


Efficient quantum interpolation of natural dataSergi Ramos-Calderer ^{*}*Quantum Research Centre, Technology Innovation Institute, Abu Dhabi, UAE
and Departament de Física Quàntica i Astrofísica and Institut de Ciències del Cosmos (ICCUB), Universitat de Barcelona,
Martí i Franquès 1, 08028 Barcelona, Spain.*

(Received 19 July 2022; accepted 5 December 2022; published 21 December 2022)

We present efficient methods to interpolate data with a quantum computer to complement uploading techniques and quantum postprocessing. The quantum algorithms make use of the efficient quantum Fourier transform (QFT) as well as classical signal and imaging processing techniques, and may facilitate quantum advantages for some relevant families of data. We showcase a QFT interpolation method, a quantum cosine transform (QCT) interpolation geared towards natural data, and we improve upon them by utilizing a quantum circuit's capabilities of processing data in superposition. A circuit for the QCT is presented. We illustrate the methods on probability distributions and quantum encoded images, and discuss the precision of the resulting interpolations.

DOI: [10.1103/PhysRevA.106.062427](https://doi.org/10.1103/PhysRevA.106.062427)**I. INTRODUCTION**

Uploading a large amount of classical data into a quantum state remains a bottleneck for quantum computation applications. Quantum states indeed offer a large Hilbert space to encode classical data, but uploading one element at a time makes a quantum approach inefficient, no matter its promises regarding processing power on the uploaded state. New strategies are needed to overcome the initial threshold of uploading data into a quantum state.

To this end, we present different methods that interpolate smooth probability distributions and natural data over a larger space that can alleviate the data uploading effort, dramatically in some cases. This method build on classical resampling techniques that employ the Fourier transform to interpolate band-limited signals. The original amplitude encoded distribution is first Fourier transformed, an ancillary register is added to the system to zero-pad higher-frequency components, and finally, an inverse Fourier transform over the larger space delivers the interpolated probability distribution. To be precise, the relevant fact that makes this approach useful when interpolating classical distributions in a quantum computer is that the quantum Fourier transform (QFT) is efficient, that is, it only needs a polynomial number of operations as a function of the number of qubits involved.

The QFT interpolation scheme was adapted with success in the context of tensor networks in Ref. [1], and introduced to enhance the result of solutions of partial differential equations in a quantum computer in Ref. [2], where the accuracy of the interpolation is studied. We expand upon this concept with more diverse efficient quantum transformations and further utilize superposition to enhance the interpolation algorithms.

These interpolation methods can also be understood as a discrete-to-discrete instance of the Nyquist-Shannon sam-

pling theorem. The sampling theorem states that all the information of a signal with finite bandwidth can be captured by samples obtained at a finite rate, known as a Nyquist rate. We incorporate these ideas when discussing the accuracy of the proposed resampling techniques. The use of Fourier-like transforms is ubiquitous when reconstructing continuous signals from samples. The QFT can introduce efficient counterparts to techniques used in classical signal processing. We present some examples in the context of probability distribution interpolation and image resampling.

Classical signal processing of natural data has been drawn towards the discrete cosine transforms (DCT) [3] over the QFT, since the real-to-real mapping yields better results. This transformation is also the basis of the widely used JPEG image compression scheme [4]. We will incorporate these ideas into the quantum regime in order to improve upon QFT interpolation for natural data.

In Sec. II we restate the QFT interpolation scheme for probability distributions and study the accuracy of the method when specifically applied to quantum states encoding probability distributions. In Sec. III we detail a quantum DCT circuit and utilize it in a quantum cosine transform (QCT) interpolation method for natural images. We discuss the power of processing data in superposition that quantum computers allow. In Sec. IV we further exploit quantum superposition for a more efficient JPEG-inspired quantum interpolation procedure. We conclude with a discussion on the practical applicability of the method and future directions.

II. EFFICIENT QFT INTERPOLATION FOR SMOOTH DISTRIBUTIONS

We showcase here the basic algorithm for interpolating a quantum-uploaded distribution to higher precision. The algorithm is detailed, then discussed in the context of enhancing quantum data uploading techniques. Afterwards, we study the accuracy of the interpolated distribution.

^{*}sergi.ramos@tii.ae




FIG. 1. Full process of interpolating a probability distribution P_n using the efficient QFT. (a) Example of a Gaussian distribution P_4 encoded in four qubits. (b) Quantum circuit that performs QFT interpolation. Here, U denotes a unitary transformation that uploads the probability distribution P_n into the q quantum register. After an initial QFT of register q , a clean ancilla register a is swapped to the first m positions after q_0 , an operation that can be done virtually by keeping track of qubit position. After applying a CNOT gate to all ancilla qubits controlled by q_0 , an inverse QFT is applied to the whole system in order to recover the interpolated probability distribution P_{n+m} in the larger space. (c) Interpolation of the initial probability distribution P_4 using $n = 4$ qubits (blue squares) and the final result P_7 over the full $n + m = 7$ qubit register (orange dots), normalized to match P_4

Given a quantum register q , with individual qubits q_i from 0 to $n - 1$, where a distribution P has been discretized and encoded into the amplitudes of the 2^n basis states of q , QFT interpolation proceeds as follows. Apply the QFT to the quantum register q . In Fourier space, the high-frequency modes of smooth distributions will be suppressed, or zero in the case of band-limited signals. Therefore, one can artificially pad the high-frequency components with quantum states at zero amplitude and not fundamentally alter the original signal. This can be achieved on a quantum circuit by adding an ancillary qubit register a , with qubits a_j from 0 to $m - 1$, between original qubits q_0 and q_1 (the first and second most significant qubits) and then applying CNOT gates, controlled by q_0 , targeting all qubits in the ancilla register a . Finally, an inverse QFT is applied to the entire quantum system.

The outcome of this circuit is the interpolation from the initial 2^n distribution to a larger 2^{n+m} space. Figures 1(a) and 1(c) showcase the interpolation of the original four-qubit distribution to a larger seven-qubit space, implemented using the quantum simulation library QIBO [5,6]. In Fig. 1(b) we present the QFT interpolation algorithm on a quantum circuit.

