




Quantics Tensor Cross Interpolation for High-Resolution, Parsimonious Representations of Multivariate Functions in Physics and Beyond

Marc K. Ritter ¹ Yuriel Núñez Fernández ² Markus Wallerberger,³

Jan von Delft ¹ Hiroshi Shinaoka ⁴ and Xavier Waintal ²

¹*Arnold Sommerfeld Center for Theoretical Physics, Center for NanoScience, and Munich Center for Quantum Science and Technology, Ludwig-Maximilians-Universität München, 80333 Munich, Germany*

²*Université Grenoble Alpes, CEA, Grenoble INP, IRIG, Pheliqs, F-38000 Grenoble, France*

³*Institute of Solid State Physics, TU Wien, 1040 Vienna, Austria*

⁴*Department of Physics, Saitama University, Saitama 338-8570, Japan*

(Dated: September 11, 2023)

Multivariate functions of continuous variables arise in countless branches of science. Numerical computations with such functions typically involve a compromise between two contrary desiderata: accurate resolution of the functional dependence, versus parsimonious memory usage. Recently, two promising strategies have emerged for satisfying both requirements: (i) The *quantics* representation, which expresses functions as multi-index tensors, with each index representing one bit of a binary encoding of one of the variables; and (ii) *tensor cross interpolation* (TCI), which, if applicable, yields parsimonious interpolations for multi-index tensors. Here, we present a strategy, *quantics TCI* (QTCI), which combines the advantages of both schemes. We illustrate its potential with an application from condensed matter physics: the computation of Brillouin zone integrals.

DOI: [10.48550/arXiv.2303.11819](https://doi.org/10.48550/arXiv.2303.11819)

Introduction.—Let f be a multivariate function of n continuous, real variables u_i ($i = 1, \dots, n$):

$$f: U \subset \mathbb{R}^n \rightarrow \mathbb{C}, \quad \mathbf{u} = (u_1, \dots, u_n) \mapsto f(\mathbf{u}). \quad (1)$$

Such functions arise in essentially all branches of science. In physics, e.g., they could stand for the fields used in classical or quantum field theories, with $\mathbf{u} = (\mathbf{x}, t)$ or $\mathbf{u} = (\mathbf{k}, \omega)$ representing space-time or momentum-frequency variables in $n = \mathcal{D} + 1$ dimensions, respectively; or for m -point correlation functions of such fields, with $\mathbf{u} = (\mathbf{x}_1, t_1, \dots, \mathbf{x}_m, t_m)$ and $n = m(\mathcal{D} + 1)$, etc.

Often such functions have structure (peaks, wiggles, divergences, even discontinuities) on length/time scales or momentum/frequency scales differing by orders of magnitude. Then, their numerical treatment is challenging due to two contrary requirements: On the one hand, accurate resolution of small-scale structures requires a fine-grained discretization grid, while large-scale structures require a large domain of definition U ; and on the other hand, memory usage should be parsimonious, hence a fine-grained grid cannot be used throughout U . In practice, compromises are needed, sacrificing resolution and/or restricting U to limit memory costs, or using nonuniform grid spacings to resolve some parts of U more finely than others.

Very recently, in different branches of physics, it was pointed out that if the structures in f exhibit *scale separation*, in a sense made precise below, they can be encoded both accurately and parsimoniously, on both small and large scales [1–5]. This is done using a representation first discussed in the context of quantum information [6–9], independently introduced in the mathematics literature by Oseledets [10], and dubbed the *quantics* representation

by Khoromskij [11]: it encodes each variable u_i through R binary digits, or bits, and expresses $f(\mathbf{u})$ as a multi-index tensor $f_{\sigma_1 \dots \sigma_{\mathcal{L}}}$, with $\mathcal{L} = nR$, where each index represents a bit. If f exhibits scale separation, this tensor is highly compressible, i.e. it can be well approximated by a tensor train (TT) of fairly low rank. These previous works found the TT via singular value decomposition (SVD) of the full tensor, demonstrating that low-rank quantics TT (QTT) representations *exist*. It remains to design more practical algorithms to find them, since the computational costs of the SVD approach grow exponentially with \mathcal{L} .

