scientific reports

Check for updates

OPEN

A quantum-inspired classification for random mixed states

Giuseppe Sergioli[1]✉, Carlo Cuccu[1], Carla Sophie Rieger[2,3], Andrés Camilo Granda Arango[1], Bikash Kumar Behera[1], Riccardo Era[1] & Roberto Giuntini[1,4]

We present a quantum-inspired classification framework designed to identify correlation structures: product, separable, and entangled, in random mixed quantum states. Building on previous work where the *Pretty-Good-Measurement* (PGM) classifier demonstrated a competitive performance on pure-state ensembles, we extend this method to the more challenging domain of mixed states. We apply our quantum-inspired classifier to randomly generated ensembles of two- and three-qubit mixed states, encompassing all possible varieties of subsystem correlations while ensuring statistical neutrality. The results indicate that learning architectures inspired by quantum state discrimination can offer scalable and physically grounded tools for the characterization of entanglement and separability even in the mixed-state regime.

Keywords Quantum-inspired machine learning, PGM classifier, Quantum states classification, Mixed states

Quantum machine learning (QML) is rapidly establishing itself as one of the most dynamic and promising research frontiers in contemporary science. By integrating the principles of quantum mechanics into the data-driven frameworks of machine learning, QML has opened new pathways for solving complex problems in high-dimensional spaces, with potential applications ranging from chemistry and material science[1] to cryptography and optimization[2].

Alongside fully quantum implementations, the development of *quantum-inspired* approaches, where classical algorithms are guided by concepts from quantum information theory, has also proven to be highly fruitful[3]. These methods, in principle, do not require quantum hardware but instead exploit structural insights from quantum theory, such as entanglement, measurement, and non-commutativity, to design novel architectures or feature embeddings that enhance classical learning tasks. Quantum-inspired models have shown benefits in tasks like kernel construction, generative modeling, and anomaly detection, often offering improved expressiveness or interpretability.

Among the most effective and physically grounded applications of both QML and quantum-inspired learning lies the analysis of microscopic systems. In this context, one typically aims to identify underlying physical properties, most notably, the presence and nature of correlations between subsystems. Correlation structures such as separability, entanglement, and non-locality are essential for quantum technologies, yet notoriously hard to detect and quantify, especially in mixed-state scenarios where experimental imperfections blur the classical–quantum boundary. This underlines the need to analyze quantum datasets proposed in recent literature[4] and e.g., large-scale quantum datasets for machine learning benchmarks in the case of one and two qubits[5].

It is within this scope that supervised state classification, whether for pure or, more realistically, for mixed states, has become a crucial benchmark for assessing an algorithm's ability to scalably distinguish product, separable, and entangled states and to uncover forms of multipartite non-locality.

In recent years, the classification of quantum correlations has been pursued from markedly different methodological perspectives. On the one hand, several contributions aim to recognize (or exploit) such correlations through quantum machine learning schemes that explicitly involve quantum algorithms and, at least in principle, execution on quantum hardware or within hybrid quantum classical pipelines. This line includes, for instance, proposals based on variants of quantum support vector machines implemented on noisy intermediate-scale

[1]Università degli Studi di Cagliari, Via Is Mirrionis 1, Cagliari 09123, Italy. [2]European Organization for Nuclear Research (CERN), Espl. des Particules 1, 1211 Meyrin, Switzerland. [3]Department of Engineering and Design, Technische Universität München, Lise-Meitner-Str. 9, Ottobrunn 85521, Germany. [4]Institute for Advanced Study (IAS), Technische Universität München, Lichtenbergstraße 2a 85748 München, Germany. ✉email: giuseppe.sergioli@gmail.com

quantum platforms (as in[6]), as well as the more general framework introduced by[7] via quantum feature maps and kernel/variational methods for supervised learning.

Alongside the "quantum computing-based" approaches, a second line of research pursues the classification of quantum correlations without requiring a quantum computer, relying on fully classical resources (in the sense of classical ML/deep learning), sometimes inspired by quantum-information principles or representations, and sometimes in hybrid forms where physical ingredients (observables, tomograms, witnesses, etc.) guide feature construction and/or labeling while the classifier itself remains classical. A recent example is the work of[8], which proposes a multipartite taxonomy grounded in how entanglement is distributed across subsystems and how "fragile" it is under partial tracing, formalized through a hierarchy of tangle-like quantities τ_i and classes $[N]_j$ (for 2, 3, and 4 qubits). A supervised classifier is then trained to map density-matrix entries to these labels, which are assigned by computing entanglement-based indicators on the generated states. While performance is very strong on pure states, it drops noticeably in the mixed-state regime and as the multipartite structure increases, revealing a significant dependence on both the class definition and the dataset distribution. A more explicitly "black-box" stance is adopted by[9], who use deep multilayer perceptrons to separate separable from entangled states (two qubits, with near-perfect reported accuracy) and, for three qubits, to assign finer families (including W- and GHZ-type classes) with typically lower performance; importantly, they also highlight a strong dependence on purity and entanglement strength, hence a sensitivity to the particular region of state space being sampled and to the labeling choices (positive partial transposition (PPT)/negativity, prescribed families, etc.).

Within the same classical-only setting but with a different trade-off between physical structure and statistical learning[10] propose classifying random states into classes such as separable/PPT entangled states/negative partial transposition (NPT) from tomograms (features given by expectation values of suitable operators), while delegating the final decision to an autoML pipeline. This avoids direct access to the density-matrix entries, yet remains strongly data-driven and depends both on dataset generation and on the criterion used to label "hard" instances (especially in the PPT region).

A line even closer to entanglement theory consists of learning a separating criterion that admits a witness interpretation. In this direction[11], represent four-qubit states as real vectors (e.g., in a Pauli basis) and train a Support Vector Machine to distinguish separable states from entangled families; the resulting hyperplane is then read as a measurable witness operator. Another often-cited reference is[12], who recast the problem as variational reconstruction of pure states via neural-network quantum states (RBMs): by imposing separability constraints (Separable Neural Network States), entanglement classification emerges from the (in)ability of constrained learners to match the fidelity achieved by an unconstrained learner. In a related spirit (though with a different target task)[13], combines RBM-like representations with learning/optimization techniques (e.g., reinforcement-learning strategies) to explore and enhance quantum-correlation signals associated with nonlocality, illustrating how classical ML architectures can be "steered" by quantum structures even when the goal is not, strictly speaking, entanglement classification.

Finally, an adjacent but distinct line aims at estimating entanglement quantifiers from limited local measurements, thereby reducing tomographic overhead[14]. train neural models that, from local Pauli expectation values (one- and two-qubit) for equilibrium states or from single-qubit time traces in nonequilibrium dynamics, infer quantities such as Rényi entropies and moments of the partial transpose, in a framework geared toward NISQ-relevant, physically structured states (ground and dynamical states of local Hamiltonians).

Overall, these contributions show that quantum-correlation classification can be successfully pursued both through QML on quantum hardware and through purely classical strategies, but they also indicate that performance and generalization often depend sensitively on the chosen class definitions, the mixing regime, and the dataset distribution. It is within this landscape that our proposal is positioned: we also work in a supervised setting on datasets of density matrices, yet the "learning" stage does not coincide with training a parametric model via statistical optimization (backpropagation, loss minimization, weight/architecture tuning, and related hyperparameterization). Instead, it consists of constructing a quantum-inspired decision rule directly from quantum-state discrimination theory, in particular, from the Pretty Good Measurement (PGM) using class operators estimated from data. In this respect, we do not claim in the present work a general superiority of our scheme over the approaches reviewed above; rather, our aim is to highlight a methodologically distinct and, in our view, novel perspective in the context of random mixed-state classification.

More specifically, we introduced in[15] a quantum-inspired classifier grounded in quantum-state discrimination theory, the PGM classifier, and showed that, on ensembles of pure states up to five qubits, it achieves robust performance comparable to the best classical models in accuracy. That study demonstrated the PGM's strong ability to separate product, separable, and entangled pure states, laying the foundations for a learning pipeline explicitly informed by quantum-information principles. The present work extends that framework to the substantially more intricate domain of mixed states, where separability is no longer a binary property and the convex geometry of density operators demands finer criteria. We show how noise reshapes the statistical distribution of classes and complicates discrimination both computationally and operationally. Specifically, we (i) analyze how different probability measures generate random mixed-state ensembles with distinct structural features; (ii) describe the dataset partitioning for two- and three-qubit systems in terms of product, separable, and entangled criteria; and (iii) apply the PGM classifier to the resulting datasets, showing that the robustness previously demonstrated for pure-state quantum-correlation classification generalizes effectively to the mixed-state scenario as well.

Accordingly, we propose our classifier as a quantum-inspired approach to mixed-state classification that is directly grounded in quantum-state discrimination. The main novelty we aim to introduce is that the decision rule is obtained constructively from discrimination theory via the PGM prescription and the class operators estimated from data rather than by positing a parametric model and learning its parameters through statistical optimization (e.g., loss minimization, gradient-based training, and extensive hyperparameter tuning). In this sense, the method does not require an iterative optimization phase to fit weights or architectures: once the class operators

are specified, the classifier is determined (as also discussed in Section Discussion). We do not claim in this work a superiority over the classical approaches reviewed above; our goal is to highlight this methodological alternative and to characterize its behavior on random mixed-state datasets. Given the heterogeneity of the above approaches in terms of class definitions, data representations, and learning paradigms, a systematic quantitative comparison with leading methods both in predictive accuracy and in overall procedural complexity lies beyond the scope of this article and will be addressed in future work.