The QFT interpolation algorithm for probability distributions is efficient. That being said, the cost of uploading the initial distribution into the small space needs to be taken into account. There are multiple proposals for uploading a probability distribution into a quantum register, some exact [7–12] and some training a quantum generator circuit [13–15]. If the cost of the initial uploading is already prohibitive, this algorithm will only provide a marginal advantage. Still, this technique opens the door for uploading methods that are effective for a small number of qubits that can later be enhanced via QFT interpolation.

We showcase one such case in Appendix A, where we use the QFT interpolation algorithm to enhance a probability distribution uploaded in the unary basis [16]. This uploading technique trades practical scalability for a device-friendly uploading method that is effective for a small number of amplitudes. After the original distribution is uploaded in unary using 2^n qubits, it is transformed into binary to n qubits,

where the full available 2^{2^n} Hilbert space is reclaimed via QFT interpolation.

In the following, we aim to establish the accuracy of the interpolated distribution when compared to an ideal uploading of the underlying distribution to the larger space. We study the effect of this procedure when applied to the amplitudes of a quantum state.

We will bound the operational distinguishability between the interpolated quantum state and the ideal one by introducing their *trace distance*,

$$\text{dist}_{\text{Tr}}(|\psi\rangle, |\phi\rangle) = \frac{1}{2} \|\psi\rangle\langle\psi| - |\phi\rangle\langle\phi|\|_1, \quad (1)$$

where $\|A\|_1 = \text{Tr} \sqrt{A^\dagger A}$. However, we first proceed by analyzing the ℓ_2 norm difference between them. The reason for this choice is twofold. Quantum states are normalized under their ℓ_2 norm, $\|\psi\rangle\|_{\ell_2} = 1$, where $|\psi\rangle$ is any pure quantum state, and the QFT preserves this norm, that is, $\|\psi\rangle\|_{\ell_2} = \|\Psi\rangle\|_{\ell_2} = 1$, with $|\Psi\rangle = \text{QFT}|\psi\rangle$ and $|\Phi\rangle = \text{QFT}|\phi\rangle$. From their ℓ_2 distance measure we can bound the trace distance between the two states by employing the Fuchs–van-de-Graaf inequality [17], which is tight for pure states [18], when we restrict ourselves to states that amplitude-encode probability distributions, see Appendix B. From here on out, all norms are assumed to be ℓ_2 unless stated otherwise.

We aim to upload a target quantum state $|\psi\rangle$ into $n + m$ qubits by using QFT interpolation on state $|\tilde{\psi}\rangle$, originally uploaded to n qubits, which will be referred to as an *n-qubit band-limited* state. An *n-qubit band-limited* state refers to a quantum state with any number of qubits that only has 2^n nonzero Fourier components. More precisely, a state with initial distribution P_n discretized over a span Δx_n , will only have nonzero Fourier components within the band $2^n / \Delta x_n$.

Analogous to the Nyquist-Shannon sampling theorem [19,20], if the target state $|\psi\rangle$ is *n-qubit band limited*, the QFT interpolation technique is able to capture all the information of the underlying distribution, and the interpolation is perfect. However, that will not always be the case. Then, the best possible initial state $|\psi\rangle$ to interpolate from will be the one that minimizes $\|\psi\rangle - |\tilde{\psi}\rangle\| = \|\Psi\rangle - |\tilde{\Psi}\rangle\|$. Classically, the

best band-limited approximation is realized by a convolution of the original signal with the sinc function [21], analogous to applying a low-pass filter to the signal, so that the high frequencies are cut off while the band-limited frequencies remain intact. This, however, is not as straightforward with quantum states, as they have to maintain their ℓ_2 norm throughout the process, and a low-pass filter is not a unitary operation. Therefore, the optimal distance between states will be

$$\begin{aligned} \|\Psi\rangle - |\tilde{\Psi}\rangle\|^2 &= \frac{2\|\Psi_{\text{out}}\|^2}{1 + \sqrt{1 - \|\Psi_{\text{out}}\|^2}}, \\ &\leq 2\|\Psi_{\text{out}}\|^2, \end{aligned} \quad (2)$$

where $|\Psi_{\text{out}}\rangle$ denotes a quantum state that contains only the Fourier components of target state $|\psi\rangle$ outside of the n -qubit band limit (see Appendix C for the detailed derivation). More explicitly, we can define both $|\Psi_{\text{out}}\rangle$ and $|\Psi_{\text{in}}\rangle$ as

$$|\Psi_{\text{out}}\rangle = \sum_{k \notin \text{band}} a_k |k\rangle \quad \text{and} \quad |\Psi_{\text{in}}\rangle = \sum_{k \in \text{band}} a_k |k\rangle, \quad (3)$$

where, up to normalization factors, a_k are the Fourier components of the real signal $|\Psi\rangle$, and *band* denotes the band limit of the uploaded quantum state $|\tilde{\Psi}\rangle$. See that we can rebuild $|\Psi\rangle$ by adding together $|\Psi_{\text{out}}\rangle$ and $|\Psi_{\text{in}}\rangle$, accounting for normalization factors.

The distance between the target state and the best possible interpolation will depend on the ℓ_2 norm of the Fourier components that are filtered out. Using this result, we can bound the trace distance between the target and interpolated state by

$$\begin{aligned} \text{dist}_{\text{Tr}}(|\psi\rangle, |\tilde{\psi}\rangle) &\leq \sqrt{2}\|\Psi_{\text{out}}\| \sqrt{1 - \frac{\|\Psi_{\text{out}}\|^2}{2}}, \\ &\leq \sqrt{2}\|\Psi_{\text{out}}\|. \end{aligned} \quad (4)$$

That being said, many applications do not have access to the filtered distribution, as only the low-resolution values of the original distribution are available. The interpolated state will then suffer from aliasing effects, where the Fourier components of subsampled distributions are mixed with its own high-frequency modes due to the periodic nature of the discrete Fourier transform. In other words, the frequency components outside the n -qubit band limit $|\Psi_{\text{out}}\rangle$ are added to the low-frequency components $|\Psi_{\text{in}}\rangle$, creating artifacts in the interpolated distribution.

