






# Q<sup>2</sup>SAR: A Quantum Multiple Kernel Learning Approach for Drug Discovery

Alejandro Giraldo  Daniel Ruiz  Mariano Caruso  Javier Mancilla  Guido Bellomo   
QNOW Technologies QNOW Technologies UGR, Granada, Spain Falcondale LLC CONICET - UBA  
Delaware, USA Delaware, USA UNIR, La Rioja, Spain Delaware, USA ICC, Argentina  
alejandro@qnow.tech daniel@qnow.tech FIDESOL, Granada, Spain javier@falcondale.pro gbellomo@icc.fcen.uba.ar  
h mcaruso@fidesol.org

**Abstract**—Quantitative Structure-Activity Relationship (QSAR) modeling is a cornerstone of computational drug discovery. This research demonstrates the successful application of a Quantum Multiple Kernel Learning (QMKL) framework to enhance QSAR classification, showing a notable performance improvement over classical methods. We apply this methodology to a dataset for identifying DYRK1A kinase inhibitors. The workflow involves converting SMILES representations into numerical molecular descriptors, reducing dimensionality via Principal Component Analysis (PCA), and employing a Support Vector Machine (SVM) trained on an optimized combination of multiple quantum and classical kernels. By benchmarking the QMKL-SVM against a classical Gradient Boosting model, we show that the quantum-enhanced approach achieves a superior AUC score, highlighting its potential to provide a quantum advantage in challenging cheminformatics classification tasks.

**Index Terms**—QSAR, classification, drug discovery, quantum machine learning, multiple kernel learning, support vector machines.

## I. INTRODUCTION

Drug discovery is an inherently complex and resource-intensive process. Computational methodologies, particularly Quantitative Structure-Activity Relationship (QSAR) modeling, are essential for the efficient evaluation and optimization of chemical compounds [1]. QSAR models establish a mathematical link between a molecule's structure and its biological activity, enabling the *in silico* prediction of properties for novel compounds, thereby prioritizing experimental efforts and reducing costs [2].

The evolution of QSAR has seen a shift from simple linear regression to sophisticated machine learning algorithms like Random Forests [3] and Support Vector Machines (SVMs) [?], which excel at capturing complex, non-linear relationships. However, the continuous growth of chemical databases like ChEMBL [4] means that molecular data is increasing in both volume and complexity. This high-dimensional feature space presents a fundamental challenge even for advanced classical algorithms.

Quantum Machine Learning (QML) offers a promising new frontier to address this challenge. By mapping classical data into the exponentially large Hilbert spaces of quantum systems, QML algorithms have the potential to identify patterns intractable for classical models [5], [6]. This work focuses on Quantum Multiple Kernel Learning (QMKL) [7], an advanced

technique that combines the strengths of various quantum and classical kernels to build a more robust and accurate classifier. We apply this framework to a real-world drug discovery use case: identifying inhibitors for the DYRK1A kinase, a critical therapeutic target for neurodegenerative disorders [8].

This paper details our complete workflow, from data curation and descriptor calculation to the implementation of the QMKL – SVM model. We provide a fair benchmark against a classical Gradient Boosting model and discuss the implications of our results, including the computational costs and future potential of this quantum-hybrid approach.

## II. BACKGROUND AND RELATED WORK

### A. Foundations of Quantitative Structure-Activity Relationships

QSAR modeling is a cornerstone methodology in computational chemistry and drug discovery. The theoretical foundation rests on the premise that similar molecular structures exhibit similar biological activities, a principle that enables the systematic prediction of pharmacological, toxicological, and physicochemical properties from structural information alone [9].

A typical QSAR workflow involves representing molecules (e.g., via SMILES strings), calculating numerical molecular descriptors, developing a model using machine learning, and validating its predictive performance [10], [11]. These descriptors can range from simple atom counts and topological indices to complex 3D features that capture the spatial arrangement of a molecule [12].