The PGM classifier

Quantum information theory, and in particular the problem of quantum state discrimination, has inspired the design of quantum-inspired classifiers that show promising performance compared to classical methods[16–19]. Among these, the PGM classifier defines a quantum-inspired multi-class classification model based on quantum theoretical principles, achieving competitive results in both accuracy and computational cost. This section is devoted to providing a brief analytical description of the PGM classifier.

Given a training set $\mathcal{S}_{tr} := \{(\mathrm{x}_1, y_1), \ldots, (\mathrm{x}_m, y_m)\}$ with $\mathrm{x}_j \in \mathbb{C}^d$ and $y_j \in \mathcal{Y} := \{0, \ldots, \ell\}$, we define the subsets $\mathcal{S}^i_{tr} = \{\mathrm{x}_j : y_j = i\}$ for each class *i*. A classifier $f : \mathbb{C}^d \to [0,1]^\ell$ assigns to each input x a vector of scores $f(\mathrm{x})$, interpreted as probabilities when normalized. The predicted label is obtained via:

$$Cl_f(\mathrm{x}) := \min \left\{ i \in \mathcal{Y} : f(\mathrm{x})_i = \max_k \{f(\mathrm{x})_k, 1 \le k \le \ell\} \right\} . \tag{1}$$

The PGM classifier encodes each classical vector x into a density matrix ρ_x, forming the quantum training set $\mathcal{S}_{Qtr} = \{(\rho_{\mathrm{x}_1}, y_1), \ldots, (\rho_{\mathrm{x}_m}, y_m)\}$. For each class *i*, the *quantum centroid* is the average state:

$$\rho_{(i)} = \frac{1}{|\mathcal{S}^i_{Qtr}|} \sum_{\mathrm{x}_j \in \mathcal{S}^i_{tr}} \rho_{\mathrm{x}_j} . \tag{2}$$

To enhance expressivity, one can employ *c*-fold tensor copies of each state, yielding generalized centroids:

$$\rho^{(c)}_{(i)} = \frac{1}{|\mathcal{S}^i_{Qtr}|} \sum_{\mathrm{x}_j \in \mathcal{S}^i_{tr}} \rho^{\otimes c}_{\mathrm{x}_j} . \tag{3}$$

The score function *f* is defined via quantum state discrimination (QSD) on the ensemble $R = \{(p_i, \rho^{(c)}_{(i)})\}$, where p_i is the class prior. Letting $\sigma = \sum_i p_i \rho^{(c)}_{(i)}$, we define:

$$E_i = (\sigma^\dagger)^{1/2} p_i \rho^{(c)}_{(i)} (\sigma^\dagger)^{1/2}, \quad F_i = E_i + \frac{1}{\ell} P_{\ker(\sigma)} ,$$

where $\sigma^\dagger$ denotes the Moore Penrose pseudoinverse of σ and $P_{\ker(\sigma)}$ is the projection onto the subspace spanned by the kernel of σ ($\ker(\sigma)$).

Since $\sigma = \sum_i p_i \rho^{(c)}_{(i)}$ is positive semidefinite, it may be rank-deficient if the image of σ ($\mathrm{im}(\sigma)$) spans a proper subspace of the Hilbert space at issue. In this case, $(\sigma^\dagger)^{1/2}$ denotes the Moore-Penrose inverse square root, i.e., the inverse square root restricted to $\mathrm{im}(\sigma)$ and acting as zero on $\ker(\sigma)$. Equivalently, if $\sigma = \sum_{k=1}^r \lambda_k |k\rangle\langle k|$ with $\lambda_k > 0$, then

$$(\sigma^\dagger)^{1/2} = \sum_{k=1}^r \lambda_k^{-1/2} |k\rangle\langle k|,$$

and $(\sigma^\dagger)^{1/2}|v\rangle = 0$ for all $|v\rangle \in \ker(\sigma)$. Therefore, the operators

$$E_i = (\sigma^\dagger)^{1/2} p_i \rho^{(c)}_{(i)} (\sigma^\dagger)^{1/2}$$

are well-defined, positive semidefinite, and satisfy $\sum_i E_i = P_{\mathrm{im}(\sigma)}$, hence they form a POV-measure on the image of σ. To obtain a *b*ona fide POV-measurement on the full space, we use the standard completion

$$F_i = E_i + \frac{1}{\ell} P_{\ker(\sigma)},$$

so that $\sum_i F_i = P_{\mathrm{im}(\sigma)} + P_{\ker(\sigma)} = I$. The added term distributes the "unseen" subspace $\ker(\sigma)$ uniformly across classes; consequently, for any test state ρ_x the kernel contribution $\frac{1}{\ell}\mathrm{Tr}(P_{\ker(\sigma)}\rho_x)$ is the same for all *i* and does not bias the $\arg\max$ decision rule.

Then the PGM classifier assigns:

$$Cl_f(\mathrm{x}) := \min_i \left\{ i \in \mathcal{Y} : p_i \mathrm{Tr}\left(F_i \rho^{(c)}_\mathrm{x}\right) = \max_k \{p_k \mathrm{Tr}\left(F_k \rho^{(c)}_\mathrm{x}\right), 1 \le k \le \ell\} \right\} . \tag{4}$$

Higher values of c can improve accuracy by better separating the centroids in Hilbert space. In[18], the authors introduced a method that yields a protocol equivalent to PGM classifier, derived via application of Kernel Trick techniques. For datasets with moderate number of samples, this approach can be particularly advantageous, as it enables the use of a high number c of copies while keeping the algorithm's complexity within tractable bounds.

To further improve the classifier's accuracy, in[20] we showed that two additional aspects can be addressed at the preprocessing stage. The first concerns the choice of a suitable *encoding*, between a real-vector representation and a density-operator representation. The second concerns the introduction of a multiplicative *rescaling factor*, applied upstream to the features of all vectors in the dataset, which yields an additional gain in accuracy. This improvement is consistent with the fact that the classification procedure is carried out within the geometric framework of the quantum state space. Hence, during the implementation of the PGM classifier it is possible to adopt a grid-search optimization procedure over the main hyperparameters mentioned above (encoding, number of copies, and rescaling factor).

Generating random mixed states

It is worth noting that the generation of random *pure* quantum states is a well-understood and technically straightforward procedure. In finite-dimensional Hilbert spaces, a pure state can be represented by a normalized vector $|\psi\rangle$ with complex entries. Sampling such a state uniformly, meaning with respect to the natural geometry of the space, amounts to drawing $|\psi\rangle$ from the unique unitarily invariant measure on the unit sphere in $\mathcal{C}^N$, known as the *Haar measure*.

In practice, this can be efficiently achieved by generating a complex vector whose entries are independent Gaussian random variables (with real and imaginary parts each drawn from a normal distribution), and then normalizing the result. This procedure guarantees uniformity over the Hilbert space without requiring any specific structure, symmetry, or decomposition of the system.

Due to its simplicity and mathematical robustness, Haar-random sampling of pure states has become a standard tool in quantum information theory, where it is routinely used in benchmarking protocols, statistical analyses, random circuit constructions[21,22], and studies of typical entanglement properties[23–25]. Its central role also extends to quantum computing and quantum machine learning, where random pure states often serve as initialization points or reference ensembles.

In contrast, the generation of an *unbiased* sample of mixed quantum states is a non-trivial task. In the quantum formalism, mixed states are represented by density operator matrices that are Hermitian, positive semidefinite, and have unit trace. The collection of all such operators on an N-dimensional Hilbert space forms a convex set, typically denoted by $\mathcal{M}_N$. This set contains all $N \times N$ Hermitian matrices that are positive semidefinite (i.e., with real and non-negative eigenvalues) and have unit trace. However, $\mathcal{M}_N$ is not a flat or homogeneous space: it is a highly curved region embedded in the real N^2-dimensional space of Hermitian matrices. This curvature reflects the geometric constraints imposed by the physical requirements of positivity and normalization. Furthermore, as the dimension N increases, the volume of $\mathcal{M}_N$ becomes exponentially concentrated near its boundary, where the density operators are nearly pure, having one dominant eigenvalue close to 1 and the rest close to zero.

This phenomenon has significant practical implications. If one uses naive methods to generate mixed states, for instance, by independently sampling the matrix elements from a Gaussian distribution, or by drawing eigenvalues uniformly from the simplex, one typically ends up producing mostly highly pure states. Such approaches fail to account for the complex geometry of $\mathcal{M}_N$ and tend to oversample peripheral regions while leaving large portions of physically relevant space (especially genuinely mixed states) virtually unexplored.

As a result, any statistical analysis or numerical experiment based on such a sample would be biased: it would reflect the peculiarities of the sampling method rather than provide an accurate and representative picture of the typical properties of mixed quantum states.

An effective strategy to overcome the aforementioned sampling bias is to employ a method grounded in symmetry principles. Instead of attempting to generate mixed states directly, one begins by sampling a pure state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ from the *Haar measure* defined on the unitary group $U(NK)$. The Haar measure is a mathematically well-defined probability measure that is both left- and right-invariant under unitary transformations. In other words, it assigns equal weights to all directions in Hilbert space, making it the natural choice for unbiased sampling of pure states[1].