The accuracy of the QFT interpolation algorithm will be worse in this approach, but we can still provide analytical bounds on the distinguishability due to the effects of aliasing. Now the distance between the target and interpolated state will be

$$\begin{aligned} \|\Psi\rangle - |\tilde{\Psi}\rangle\|^2 &= \frac{2\|\Psi_{\text{out}}\|^2}{N} - \frac{(N-1)^2}{N} \\ &\leq 2\|\Psi_{\text{out}}\|^2, \end{aligned} \quad (5)$$

where $N \geq 1$ is the normalization factor of the target state $|\tilde{\Psi}\rangle$ due to the effects of aliasing (refer to Appendix D for details). Therefore, the upper bound on the trace distance under aliasing effects remains the same as in Eq. (4).

We have shown how the QFT interpolation algorithm approaches the target quantum state for a probability distribution. The smaller the norm of the high-frequency components,

the better the interpolation will be. When the distribution is n -qubit band limited, that is, $\|\Psi_{\text{out}}\| = 0$, it captures all the information of the underlying function and can be interpolated with as many qubits as needed. Even for non-band-limited distributions, we show in Fig. 1 that with an initial uploading of four qubits, a Gaussian distribution can be interpolated with high fidelity to an exponentially larger space.

This technique can then be generalized and improved with more complex transformations for probability distributions over more dimensions. In order to expand on the effect of interpolation for multidimensional data, we will display interpolation of images, as it provides a useful tool set and clear visual representation that general distributions lack. However, everything can be extended to distributions over multidimensional data, which can yield further benefits for quantum algorithms that require them.

III. EFFICIENT QCT RESAMPLING OF NATURAL IMAGES

Interpolation methods, also referred to as resampling, are very common techniques in image processing. Moreover, natural images tend to have suppressed high-frequency components, and in particular, algorithms designed to process natural images tend to utilize the DCT [3]. This real-to-real variant of the Fourier transform is particularly suited for this type of signal, so much so that some of the most prominent image processing techniques, like the image compression scheme JPEG [4], use this transformation as the basis of the algorithm. The field of quantum image processing has experienced a lot of progress during the recent years [22–26] due to the efficiency of such transformations in quantum computing. Therefore, resampling quantum encoded images using a QCT interpolation algorithm seems a natural step forward to generalize the technique to more dimensions.

Encoding images into a quantum state is a costly task. Most algorithms that attempt to upload a classical image into a quantum state [23,24] inevitably scale with the pixel count of the image. These encoding techniques quickly become prohibitive for large images needing to be uploaded. The proposed algorithm can alleviate part of this complexity by applying QCT interpolation to resample an image in an efficient way after an initial uploading, showcased using two examples in Fig. 2. This method also supports other uploading methods that rely on machine learning or quantum sensing techniques. This interpolation technique, however, cannot enhance the data retrieval process; therefore, quantum imaging techniques that require the pixel data might still be unfeasible regardless of improvements in data encoding.

An efficient quantum implementation of the DCT is needed. In particular, we are interested in the type-II DCT, defined as

$$\text{DCT-II} : X_k = \alpha_k \sum_{i=0}^{I-1} x_i \cos\left(\frac{(2i+1)k\pi}{2I}\right), \quad (6)$$

where α_k are normalization constants for $k = 0, \dots, I-1$.

There have been some works addressing quantum circuits that implement the DCT: Ref. [27] proposes an efficient circuit that computes both the DCT and the discrete sine




FIG. 2. Example of the QCT resampling technique on 512 X 512 pixel input images. (Left) Grayscale image enlarged four times using the two-dimensional QCT resampling technique. (Right) RGB image enlarged four times using the two-dimensional QCT resampling technique; ancillary qubits are used to label each color layer. Both quantum interpolation circuits require the same number of gates regardless of the layers of the image.

transform (DST) in parallel, and Ref. [28] achieves a quadratic speedup over the classical DCT for data encoded into the quantum register instead of its amplitudes. However, none are suitable for this implementation. Therefore, we need introduce a novel method of implementing the type-II DCT efficiently on a quantum computer. There is a direct mapping between the DCT-II and a classical Fourier transform on a larger, symmetric space. Specifically, one needs to upload the original (real) signal x_i into a $4I$ space in the following way. The signal x_i is uploaded in the odd-indexed elements of the state in the first $2I$ entries. The inputs from $2I$ to $4I$ will mirror the first half. This can be mapped to direct operations on a quantum computer. Two ancilla qubits are added to the system, one as the most significant qubit and the other as the least significant one. A Hadamard gate on the first ancilla, followed by CNOT gates from that qubit to the rest of the signal-encoding qubits, will copy the initial signal in a symmetric way over a $2I$ space. The second ancilla qubit only requires a single X gate in order to encode the symmetric signal on the odd numbered elements of the quantum state. A QFT applied on the $n + 2$ qubit system will yield the desired transformation, up to an overall scale factor. This transformation needs only to be inverted before the data is recovered. A circuit implementing the QCT is illustrated in Fig. 3.

The QCT interpolation algorithm, understood as substituting the QFT by the QCT in the scheme presented in Sec. II, can be extended to two, and up to any, dimensions. The image data that we want to process requires the following amplitude encoding. For two dimensions, values labeled by $\{x, y\}$ will be encoded into the amplitude of the quantum state $|x\rangle|y\rangle$ in the computational basis of registers q_x and q_y . If this encoding is realized, applying the QCT circuit to both q_x and q_y registers separately (and simultaneously) achieves the two-dimensional discrete cosine transform over the data.

A. Grayscale images

The quantum resource needed to encode and interpolate a grayscale image with $2^n \times 2^n$ pixels using the proposed technique is $2n$ qubits, plus constant ancilla depending on the uploading method, $2m$ ancillary qubits used to enlarge the image to $2^{n+m} \times 2^{n+m}$ pixels, and two ancilla qubits for each dimension the QCT needs to be applied to. An example of an interpolated image using this method is displayed in Fig. 2 (left). A 512×512 image (top left) has been enlarged four times in both axes. This implementation requires $2 \times 9 + 2 \times 2 + 4 = 26$ qubits, omitting encoding ancillas, to transform the original image to the enlarged space, and has been simulated using the quantum simulation library QIBO using images available in the image processing library SCIKIT-IMAGE [29].