In an unrelated very recent development [12], TT representations were used for multivariate correlation functions arising in diagrammatic Monte Carlo methods (albeit without using the quantics encoding). It was found that these TTs are not only highly compressible, but that the compression can be achieved very efficiently using the *tensor cross interpolation* (TCI) algorithm. This technique, pioneered by Oseledets and coworkers [13–15] and improved by Dolgov and Savostyanov [16, 17], is computationally exponentially cheaper than SVDs (albeit theoretically less optimal, though with controlled errors [16]).

The purpose of this paper is to point out that quantics TTs and TCI can be combined. This leads to a strategy that we call *quantics tensor cross interpolation* (QTCI). It has several highly desirable properties: (i) Arbitrary resolution via an exponentially large grid with 2^R points for each variable, obtained at a cost linear in R . (ii) Efficient construction of the QTCI at cost linear in $\mathcal{L} = nR$. (iii) Access to many ultrafast algorithms once the QTCI has been obtained [1, 5, 18–20]; for instance, integrals,

convolutions and Fourier transforms can be computed at $\mathcal{O}(\mathcal{L})$ costs (i.e. exponentially cheaper than standard fast Fourier transforms) [21]. (iv) More generally, TT representations yield access to a whole range of matrix product states/matrix product operators (MPS/MPO) algorithms which were devised in the context of many-body physics [18] and have spawned the mathematical field of TTs.

We illustrate the power of QTCI by using it to resolve the momentum dependence of functions defined on the Brillouin zone of the celebrated Haldane model [28]. We construct QTCIs for its noninteracting Green’s function and Berry curvature, and compute the Chern number of a band with topological properties.

Quantics tensor trains.—For the quantics representation of $f(\mathbf{u})$, each variable u_i is rescaled to lie within the unit interval $I = [0, 1)$, then discretized on a grid of 2^R points and expressed as a tuple of R bits [10, 11],

$$u_i = \sum_{b=1}^R \frac{\sigma_{ib}}{2^b} \mapsto (\sigma_{i1} \dots \sigma_{iR}), \quad \sigma_{ib} \in \{0, 1\}. \quad (2)$$

Here, σ_{ib} resolves the variable u_i at the scale 2^{-b} . Arbitrarily high resolution can be achieved by choosing R sufficiently large. Thus, \mathbf{u} is represented by a tuple of $\mathcal{L} = nR$ bits. To facilitate scale separation, the bits are relabeled [10] as $\sigma_\ell = \sigma_{ib}$, using a single index $\ell = i + (b-1)n \in 1, \dots, \mathcal{L}$. This interleaves them such that all bits σ_{ib} describing the same scale 2^{-b} have contiguous σ_ℓ labels. Then, f can be viewed and graphically depicted as tensor of degree \mathcal{L} :

$$f : \{0, 1\}^{\mathcal{L}} \rightarrow \mathbb{C}, \quad \sigma = (\sigma_1, \dots, \sigma_{\mathcal{L}}) \mapsto f_\sigma = f(\mathbf{u}). \quad (3)$$

$$\begin{aligned} f_\sigma &= \overbrace{\sigma_{11} \dots \sigma_{n1} \sigma_{12} \dots \sigma_{n2} \dots \sigma_{1R} \dots \sigma_{nR}}^{\mathcal{L}} \\ &= \overbrace{\sigma_1 \sigma_2 \dots \sigma_\ell \dots \sigma_{\mathcal{L}}}^{\mathcal{L}}. \end{aligned}$$

Alternatively, all same-scale bits can be *fused* together as $\tilde{\sigma}_b = \sum_{i=1}^n 2^{i-1} \sigma_{ib} \in \{0, \dots, 2^n - 1\}$, yielding the *fused* representation $f_{\tilde{\sigma}} = f(\mathbf{u})$, $\tilde{\sigma} = (\tilde{\sigma}_1, \dots, \tilde{\sigma}_R)$. It employs only $\tilde{\mathcal{L}} = R$ indices, each of dimension $d = 2^n$ [29].