Once a pure state $|\psi\rangle$ has been drawn from the Haar distribution, a corresponding mixed state ρ is obtained by performing a partial trace over one of the two subsystems, typically interpreted as tracing out an unobserved environment of dimension K: $\rho = \mathrm{Tr}_B\,(|\psi\rangle\langle\psi|)$. This procedure produces an ensemble of mixed states whose statistical properties inherit the uniformity and symmetry of the underlying Haar distribution. Crucially, it avoids the geometric biases associated with naive sampling techniques and yields a well-behaved distribution over the space of density operators that respects the natural structure of the Hilbert space.

Because the Haar measure distributes probability uniformly over the entire space of pure states in the composite system $\mathcal{H}_A \otimes \mathcal{H}_B$, it inherently *ignores the internal tensor structure* that is, it does not distinguish between the subsystems A and B. As a result, the statistical properties of the reduced mixed state ρ, obtained by tracing out one subsystem, depend only on the total dimension NK of the composite Hilbert space, and not on how that dimension is partitioned between the subsystems. A particularly interesting case arises when the bipartition is *balanced*, i.e., when both subsystems have the same dimension ($K = N$). In this setting, the distribution of the resulting mixed states coincides with the well-known *Hilbert–Schmidt ensemble*, a standard random matrix model for quantum density operators. In the limit $N \to \infty$, the eigenvalue distribution of such states converges to the *Marchenko–Pastur law*, a fundamental result from random matrix theory describing the spectral density of large random

[1] Formally, if μ_{H} denotes the Haar measure and $U, V \in U(NK)$, then $\mu_{\mathrm{H}}(U\psi) = \mu_{\mathrm{H}}(\psi) = \mu_{\mathrm{H}}(V\psi)$.

positive matrices[26]. We refer to this family of states as the *structureless ensemble* because it preserves no information about the subsystem decomposition. In this ensemble, no subsystem is distinguished or favored, and any quantum correlations that arise, such as entanglement, are merely incidental rather than purposefully embedded or controlled.

To preserve the information encoded in the subsystem structure of a composite quantum system, one relaxes the requirement of invariance under global unitaries $U(NK)$ which treat the entire Hilbert space $\mathcal{H}_A \otimes \mathcal{H}_B$ as a single undifferentiated entity and instead adopts invariance under the action of *local* unitaries $U(N) \times U(K)$, which act independently on each subsystem. This shift in symmetry has important consequences. Measures that are invariant under $U(N) \times U(K)$ define what are known as *structured ensembles*. These ensembles are physically more realistic, as they reflect the operational constraints inherent to composite systems, where only local transformations are typically allowed.

Unlike structureless ensembles, which effectively erase the distinction between subsystems and thus obscure features like separability and entanglement, structured ensembles preserve these distinctions. By assigning statistical weight uniformly with respect to local operations, they maintain the relevant features of quantum correlations. In this way, structured ensembles allow one to explore the space of mixed quantum states in a way that retains information about non-classical correlations. This makes them especially valuable for tasks where the nature and strength of quantum entanglement must be controlled or assessed. They provide a balanced and tunable framework for sampling states with varying degrees of entanglement, ensuring that correlation-dependent properties are not washed out by overly strong symmetry assumptions. A major advantage of structured ensembles is their flexibility: by choosing suitable methods of construction, it becomes possible to generate quantum states that span a wide range of correlation profiles from nearly uncorrelated (product-like) states to strongly entangled ones while avoiding the geometric distortions discussed earlier.

Several families of structured ensembles have been developed precisely for this purpose. For example, one can consider *superpositions of maximally entangled vectors*, which allow for fine-grained control over the degree of entanglement by adjusting the relative weights in the superposition. Alternatively, ensembles constructed via *Bures rotations*, which are based on geodesic flows in state space with respect to the Bures metric, yield a natural interpolation between classical and quantum correlations while preserving operational relevance. Another powerful method involves the use of *products of Ginibre matrices*, whose statistical properties enable the construction of mixed states with tunable entanglement spectra. Through these and similar constructions, one can define ensembles that vary continuously along a meaningful axis of entanglement, all while maintaining statistical neutrality and avoiding the sampling biases introduced by naive or unstructured approaches[27].

Let us emphasize that the aim of this article is to evaluate how effectively the PGM classifier recognizes and exploits quantum correlations. For a meaningful benchmarking study, the training and test data must *exhibit* those correlations in a controlled manner. Accordingly, all numerical experiments reported below are based on structured ensembles. This choice preserves the statistical neutrality afforded by Haar symmetry now applied locally while ensuring that the datasets contain a balanced and representative spectrum of separable, weakly entangled, and strongly entangled states.

In the following section, we will consider random density operators of dimension $N = 2^n$, where n denotes the number of qubits. Based on the considerations outlined above, in principle, two complementary strategies to generate *unbiased* sets of random density matrices could be adopted, in line with the framework of *structured ensembles* introduced earlier.

The first strategy proceeds as follows: to obtain a random density matrix of dimension 2^n, we begin with a pure state defined on a Hilbert space of dimension 2^{2n+1}, and perform a partial trace over an environment of dimension 2^{n+1}. This yields a reduced density matrix of the desired size 2^n.

The second strategy fixes the total pure state dimension at 2^7, and generates random density matrices by tracing out environments of dimension 2^4, 2^5, or 2^6. This produces, respectively, reduced states of dimension 2^3, 2^2, and 2^1.

In both approaches, the environment dimension is deliberately chosen to be larger than that of the retained subsystem. This design choice is consistent with the structured ensemble approach and ensures that the resulting mixed states are sufficiently entropic. In particular, it guarantees that the reduced density matrices reflect typical features of mixed states, such as low purity and a broad eigenvalue spectrum, while avoiding statistical distortions that would otherwise arise from an environment that is too small. This allows for representative and physically meaningful sampling across a wide range of correlation structures.

Nevertheless, in the following experiments, we restrict ourselves to using only the first sampling strategy, as the second one introduces a significant complication. In fact, starting from pure states in dimension 2^7 and tracing out large subsystems of dimension 2^5 yields reduced states of dimension 2^2 that are extremely mixed, often approaching the maximally mixed state. This high entropy substantially reduces the likelihood of obtaining entangled states[2] as most of the resulting density matrices lie deep within the separable region of state space. Therefore, for the purpose of generating balanced datasets with a non-negligible proportion of entangled states, the most reliable approach remains tracing out from pure states of dimension 2^5, as in the first strategy described above.

To make the above point quantitative, we performed an auxiliary numerical study aimed at characterizing how the purity distribution of reduced states changes as the size of the traced-out environment increases, and how this affects the likelihood of obtaining entangled samples. Concretely, we fixed the reduced system to three qubits (dimension 2^3) and generated reduced density operators by tracing out subsystems from larger n-qubit states (dimension 2^n), with $n \in \{4, \ldots, 10\}$.

For each value of n, we repeated the procedure 20000 times and computed the purity $\mathrm{tr}(\rho^2)$ of the resulting three-qubit reduced state. The results are summarized in Fig. 1, which reports in black the mean purity (with error

[2]Where, in this class, we include states that fail the PPT test for all possible bipartitions, as largely discussed at the end of Section Analytical description.

bars) as a function of n. As n increases, the purity decreases monotonically and concentrates around values close to the maximally mixed state $\frac{1}{8}I$, showing that larger environments induce strongly mixed reduced states. This provides a direct, quantitative comparison between the regimes underlying our two sampling strategies: the second strategy corresponds precisely to tracing out significantly larger subsystems, and the resulting purity distribution shifts towards highly mixed states. The first strategy, instead, considers a global system defined on a Hilbert space of dimension markedly larger than that of the subsystem of interest (from 2^{2n+1} to 2^n), but without resorting to an *excessively large* enlargement. As a result, the reduced states obtained via partial trace are certainly mixed, while preventing the procedure from degenerating into the generation of overly mixed states, i.e., states that lie too close to the maximally mixed state.

In addition, we quantified the associated loss of quantum correlations by computing, for the same reduced states, the negativity (as an NPT *-negative partial transpose* - entanglement witness) across a three-qubit bipartition. Fig. 1 (blue line) shows that the average negativity decreases rapidly with n and becomes essentially zero for sufficiently large n. This confirms that the "large-environment" regime produces reduced states that are not only highly mixed, but also overwhelmingly PPT, hence yielding an increasingly negligible proportion of NPT-entangled samples. Obviously, the analysis carried out here for obtaining density operators of dimension 2^3 applies, in general, to the generation of density operators of an arbitrary dimension 2^m (that is, for m both smaller and larger than 3).

Summarizing, this quantitative analysis allows us to understand that the second strategy tends to generate 2^2 -dimensional density operators and, especially, 2^1-dimensional density operators that are excessively mixed (due to tracing out a very large environment, starting from a 2^7-dimensional space). This shifts the distributions toward predominantly PPT states (with negativity close to zero), thereby making the dataset imbalanced against states that fail the PPT test (i.e., NPT-entangled states), which become rare (obviously, for $n \geq 2$). By contrast, the first strategy allows one to achieve a high but not extreme level of mixedness, preserving a non-negligible fraction of NPT states and thus yielding an overall more balanced sample generation.

In the next Section, we detail the sampling routines used for two- and three-qubit systems and describe the criteria by which each state is labeled according to its correlation content.