B. Multilayer images

We highlight a genuinely quantum speedup that arises when working with multilayer data. By properly encoding the




FIG. 3. Circuit example that implements a QCT on the quantum state encoded in register $q = q_0, \dots, q_{n-1}$ using two ancilla qubits.

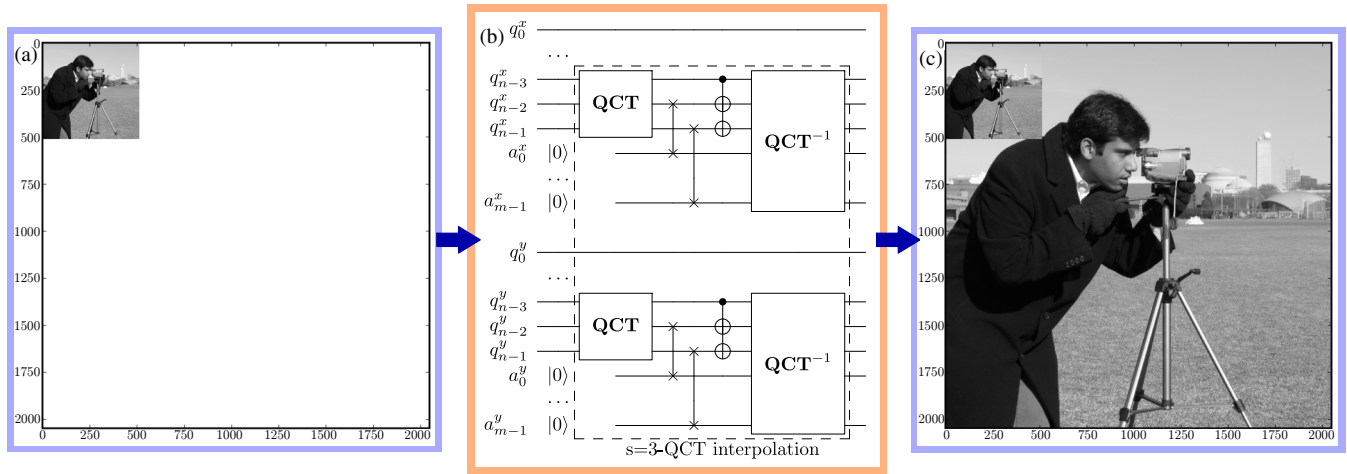


FIG. 4. Interpolation of a grayscale 512×512 pixel image using the JPEG-inspired $s = 3$ -QCT method. (a) Original image embedded in the larger subspace. (b) Quantum circuit that performs $s = 3$ -QCT interpolation for two-dimensional data. The decomposition of the QCT blocks are shown in Fig. 3, and the two ancilla needed are implied. The interpolation algorithm for two dimensions can be applied simultaneously, therefore not incurring extra depth cost. The algorithm is independent of the original system size, as it only acts on the subspace s and the added ancilla m . (c) Result of the interpolated image using $m = 2$ ancilla for each of the two dimensions. The original image is shown for scale in the bigger 2048×2048 pixel space.

multiple layers, we can process all the quantum data with a single call to the QCT interpolation algorithm. Starting from the same encoding technique used for a single image, each layer l_i of the image will be labeled by state $|l_i\rangle$ of a new label quantum register q_l , with $\lceil \log(l) \rceil$ qubits, where l is the total number of layers. That is, the value of pixel $\{x, y\}$ of layer l_i will be encoded in the amplitude of the quantum state $|x\rangle |y\rangle |l_i\rangle$. This will require a more complex encoding of the image, but it will become useful in the following steps.

In this encoding, applying the QCT interpolation algorithm in the same way as one would for grayscale images, acting on quantum registers q_x and q_y only, will perform the interpolation to all layers of the image in superposition at no extra quantum cost. Since every pixel value is now entangled with its label, the amplitude interference that makes the QFT possible (and efficient) will only act on the amplitudes of pixel states that share the same label state. We showcase an example in Fig. 2 (right), where we interpolate an RGB image using two qubits to keep track of the color channel. Since the pixel dimensions are the same as Fig. 2 (left), the gate count of the quantum circuit required for interpolation is the same for both instances.

The speedup provided by acting on all layers in superposition is made more apparent the more layers, or data instances sharing the same shape, that are encoded in the proposed way. Additionally, this kind of speedup can be enhanced when implementing quantum transformations to all subsets within an image by acting only once in a smaller set of qubits. Beyond that, this can be extended to other types of data uploaded in superposition. If the initial cost of uploading data using a label register is feasible, any quantum processing using a transformation such as the QFT or QCT can be applied to all members of the superposition in singular cost. However, we need to keep in mind that in the case where the data in superposition needs to then be extracted via measurements, the advantage that the parallel processing power introduces will be lost.

IV. SUBSPACE-QCT RESAMPLING OF NATURAL IMAGES

The power of parallel QFT computation can be extended beyond actions on different sets of data encoded in superposition. Subsets of a single data entry can also be processed in parallel.

Correlations present in natural data tend to be localized within short distances. Indeed, natural images are usually built from large structures (in number of pixels) that do not necessarily correlate with the rest of the picture. The JPEG protocol for image compression [4] exploits this fact by performing DCT onto 8×8 blocks of the image. This way, the small DCT only capture short range correlations, and the following steps of the protocol can be described using 8×8 matrices.

A similar procedure can be applied in a quantum circuit. All the 8×8 pixel matrices of the image encoded in the proposed way are stored in the superposition of the three least significant qubits. The rest of the system can be understood in the same way we considered the label register for multilayer images. Applying a transformation to only the three least significant qubits, a QCT in this case, achieves in constant gate count the transformation of all subspaces of the encoded data.