### B. Contemporary Relevance and Industrial Impact

The relevance of QSAR has intensified due to the exponential growth of chemical data. QSAR offers three fundamental advantages in modern drug development:

- **Predictive Efficiency:** Once validated, models allow for instantaneous property prediction, accelerating the assessment of large compound libraries.
- **Scalability:** Modern frameworks support the automated, high-throughput computational screening of millions of virtual compounds.
- **Regulatory Acceptance:** QSAR models are accepted by regulatory agencies like the OECD, EPA, and EMA for

toxicity and risk assessment, providing a path from computation to clinical application.

### C. Quantum Frontiers in Cheminformatics

Despite advances, classical machine learning methods face limitations in exploring high-dimensional feature spaces. The emergence of QML in cheminformatics represents a paradigm shift, aiming to harness quantum principles for enhanced molecular property prediction [13]. By mapping classical molecular data into quantum states, quantum algorithms can potentially access richer representational spaces and capture subtle structure-activity relationships that classical methods might overlook [14].

## III. METHODOLOGY

Our pipeline consists of three main stages: (1) data curation and feature engineering, (2) classification using a QMKL-SVM model, and (3) benchmarking against a classical Gradient Boosting model. Predictive performance is assessed using standard metrics, with a focus on the Area Under the Receiver Operating Characteristic Curve (ROC – AUC).

### A. Dataset Curation and Preprocessing

This study is based on a high-quality dataset of DYRK1A inhibitors, provided and curated by ProtoQSAR SL from the ChEMBL database [4]. DYRK1A is a critical therapeutic target for neurodegenerative disorders like Alzheimer’s disease [8].

The raw data for 1291 molecules underwent rigorous curation: duplicates were removed, chemical structures were validated, and only assays labeled as “single protein assay” were retained to ensure direct target-compound interaction measurements. The final dataset comprised 354 unique molecules. For binary classification, a standard medicinal chemistry threshold was used: molecules with  $pIC_{50} \geq 6.0$  (corresponding to  $IC_{50} \leq 1 \mu M$ ) were labeled “active,” and all others were labeled “inactive.” The dataset was pre-divided into a training set (283 molecules) and a test set (71 molecules).

1) *Molecular Descriptors*: A comprehensive set of molecular descriptors was calculated using RDKit and validated to ensure they fell within typical drug-like ranges. Key descriptors included:

- **Shape/Geometric**: Principal Moments of Inertia (e.g., **PMI1**) and Normalized Principal Ratios (e.g., **NPR1**).
- **Electronic/Topological**: Molecular LogP (**MolLogP**) and Topological Polar Surface Area (**TPSA**), where values  $< 140 \text{ \AA}^2$  often correlate with good blood-brain barrier permeability.
- **Connectivity**: Bertz Complexity Index (**BertzCT**) [15] and Kier-Hall Connectivity Indices (e.g., **Chi0n**).
- **Pharmacophore/Functional Group**: Counts of Hydrogen Bond Donors/Acceptors (**NumHBD**, **NumHBA**) and the fraction of  $sp^3$  hybridized carbons (**FractionCsp3**).

This high-dimensional descriptor matrix provides a comprehensive mathematical representation of each molecule’s structure.

2) *Data Processing Pipeline*: A multi-step preprocessing pipeline was implemented:

- 1) **Numerical Feature Extraction**: Each SMILES string was converted into a vector of 200 physicochemical descriptors using RDKit.
- 2) **Standardization**: Features were scaled to have a mean of 0 and a standard deviation of 1 using Scikit-learn’s `StandardScaler`, fitted only on training data.
- 3) **Dimensionality Reduction**: To prepare the data for quantum simulation, we applied Principal Component Analysis (PCA). This aggressive reduction to four components, which capture 37.1% of the total variance, was necessary to ensure compatibility with quantum simulators, despite the substantial loss of information from the original descriptor space. These four components serve as the final input features for all models.

### B. Quantum Multiple Kernel Learning

SVMs use a kernel function,  $K(\mathbf{x}_i, \mathbf{x}_j)$ , to compute similarity in a high-dimensional feature space. Our approach extends this with QMKL.