Any tensor can be *unfolded* as a TT [11, 13, 14, 30], graphically depicted as a chain of ℓ sites connected by bonds representing sums over repeated indices:

$$\begin{aligned} f_\sigma &= \prod_{\ell=1}^{\mathcal{L}} M_\ell^{\sigma_\ell} = [M_1]_{1\alpha_1}^{\sigma_1} [M_2]_{\alpha_1\alpha_2}^{\sigma_2} \dots [M_{\mathcal{L}}]_{\alpha_{\mathcal{L}-1}1}^{\sigma_{\mathcal{L}}} \quad (4) \\ &= \begin{array}{c} \times \quad 1 \quad M_1 \quad D_1 \quad M_2 \quad D_2 \quad \dots \quad M_\ell \quad D_\ell \quad \dots \quad M_{\mathcal{L}} \quad 1 \quad \times \\ \circ \quad \alpha_1 \quad \circ \quad \alpha_2 \quad \dots \quad \circ \quad \alpha_\ell \quad \dots \quad \circ \quad \alpha_{\mathcal{L}} \\ \sigma_1 \quad \quad \quad \sigma_2 \quad \quad \quad \quad \quad \sigma_\ell \quad \quad \quad \sigma_{\mathcal{L}} \end{array} \end{aligned}$$

Each site ℓ hosts a three-leg tensor M_ℓ with elements $[M_\ell]_{\alpha_{\ell-1}\alpha_\ell}^{\sigma_\ell}$. Its “local” and “virtual” bond indices, σ_ℓ and $\alpha_{\ell-1}, \alpha_\ell$, have dimensions $d=2$ and $D_{\ell-1}, D_\ell$, respectively, with $D_0 = D_{\mathcal{L}} = 1$ for the outermost (dummy) bonds.

If f_σ is full rank, exact TT unfoldings have exponential bond growth towards the chain center, $D_\ell = 2^{\min\{\ell, \mathcal{L}-\ell\}}$,

implying exponential memory costs, $\mathcal{O}(2^{\mathcal{L}/2})$. However, tensors f_σ with lower information content admit accurate TT unfoldings with lower virtual bond dimensions. Such unfoldings are obtained via iterative factorization and truncation of bonds with low information content. Usually, this is done using a sequence of SVDs, discarding all singular values smaller than a specified *truncation threshold* ϵ . The largest D_ℓ value so obtained, D_{\max} , is the rank of the ϵ -truncated TT. SVD truncation is provably optimal [14], yielding the smallest possible D_{\max} for specified ϵ . If $D_{\max} \ll 2^{\mathcal{L}/2}$, f_σ is *strongly compressible*, implying that it has internal structure. Building on the pioneering studies of Oseledets [13, 30] and Khoromskij [11], Refs. [2–5] argued that for quantics tensors $f_\sigma = f(\mathbf{u})$, strong compressibility reflects scale separation: structures in $f(\mathbf{u})$ occurring on different scales are only “weakly entangled,” in that the virtual bonds connecting the corresponding sites in the TT do not require large dimensions.

This perspective is informed by the study of one-dimensional quantum lattice models using matrix product states (MPSs)—many-body wave-functions of the form (4) [18]. In that context, σ_ℓ labels physical degrees of freedom at site ℓ , and the entanglement of sites ℓ and $\ell+1$ is characterized by an entanglement entropy bounded by 2^{D_ℓ} . By analogy, if a quantics TT is strongly compressible, requiring only small D_{\max} , the sites representing different scales are not strongly entangled—indeed, D_{\max} quantifies the degree of scale separation inherent in $f(\mathbf{u})$.

The SVD unfolding strategy requires knowledge of the full tensor f_σ : it uses $2^{\mathcal{L}}$ function calls, implying exponentially long runtimes, even if f_σ is strongly compressible. Thus, this strategy is optimally accurate but exponentially inefficient: it uncovers structure in f_σ , but does not exploit it already while constructing the unfolded TT.

Tensor cross interpolation.—The TCI algorithm [12–14, 17] solves this problem. It serves as a *black box* that samples f_σ at some clever choices of σ and iteratively constructs the TT from the sampled values. TCI is slightly less accurate than SVD unfoldings, requiring a slightly larger D_{\max} for a specified error tolerance ϵ . But it is exponentially more efficient, needing at most $\mathcal{O}(D_{\max}^2 \mathcal{L})$ function evaluations and a run-time of at most $\mathcal{O}(D_{\max}^3 \mathcal{L})$.