To be precise, this JPEG-inspired interpolation algorithm using the DCT proceeds as follows. Given an image encoded into a quantum state in registers $|x\rangle |y\rangle |l\rangle$, perform the QCT on only the last three qubits of registers $|x\rangle$, and simultaneously in only the last three qubits of register $|y\rangle$. Introduce the ancilla qubits needed for the interpolation after the third least significant qubit on both registers. Undo the QCT transformation on the extended space. This achieves image interpolation at a gate count independent of the original system size. Given a fixed subspace 2^s , $s = 3$ in the JPEG-inspired procedure, this algorithm resamples an 2^n signal into a 2^{n+m} space with a gate count of $\mathcal{O}[(s+m)^2]$, it no longer depends on the size of the original image, only on the subspace of interest. A circuit depicting this algorithm is sketched in Fig. 4(b).

TABLE I. Comparison of the peak signal-to-noise ratio (PSNR) and structural SIMilarity (SSIM) of the grayscale *camera* image, after different methods of interpolation. The image is downsized to half its original size ($m = 1$) via a classical algorithm using the pixel area relation of the image implemented in `OPENCV`. The image is interpolated to the original size and compared to the base image. The classical bicubic interpolation scheme yield the best PSNR and SSIM values, closely followed by the $s = 3$ -QCT. From the quantum algorithms, the best choice is the JPEG-inspired $s = 3$ -QCT method, as it provides the best results and is the most efficient option.

$m=1$	Bicubic	QFT	n -QCT	3-QCT
PSNR	30.095	27.395	29.930	29.988
SSIM	0.880	0.829	0.871	0.878

This algorithm can be generalized to any subspace s , and can introduce improvements depending on the underlying structure of the signal. In particular, an s -QCT with $s = n$ is equivalent to the full QCT approach presented in the previous section.

We compare the performance of the presented algorithms for image interpolation, as the field contains useful metrics for this type of comparison. A trial image is first downscaled using the recommended classical algorithm as provided by the `OPENCV` library [30]. Then, the downscaled image is interpolated back into its original size, and compared to the base image. Two metrics are commonly used in order to assess the accuracy of image interpolation [31]: the peak signal-to-noise ratio (PSNR) and structural SIMilarity (SSIM) [32], detailed in Appendix E. The interpolation is performed using four algorithms: classical bicubic interpolation, QFT interpolation, QCT interpolation, and $s = 3$ -QCT interpolation.

In Table I we compare the two metrics for the different interpolation methods. Bicubic interpolation is one of the go-to methods for classical image interpolation; it uses the adjacent 4×4 pixels to compute each new value. This method achieves the best results by a slight margin when compared to the better of the quantum algorithms presented. In particular, the interpolation incurs a computational cost of $\mathcal{O}(N_p^2)$ [33], where N_p is the number of pixels in the image, exponentially more expensive than the presented quantum algorithms for interpolation. Alternatively, the best performing quantum interpolation method is the $s = 3$ -QCT, which is also the least expensive in terms of computational complexity. By focusing on small subsets of the image, the algorithm is both faster, due to processing data in superposition, and more accurate in reproducing the original data.

By using techniques of classical image processing we have enhanced the quantum interpolation algorithm, in terms of both accuracy of reconstruction and complexity. The ideas presented in the previous sections about computing in superposition have been extended to small subsets of the same image. By fixing the subspace where the QCT acts on three qubits, the interpolation algorithm no longer scales with the size of the original image, allowing for a very efficient circuit for interpolation of natural data in a quantum computer. This, however, only pertains to the interpolation part of the algorithm, as mentioned before; if the uploading or further

processing is prohibitive, this will only marginally aid in the overall cost of the task.

V. CONCLUSION

We have proposed efficient quantum techniques to interpolate distributions, all exploiting the quantum Fourier transform. We show that this QFT interpolation technique achieves favorable results on data with negligible high-frequency components, as well as in practical examples using natural images as a first visual example of multidimensional distributions.

This technique can be used to enhance current uploading algorithms by focusing on smaller scale and accurate uploading techniques that can then be efficiently resampled via interpolation algorithms. It can also be extended to any uploading technique that deals with nonclassical band-limited quantum signals. Additionally, this interpolation technique can be used on genuinely quantum states, but the interpolation will be extended along computational basis states which might fail to capture correlations that go beyond that.

We have also showcased the power of processing quantum data in superposition. Encoding different sets of data using label ancillas allows for a single implementation of a quantum transformation to act on all the encoded data in parallel. Furthermore, by extending native ideas of natural image processing, used in the JPEG procedure for image compression, we can employ this parallel processing power of a quantum circuit for a substantial gain. By processing subspaces of an image simultaneously using quantum superposition, we can perform interpolation using a number of gates that are constant with the original size of the image.

We would like to further highlight the implementation of efficient quantum transformations in order to unlock quantum advantages to other algorithms that might not initially achieve them due to data uploading costs. After all, at the core of the efficiency of Shor's factorization algorithm [34] is the use of the QFT. Looking into areas where these types of transformations are extensively used may result in further avenues where quantum algorithms can be enhanced. Natural image and signal processing are fields that have evolved around such transformations, and further work might borrow from such well understood fields in order to explore interpolation techniques with even more advanced transformations. Compression techniques that rely on these transformations might also prove useful.

The code used to simulate the quantum circuits presented is available online [35].

ACKNOWLEDGMENTS

The author would like to acknowledge J. I. Latorre, I. Roth, and N. Aburaed for fruitful discussions.

APPENDIX A: ENHANCING A UNARY UPLOADING

We illustrate how the QFT resampling algorithm can enhance uploading techniques using the amplitude distribution in the unary basis presented in Ref. [16]. This proposal requires only nearest-neighbor connectivity and employs gates




FIG. 5. Circuit example that encodes an 8 – to – 3 basis change from the unary basis to the binary one. The u register depicts the unary register where u_0 is the most significant qubit in unary representation. On the left-hand side the labels for the binary qubits n and the ancilla qubits m are depicted. The final SWAP gates can be neglected if we keep track of the qubit order and perform classical reshuffling.

that are well suited for implementation on near-term quantum devices. The distribution is uploaded in the unary basis where only one qubit is in state $|1\rangle$ while the others remain at $|0\rangle$, i.e., $|10000\rangle, |01000\rangle, \dots, |00001\rangle$. This, however, reduces the Hilbert space available for computation and requires linear depth with the number of amplitudes needed to upload, limiting its usability. While this basis helps with gate application and control of the device, one would want to exploit the exponential Hilbert space that qubits support.

