

Quantum-Inspired Machine Learning for Bioimaging: Two State-Discrimination Case Studies

Giuseppe Sergioli¹, Andrés Camilo Granda Arango¹, Carlo Cuccu¹, and Roberto Giuntini¹

University of Cagliari, Italy

Abstract. Bioimaging offers a concrete stress test for quantum machine learning (QML): data are high dimensional, labels are limited, acquisition protocols vary, and clinically relevant classes often overlap. We summarize two biomedical applications of a quantum-inspired framework in which classification is derived from quantum state discrimination. Image or radiomics vectors are encoded as density operators, class prototypes are mixed states, and decisions are produced by Positive Operator-Valued Measures (POVMs). The first case uses the Helstrom Quantum Classifier (HQC) for clonogenic-assay segmentation; the second uses the Pretty Good Measurement (PGM) classifier for computed-tomography (CT) and positron-emission-tomography/computed-tomography (PET/CT) radiomics. The main point is methodological: state-discrimination principles yield hardware-independent QML algorithms that are interpretable, data-efficient and promising for practical bioimaging.

Keywords: Quantum-inspired machine learning · Pretty Good Measurement · Helstrom measurement · Bioimaging · Radiomics

1 Quantum state discrimination as a classifier

We show that quantum state discrimination can be used as a constructive route to quantum-inspired learning algorithms: the quantum formalism suggests how to encode data, aggregate examples into states and turn class overlap into a measurement, while the resulting classifier remains executable on classical hardware. This is particularly relevant for bioimaging, where practical adoption requires robust, interpretable and data-efficient models rather than immediate access to quantum devices.

The framework treats supervised learning as a measurement problem rather than only as empirical risk minimization [7, 8, 1]. Given labelled samples $S = \{(x_j, \lambda_j)\}$, each feature vector is mapped to a density operator ρ_x . Samples of class i define a mixed centroid $\rho^{(i)} = |S_i|^{-1} \sum_{x_j \in S_i} \rho_{x_j}$ with prior p_i . In binary tasks, HQC measures the sign structure of $p_1 \rho^{(1)} - p_2 \rho^{(2)}$. In multi-class tasks, PGM builds one POVM from $\sigma = \sum_i p_i \rho^{(i)}$ via $F_i = \sigma^{-1/2} p_i \rho^{(i)} \sigma^{-1/2}$, with completion on the kernel of σ ; prediction maximizes $\text{tr}(F_i \rho_x)$.

This formulation is directly connected to QTML topics. Density-matrix embeddings and tensor-copy maps $\rho_x^{\otimes n}$ act as quantum-inspired feature maps related to quantum kernels and Hilbert-space encodings [3, 2]. The POVM returns Born-like scores, hence a natural confidence profile over competing classes. Fixed Helstrom or PGM measurements also provide baselines and initial geometries for trainable POVM layers, quantum kernels or variational circuits.

2 Case study I: clonogenic-assay evaluation

A clonogenic assay estimates whether treated cells retain the ability to proliferate. After irradiation or drug exposure, surviving cells form stained colonies; colonies with at least 50 densely packed cells are counted to estimate surviving fractions in radiobiology. The image-analysis task is therefore a pixel-level separation of biologically meaningful colony material from background, stains and merged structures, with direct impact on dose-response curves.

The reference study applies HQC to scanner-acquired well-plate images from four tumour cell lines: MDA-MB-231 and MCF7 breast-cancer cells, and U87-MG and U251 glioblastoma cells [9]. These lines differ in morphology, staining and background artifacts. For each line, 30 images were acquired at 800 dpi and 24-bit colour depth; one 301×301 pixel region was analysed per image. Biologist-validated masks were obtained through spatial fuzzy C-means on entropy maps and post-processing that removed components below the colony-size criterion.

Six image representations were compared: RGB, perceptual CIE $L^*u^*v^*$, and four Haralick texture maps from the Gray-Level Co-occurrence Matrix (GLCM)—contrast, correlation, energy and homogeneity. The GLCM counts how often neighbouring grey levels occur together; homogeneity is high for compact, smoothly stained regions, while contrast highlights local intensity changes. Thus the experiment also evaluates which descriptors carry biological information.

Overall, $4 \times 30 \times 6 = 720$ datasets were generated. Each used only 0.2% of pixels (181 samples), with standardization, development/test splitting and 5-fold cross-validation. HQC hyperparameters included encoding, rescaling, tensor-copy number and class weights, and the comparison involved 18 classical classifiers for balanced accuracy. HQC reached balanced accuracy 0.959 ± 0.036 on MDA-MB-231, 0.919 ± 0.050 on U87-MG, 0.965 ± 0.033 on MCF7 and 0.979 ± 0.029 on U251. GLCM homogeneity was best for three cell lines, while $L^*u^*v^*$ was best for MCF7. The QML relevance is that class imbalance, few labelled pixels and texture uncertainty are combined in one operator-based decision rule.

3 Case study II: PGM radiomics in CT and PET/CT

Radiomics converts medical images into descriptors of shape, intensity and texture [4]; operationally, it turns a tumour volume into a high-dimensional fingerprint. This is naturally close to QML: each patient is a noisy sample, each clinical group is a statistical mixture, and PGM compares all mixtures through a native multi-class measurement.