We refer to [12] for details about TCI in general and its actual implementation. Here, we just sketch the main idea. TCI achieves the factorization needed for unfolding by employing matrix cross interpolation (MCI) rather than SVD. Given a matrix A , the MCI formula approximates it as $A \approx CP^{-1}R = \tilde{A}$, graphically depicted as follows:

$$A = \begin{pmatrix} \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \end{pmatrix} \approx \begin{pmatrix} \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet \end{pmatrix}^{-1} \begin{pmatrix} \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \\ \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet & \bullet \end{pmatrix},$$

or $\begin{array}{c} \square \approx \square \diamond \square \end{array}$.

Here, the column, row and pivot matrices C , R , and P , are all constructed from elements of A : C contains D columns (red), R contains D rows (blue), and P their intersections, the *pivots* (purple). The resulting \tilde{A} exactly

reproduces all elements of A contained in C and R ; the remaining elements are in effect interpolated from the “crosses” formed by these (hence “cross interpolation”). The accuracy of the interpolation depends on the number and choice of pivots; it can be improved systematically by adding more pivots. If $D = \text{rank}(A)$, one can obtain an exact representation of the full matrix, $A = \tilde{A}$ [14].

Tensors can be unfolded by iteratively using MCI while treating multiple indices (e.g. $\sigma_2 \dots \sigma_{\mathcal{L}}$) as a single composite index, e.g. $f_{\sigma_1 \sigma_2 \dots \sigma_{\mathcal{L}}} \approx [C_1]_{1\beta_1}^{\sigma_1} [P_1^{-1}]_{\beta_1 \alpha_1} [R_1]_{\alpha_1}^{\sigma_2 \dots \sigma_{\mathcal{L}}}$. Iterative application of MCI to each new tensor on the right ultimately yields a fully unfolded TT, $f_{\sigma} \approx f_{\sigma}^{\text{QTCl}}$:

$$\overbrace{\sigma_1 \sigma_2 \dots \sigma_{\mathcal{L}}} \approx \overbrace{\sigma_1 \sigma_2 \dots \sigma_{\mathcal{L}}} \approx \dots \approx \overbrace{\sigma_1 \sigma_2 \dots \sigma_{\mathcal{L}}} \cdot$$

[A TT of the form (4) is obtained by defining $M_{\ell} = C_{\ell} P_{\ell}^{-1}$, $\diamond = \square \diamond \square$.] This naive algorithm is inefficient, but illustrates how the interpolation properties of MCI carry over to TCI. In practice, it is more efficient to use a sweeping algorithm, successively sampling more function values f_{σ} and adding pivots to each tensor until the relative error ε , which decreases during sweeping, drops below a specified tolerance ϵ [21]. We define ε as $\max_{\sigma \in S} |f_{\sigma}^{\text{QTCl}} - f_{\sigma}| / \max_{\sigma \in S} |f_{\sigma}|$, where S is the set of all σ index values sampled while constructing f_{σ}^{QTCl} .

Integration.—The integral over a function f in QTT form is easily accessible in $\mathcal{O}(D_{\max}^2 \mathcal{L})$ steps [12, 14]. It can be approximated by a Riemann sum since the quantics grid is exponentially fine, and all σ_{ℓ} sums can be performed independently due to the TT’s factorized form:

$$\int_{I^n} d^n \mathbf{u} f(\mathbf{u}) \approx \frac{1}{2^{\mathcal{L}}} \sum_{\sigma} f_{\sigma} \approx \frac{1}{2^{\mathcal{L}}} \prod_{\ell} \sum_{\sigma_{\ell}} [M_{\ell}]^{\sigma_{\ell}}. \quad (5)$$

1D example.—We first demonstrate QTCl for computing the integral $I[f] = \int_0^{\ln(20)} dx f(x)$ of the function $f(x) = \cos(\frac{x}{B}) \cos(\frac{x}{4\sqrt{5}B}) e^{-x^2} + 2e^{-x}$ with $B = 2^{-30} \approx 10^{-9}$. This function, shown in Fig. 1(a), involves structure on widely different scales: rapid, incommensurate oscillations and a slowly decaying envelope. A standard representation thereof on an equidistant mesh would require much more than $\mathcal{O}(1/B)$ sampling points, as would the computation of the integral $I[f] = \frac{19}{10} + \mathcal{O}(e^{-1/(4B^2)})$. By contrast, for a quantics representation, it suffices to choose R somewhat larger than 30 (ensuring $2^{-R} \ll B$); and since the information content of $f(x)$ is not very high, f_{σ} is strongly compressible. We unfolded it using QTCl with $\epsilon = 10^{-8}$ and $R = 50$ (quite a bit larger than 30, just to demonstrate the capabilities of TCI). Figure 1(b) shows the resulting profile of D_{ℓ} vs ℓ , revealing the scale separation inherent in $f(x)$: the initial growth of the bond dimension, $D_{\ell} \sim e^{\ell}$, quickly stops at a fairly small maximum, $D_{\max} = 15$, confirming strong compressibility; thereafter, D_{ℓ} decreases steadily with ℓ , becoming $\mathcal{O}(1)$ for ℓ larger than 30, since $f(x)$ has very little structure at scales below 2^{-30} . Remarkably, although f_{σ} has