In order to reclaim the lost Hilbert space using QFT interpolation, we propose the following. Upload a probability distribution into the unary basis over a 2^n qubit register. Then, perform a unitary change of basis from unary to binary basis [36], as detailed in Algorithm 1 and illustrated with a small example in Fig. 5 for $n = 3$. After the change of basis, the quantum state contains n qubits storing the superposition and $2^n - n$ clean ancillas at state $|0\rangle$. This allows the implementation of the QFT interpolation algorithm using the clean ancilla register to encode the high-frequency components. Now the full extent of the available Hilbert space is used to encode the interpolated probability distribution. Shown in Fig. 6 are the simulation results of a 16-qubit total QFT interpolation with unary uploading.

Algorithm 1.Unary to binary encoding.

```

1 Unary2Binary ( $n$ )
2  $c \leftarrow \text{Circuit}(2^n)$ 
3 Ensure  $q = 0$ 
4 for  $i \leftarrow 0$  to  $n - 1$  do
5    $qq \leftarrow 2^{n-i-1}$ 
6    $c.add[\text{CNOT}(q, qq)]$ 
7   for  $j \leftarrow 1$  to  $qq - 1$  do
8      $c.add[\text{CNOT}(q + j, q)]$ 
9   for  $j \leftarrow 1$  to  $qq - 1$  do
10     $c.add[\text{SWAP}(q + j, q + j + qq).controlled\_by(q)]$ 
11     $q \leftarrow q + qq$ 
12  $c.add[X(2^n - 1)]$ 
13 return
    Binary basis in qubits  $\{2^n - 2^i\}$  with  $i \leftarrow n$  to 1
  
```




FIG. 6. Comparison of a Gaussian probability distribution uploaded using the unary basis to 16 qubits P_{16} that is enlarged into P_4 using the QFT interpolation scheme to span the full 2^{16} Hilbert space.

APPENDIX B: TRACE DISTANCE FOR PURE STATES ENCODING REAL PROBABILITY DISTRIBUTIONS

The trace distance for pure quantum states is defined as

$$\text{dist}_{\text{Tr}}(|\psi\rangle, |\phi\rangle) = \frac{1}{2} \|\ |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1. \quad (\text{B1})$$

At the same time, for pure quantum states, the Fuchs–van-de-Graaf inequality is tight, meaning that

$$\|\ |\psi\rangle\langle\psi| - |\phi\rangle\langle\phi| \|_1 = 2\sqrt{1 - |\langle\psi|\phi\rangle|^2}. \quad (\text{B2})$$

Recall then that the ℓ_2 norm of the same pure states can be decomposed as

$$\|\ |\psi\rangle - |\phi\rangle \|^2 = \|\ |\psi\rangle \|^2 + \|\ |\phi\rangle \|^2 - 2\text{Re}(\langle\psi|\phi\rangle). \quad (\text{B3})$$

In particular, when the pure states being considered are amplitude encodings of probability distributions, that is, all amplitudes are real valued, the ℓ_2 norm can be simplified as

$$\|\ |\psi\rangle - |\phi\rangle \|^2 = 2(1 - |\langle\psi|\phi\rangle|). \quad (\text{B4})$$

Therefore, one can substitute the value

$$|\langle\psi|\phi\rangle| = \left(1 - \frac{\|\ |\psi\rangle - |\phi\rangle \|^2}{2}\right) \quad (\text{B5})$$

in Eq. (B2) in order to recover the trace distance from the ℓ_2 distance in this particular case. Precisely,

$$\text{dist}_{\text{Tr}}(|\psi\rangle, |\phi\rangle) = \|\ |\psi\rangle - |\phi\rangle \| \sqrt{1 - \frac{\|\ |\psi\rangle - |\phi\rangle \|^2}{4}}. \quad (\text{B6})$$

APPENDIX C: DISTANCE FOR BAND-FILTERED PURE STATES

The *bra-ket* notation will be dropped in the derivation section for clarity.

The band-limited version of target state Ψ can be defined as

$$\tilde{\Psi} = \frac{\Psi_{\text{in}}}{\|\Psi_{\text{in}}\|}, \quad (\text{C1})$$

where Ψ_{in} are the components of the target Ψ that fall within the n -qubit band-limited space. Alternatively, Ψ_{out} are the components outside of the n -qubit band limit, and $\Psi_{\text{in}} + \Psi_{\text{out}} = \Psi$, see Eq. (3).

Therefore, the difference between the target state and the initial state used in the interpolation is

$$\|\Psi - \tilde{\Psi}\|^2 = \|\Psi_{\text{out}}\|^2 + \left(1 - \frac{1}{\|\Psi_{\text{in}}\|}\right)^2 \|\Psi_{\text{in}}\|^2. \quad (\text{C2})$$

Developing the square and using the relation $\|\Psi_{\text{in}}\|^2 + \|\Psi_{\text{out}}\|^2 = 1$ results in

$$\|\Psi - \tilde{\Psi}\|^2 = 2 + 2 \frac{\|\Psi_{\text{out}}\|^2 - 1}{\sqrt{1 - \|\Psi_{\text{out}}\|^2}} \quad (\text{C3})$$

that will only depend on the norm of the high-frequency components. By combining the terms into

$$\|\Psi - \tilde{\Psi}\|^2 = 2 \frac{\sqrt{1 - \|\Psi_{\text{out}}\|^2} - 1 + \|\Psi_{\text{out}}\|^2}{\sqrt{1 - \|\Psi_{\text{out}}\|^2}} \quad (\text{C4})$$

and multiplying both terms in the fraction by $\sqrt{1 - \|\Psi_{\text{out}}\|^2} + 1$, we reach the simplified form shown in the main text,

$$\|\Psi - \tilde{\Psi}\|^2 = \frac{2\|\Psi_{\text{out}}\|^2}{(1 + \sqrt{1 - \|\Psi_{\text{out}}\|^2})}. \quad (\text{C5})$$

APPENDIX D: DISTANCE FOR ALIASED PURE STATES

The derivation of the ℓ_2 distance between the target and aliased quantum state will proceed similarly to the band-limited one, with more detail on the normalization constant N .