Lung-cancer subtyping. The first radiomics task concerns non-small-cell lung carcinoma (NSCLC). Computed tomography (CT) scans from two public datasets include 466 patients: squamous cell carcinoma (SCC, $N = 152$), adenocarcinoma (ADC, $N = 150$), large cell carcinoma (LCC, $N = 106$), and not-otherwise-specified tumours (NOS, $N = 58$) [6]. SCC, ADC and LCC are distinct tissue subtypes; NOS is a heterogeneous residual category. Three tasks were studied: SCC vs. ADC, SCC/ADC/LCC, and




Fig. 1. Bioimaging inputs and PGM scores.

SCC/ADC/LCC/NOS. CT volumes were resampled to $1 \times 1 \times 1 \text{ mm}^3$ and harmonized to reduce scanner/site effects. PGM reached test accuracy 0.755 ± 0.041 and ADC area under the curve (AUC) 0.812 ± 0.033 in the binary task. In the 3-class task, accuracy was 0.532 ± 0.035 ; in the 4-class task it decreased to 0.424 ± 0.086 , although NOS retained one-vs-rest AUC 0.832. This is informative rather than merely negative: radiomic overlap appears as overlap between class states, especially for the ambiguous NOS group.

Prostate-cancer risk stratification. The second task concerns prostate cancer (PCa), where the aim is to distinguish treatment-naïve high-risk from low-risk patients. The modality is [^{18}F]PSMA-1007 PET/CT: positron-emission tomography (PET) highlights prostate-specific membrane antigen (PSMA) uptake, linked to tumour activity, while CT supplies anatomy. The cohort contains 143 patients, balanced by post-surgical Gleason grade—a histological score of aggressiveness—into 72 high-risk and 71 low-risk subjects [5]. Lesion volumes were manually delineated. PET intensities were normalized as standardized uptake value by body weight (SUVbw), resampled to $2 \times 2 \times 2 \text{ mm}^3$, discretized with bin width 0.25, and described by original, Laplacian-of-Gaussian (LoG) and wavelet features. Robust subsets were selected by Least Absolute Shrinkage and Selection Operator (LASSO) analysis. On the same 30 stratified 80/20 splits and 13 feature subsets used by the reference ensemble model, PGM obtained best average accuracy 72.62%, best AUC 80.10%, and sensitivity/specificity around 72–73%. Sensitivity and specificity matter clinically because they balance missed high-risk disease against overtreatment; in PGM this trade-off can be tuned through priors, thresholds or modified measurement effects.

4 QTML perspective

These case studies show that quantum state discrimination is not only a quantum-information primitive, but also a design principle for practical machine-learning algorithms. HQC and PGM retain the operational structure of state preparation, priors, mixed states, POVMs and Born-like decision scores, while remaining hardware-independent. This makes them useful intermediate models between classical radiomics pipelines and future circuit-based QML.

The biomedical relevance follows from the nature of bioimaging data. Colonies versus background, lung-cancer subtypes, and prostate-cancer risk groups are not separated by one deterministic marker; they are overlapping distributions affected by noise, scanner variability, biological heterogeneity and limited labels. A state-discrimination view makes this overlap explicit and turns it into measurable geometry. The resulting algorithms are interpretable, adjustable to asymmetric costs, and immediately deployable on classical hardware. For QTML, bioimaging is therefore not merely an application area: it is a realistic benchmark where density embeddings, quantum kernels, measurement design and uncertainty quantification can be tested against clinically meaningful constraints.

References

1. Giuntini, R., Granda Arango, A.C., Freytes, H., Holik, F.H., Sergioli, G.: Multi-class classification based on quantum state discrimination. *Fuzzy Sets and Systems* **467**, 108509 (2023)
2. Giuntini, R., Holik, F., Park, D.K., Freytes, H., Blank, C., Sergioli, G.: Quantum-inspired algorithm for direct multi-class classification. *Applied Soft Computing* **134**, 109956 (2023)
3. Havlicek, V., Corcoles, A.D., Temme, K., Harrow, A.W., Kandala, A., Chow, J.M., Gambetta, J.M.: Supervised learning with quantum-enhanced feature spaces. *Nature* **567**(7747), 209–212 (2019)
4. Lambin, P., Rios-Velazquez, E., Leijenaar, R., Carvalho, S., van Stiphout, R.G.P.M., Granton, P., Zegers, C.M.L., Gillies, R., Boellard, R., Dekker, A., Aerts, H.J.W.L.: Radiomics: Extracting more information from medical images using advanced feature analysis. *European Journal of Cancer* **48**(4), 441–446 (2012)
5. Pasini, G., Stefano, A., Mantarro, C., Richiusa, S., Comelli, A., Russo, G.I., Sabini, M.G., Cosentino, S., Ippolito, M., Russo, G.: A robust [18f]-psma-1007 radiomics ensemble model for prostate cancer risk stratification. *Journal of Imaging Informatics in Medicine* **38**(3), 1388–1402 (2025)
6. Pasini, G., Stefano, A., Russo, G., Comelli, A., Marinozzi, F., Bini, F.: Phenotyping the histopathological subtypes of non-small-cell lung carcinoma: How beneficial is radiomics? *Diagnostics* **13**(6), 1167 (2023)
7. Sergioli, G., Bosyk, G.M., Santucci, E., Giuntini, R.: A quantum-inspired version of the classification problem. *International Journal of Theoretical Physics* **56**(12), 3880–3888 (2017)
8. Sergioli, G., Giuntini, R., Freytes, H.: A new quantum approach to binary classification. *PLOS ONE* **14**(5), 1–14 (2019)
9. Sergioli, G., Militello, C., Rundo, L., Minafra, L., Torrisi, F., Russo, G., Chow, K.L., Giuntini, R.: A quantum-inspired classifier for clonogenic assay evaluations. *Scientific Reports* **11**(1), 2830 (2021)