1) *Quantum Kernels*: A quantum kernel is created by encoding classical data  $\mathbf{x}$  into quantum states  $|\psi(\mathbf{x})\rangle$  via a parameterized quantum circuit (a feature map)  $U(\mathbf{x})$  [16]. The kernel value is the fidelity between the states:

$$\mathcal{K}(\mathbf{x}_i, \mathbf{x}_j) = |\langle \psi(\mathbf{x}_i) | \psi(\mathbf{x}_j) \rangle|^2 = |\langle 0 | U^\dagger(\mathbf{x}_i) U(\mathbf{x}_j) | 0 \rangle|^2. \quad (1)$$

This process can map data into a richer feature space [17]. We used a feature map of angle-embedding rotations and CNOT entangling gates, implemented in PennyLane [18].

2) *Multiple Kernel Learning and Weight Optimization*: Rather than relying on a single kernel, MKL seeks to find the best linear combination of a pre-defined set of basis kernels  $\{K_m\}$ :

$$K_{comb} = \sum_{m=1}^M w_m K_m, \quad \text{with } w_m \geq 0, \quad \sum_{m=1}^M w_m = 1. \quad (2)$$

Our basis set included 8 distinct quantum fidelity kernels (from varied feature maps) and 3 classical Radial Basis Function (RBF) kernels. To determine the optimal weights  $w_m$ , we implemented and compared multiple strategies within our framework. The two primary methods evaluated were: 1) a simple Averaging approach, where each of the  $M$  kernels is given an equal weight ( $w_m = 1/M$ ), and 2) a selection strategy based on Centered Kernel Alignment (CKA). The CKA method measures the similarity between each base kernel  $K_m$  and an ideal target kernel  $K_y = yy^T$  (where  $y$  is the vector of class labels), and our implementation selects the single kernel with the highest alignment score.

Upon evaluation, the CKA selection strategy consistently yielded the best-performing model. Therefore, the results reported in this paper correspond to the SVM trained on the single, best-aligned kernel identified through this process.

### C. Classical Benchmark Model

A classical `GradientBoostingClassifier` from Scikit-learn was trained and evaluated on the exact same 4-dimensional, PCA-processed data to ensure a fair comparison.

## IV. RESULTS AND DISCUSSION

The performance of the final QMKL – SVM model and the classical Gradient Boosting model was evaluated on the unseen test set. Table I summarizes the performance across multiple metrics.

Our proposed Q<sup>2</sup>SAR model showed superior overall performance, outperforming the baseline in the key metrics of AUC, Accuracy, Recall, and F1-Score. While the classical model achieved slightly higher precision (0.8333 vs. 0.8070), the significant gain in recall (sensitivity) from the quantum-enhanced model (0.9200 vs. 0.8000) indicates a much stronger ability to correctly identify true active compounds, which is often a primary goal in initial drug screening campaigns.

TABLE I

COMPREHENSIVE PERFORMANCE COMPARISON ON THE DYRK1A TEST SET

Metric	Gradient Boosting	Q <sup>2</sup> SAR Model
AUC	0.8037	<b>0.8750</b>
Accuracy	0.8000	<b>0.8333</b>
Precision	<b>0.8333</b>	0.8070
Recall (Sensitivity)	0.8000	<b>0.9200</b>
F1-Score	0.8163	<b>0.8598</b>

Figure 1 provides a visual comparison of the ROC curves. The plot clearly illustrates the superior discriminative power of the QMKL – SVM, as its curve consistently stays above the Gradient Boosting curve, encompassing a larger area. The results suggest that the optimally chosen kernel was able to capture complex, non-linear relationships within the data that were advantageous for identifying active molecules, even at the expense of slightly lower precision.

