press, Princeton, NJ, 1997).
- [128] The convex-linear functional sending states $\sigma \in \mathcal{S}(\mathcal{H})$ to $\text{Tr}[\sigma O]$ defines a hyperplane $\{\sigma \in \mathcal{S}(\mathcal{H}) : \text{Tr}[\sigma O] = 0\}$ in the space of states, which is one of the reasons it is often used for classification tasks in QML [11].
- [129] G. Strang, *Introduction to Linear Algebra*, 6th ed. (Wellesley-Cambridge Press, Wellesley, MA, 2022).
- [130] An *equivalence* of categories is a technical term whose formal definition we omit. Intuitively, however, two categories are said to be *equivalent* if they are effectively the same.
- [131] J. F. Adams, *Infinite Loop Spaces*, Annals of Mathematics Studies (Princeton University Press, Princeton, NJ, 1978), Vol. 90.
- [132] A. Pressley and G. B. Segal, *Loop Groups*, Oxford Mathematical Monographs (Clarendon Press, Oxford, UK, 1986).
- [133] J.-L. Brylinski, *Loop Spaces, Characteristic Classes and Geometric Quantization*, Modern Birkhäuser Classics (Birkhäuser Boston, Inc., Boston, MA, 2008).
- [134] A. J. Parzygnat, Gauge invariant surface holonomy and monopoles, *Theory Appl. Categ.* **30**, 1319 (2015).
- [135] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, Data re-uploading for a universal quantum classifier, *Quantum* **4**, 226 (2020).
- [136] P. R. Halmos, *Finite-Dimensional Vector Spaces*, Undergraduate Texts in Mathematics (Springer, New York, 1958).
- [137] A. J. Parzygnat, Two-dimensional algebra in lattice gauge theory, *J. Math. Phys.* **60**, 043506 (2019).