Generating a dataset of mixed states

Analytical description

In this section, we describe the preparation of datasets of mixed density operators, which are subsequently used in order to be evaluated by a classification algorithm. Our aim is to investigate whether the PGM classifier can identify quantum correlations based on the matrix structure of the respective density operators. Therefore, it is crucial to generate well-balanced classes of density operators, each corresponding to a distinct type of correlation, while ensuring a minimal bias in their construction. The first subsection focuses on the two-qubit case, while the second addresses the three-qubit scenario.

Two-Qubit case

In the case of two qubits, the set of all possible density operators can be partitioned into the following three classes:

1. **Product states (factorized states)**: This class consists of density matrices of the form $\rho = \rho_1 \otimes \rho_2$, where each ρ_i is a 2×2 random density matrix generated according to the procedure described in the previous section. These states exhibit no quantum correlations and are fully factorized.

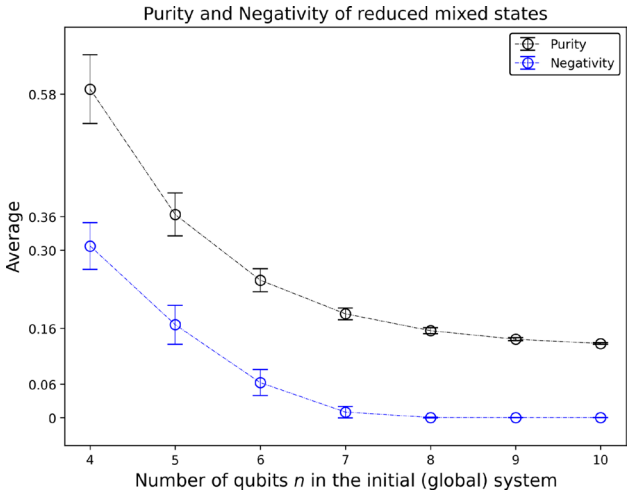

Fig. 1. Average purity (black) and average negativity (blue) of three-qubit reduced states obtained by tracing out an $(n-3)$-qubit environment from random pure n-qubit states, for $n \in \{4, \ldots, 10\}$. For each n, 2×10^4 samples were generated; markers report sample means and error bars indicate the statistical uncertainty. As the environment size increases, the reduced states become increasingly mixed (purity approaching the maximally mixed value 1/8) and the negativity rapidly vanishes, indicating that NPT entanglement becomes negligible in the large-environment regime.

2. **Separable (but non-product) states**: These are states that are not strictly of product form, yet are still separable. That is, they can be expressed as convex combinations of product states. Formally, they are of the form

$$\rho = \sum_i \lambda_i \, \rho_{i_1} \otimes \rho_{i_2},$$

where $\lambda_i \geq 0$, $\sum_i \lambda_i = 1$, and each ρ_{i_j} is a 2×2 random density matrix. The indices *i* range over a set of 2 to 4^2 elements, ensuring sufficient variety while maintaining separability by the Carathéodory Theorem[28].

3. **Entangled states**: This class includes generic 4×4 random density matrices. Since the random generation of 4×4 density matrices typically results in entangled states with high probability, we introduce an additional filtering step to eliminate any bias. Specifically, we retain only those states that fail the PPT (Positive Partial Transpose) test, thereby guaranteeing that all states in this class are genuinely entangled.

Three-Qubit case

1. **Fully Product States**: States of the form $\rho_1 \otimes \rho_2 \otimes \rho_3$, where each ρ_i is a random 2×2 mixed density matrix.
2. **Separable (but not fully product) States**

   (a) **Tripartite Separable States**: This subclass consists of convex combinations of fully product states, that is, states of the form $\sum_i \lambda_i(\rho_1 \otimes \rho_2 \otimes \rho_3)$, where each ρ_i is a 2×2 random mixed density matrix, and λ_i are real coefficients such that $i \in [2, 8^2]$.

   (b) **Bipartite Product States (Three-Qubit Case)**: This class includes states of the form $\rho \otimes \sigma$, where ρ is a random 2×2 density matrix and σ is a random 4×4 density matrix that fails the PPT test and is therefore entangled. These represent three-qubit states in which the first qubit is uncorrelated with the remaining two, which are entangled. We denote this configuration as 1; 2–3, indicating that qubit 1 is separable from qubits 2 and 3, which are mutually entangled. Our aim is to investigate whether the PGM classifier can capture physical features of quantum systems from their matrix structure. To that end, the dataset must include all possible configurations, equally represented. Therefore, this subclass must also include states where the first and third qubits are entangled, denoted 1–3; 2, obtained by applying the operator $Swap \otimes I$ to the previous configuration, where I is the 2×2 identity matrix. Another relevant configuration is the one in which the first two qubits are entangled and separable from the third, denoted 1–2; 3. This is obtained from the previous configuration by applying the operator $I \otimes Swap$. Consequently, any state in this class will be generated in the initial 1; 2–3 configuration and then randomly transformed by applying one of the following operations:

- $I^{\otimes 3}$: the 8×8 identity matrix (preserves the initial configuration);
- $Swap \otimes I$: transforms the state into the 1–3; 2 configuration;
- $(I \otimes Swap) \cdot (Swap \otimes I)$: yields the 1–2; 3 configuration.

   (c) **Bipartite Separable Mixed States**: This subclass contains convex combinations of states from the previous class, i.e., $\sum_i \lambda_i(\rho \otimes \sigma)$, where ρ is a random 2×2 density matrix and σ is a random 4×4 density matrix that fails the PPT test. As before, $i \in [2, 2^6]$. These states also undergo, prior to dataset inclusion, one of the three transformations I, $Swap \otimes I$, or $(I \otimes Swap) \cdot (Swap \otimes I)$ randomly applied.
3. **Entangled States**: Random 8×8 density matrices that fail the PPT test.

Let us clarify that in the present work we adopt an operational notion of entanglement based on the Peres Horodecki criterion: for three-qubit states we label as entangled those that are NPT (negative partial transposition) with respect to each of the three one-vs-two possible bipartitions, i.e., those for which the partial transpose has at least one negative eigenvalue for some of the possible bipartitions. This choice yields a set of entangled samples that can be certified unambiguously and with high numerical reliability, thereby reducing the risk of mislabeling and ensuring a clear separation from the classes defined by construction. We note that this class forms a subset of the set of entangled states in general: in particular, it does not include bound entangled states, nor states that are NPT only on a strict subset of the bipartitions.

This restriction is motivated by concrete requirements of supervised benchmarking. First, supervised classification requires labels that can be verified through procedures that are easily reproducible and endowed with explicit acceptance/rejection criteria: the NPT test provides a certifying and readily automatable condition, which can be applied uniformly to large datasets without introducing *ad hoc* choices or heuristic dependencies. Second, extending the entangled class to include "marginal" regions (e.g. bound entangled states, or states for which NPT entanglement is extremely weak, with the minimum eigenvalue of the partial transpose lying very close to zero) would, in order to avoid ambiguities, require more demanding verification procedures (semidefinite programming hierarchies, optimized witnesses, etc.) and would inevitably introduce a grey zone in which labeling becomes more sensitive to numerical tolerances and implementation details. By contrast, selecting states that are NPT across all bipartitions allows one to construct a benchmark in which (i) the entangled class is defined by a physically motivated and computationally simple criterion, (ii) overlap with separable/biseparable classes is minimized by construction, and (iii) the classifier's performance metrics can be interpreted more stably, since they refer to categories whose ground truth can be controlled with high reliability.

A brief remark may be useful here. It is natural to observe that, as the number n of qubits increases, the classification of quantum-correlation properties tends to involve a proliferation of classes, due both to the wider variety of separability structures that become available and to the presence of multiple non-equivalent partitions of the system (e.g., different bipartitions and, more generally, partitions into multiple blocks). However, by construction, the PGM is a multi-class discrimination procedure and extends without conceptual modification to the case in which the dataset is split into an arbitrary number of classes. Therefore, although the experimental analysis in the present work is limited to the two- and three-qubit cases, the proposed methodology is directly scalable, at a conceptual level, to classification tasks involving $n \geq 3$ qubit systems, where the number of classes naturally grows as a function of the correlation patterns and of the partitions under consideration.

Setup of the experiment

We provide a detailed description of the experimental setup used in our study. For each of the scenarios outlined in the previous section specifically, systems comprising 2 and 3 qubits we generated random quantum mixed states to serve as input data. Classification in the 2-qubit case involves three classes. In the 3-qubit case, there are a total of five classes, three of which correspond to different kinds of separable states. Therefore, in the three-qubit case, we first present the data by aggregating all separable states into a single class, and then we provide a finer classification by distinguishing the various types of separable states into separate classes. Classification experiments were conducted across multiple dataset sizes: 10,000, 20,000, 30,000, and 40,000 samples. Balanced class distributions were ensured by the quantum state generation method described in Section Analytical description, which provides direct control over the number of samples per class.