## V. COMPUTATIONAL COST AND SCALABILITY

A fundamental challenge for kernel methods is scalability. Training a conventional SVM scales polynomially (typically  $O(n^2)$  to  $O(n^3)$ ) with the number of samples  $n$ , due to the need to compute the  $n \times n$  kernel matrix. For quantum kernels, this is exacerbated, as building the matrix requires  $O(n^2)$  quantum circuit executions, a significant bottleneck for near-term quantum hardware. It is critical to note that all experiments in this study were performed on a classical simulator. On real quantum devices, challenges such as hardware noise and decoherence would further impact the fidelity of the kernel computation and must be addressed with error mitigation techniques.

Strategies to mitigate this exist, such as quantum kernel estimation, using representative data centroids (e.g., QuACK [19]), or reformulating the SVM optimization to scale linearly (e.g., PegasosQSVC [20]). These approaches will be crucial for achieving practical quantum advantage on large-scale problems.




Fig. 1. Comparison of ROC curves for the QMKL – SVM and Gradient Boosting models on the DYRK1A test set, with performance metrics detailed in Table I.

## VI. CONCLUSION

In this work, we developed a complete QSAR pipeline demonstrating that a quantum-enhanced model can outperform a strong classical baseline. Our QMKL – SVM approach achieved a superior AUC score (0.875 vs. 0.804) on the DYRK1A inhibitor dataset, highlighting the potential of quantum feature spaces in drug discovery. The model’s high recall, in particular, suggests its utility for hit-finding in early-stage discovery, where minimizing false negatives is paramount.

We acknowledge the limitations of this study, including the reliance on a single train-test split and the inherent “black-box” nature of kernel models. Future work should employ more rigorous validation schemes like cross-validation and explore model interpretability with techniques such as Kernel-SHAP to understand which chemical features drive the predictions.

As quantum hardware matures, overcoming the challenges of noise and scalability will be paramount. When that happens, frameworks like Q<sup>2</sup>SAR could become impactful tools for accelerating the discovery of new medicines.

## ACKNOWLEDGMENT

This work was supported by the project ECO – 20241014 QUORUM funded by the Ministerio de Ciencia, Innovación y Universidades, through CDTI. We gratefully acknowledge ProtoQSAR SL and the team led by Eva Serrano-Candelas, Laureano E. Carpio, and Rafael Gozalbes for their fundamental contribution in providing, curating, and validating the dataset used in this study. For more information, refer to *Computational Modeling of DYRK1A Inhibitors as Potential Anti-Alzheimer Agents* [12].

## REFERENCES

- [1] C. Hansch and T. Fujita, “ $\rho$ - $\sigma$ - $\pi$  analysis. a method for the correlation of biological activity and chemical structure,” *Journal of the American Chemical Society*, vol. 86, no. 8, pp. 1616–1626, 1964. [Online]. Available: <https://doi.org/10.1021/ja01062a035>