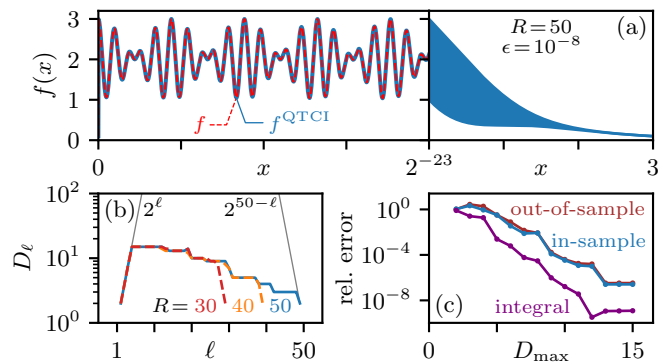


FIG. 1. QTCl representation of a rapidly oscillating function, for $\mathcal{L} = R = 50$ with tolerance $\epsilon = 10^{-8}$. (a) Plot of $f(x)$ (see text). Left: the interval $x \in [0; 2^{-23}]$; red dashed: the actual function, blue: its QTCl representation. Right: the envelope structure up to $x = \log(20) \approx 3$; the rapid oscillations are not resolvable on this scale. (b) Virtual bond dimensions D_{ℓ} of the QTT, for $R = 30, 40, 50$. Gray lines indicate how $D_{\ell}^{R=50}$ would grow without any truncation. (c) Relative error estimates as a function of $D_{\max} = \max(D_{\ell})$, for $R = 50$.

$2^{50} \approx 10^{15}$ elements, the TCI algorithm finds the relevant structure using only 8706 samples, i.e. roughly 1 sample per 59 000 oscillations. Nevertheless, it yields an accurate representation of $f(x)$: the in-sample error, the out-of-sample error (defined as maximum error over 2000 random samples) [31], and the error for the integral $I[f]$, computed via Eq. (5), all decrease exponentially with D_{\max} [Fig. 1(c)]. The runtimes for computing $I[f]$ using QTCl or adaptive Gauss-Kronrod quadrature are 44 ms vs 6 h on an Intel Xeon W-2245 processor, illustrating the efficiency of QTCl vs conventional approaches.

Haldane model.—As an example with relevance in physics, we apply QTCl to the Green’s function and Berry curvature of the well-known Haldane model [28]. It is one of the simplest models with topological properties, yet produces nontrivial structure with multiple peaks and sign changes in reciprocal space. Its Bloch Hamiltonian is

$$H(\mathbf{k}) = \sum_{i=1}^3 \left[\sigma^1 \cos(\mathbf{k} \cdot \mathbf{a}_i) + \sigma^2 \sin(\mathbf{k} \cdot \mathbf{a}_i) \right] + \sigma^3 \left[m - 2t_2 \sum_{i=1}^3 \sin(\mathbf{k} \cdot \mathbf{b}_i) \right], \quad (6)$$

where σ^{μ} are Pauli matrices, $\mathbf{k} = (k_x, k_y)$, while $\mathbf{a}_{1,2,3}$ connect nearest neighbors and $\mathbf{b}_{1,2,3}$ next-nearest neighbors of a honeycomb lattice. Compared to Haldane’s more general version of the model, we fix his parameters $\phi = \frac{\pi}{2}$, $t_1 = 1$ and set $t_2 = 0.1$. The parameter m tunes the model through two phase transitions: $|m| < m_c = t_2 3\sqrt{3}$ yields a Chern insulator with Chern number $\mathcal{C} = -1$, and $|m| > m_c$ a trivial phase with $\mathcal{C} = 0$ [28]. At $m = \pm m_c$, a single Dirac point appears at $\mathbf{k} = (\mp \frac{1}{3}\pi, 0)$ and symmetry-related \mathbf{k} ; there, the Chern number is $\mathcal{C} = -\frac{1}{2}$ [33].

