PROCEEDINGS OF SCIENCE

A Quantum Adaptive Importance Sampling Algorithm for Multidimensional Integration

Konstantinos Pyretzidis[a,*]

[a] *Instituto de Física Corpuscular, Universitat de València – Consejo Superior de Investigaciones Científicas, Parc Científic, E-46980 Paterna, Valencia, Spain.*

E-mail: k.pyretzidis@ific.uv.es

This work presents Quantum Adaptive Importance Sampling (QAIS), a hybrid quantum-classical algorithmic workflow that performs Monte Carlo numerical integration of multivariate functions. This workflow extends standard Importance Sampling to a quantum computational system. Its primary purpose is to produce an unbiased estimate of integrals, with a limited number of measurements. By using the exponentially sized Hilbert space of a Parameterized Quantum Circuit (PQC), we manipulate the Probability Density Function (PDF) over a grid in its entirety. With adequate expressivity and entanglement, we effectively capture correlations among different integrand variables, thus bypassing the separable PDF assumption, as implemented in VEGAS. By optimizing the PQC, we adapt and ultimately load a PDF that approximates the integrand's behavior. Direct sampling from it, allocates samples to the important regions, allowing a highly accurate integral estimation. As an application, we study a benchmark two-dimensional integral, and compare the precision of the best grid of VEGAS against that of the quantum-generated proposal PDF, as a function of the number of samples.

The European Physical Society Conference on High Energy Physics (EPS-HEP2025)
7-11 July 2025
Marseille, France

*Speaker

PoS(EPS-HEP2025)638

1. Introduction

Perturbative Quantum Field Theory is central in performing accurate theoretical predictions of observables at high-energy colliders. Fundamental ingredients in this framework, such as loop Feynman diagrams and the phase-space, involve evaluating multidimensional integrals that are computationally intensive due to complex mathematical structures. To address this challenge and improve the precision of such theoretical predictions, the standard approach relies on Adaptive Importance Sampling. The most notable example is the VEGAS [1–3] algorithm, which adaptively updates a multidimensional grid. However, because the computational cost of handling each grid cell separately grows exponentially with the number of the integral's dimensions, the model for the Probability Density Function (PDF) is simplified to a separable product of PDFs, factorized along each integration variable's axis.

As classical integrators approach their practical limits, attention has shifted towards addressing integration through deep generative models, in particular normalizing flows [4–10], and quantum computing, mainly through the quadratic reduction in Monte Carlo query complexity [11, 12], with recent applications [13–21]. This work introduces Quantum Adaptive Importance Sampling (QAIS) [22] as a sampling-centric NISQ oriented quantum integrator for multi-dimensional integration. QAIS encodes a non-separable proposal PDF on a grid, encoded in the amplitudes of a Parametrized Quantum Circuit (PQC) and allocates samples via computational basis measurements. In contrast to VEGAS, which adapts a separable PDF, QAIS captures cross-dimensional correlations through entanglement and samples from a non-factorizable proposal PDF. The method is illustrated on a two-dimensional benchmark, highlighting the efficient sampling in the truly important regions.

2. Numerical Integration through Importance Sampling

Let $\Omega \subset \mathbb{R}^d$ be the integration domain and $f : \Omega \to \mathbb{R}$ the integrand. In plain Monte Carlo integration one draws N points x_i uniformly in Ω and approximates the integral by the corresponding sample average. A standard variance reduction technique is Importance Sampling (IS). By utilizing a proposal PDF that is efficient to sample from and approximates the shape of $f(x)$, IS rewrites the integral as $I = \int_\Omega f(x)\,dx = \int_\Omega \frac{f(x)}{q(x)}\,q(x)\,dx$ and uses samples $x_i \sim q$ to form

$$\hat{I}_N^{(\mathrm{IS})} = \frac{1}{N}\sum_{i=1}^{N}\frac{f(\mathbf{x}_i)}{q(\mathbf{x}_i)}, \qquad \left(\hat{\sigma}_N^{(\mathrm{IS})}\right)^2 = \frac{1}{N-1}\left(\frac{1}{N}\sum_{i=1}^{N}\frac{f(\mathbf{x}_i)^2}{q(\mathbf{x}_i)^2} - \left(\hat{I}_N^{(\mathrm{IS})}\right)^2\right). \tag{1}$$

VEGAS implements Adaptive IS by sampling uniformly in a unit hypercube $y \in [0,1]^d$ and then mapping the samples to the integration domain of interest $x(y)$ through the factorized Jacobian $J(y) = \prod_{k=1}^{d} J_k(y_k)$, that is adapted along different iterations. Through this separability, the algorithm effectively reduces the degrees of freedom from exponentially growing in the number of integration dimensions d to linear. The trade-off of this robustness is that the algorithm cannot encode cross-dimensional correlations, which is the root cause of phantom structures on multi-modal or other complicated targets.

