Quantum Feature Maps for Riemann Zeta Zeros: A Mathematically-Driven Approach

Gabriel D. Azevedo, Demerson N. Gonçalves and Tharso D. Fernandes

***Abstract*— This work explores whether quantum embeddings informed by the mathematics of the Riemann zeta function can provide more effective kernels for predicting deviations between zeros and Gram points. In contrast to generic approaches such as the *PauliFeatureMap*, we propose a feature map that incorporates the analytic structure of the problem, including phase, magnitude, and oscillatory components, as well as known functional couplings. The design also leverages domain-guided compression to remain feasible on near-term hardware. We outline its principles and an evaluation protocol comparing classical SVR, QSVR with *PauliFeatureMap*, and QSVR with the proposed construction using Odlyzko's datasets. Empirical validation is part of ongoing work and will be presented in future results.**

***Keywords*— Quantum feature maps, Riemann zeta function, QSVR, kernel methods, PauliFeatureMap**

I. Introduction

The Riemann zeta function $\zeta(s)$ occupies a central place in analytic number theory. The Riemann Hypothesis asserts that all nontrivial zeros lie on the critical line $\Re(s) = \frac{1}{2}$, a statement deeply linked to the distribution of prime numbers. High-precision studies of these zeros make use of structures such as the Riemann–Siegel theta function $\theta(t)$, Hardy's real-valued function $Z(t) = e^{i\theta(t)}\zeta(\frac{1}{2} + it)$, and Gram points g_n defined by $\theta(g_n) = (n-1)\pi$.

Recent work has explored the use of machine learning to support this line of investigation. The authors in [1] introduced a large pool of engineered features, including values of Z at Gram and related points, terms of the Riemann–Siegel formula, and lagged sequences, and showed that Support Vector Regression (SVR) could learn local patterns in the distribution of zeros. Ref. [2] took a first step into the quantum setting, applying Quantum SVR with Qiskit's *PauliFeatureMap* as a baseline quantum kernel. These studies indicate that data-driven approaches are viable, but they also reveal a limitation: existing quantum embeddings are generic and do not take advantage of the specific mathematical relationships among θ, Z, Gram sequences, and series contributions.

In this work we outline a different approach. Instead of relying on generic embeddings, we seek to design a quantum feature map that reflects the mathematical semantics of the problem itself. The idea is to encode phase-related quantities, oscillatory components, and known analytic couplings in a way that mirrors their mathematical roles, while also compressing high-dimensional classical features into a feasible number of qubits through domain-guided strategies. The expectation is that such a construction may lead to kernels more closely aligned with the structure of the data, potentially improving predictive accuracy compared to both classical SVR and quantum SVR with *PauliFeatureMap*. Our research is ongoing, and the present paper describes the motivation, design principles, and planned evaluation.

D. N. Gonçalves, Collegiate of Mathematics, CEFET-RJ, Petrópolis-RJ (demerson.goncalves@cefet-rj.br). G. D. Azevedo, Computer Engineering, CEFET-RJ, Petrópolis-RJ (gabriel.dias@aluno.cefet-rj.br). T. D. Fernandes, Pure and Applied Mathematics, UFES, Alegre-ES (tharso.fernandes@ufes.br).

II. Background and Feature Engineering

The definition of the extended Riemann zeta function is given by $\zeta(s) = 2^s\pi^{s-1}\sin\left(\frac{\pi s}{2}\right)\Gamma(1-s)\zeta(1-s)$, where Γ is the Gamma function, a factorial extension defined as $\Gamma(s) = \int_0^\infty x^{s-1}e^{-x}dx$. Thus, we can define the Hardy Z-function:

$$Z(t) = e^{i\theta(t)}\zeta\left(\tfrac{1}{2} + it\right), \tag{1}$$

with $\theta(t) = \arg\Gamma\left(\frac{1}{4} + \frac{it}{2}\right) - \frac{\ln\pi}{2}\,t$ and Gram points g_n defined by $\theta(g_n) = (n-1)\pi$. The target in the regression setting is

$$\delta_n = \gamma_n - g_{n-1}, \tag{2}$$

where γ_n is the ordinate of the n-th zero on the critical line.

The feature pool proposed in [1] includes: (i) *core analytic quantities* $(g_n, \theta(g_n), Z(g_n))$; (ii) *complementary sequences*, such as co-Gram points and Z evaluated at successive integers; (iii) *Riemann–Siegel contributions* (terms 2–10 of the truncated Riemann–Siegel formula evaluated at Gram points),

$$Z(t) \approx 2\sum_{n=1}^{N}\frac{\cos(\theta(t) - t\ln n)}{\sqrt{n}} + R(t), \quad N \approx \sqrt{\tfrac{t}{2\pi}}, \tag{3}$$

and (iv) *temporal context* via lagged versions of Gram sequences, Z-values, and distances. These quantities are motivated by the fact that sign changes of Z across Gram intervals often bracket zeros, while oscillatory terms and lags help capture local dynamics and spacing patterns relevant to δ_n.

A direct quantum encoding of all 40–94 engineered features would be infeasible on near-term devices. Our plan is therefore to adopt compression strategies within each mathematical group rather than globally. For example, the 19 terms of the Riemann–Siegel expansion could be summarized by a few principal components, while the core Gram/Zeta quantities $(g_n, \theta(g_n), Z(g_n))$ would remain explicit due to their direct phase–magnitude semantics. Features related to prime distributions or zero spacings may also be aggregated into a compact representation. This type of domain-guided compression is expected to preserve mathematical interpretability while reducing the effective qubit requirement to about 8–12, a range more suitable for NISQ hardware [3].