For each dataset size, the data was randomly split into a training set (80%) and a test set (20%). Model selection was carried out on each training set using a grid search, whose research spaces are specified in Table S1 (in Supplementary Material), combined with 5-fold cross-validation, with balanced accuracy as the evaluation metric. The model achieving the highest balanced accuracy on the validation folds was then evaluated on the corresponding

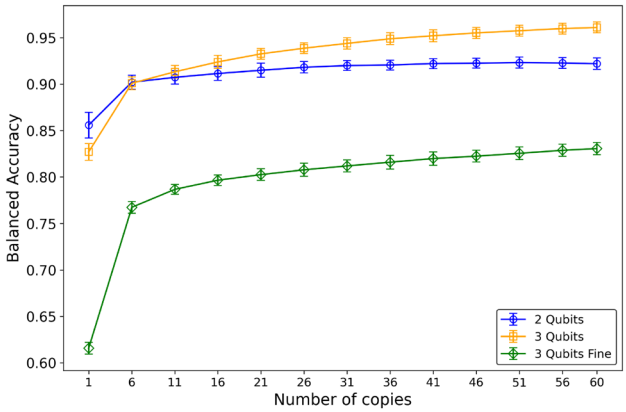

Fig. 2. Average performance of the PGM as a function of the number of copies for a dataset of 10,000 samples. For each number of copies, the experiment was repeated ten times using the same set of random seeds to reproduce identical 80%/20% train–test splits.

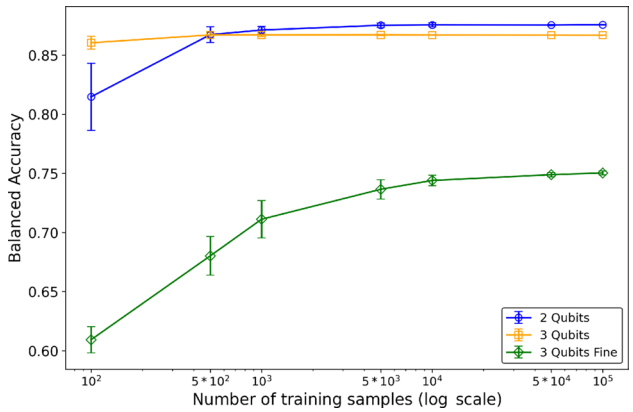

Fig. 3. Relationship between balanced accuracy and the number of training samples. The analysis was conducted using a fixed test set of 50,000 samples. For each training set size, the experiment was repeated ten times using an identical set of random seeds to ensure the reproduction of the same samplings, and the reported results were averaged over these repetitions.

test set. This entire procedure was repeated 10 times with different random seeds to ensure statistical robustness. For each dataset size, the final balanced accuracy reported for each model was obtained by averaging the test balanced accuracy scores across the repetitions.

Given that 10 splits were performed for each dataset size, multiple hyperparameter configurations may have been identified as optimal for each classifier. Therefore, in Table S2 (in Supplementary Material), the reported hyperparameter configuration corresponds to the most frequently selected one across all training set sizes. In cases of equal frequency, configurations identified in the largest datasets (30,000 or 40,000) were prioritized.

In more details, as mentioned at the end of Section The PGM classifier, a particularly useful preprocessing step for the PGM consists in the near-optimal tuning of three parameters, namely the choice of the *encoding*, the *rescaling factor* and the *number of copies*, which allow one to achieve further significant improvements in classification accuracy. Previous empirical investigations[3,15,19] show that, as the number of copies increases, accuracy typically improves up to a plateau, beyond which the gains become marginal. Anticipating the results discussed in detail in the next section, Fig. 2 reports the accuracy trend as a function of the number of copies, highlighting both the improvement and the subsequent saturation. Interestingly, this saturation also depends on the choice of the rescaling factor. For this reason, we perform a *GridSearch* over the *rescaling factor*, with values specified in Table S1 of Supplementary Material, in order to identify the minimum number of copies (together with the best choice of the encoding) sufficient to yield a genuinely significant increase in accuracy.

Moreover, as already noted, in the experimental campaign presented in the next Section, each classification is performed on datasets of different cardinalities. Figure 3 shows that increasing the dataset cardinality yields a further improvement in the classification accuracy. It is also evident that this increase reaches a barren plateau for training-set sample counts in the range immediately above 10^4. For this reason, each experiment was repeated on datasets with cardinalities of 10.000, 20.000, 30.000, and 40.000, so that the number of samples in the corresponding training sets (80% of the initial dataset) falls precisely within that range.

Overall, these considerations indicate that the PGM protocol, besides offering a natural scalability with respect to the number of qubits (as discussed at the end of Section Analytical description), also exhibits scalability with respect to both the number of copies and samples. Further, thanks to the Kernel Trick techniques[18] mentioned at the end of Section The PGM classifier, in the regime where the dataset cardinality is not too large, these choices do not entail a prohibitive increase in computational cost.

In the following section, we present the experimental results in detail, comparing the performance of the PGM with that of other well-established classical classifiers. The performance indicators adopted in our evaluation are *balanced accuracy*, *precision*, *recall*, and *F1-score*.

Results

In this section, we present a quantitative comparison between the PGM classifier and a set of classical baseline classifiers optimised through hyper-parameter tuning and applied to different sets of mixed states, as discussed in Section Analytical description. In particular, the analysis is performed for three quantum-system configurations: 2-qubit (Table 1, Figs. 4, 5, 6), 3-qubit (Table 2, Figs. 7, 8, 9), and 3-qubit-with-five-labels (Table 3, Figs. 10, 11, 12). For each configuration, a dedicated table reports the mean balanced accuracy together with macro-averaged precision, recall, and F1-score, aggregated over all four training-set sizes (10k, 20k, 30k, and 40k samples). In addition, confusion matrices for the PGM classifier are reported for the three largest datasets (20k, 30k, and 40k samples), providing a detailed view of class-specific misclassification patterns for each problem setting. For all qubit configurations, the balanced accuracy, the Precision, the Recall, and the F1-score of the PGM and the classical baselines are reported as mean ± standard deviation over ten independent runs with different random seeds. The dataset size indicated in each table corresponds to the total number of samples before applying the random 80%/20% train–test split. Detailed per-class precision, recall, and F1-score values for the PGM classifier in the 2-qubit, 3-qubit, and 3-qubit-with-five-labels experiments are reported in the Supplementary Material (Tables S3, S4, and S5).

Two-Qubit case

See Table 1, Figs. 4, 5, 6.

Classifier	Bal. Accuracy	Precision	Recall	F1
Quadratic Discriminant Analysis	0.344 ± 0.012	0.228 ± 0.107	0.344 ± 0.012	0.255 ± 0.082
Linear Discriminant Analysis	0.365 ± 0.011	0.364 ± 0.011	0.365 ± 0.011	0.363 ± 0.011
Bernoulli Naive Bayes	0.608 ± 0.007	0.579 ± 0.008	0.608 ± 0.007	0.585 ± 0.008
Gaussian Naive Bayes	0.826 ± 0.004	0.824 ± 0.004	0.826 ± 0.004	0.824 ± 0.004
Nearest Neighbors	0.855 ± 0.015	0.871 ± 0.012	0.855 ± 0.015	0.857 ± 0.016
Extra Tree	0.911 ± 0.009	0.911 ± 0.009	0.911 ± 0.009	0.911 ± 0.009
Random Forest	0.923 ± 0.007	0.924 ± 0.007	0.923 ± 0.007	0.923 ± 0.008
PGM	$\mathbf{0.931 \pm 0.005}$	$\mathbf{0.933 \pm 0.006}$	$\mathbf{0.932 \pm 0.006}$	$\mathbf{0.932 \pm 0.006}$
Multilayer Perceptron	0.943 ± 0.012	0.944 ± 0.012	0.943 ± 0.012	0.943 ± 0.012
Support Vector Machine	0.945 ± 0.005	0.946 ± 0.005	0.945 ± 0.005	0.945 ± 0.005

Table 1. Mean balanced accuracy of each classifier evaluated across all dataset sizes for the 2-qubit case. All reported precision, recall, and F1-score values are macro-averaged across the classes.

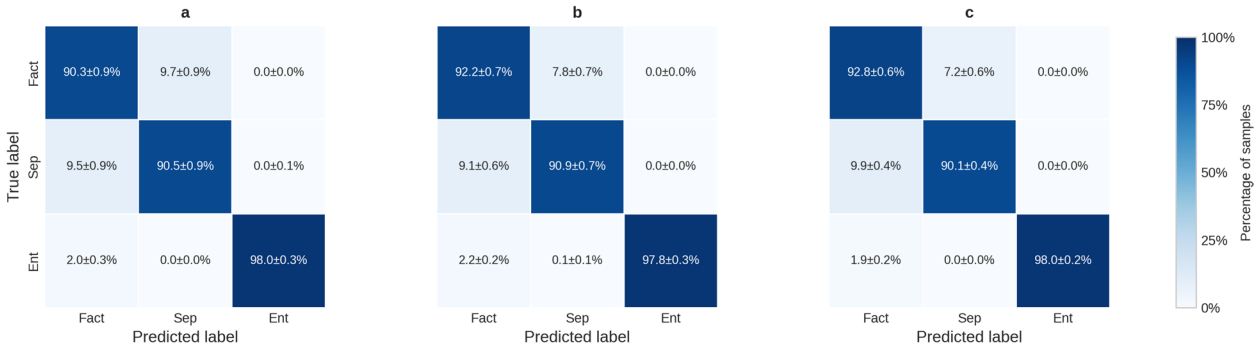

Fig. 4. Confusion matrices of the PGM classifier for the 2-qubit mixed system, averaged element-wise over 10 independent runs. Each matrix is computed on a test set comprising 20% of the samples from datasets of size (**a**) 20,000, (**b**) 30,000, and (**c**) 40,000.