PoS(EPS-HEP2025)638

3. Quantum Adaptive Importance Sampling

The full computational framework of QAIS, is presented in detail in [22]. In summary, it consists of three main parts. First, there is the Encoding, where each dimension i of the integration domain $\Omega = \prod_{l=1}^{d}[a_l, b_l]$ is discretized with q_l qubits, giving 2^{q_l} bins of width $\Delta_l = (b_l - a_l)/2^{q_l}$. With $n = \sum_l q_l$ qubits, a computational basis state $|i\rangle$ labels a d-dimensional cell $\Omega^{(i)}$. A PQC prepares

$$|\psi\rangle = \sum_i c_i |i\rangle, \qquad p_i = |c_i|^2, \tag{2}$$

so that the probabilities p_i define the cell heights of a piecewise-constant, non-separable proposal PDF on Ω.

Second is the State Preparation. The objective is to prepare a discretized state whose probabilities approximate the target integrand on the grid. In practice, this is achieved by training a Quantum Circuit Born Machine. We use an all-to-all entangling Ansatz with two-qubit interaction gates and single-qubit rotations. The number of parameters per layer scales as $O(n^2)$. The loss function used is the discretized Kullback–Leibler divergence between the target integrand and the PQC's probability mass function and the optimization is carried out with COBYLA on a noiseless state-vector simulator.

Third, there is the Integral Estimation. Given an optimized PQC, N measurements are performed and M distinct basis states are observed, corresponding to a subset $\Omega^- \subset \Omega$ of visited cells. For each visited cell $|i\rangle \in \Omega^-$ that has been measured N_i times, we draw N_i random continuous points $x_j^{(i)} \in \Omega^{(i)}$ and use these in an IS estimator weighted by the learned proposal. A well-trained proposal PDF concentrates probability on a small subset of all possible cells, so many cells remain unobserved at finite N. Simply discarding unobserved cell's contribution would induce a systematic bias in the estimate. QAIS addresses this missing mass problem by a classical tiling procedure. To correct this missing mass without extra calls to f, Ω^- is partitioned into Important cells Ω_I and a set made of Boundary and Noise cells, Ω_B and Ω_N. The Tiling algorithm then covers the non-Important region $\Omega_{N-I} = (\Omega \setminus \Omega^-) \cup \Omega_B \cup \Omega_N$ with a small collection of enlarged hyper-rectangles that are greedily constructed by the Boundary and Noise cells. The final QAIS estimator is written as:

$$\hat{I}_N^{(\mathrm{QAIS})} = \frac{1}{N}\sum_{i\in\Omega_I} w^{(i)} \sum_{j=1}^{N_i} f(\mathbf{x}_j^{(i)}) + \frac{1}{N}\sum_{k\in\Omega_{N-I}} w^{(k)} \sum_{j=1}^{N_k} f(\mathbf{x}_j^{(k)}), \qquad w^{(l)} = \frac{|\Omega^{(l)}|N}{N_l}. \tag{3}$$

4. Demonstration of the Method

For the demonstration of QAIS, a two-dimensional two-peak Gaussian benchmark with peaks along the diagonal is chosen. The target function is:

$$f(x) = \sum_{i=0}^{1} e^{-200\,\|x-r_i\|^2}, \qquad \begin{cases} r_0 = (0.23, 0.23) \\ r_1 = (0.74, 0.74) \end{cases}. \tag{4}$$

This function is integrated on $[0,1]^2$. The discretization used is five qubits per dimension ($n_{\mathrm{qubits}} = 10$). For the training, 5×10^3 COBYLA iterations have been used. The comparison focuses on the trained QAIS proposal PDF against the best grid of VEGAS integration accuracy. The results