- [2] R. Natarajan, G. S. Natarajan, and S. C. Basak, "Quantitative structure–activity relationship (QSAR) modeling of chiral CCR2 antagonists with a multidimensional space of novel chirality descriptors," *Molecules*, vol. 30, no. 2, p. 307, 2025. [Online]. Available: <https://doi.org/10.3390/molecules30020307>
- [3] V. Svetnik, A. Liaw, C. Tong, J. C. Culberson, R. P. Sheridan, and B. P. Feuston, "Random forest: A classification and regression tool for compound classification and qsar modeling," *Journal of Chemical Information and Computer Sciences*, vol. 43, no. 6, pp. 1947–1958, 2003.
- [4] A. Gaulton *et al.*, "The chembl database in 2023: a drug discovery platform spanning multiple bioactivity data types and time periods," *Nucleic Acids Research*, 2023.
- [5] M. Schuld and N. Killoran, "Is quantum advantage the right goal for quantum machine learning?" *PRX Quantum*, vol. 3, no. 3, p. 030101, 2022. [Online]. Available: <https://link.aps.org/doi/10.1103/PRXQuantum.3.030101>
- [6] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, no. 7747, pp. 209–212, 2019. [Online]. Available: <https://doi.org/10.1038/s41586-019-0980-2>
- [7] S. Miyabe, B. Quanz, N. Shimada, A. Mitra, T. Yamamoto, V. Rastunkov, D. Alevras, M. Metcalf, D. J. M. King, M. Mamouei, M. D. Jackson, M. Brown, P. Intallura, and J.-E. Park, "Quantum multiple kernel learning in financial classification tasks," 2023. [Online]. Available: <https://arxiv.org/abs/2312.00260>
- [8] G. V. W. Johnson and W. H. Stoothoff, "Tau phosphorylation in neuronal cell function and dysfunction," *Journal of Cell Science*, vol. 117, no. 24, pp. 5721–5729, 2004. [Online]. Available: <https://doi.org/10.1242/jcs.01558>
- [9] R. R. Neubig, "International union of pharmacology committee on receptor nomenclature and drug classification. xxxviii. update on terms and symbols in quantitative pharmacology," *Pharmacological Reviews*, vol. 55, no. 4, pp. 597–606, 2003. [Online]. Available: <https://doi.org/10.1124/pr.55.4.4>
- [10] L. E. Carpio, Y. Sanz, R. Gozalbes, and S. J. Barigye, "Computational strategies for the discovery of biological functions of health foods, nutraceuticals and cosmeceuticals: a review," *Molecular Diversity*, vol. 25, pp. 1425–1438, 2021. [Online]. Available: <https://doi.org/10.1007/S11030-021-10277-5>
- [11] A. Tropsha, "Best practices for qsar model development, validation, and exploitation," *Molecular Informatics*, vol. 29, pp. 476–488, 2010. [Online]. Available: <https://doi.org/10.1002/MINF.201000061>
- [12] E. Serrano-Candelas, L. E. Carpio, and R. Gozalbes, "Computational modeling of dyrk1a inhibitors as potential anti-alzheimer agents," in *Computational Modeling of Drugs Against Alzheimer's Disease*, ser. Neuromethods, K. Roy, Ed. New York, NY: Springer US, 2023, vol. 203, pp. 295–324. [Online]. Available: [https://doi.org/10.1007/978-1-0716-3311-3\\_10](https://doi.org/10.1007/978-1-0716-3311-3_10)
- [13] T. Suzuki, T. Hasebe, and T. Miyazaki, "Quantum support vector machines for classification and regression on a trapped-ion quantum computer," 2023. [Online]. Available: <https://arxiv.org/abs/2307.02091>
- [14] K.-C. Chen, X. Xu, H. Makhanov, H.-H. Chung, and C.-Y. Liu, "Quantum-enhanced support vector machine for large-scale stellar classification with GPU acceleration," 2023. [Online]. Available: <https://arxiv.org/abs/2311.12328>
- [15] S. H. Bertz, "The first general index of molecular complexity," *Journal of the American Chemical Society*, vol. 103, no. 12, pp. 3599–3601, 1981. [Online]. Available: <https://doi.org/10.1021/ja00402a071>
- [16] P. Rebentrost, M. Mohseni, and S. Lloyd, "Quantum support vector machine for big data classification," *Physical Review Letters*, vol. 113, no. 13, p. 130503, 2014. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.113.130503>
- [17] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, "Power of data in quantum machine learning," *Nature Communications*, vol. 12, no. 1, p. 2508, 2021. [Online]. Available: <https://doi.org/10.1038/s41467-021-22539-9>
- [18] V. Bergholm, J. Izaac, M. Schuld, C. Gogolin, C. Blank, K. McKiernan, N. Killoran *et al.*, "PennyLane: Automatic differentiation of hybrid quantum-classical computations," *arXiv preprint arXiv:1811.04968*, 2018. [Online]. Available: <https://arxiv.org/abs/1811.04968>
- [19] K. Tschärke, S. Issel, and P. Debus, "QUACK: Quantum aligned centroid kernel," *arXiv preprint arXiv:2405.00304*, 2024, arXiv:2405.00304.
- [20] S. Shalev-Shwartz, Y. Singer, and N. Srebro, "Pegasos: Primal estimated sub-gradient solver for SVM," in *Proceedings of the 24th International Conference on Machine Learning (ICML)*, 2007, pp. 807–814.