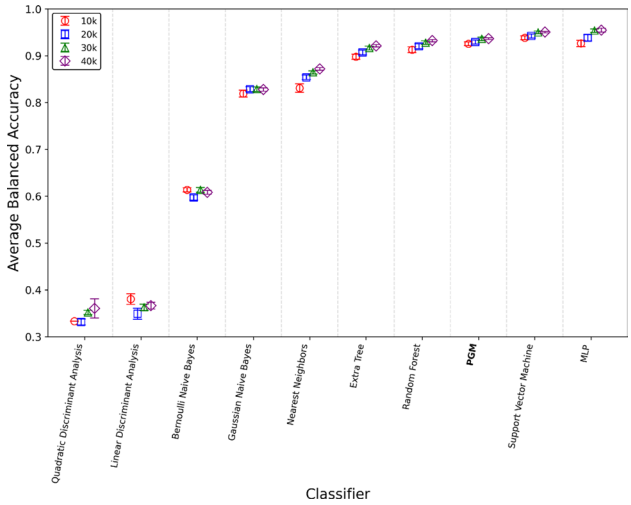

Fig. 5. "Balanced accuracy of hyperparameter-optimized classical and quantum classifiers for 2 qubits. We display the mean and standard deviation over 10 individual runs with distinct seeds. The dataset size is given before an 80%/20% train/test split is applied.

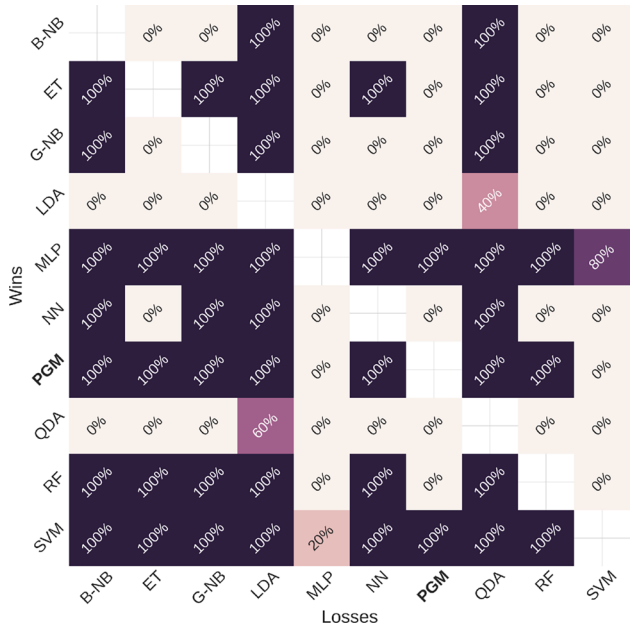

Fig. 6. Biclustering map, when a classifier A outperforms a classifier B according to the Balanced Accuracy. A lighter color indicates a lower percentage.

Three-Qubit case

See Table 2, Figs. 7, 8, 9.

Classifier	Bal. Accuracy	Precision	Recall	F1
Bernoulli Naive Bayes	0.346 ± 0.009	0.347 ± 0.01	0.346 ± 0.009	0.345 ± 0.009
Linear Discriminant Analysis	0.357 ± 0.005	0.357 ± 0.005	0.357 ± 0.005	0.357 ± 0.005
Quadratic Discriminant Analysis	0.389 ± 0.018	0.362 ± 0.02	0.389 ± 0.018	0.326 ± 0.012
Nearest Neighbors	0.611 ± 0.007	0.714 ± 0.065	0.611 ± 0.007	0.507 ± 0.007
Gaussian Naive Bayes	0.837 ± 0.001	0.835 ± 0.002	0.837 ± 0.001	0.833 ± 0.001
Extra Tree	0.874 ± 0.003	0.893 ± 0.001	0.874 ± 0.003	0.865 ± 0.00
Random Forest	0.949 ± 0.011	0.952 ± 0.01	0.949 ± 0.011	0.948 ± 0.012
Multilayer Perceptron	0.962 ± 0.012	0.963 ± 0.011	0.962 ± 0.012	0.962 ± 0.012
PGM	$\mathbf{0.971 \pm 0.011}$	$\mathbf{0.972 \pm 0.006}$	$\mathbf{0.971 \pm 0.006}$	$\mathbf{0.971 \pm 0.006}$
Support Vector Machine	0.978 ± 0.008	0.978 ± 0.007	0.978 ± 0.008	0.978 ± 0.008

Table 2. Mean balanced accuracy of each classifier evaluated across all dataset sizes for the 3-qubit case. All reported precision, recall, and F1-score values are macro-averaged across the classes.

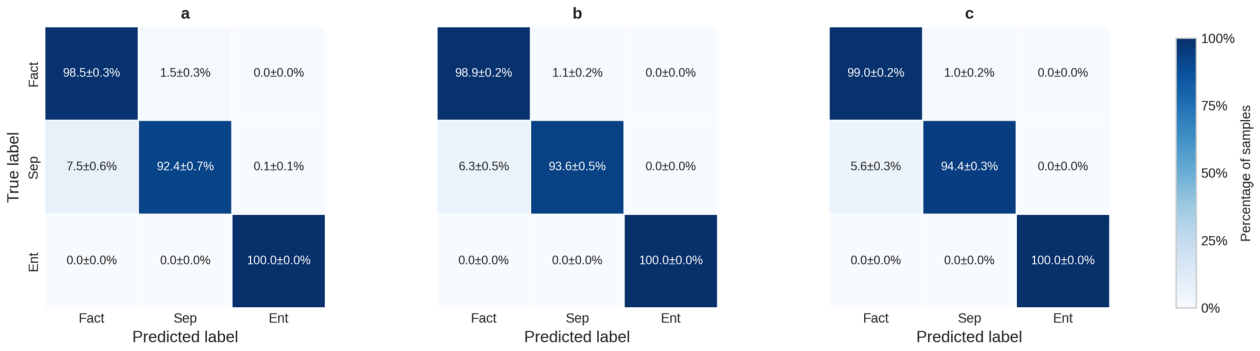

Fig. 7. Confusion matrices of the PGM classifier for the 3-qubit mixed system, averaged element-wise over 10 independent runs. Each matrix is computed on a test set comprising 20% of the samples from datasets of size (**a**) 20,000, (**b**) 30,000, and (**c**) 40,000.

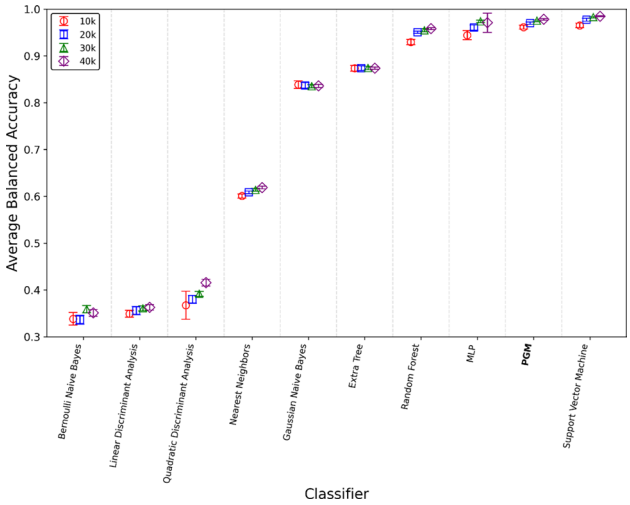

Fig. 8. Balanced accuracy of hyperparameter-optimized classical and quantum classifiers for 3 qubits. We display the mean and standard deviation over 10 individual runs with distinct seeds. The dataset size is given before an 80%/20% train/test split is applied.

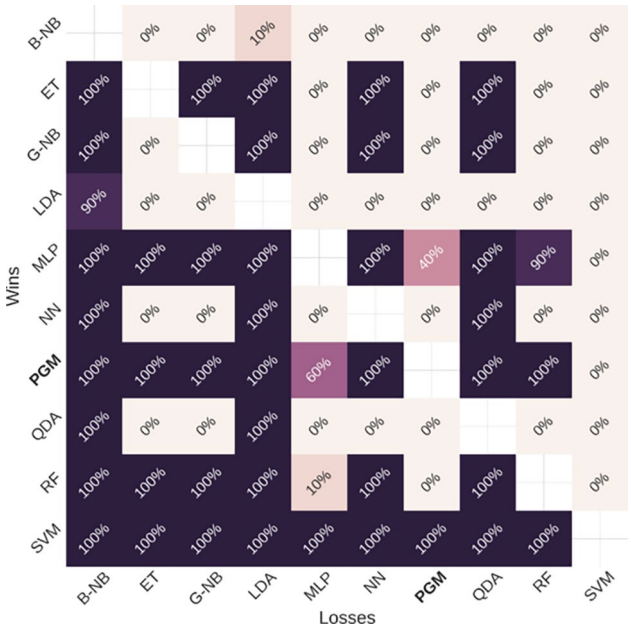

Fig. 9. Biclustering map, when a classifier A outperforms a classifier B according to the Balanced Accuracy. A lighter color indicates a lower percentage.

Three-Qubit case extended - from three to five classes via refinement of the separable class

See Table 3, Figs. 10, 11, 12.