PoS(EPS-HEP2025)638

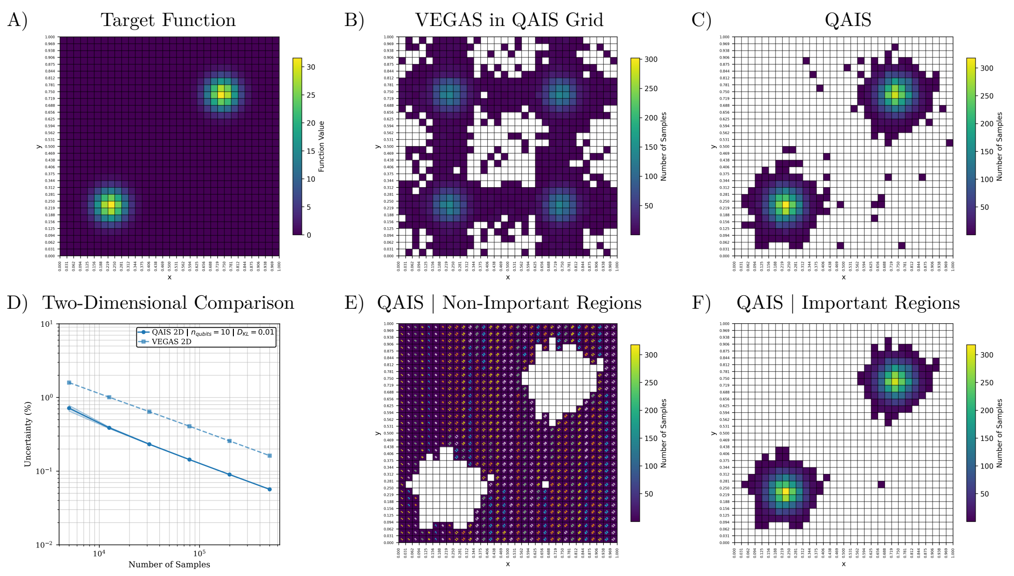

Figure 1: Two-dimensional two-peak Gaussian benchmark. (A) Target integrand. (B) VEGAS samples projected onto the QAIS grid ($N = 10^4$ samples). (C) QAIS sample allocation ($N = 10^4$ samples). (D) Relative uncertainty (σ/I) for different number of samples. The band represents one standard deviation from 100 integrations with the same proposal PDF. (E) Non-Important region with tiles. (F) Important region.

are shown in Fig. 1, together with the methods' explicit sample allocation in the two-dimensional domain as well as an explicit illustration of the Tiling debiasing. Projecting VEGAS onto the QAIS grid reveals the phantom peak regions that lay away from the true modes, reflecting the limitations of it's separable proposal on correlated multi-modal structure. In contrast, QAIS concentrates samples around the two peaks, while the tiling step reconstructs the coverage in the full domain. As a result, QAIS attains improved relative uncertainty at fixed sample budget in this benchmark.

5. Conclusions and Outlook

QAIS [22] is a hybrid quantum–classical algorithm designed to perform numerical integration of multivariate functions. It works by adapting a PQC whose Born probabilities drive a proposal PDF over a discretized multidimensional grid, allowing sampling from a non-separable model encoded directly in the circuit's amplitudes. Samples are generated by computational-basis measurements and are promoted to the continuum by drawing random points within the cell. A Tiling debiasing step corrects for unobserved cells at finite shots.

The main bottleneck is state preparation, since training becomes challenging as circuit size grows, consistent with known trainability limitations such as barren plateaus [23, 24]. Improving trainability and incorporating better initializations or priors are key directions toward higher-dimensional collider applications.

Acknowledgements: This work is supported by the Spanish Government and ERDF/EU - Agencia Estatal de Investigación (MCIU/AEI/10.13039/501100011033), Grants No. PID2020-114473GB-I00, No. PID2023-146220NB-I00, and No. CEX2023-001292-S; and Generalitat Valenciana, Grant No. ASFAE/2022/009 (Planes Complementarios de I+D+i, NextGenerationEU).

References

[1] G. Peter Lepage. A New Algorithm for Adaptive Multidimensional Integration. *J. Comput. Phys.*, 27:192, 1978.

[2] G. Peter Lepage. Adaptive multidimensional integration: VEGAS enhanced. *J. Comput. Phys.*, 439:110386, 2021.

[3] T. Hahn. CUBA: A Library for multidimensional numerical integration. *Comput. Phys. Commun.*, 168:78–95, 2005.

[4] Christina Gao, Joshua Isaacson, and Claudius Krause. i-flow: High-dimensional integration and sampling with normalizing flows. *Machine Learning: Science and Technology*, 1(4):045023, November 2020.