III. Design of the Quantum Feature Map

Our proposal organizes the encoding into three functional blocks, shown schematically in Fig. 1. Each block corresponds to a distinct group of mathematical features, with couplings drawn as arrows indicating how information flows across them.

(1) Gram–Zeta block (blue, left, 3 qubits). This block encodes the most fundamental analytic quantities: the Gram point g_n, the Riemann–Siegel phase $\theta(g_n)$, and the Hardy Z-function. Each variable is assigned to a different rotation axis according to its role: g_n to R_z (phase accumulation), $\theta(g_n)$ to R_y (angular displacement), and $Z(g_n)$ to R_x (oscillatory amplitude). The arrow from Gram–Zeta to the Riemann–Siegel block in Fig. 1 represents controlled-phase gates coupling g_n and $\theta(g_n)$, reflecting their definitional link in $Z(t) = e^{i\theta(t)}\zeta(1/2 + it)$.

(2) Riemann–Siegel block (green, top right, 5 qubits). The oscillatory terms of the truncated Riemann–Siegel expansion are first reduced to a handful of principal components. These compressed features are then encoded by R_y rotations, chained together with linear entanglement using controlled gates (CX, CZ). In Fig. 1, this block sits to the right of Gram–Zeta, receiving input through controlled gates that establish coherence between the fundamental quantities and the oscillatory structure.

(3) Primes/spacing block (orange, bottom right, 2 qubits). A compact block encodes proxies for prime density and zero spacing, believed to influence zero distribution. In the diagram, it appears below the Riemann–Siegel block. Solid arrows indicate weak couplings (CRX, CZ) linking it to the series, while the dashed arrow shows sparse cross-entanglement with Gram–Zeta, providing interaction without excessive circuit depth.

Overall, the map repeats this three-block structure over a small number of layers L. The design is heterogeneous and semantics-aware: each block corresponds to a mathematically distinct feature group, and the couplings in Fig. 1 represent known analytic relations. This contrasts with the generic *PauliFeatureMap*, which applies uniform rotations and entanglement without regard to the mathematical role of each variable.

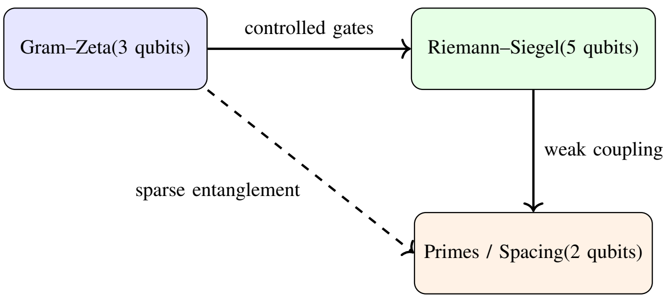

Fig. 1. **Proposed quantum feature map.** Each block encodes a mathematically distinct group of features, with couplings indicating known analytic relationships.

IV. Methodology for Evaluation

We will rely on publicly available datasets of Riemann zeta zeros such as Odlyzko's tables [4]. From these, we intend to compute the relevant analytic features, including Gram points, the Riemann–Siegel phase, Z-function values, truncated series contributions, and lagged sequences, with the prediction target defined as the deviation δ_n.

Our evaluation will compare three approaches: (i) classical SVR with RBF kernel; (ii) QSVR with *PauliFeatureMap*, serving as a generic quantum baseline; and (iii) QSVR with the proposed mathematically-driven map, which incorporates domain-guided compression to remain feasible on NISQ hardware.

Performance will be assessed through standard regression metrics such as MSE, RMSE, R^2, and through kernel alignment to examine how well the quantum kernel captures structural relations among zeros. We also plan controlled variations in the feature map design, for example disabling cross-block couplings, altering rotation assignments, or adjusting qubit counts, in order to understand which components contribute most to predictive power.

V. Expected Results

We anticipate that a mathematically-structured feature map can achieve better kernel alignment and lower prediction error than generic baselines, provided enough qubits are available to encode both phase–magnitude information and oscillatory components. By applying compression within coherent feature groups, the approach is expected to remain interpretable and relatively stable under small perturbations. At very small scales, however, we recognize that simpler maps such as *PauliFeatureMap* may perform comparably due to lower circuit depth and variance.

VI. Conclusion

We proposed a quantum embedding strategy informed by the mathematical structure of zero localization in the Riemann zeta function. Phase variables are associated with phase rotations, oscillatory or magnitude terms with amplitude rotations, and known couplings with controlled entanglement, while compression is applied group-wise to reduce dimensionality. This framework provides a principled alternative to generic maps such as *PauliFeatureMap*. A complete empirical evaluation, which addresses kernel quality, predictive accuracy, and the role of each design choice, remains for future work.

Acknowledgment

Gabriel D. Azevedo acknowledges the support of CNPq through an undergraduate research scholarship under the PIBIC program at CEFET/RJ.

Referências

[1] J. Kampe, A. Vysogorets. Predicting zeros of the Riemann Zeta function using machine learning. 2018. https://www.sci.sdsu.edu/main-reu/2018-2.pdf

[2] T. D. Fernandes, D. N. Gonçalves, J. T. Dias. Quantum Support Vector Regression for Predicting Zeros of the Riemann Zeta Function. WECIQ, 2024.

[3] M. Schuld, F. Petruccione. Supervised learning with quantum computers. Springer, 2018.

[4] A. M. Odlyzko. Tables of zeros of the Riemann zeta function. https://www-users.cse.umn.edu/~odlyzko/zeta_tables/

[5] V. Havlíček et al. Supervised learning with quantum-enhanced feature spaces. *Nature* 567(7747), 209–212, 2019.