The difference now can be written as

$$\|\Psi - \tilde{\Psi}\|^2 = \left\| \Psi - \frac{1}{N}(\Psi_{\text{in}} + \Phi) \right\|^2, \quad (\text{D1})$$

where $N = \|\Psi_{\text{in}} + \Phi\|$ is the normalization factor of the interpolated state $\tilde{\Psi}$ and Φ is the aliasing effect of the off band-limit components which satisfies $\|\Phi\| = \|\Psi_{\text{out}}\|$. Notice that $N \geq 1$, since

$$\|\Psi_{\text{in}} + \Phi\|^2 \geq \|\Psi_{\text{in}}\|^2 + \|\Phi\|^2 = \|\Psi_{\text{in}}\|^2 + \|\Psi_{\text{out}}\|^2 = 1, \quad (\text{D2})$$

due to Ψ_{in} and Ψ_{out} composing to reconstruct Ψ , which is a normalized quantum state.

The norm outside the band limit remains the same, leaving

$$\|\Psi - \tilde{\Psi}\|^2 = \|\Psi_{\text{out}}\|^2 + \left\| \left(1 - \frac{1}{N}\right)\Psi_{\text{out}} - \frac{1}{N}\Phi \right\|^2. \quad (\text{D3})$$

The second term can be decomposed, making the distance then

$$\begin{aligned} \|\Psi - \tilde{\Psi}\|^2 &= \|\Psi_{\text{out}}\|^2 + \left(1 - \frac{1}{N}\right)^2 \|\Psi_{\text{in}}\|^2 + \frac{1}{N^2} \|\Phi\|^2 \\ &\quad - 2 \left(1 - \frac{1}{N}\right) \frac{1}{N} \text{Re}(\langle \Psi_{\text{in}} | \Phi \rangle). \end{aligned} \quad (\text{D4})$$

In order to get rid of the last term, we recall that the norm N of the aliased state is

$$N^2 = \|\Psi_{\text{in}} + \Phi\|^2 = \|\Psi_{\text{in}}\|^2 + \|\Phi\|^2 + 2 \text{Re}(\langle \Psi_{\text{in}} | \Phi \rangle), \quad (\text{D5})$$

and since $\|\Psi_{\text{in}}\|^2 + \|\Phi\|^2 = \|\Psi_{\text{in}}\|^2 + \|\Psi_{\text{out}}\|^2 = 1$, we can substitute

$$2 \text{Re}(\langle \Psi_{\text{in}} | \Phi \rangle) = N^2 - 1 \quad (\text{D6})$$

into Eq. (D4) to reach

$$\begin{aligned} \|\Psi - \tilde{\Psi}\|^2 &= \|\Psi_{\text{out}}\|^2 + \left(1 - \frac{1}{N}\right)^2 (1 - \|\Psi_{\text{out}}\|^2) \\ &\quad + \frac{1}{N^2} \|\Psi_{\text{out}}\|^2 - \left(1 - \frac{1}{N}\right) (N^2 - 1) \frac{1}{N}. \end{aligned} \quad (\text{D7})$$

This equation simplifies to the form shown in the main text,

$$\|\Psi - \tilde{\Psi}\|^2 = \frac{2\|\Psi_{\text{out}}\|^2}{N} - \frac{(N-1)^2}{N}. \quad (\text{D8})$$

APPENDIX E: PSNR AND SSIM FOR IMAGE INTERPOLATION

The PSNR between the original image f and the interpolation result g is defined as

$$\text{PSNR}(f, g) = 10 \log_{10} [255^2 / \text{MSE}(f, g)], \quad (\text{E1})$$

with the mean squared error (MSE) being

$$\text{MSE}(f, g) = \frac{1}{NM} \sum_{i=1}^M \sum_{j=1}^N (f_{ij} - g_{ij})^2 \quad (\text{E2})$$

for $M \times N$ images. The higher the PSNR value, the closer the images are in terms of numerical value. Generally, one aims at a PSNR of 30 or above.

The SSIM index is another image quality metric that is correlated with how humans perceive images. It compares three metrics—the luminance l , the contrast c , and the structure s —of two images, f and g . Each element is defined as

$$l(f, g) = \frac{2\mu_f \mu_g + c_1}{\mu_f^2 + \mu_g^2 + c_1}, \quad (\text{E3})$$

$$c(f, g) = \frac{2\sigma_f \sigma_g + c_2}{\sigma_f^2 + \sigma_g^2 + c_2}, \quad (\text{E4})$$

$$s(f, g) = \frac{\sigma_{fg} + c_3}{\sigma_f \sigma_g + c_3}, \quad (\text{E5})$$

where $\mu_{f,g}$ is the average of the image, $\sigma_{f,g}^2$ is the variance, σ_{fg} is the covariance, and $c_{1,2,3}$ are constants to avoid 0 on the denominator. The SSIM is the product of the three metrics,

$$\text{SSIM}(f, g) = l(f, g) c(f, g) s(f, g), \quad (\text{E6})$$

and is equal to 1 when the two images are the same.