[5] Christina Gao, Stefan Höche, Joshua Isaacson, Claudius Krause, and Holger Schulz. Event generation with normalizing flows. *Physical Review D*, 101(7), April 2020.

[6] Enrico Bothmann, Timo Janßen, Max Knobbe, Tobias Schmale, and Steffen Schumann. Exploring phase space with Neural Importance Sampling. *SciPost Phys.*, 8(4):069, 2020.

[7] Ramon Winterhalder, Vitaly Magerya, Emilio Villa, Stephen P. Jones, Matthias Kerner, Anja Butter, Gudrun Heinrich, and Tilman Plehn. Targeting multi-loop integrals with neural networks. *SciPost Phys.*, 12(4):129, 2022.

[8] Theo Heimel, Ramon Winterhalder, Anja Butter, Joshua Isaacson, Claudius Krause, Fabio Maltoni, Olivier Mattelaer, and Tilman Plehn. Madnis - neural multi-channel importance sampling. *SciPost Physics*, 15(4), October 2023.

[9] Theo Heimel, Nathan Huetsch, Fabio Maltoni, Olivier Mattelaer, Tilman Plehn, and Ramon Winterhalder. The madnis reloaded. *SciPost Physics*, 17(1), July 2024.

[10] Theo Heimel, Olivier Mattelaer, Tilman Plehn, and Ramon Winterhalder. Differentiable madnis-lite. *SciPost Physics*, 18(1), January 2025.

[11] Gilles Brassard, Peter Høyer, Michele Mosca, and Alain Tapp. Quantum amplitude amplification and estimation, 2002.

[12] Ashley Montanaro. Quantum speedup of monte carlo methods. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 471(2181):20150301, September 2015.

[13] Gabriele Agliardi, Michele Grossi, Mathieu Pellen, and Enrico Prati. Quantum integration of elementary particle processes. *Phys. Lett. B*, 832:137228, 2022.

[14] Ifan Williams and Mathieu Pellen. A general approach to quantum integration of cross sections in high-energy physics. 2 2025.

[15] Jorge J. Martínez de Lejarza, David F. Rentería-Estrada, Michele Grossi, and Germán Rodrigo. Quantum integration of decay rates at second order in perturbation theory. *Quantum Sci. Technol.*, 10(2):025026, 2025.

[16] Jorge J. Martínez de Lejarza, Leandro Cieri, Michele Grossi, Sofia Vallecorsa, and Germán Rodrigo. Loop Feynman integration on a quantum computer. *Phys. Rev. D*, 110(7):074031, 2024.

[17] Steven Herbert. Quantum monte carlo integration: The full advantage in minimal circuit depth. *Quantum*, 6:823, September 2022.

[18] Juan M. Cruz-Martinez, Matteo Robbiati, and Stefano Carrazza. Multi-variable integration with a variational quantum circuit. *Quantum Sci. Technol.*, 9(3):035053, 2024.

[19] Jorge J. Martínez de Lejarza, Michele Grossi, Leandro Cieri, and Germán Rodrigo. Quantum Fourier Iterative Amplitude Estimation. In *2023 International Conference on Quantum Computing and Engineering*. IEEE, 5 2023.

[20] Ismail Yunus Akhalwaya, Adam Connolly, Roland Guichard, Steven Herbert, Cahit Kargi, Alexandre Krajenbrink, Michael Lubasch, Conor Mc Keever, Julien Sorci, Michael Spranger, and Ifan Williams. A Modular Engine for Quantum Monte Carlo Integration. 8 2023.

[21] Kirill Plekhanov, Matthias Rosenkranz, Mattia Fiorentini, and Michael Lubasch. Variational quantum amplitude estimation. *Quantum*, 6:670, March 2022.

[22] Konstantinos Pyretzidis, Jorge J. Martínez de Lejarza, and Germán Rodrigo. Unlocking Multi-Dimensional Integration with Quantum Adaptive Importance Sampling. 6 2025.

[23] Manuel S. Rudolph, Sacha Lerch, Supanut Thanasilp, Oriel Kiss, Oxana Shaya, Sofia Vallecorsa, Michele Grossi, and Zoë Holmes. Trainability barriers and opportunities in quantum generative modeling. *npj Quantum Inf.*, 10(1):116, 2024.

[24] Martín Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo. Barren plateaus in variational quantum computing. *Nature Reviews Physics*, 7(4):174–189, March 2025.