Classifier	Bal. Accuracy	Precision	Recall	F1
Quadratic Discriminant Analysis	0.229 ± 0.007	0.237 ± 0.01	0.229 ± 0.007	0.221 ± 0.009
Linear Discriminant Analysis	0.247 ± 0.009	0.251 ± 0.009	0.247 ± 0.009	0.244 ± 0.009
Bernoulli Naive Bayes	0.407 ± 0.001	0.408 ± 0.002	0.407 ± 0.001	0.406 ± 0.001
Nearest Neighbors	0.595 ± 0.028	0.717 ± 0.034	0.595 ± 0.028	0.539 ± 0.027
Gaussian Naive Bayes	0.776 ± 0.002	0.797 ± 0.001	0.776 ± 0.002	0.768 ± 0.002
Support Vector Machine	0.865 ± 0.021	0.874 ± 0.022	0.865 ± 0.021	0.859 ± 0.024
Extra Tree	0.871 ± 0.020	0.874 ± 0.018	0.871 ± 0.020	0.869 ± 0.021
Random Forest	0.895 ± 0.014	0.893 ± 0.014	0.895 ± 0.014	0.893 ± 0.014
PGM	$\mathbf{0.895 \pm 0.012}$	$\mathbf{0.897 \pm 0.011}$	$\mathbf{0.895 \pm 0.013}$	$\mathbf{0.894 \pm 0.013}$
Multilayer Perceptron	0.902 ± 0.03	0.902 ± 0.03	0.902 ± 0.03	0.902 ± 0.03

Table 3. Mean balanced accuracy of each classifier evaluated across all dataset sizes for the extended 3-qubit case. All reported precision, recall, and F1-score values are macro-averaged across the classes.

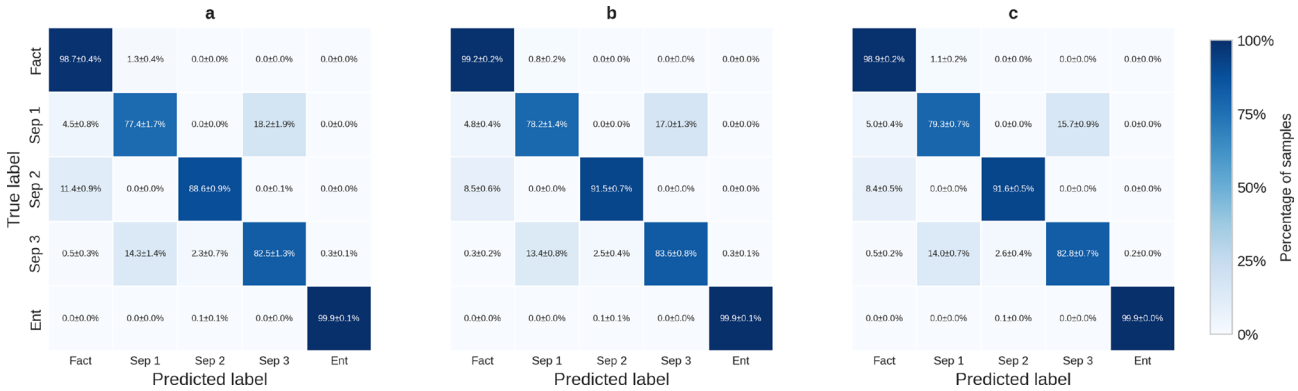

Fig. 10. Confusion matrices of the PGM classifier for the 3-qubit mixed system with five classification labels, averaged element-wise over 10 independent runs. Each matrix is computed on a test set comprising 20% of the samples from datasets of size (**a**) 20,000, (**b**) 30,000, and (**c**) 40,000.

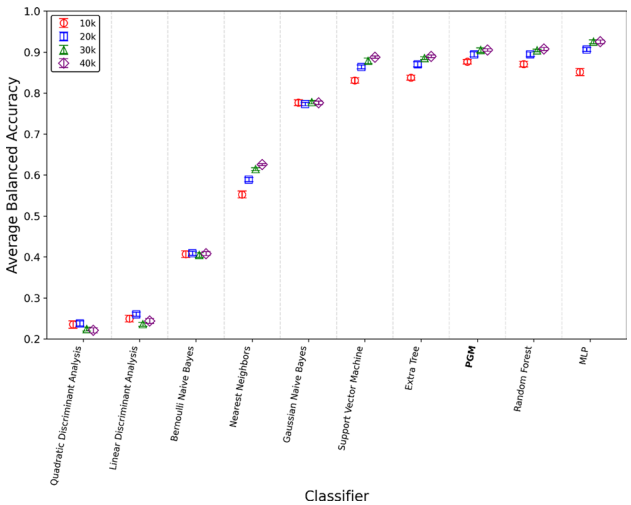

Fig. 11. Balanced accuracy of hyperparameter-optimized classical and quantum classifiers for 3 qubits. We display the mean and standard deviation over 10 individual runs with distinct seeds. The dataset size is given before an 80%/20% train/test split is applied.

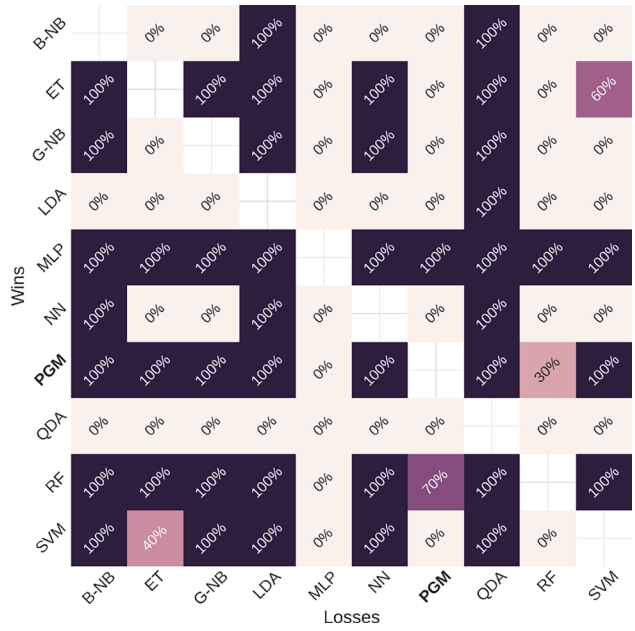

Fig. 12. Biclustering map, when a classifier A outperforms a classifier B according to the Balanced Accuracy. A lighter color indicates a lower percentage.

Discussion

The results presented in this work confirm the strong and consistent performance of the PGM classifier, extending its effectiveness from pure states to the more complex domain of mixed quantum states. The comparative evaluation against a wide range of classical classifiers such as Support Vector Machine and Neural Networks shows that the PGM classifier performs on par with, or in some cases slightly below, the most accurate models.

Both in the two-qubit and in the three-qubit scenario, the classifier achieves a balanced accuracy that generally exceeds 90% across all correlation classes product, separable, and entangled.

As shown in Fig. 3, although increasing the number of samples generally leads to an improvement in classification performance, the gain becomes progressively smaller and eventually negligible for training-set sizes that fall in the range immediately above 10.000. This behavior also extends to the three-qubit case, and is clearly confirmed by the previous tables (Tables 1,2,3). Moreover, the low values of the standard deviation indicate that this asymptotic trend observed for the balanced accuracy also holds for the other metrics considered (Precision, Recall and F1-score).

A more nuanced analysis of the fine-grained separable classes reveals that tripartite separable states (Sep 1) are often misclassified as bipartite separable mixed states (Sep 3), and vice-versa. This ambiguity can be physically motivated by the overlapping statistical signatures between convex mixtures of product states and bipartite-entangled marginal structures: both can lead to intermediate purities and partial correlations that are not easily distinguishable based solely on second-order features. Moreover, certain tripartite separable states may exhibit bipartite entanglement-like behavior under partial traces, complicating their identification within a classification framework. Moreover, it is quite reasonable to note that a non-negligible fraction of bipartite product states (Sep 2) are sometimes misclassified as fully product (factorized) states: indeed, the presence of a tensor-product structure in both cases (not mitigated by any convex combination) can sometimes make it harder to distinguish between the two classes.

Overall, the PGM classifier emerges as a powerful, theoretically principled, and computationally tractable tool for learning correlation structures in quantum systems. Its performance across different correlation classes and dimensional regimes highlights its potential as a benchmark method for quantum-inspired classification tasks in the mixed-state regime.

Further developments of this line of research could involve extending the method to higher-dimensional systems, both in terms of qubit number and degree of state mixing, with the aim of assessing the robustness of the PGM classifier under realistic conditions of noise, decoherence, and experimental imperfections. Another particularly promising direction lies in integrating the classifier with advanced machine learning architectures, such as deep neural networks or variational autoencoders, in order to explore adaptive configurations in supervised or unsupervised learning settings. In particular, the combination of PGM-based approaches with variational optimization techniques and automated feature extraction could enhance the efficiency and scalability of quantum-inspired classifiers in high-dimensional environments.

Moreover, the use of the PGM classifier in concrete applications represents a highly promising perspective. Examples include quantum device benchmarking, where accurate classification of state correlation properties may provide reliable indicators of operational quality and coherence; the efficient detection of quantum resources such as entanglement or non-locality in distributed computation or quantum communication scenarios; and its use as a preprocessing stage for more complex quantum machine learning algorithms, where fast classification of quantum inputs can dynamically inform the selection of circuits or learning strategies.

These directions, in addition to strengthening the theoretical significance of the approach, also highlight its potential for practical impact, making the PGM classifier a compelling candidate within the rapidly evolving landscape of quantum-enhanced machine learning.

Finally, it is worth emphasizing that the experimental results presented here are not meant to establish a general hierarchy among classification approaches for characterizing quantum correlations in mixed states. Rather, they indicate that a classifier grounded in state-discrimination principles specifically, in the form of a Pretty Good Measurement achieves competitive performances for the classification of different correlation regimes in mixed quantum states, where the geometric and statistical structure of the state space poses particular challenges.