- [1] J. J. García-Ripoll, Quantum-inspired algorithms for multivariate analysis: From interpolation to partial differential equations, *Quantum* **5**, 431 (2021).
- [2] P. García-Molina, J. Rodríguez-Mediavilla, and J. J. García-Ripoll, Quantum Fourier analysis for multivariate functions and applications to a class of Schrödinger-type partial differential equations, *Phys. Rev. A* **105**, 012433 (2022).
- [3] N. Ahmed, T. Natarajan, and K. R. Rao, Discrete cosine transform, *IEEE Trans. Comput.* **100**, 90 (1974).
- [4] G. K. Wallace, The JPEG still picture compression standard, *IEEE Trans. Consumer Electron.* **38**, xviii (1992).
- [5] S. Efthymiou, S. Ramos-Calderer, C. Bravo-Prieto, A. Pérez-Salinas, D. García-Martín, A. Garcia-Saez, J. I. Latorre, and S. Carrazza, Qibo: A framework for quantum simulation with hardware acceleration, *Quantum Sci. Technol.* **7**, 015018 (2021).
- [6] S. Efthymiou, S. Carrazza, C. Bravo-Prieto, A. P. Salinas, S. Ramos, D. García-Martín, M. Lazzarin, N. Zattarin, A. Pasquale, Paul, and J. Serrano, *qiboteam/qibo: Qibo 0.1.7* (Zenodo, 2022), doi: [10.5281/zenodo.7307979](https://doi.org/10.5281/zenodo.7307979).
- [7] L. Grover and T. Rudolph, Creating superpositions that correspond to efficiently integrable probability distributions, [arXiv:quant-ph/0208112](https://arxiv.org/abs/quant-ph/0208112) (2002).
- [8] V. V. Shende, S. S. Bullock, and I. L. Markov, Synthesis of quantum-logic circuits, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **25**, 1000 (2006).
- [9] A. Kitaev and W. A. Webb, Wavefunction preparation and resampling using a quantum computer, [arXiv:0801.0342](https://arxiv.org/abs/0801.0342) (2008).
- [10] M. Plesch and Č. Brukner, Quantum-state preparation with universal gate decompositions, *Phys. Rev. A* **83**, 032302 (2011).
- [11] A. Holmes and A. Matsuura, Efficient quantum circuits for accurate state preparation of smooth, differentiable functions, in *2020 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2020) pp. 169–179.
- [12] A. G. Rattew, Y. Sun, P. Minssen, and M. Pistoia, The efficient preparation of normal distributions in quantum registers, *Quantum* **5**, 609 (2021).
- [13] S. Lloyd and C. Weedbrook, Quantum Generative Adversarial Learning, *Phys. Rev. Lett.* **121**, 040502 (2018).
- [14] P.-L. Dallaire-Demers and N. Killoran, Quantum generative adversarial networks, *Phys. Rev. A* **98**, 012324 (2018).
- [15] C. Zoufal, A. Lucchi, and S. Woerner, Quantum generative adversarial networks for learning and loading random distributions, *npj Quantum Inf.* **5**, 103 (2019).
- [16] S. Ramos-Calderer, A. Pérez-Salinas, D. García-Martín, C. Bravo-Prieto, J. Cortada, J. Planagumà, and J. I. Latorre, Quantum unary approach to option pricing, *Phys. Rev. A* **103**, 032414 (2021).
- [17] C. A. Fuchs and J. Van De Graaf, Cryptographic distinguishability measures for quantum-mechanical states, *IEEE Trans. Inf. Theory* **45**, 1216 (1999).
- [18] M. Kliesch and I. Roth, Theory of quantum system certification, *PRX Quantum* **2**, 010201 (2021).
- [19] H. Nyquist, Certain topics in telegraph transmission theory, *Trans. Am. Inst. Electr. Eng.* **47**, 617 (1928).
- [20] C. E. Shannon, Communication in the presence of noise, *Proc. IRE* **37**, 10 (1949).
- [21] E. T. Whittaker, XVIII.—On the functions which are represented by the expansions of the interpolation-theory, *Proc. R. Soc. Edinburgh* **35**, 181 (1915).
- [22] J. I. Latorre, Image compression and entanglement, [arXiv:quant-ph/0510031](https://arxiv.org/abs/quant-ph/0510031) (2005).
- [23] P. Q. Le, F. Dong, and K. Hirota, A flexible representation of quantum images for polynomial preparation, image compression, and processing operations, *Quant. Info. Proc.* **10**, 63 (2011).
- [24] Y. Zhang, K. Lu, Y. Gao, and M. Wang, NEQR: A novel enhanced quantum representation of digital images, *Quant. Info. Proc.* **12**, 2833 (2013).
- [25] Y. Zhang, K. Lu, and Y. Gao, QSobel: A novel quantum image edge extraction algorithm, *Sci. China Inf. Sci.* **58**, 1 (2015).
- [26] X.-W. Yao, H. Wang, Z. Liao, M.-C. Chen, J. Pan, J. Li, K. Zhang, X. Lin, Z. Wang, Z. Luo *et al.*, Quantum Image Processing and Its Application to Edge Detection: Theory and Experiment, *Phys. Rev. X* **7**, 031041 (2017).
- [27] A. Klappenecker and M. Rotteler, Discrete cosine transforms on quantum computers, in *ISPA 2001. Proceedings of the 2nd International Symposium on Image and Signal Processing and Analysis. In Conjunction with 23rd International Conference on Information Technology Interfaces (IEEE Cat. No.01EX480)* (IEEE, New York, 2001) pp. 464–468.
- [28] C. Y. Pang, Z. W. Zhou, and G. C. Guo, Quantum discrete cosine transform for image compression, [arXiv:quant-ph/0601043](https://arxiv.org/abs/quant-ph/0601043) (2006).
- [29] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, scikit-image: Image processing in Python, *PeerJ* **2**, e453 (2014).
- [30] G. Bradski, The OpenCV Library, *Dr. Dobb's Journal of Software Tools* (2000), https://www.bibsonomy.org/bibtex/25fa6cccd386170d95ae17fd29a42303d/ross_mck.
- [31] A. Hore and D. Ziou, Image quality metrics: PSNR vs. SSIM, in *2010 20th International Conference on Pattern Recognition* (IEEE, 2010) pp. 2366–2369.
- [32] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, Image quality assessment: From error visibility to structural similarity, *IEEE Trans. Image Process.* **13**, 600 (2004).
- [33] T. Acharya and P.-S. Tsai, Computational foundations of image interpolation algorithms, *Ubiquity* **2007**, 4 (2007).
- [34] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM Rev.* **41**, 303 (1999).
- [35] S. Ramos-Calderer, <https://github.com/qiboteam/qft-resampling> (2022).
- [36] Brought to my attention in a post by Craig Gidney in <https://quantumcomputing.stackexchange.com/questions/5526/garbage-free-reversible-binary-to-unary-decoder-construction>.