Moreover, the practical significance of the PGM approach also concerns concrete operational features. First, the PGM classifier does not require parametric optimization procedures: the measurement construction follows deterministically from representative class states and their *a priori* probabilities, avoiding the extensive tuning of architectures, learning rates, regularization schemes, or other optimization processes, typical of neural models. These features make the PGM classifier particularly suitable in regimes where reproducibility, physical transparency, and label certifiability are prioritized over marginal gains in predictive accuracy. Second, the decision rule is expressed directly in terms of measurement operators $\{F_i\}$ and observable quantities of the form $\mathrm{tr}(F_i\rho)$, providing a level of transparency and controllability that is particularly valuable when the goal is to relate the classifier's output to quantities and structures native to quantum information theory.

Finally, and most importantly, the relevance of the PGM approach is further motivated by its potential implementability on quantum hardware. Indeed, the testing phase of the classifier amounts to applying a POVM and can naturally be realized on a quantum device via a projective measurement on an extended Hilbert space, by directly exploiting the Naimark dilation theorem, as we have already illustrated in some recent simple applications[3]. Current work is devoted to the quantum implementation of the training phase (i.e., the efficient construction of the measurement operators from the training set). Looking ahead, this research direction may lead to an algorithm that, while "quantum-inspired" in its formulation, can be executed natively on quantum hardware, with the possibility of combining the high-accuracy performances associated with the PGM protocol with potential computational-complexity advantages enabled by implementation on a quantum device. Naturally, these developments should be regarded as work in progress and constitute one of the main future directions opened by the present study.

Data availability

The datasets generated and analyzed during the current study are available in the Zenodo repository: https://zenodo.org/records/18229713[29].

Received: 30 September 2025; Accepted: 9 March 2026
Published online: 30 March 2026

References

1. Motta, M. & Rice, J. E. Emerging quantum computing algorithms for quantum chemistry. *WIREs Comput. Mol. Sci.* **12**, e1580. https://doi.org/10.1002/wcms.1580 (2022).
2. Huang, H.-Y. et al. Power of data in quantum machine learning. *Nat. Commun.* **12**, 2631. https://doi.org/10.1038/s41467-021-22539-9 (2021).
3. Giuntini, R., Granda Arango, A. C., Freytes, H., Holik, F. H. & Sergioli, G. Multi-class classification based on quantum state discrimination. *Fuzzy Sets Syst.* **467**, 108509. https://doi.org/10.1016/j.fss.2023.03.012 (2023).
4. Schatzki, L., Arrasmith, A., Coles, P. J. & Cerezo, M. Entangled datasets for quantum machine learning. *arXiv* 2109.03400. https://doi.org/10.48550/arXiv.2109.03400 (2021).
5. Perrier, E., Youssry, A. & Ferrie, C. QDataSet: Quantum datasets for machine learning. *Sci. Data* **9**, 582. https://doi.org/10.1038/s41597-022-01639-1 (2022).
6. Mahdian, M. & Mousavi, Z. Entanglement detection with quantum support vector machine on near-term quantum devices. *Sci. Rep.* **15**, 11931. https://doi.org/10.1038/s41598-025-95897-9 (2025).
7. Havlíček, V. et al. Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212. https://doi.org/10.1038/s41586-019-0980-2 (2019).
8. El Ayachi, F., Ait Mansour, H. & El Baz, M. Classification of entanglement distribution using machine learning. *Commun. Theor. Phys.* **77**, 065104. https://doi.org/10.1088/1572-9494/ad9f47 (2025).
9. Ureña, J. et al. Entanglement detection with classical deep neural networks. *Sci. Rep.* **14**, 18109. https://doi.org/10.1038/s41598-024-68213-0 (2024).
10. Goes, C. B. D., Canabarro, A., Duzzioni, E. I. & Maciel, T. O. Automated machine learning can classify bound entangled states with tomograms. *Quantum Inf. Process.* **20**, 99. https://doi.org/10.1007/s11128-021-03037-9 (2021).
11. Vintskevich, S. V., Bao, N., Nomerotski, A., Stankus, P. & Grigoriev, D. A. Classification of four-qubit entangled states via machine learning. *Phys. Rev. A* **107**, 032421. https://doi.org/10.1103/PhysRevA.107.032421 (2023).
12. Harney, C., Pirandola, S., Ferraro, A. & Paternostro, M. Entanglement classification via neural network quantum states. *New J. Phys.* **22**, 045001. https://doi.org/10.1088/1367-2630/ab783d (2020).
13. Deng, D.-L. Machine learning detection of Bell nonlocality in quantum many-body systems. *Phys. Rev. Lett.* **120**, 240402. https://doi.org/10.1103/PhysRevLett.120.240402 (2018).
14. Huang, Y. et al. Direct entanglement detection of quantum systems using machine learning. *npj Quantum Inf.* **11**, 29. https://doi.org/10.1038/s41534-025-00970-w (2025).
15. Sergioli, G. et al. Classification of quantum correlations via quantum-inspired machine learning. *Quantum Mach. Intell.* **7**, 112. https://doi.org/10.1007/s42484-025-00339-4 (2025).
16. Bae, J. & Kwek, L.-C. Quantum state discrimination and its applications. *J. Phys. A Math. Theor.* **48**, 083001 (2015).
17. Barnett, S. M. & Croke, S. *Quantum state discrimination. Adv. Opt. Photon.* **1**, 238–278 (2009).
18. Cruzeiro, E. Z., De Mol, C., Massar, S. & Pironio, S. Quantum-inspired classification based on quantum state discrimination. *Quantum Mach. Intell.* **6**, 1–18. https://doi.org/10.1007/s42484-024-00216-6 (2024).
19. Giuntini, R. et al. Quantum-inspired algorithm for direct multi-class classification. *Appl. Soft Comput.* **134**, 109956. https://doi.org/10.1016/j.asoc.2022.109956 (2023).
20. Sergioli, G., Bosyk, G. M., Santucci, E. & Giuntini, R. A quantum-inspired version of the classification problem. *Int. J. Theor. Phys.* **56**, 3880–3888. https://doi.org/10.1007/s10773-017-3371-1 (2017).

21. Harrow, A. W. & Low, R. A. Random quantum circuits are approximate 2-designs. *Commun. Math. Phys.* **291**, 257–302. https://doi.org/10.1007/s00220-009-0873-6 (2009).
22. Choi, J. et al. Preparing random states and benchmarking with many-body quantum chaos. *Nature* **613**, 468–473. https://doi.org/10.1038/s41586-022-05442-1 (2023).
23. Datta, A. Negativity of random pure states. *Phys. Rev. A* **81**, 052312. https://doi.org/10.1103/PhysRevA.81.052312 (2010).
24. Dahlsten, O. C. O., Lupo, C., Mancini, S. & Serafini, A. Entanglement typicality. *J. Phys. A* **47**, 363001. https://doi.org/10.1088/1751-8113/47/36/363001 (2014).
25. Singh, U., Zhang, L. & Pati, A. K. Average coherence and its typicality for random pure states. *Phys. Rev. A* **93**, 032125. https://doi.org/10.1103/PhysRevA.93.032125 (2016).
26. Życzkowski, K. & Sommers, H.-J. Induced measures in the space of mixed quantum states. *J. Phys. A* **34**, 7111–7125 (2001).
27. Życzkowski, K., Penson, K. A., Nechita, I. & Collins, B. Generating random density matrices. *J. Math. Phys.* **52**, 062201. https://doi.org/10.1063/1.3595693 (2011).
28. Bengtsson, I. & Życzkowski, K. *Geometry of Quantum States: An Introduction to Quantum Entanglement* 2nd edn. (Cambridge University Press, 2017).
29. Cuccu, C., Rieger, C. S., Sergioli, G. & Giuntini, R. Mixed quantum states datasets for entanglement classification. *Zenodo* https://doi.org/10.5281/zenodo.18229713 (2025).

Author contributions

Giuseppe Sergioli, Carlo Cuccu, Carla Sophie Rieger, Andrés Camilo Granda Arango, Bikash K. Behera, Riccardo Era, and Roberto Giuntini all contributed equally and collaboratively to all stages of the work. The conception of the study, the development of methods, the implementation of experiments, the analysis of results, and the writing and revision of the manuscript were carried out in a transversal and joint manner by all authors.

Funding

Giuseppe Sergioli and Roberto Giuntini acknowledge financial support from the PRIN-PNRR project "*Quantum Models for Logic, Computation and Natural Processes* – Qm4Np" (code: P2022A52CR), from the PRIN-2022 project "*The COst of Reasoning: Theory and EXperiments* – CORTEX" (code: 2022ZLLR3T) and from the Fondazione di Sardegna project "*An algebraic approach to hyperintensionality*" (code: F23C25000360007). Roberto Giuntini is partially funded by the TÜV SÜD Foundation, the Federal Ministry of Education and Research (BMBF) and the Free State of Bavaria under the Excellence Strategy of the Federal Government and the Länder, as well as by the Technical University of Munich-Institute for Advanced Study. Carla Rieger has been sponsored by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 13E18CHA) and CERN through the CERN Quantum Technology Initiative.

Declarations

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1038/s41598-026-44068-5.

Correspondence and requests for materials should be addressed to G.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by-nc-nd/4.0/.

© The Author(s) 2026