Green’s function in reciprocal space.—To illustrate QTCl for the Haldane model, we study the momentum dependence of the Green’s function, $G(\mathbf{k}, i\omega_0) = \text{Tr}[(i\omega_0 - H(\mathbf{k}) + \mu)^{-1}]$, with $\omega_0 = \pi/\beta$ the lowest fermionic

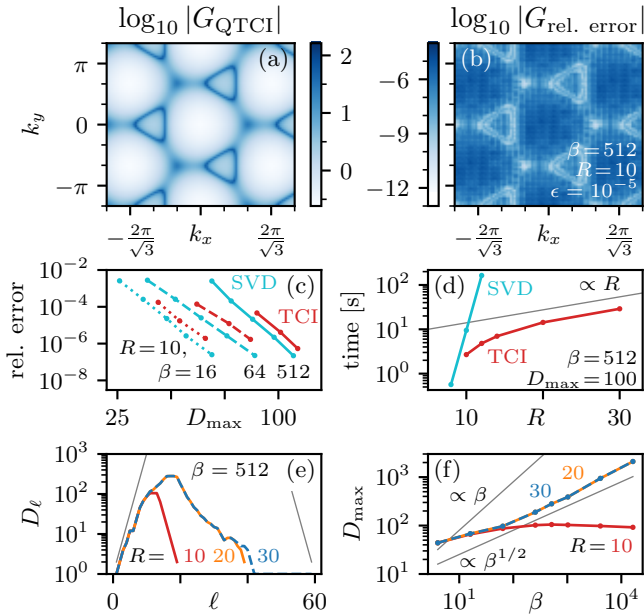


FIG. 2. Green's function $G(\mathbf{k})$ of the Haldane model, computed with error tolerance $\epsilon = 10^{-5}$ throughout. (a,b) QTCI of the $|G|$ and its relative error, $|G_{\text{QTCI}} - G|/|G|$, for $\beta = 512$, $R = 10$. (c,d) Comparison of QTT unfoldings via SVD and TCI, showing (c) the relative error vs the maximum bond dimension for $R = 10$ and $\beta = 16, 64, 512$; and (d) runtimes vs R for $\beta = 512$. (e,f) QTCI bond dimensions for $R = 10, 20, 30$, showing (e) D_ℓ vs ℓ for $\beta = 512$; and (f) D_{max} vs β .

Matsubara frequency and Tr traces over the 2×2 space of H .

Figure 2(a) shows an intensity plot of the QTCI representation of G in reciprocal space; Fig. 2(b) shows that the relative error with respect to the exact value is below 10^{-5} throughout, hence the momentum dependence is captured accurately. There are small Fermi surfaces around $\mathbf{k} = (-\frac{4}{3}\pi, 0)$ and symmetry-related \mathbf{k} . To construct QTTs, we define $f_\sigma = G(\mathbf{k}, i\pi/\beta)$, where σ encodes \mathbf{k} and β is fixed. Figure 2(c) shows the relative in-sample error as a function of D_{max} for TTs constructed with $R = 10$ for $\beta = 16, 64, 512$, using either SVD or TCI. For both, the error decreases exponentially as D_{max} increases. Moreover, TCI is nearly optimal, achieving the same error as SVD for a D_{max} that is only a few percent larger.

Figure 2(d) shows how SVD and TCI runtimes depend on the number of bits, R , for a fixed D_{max} at large $\beta = 512$, where the features in G are sharp. The times, including function evaluations, were measured on a single CPU core of AMD EPYC 7702P. The SVD runtimes become prohibitively large for $R > 10$ due to exponential scaling; by contrast, the TCI runtimes depend only mildly on R .

Figure 2(e) shows how TCI profiles of D_ℓ vs ℓ depend on R , for $\beta = 512$ and a specified error tolerance $\epsilon = 10^{-5}$. The bond dimension initially grows as $D_\ell \sim 2^\ell$, reaches a maximum near $\ell \approx 20$, then decreases back to 1. The curves for $R=20$ and 30 almost coincide, indicating that a

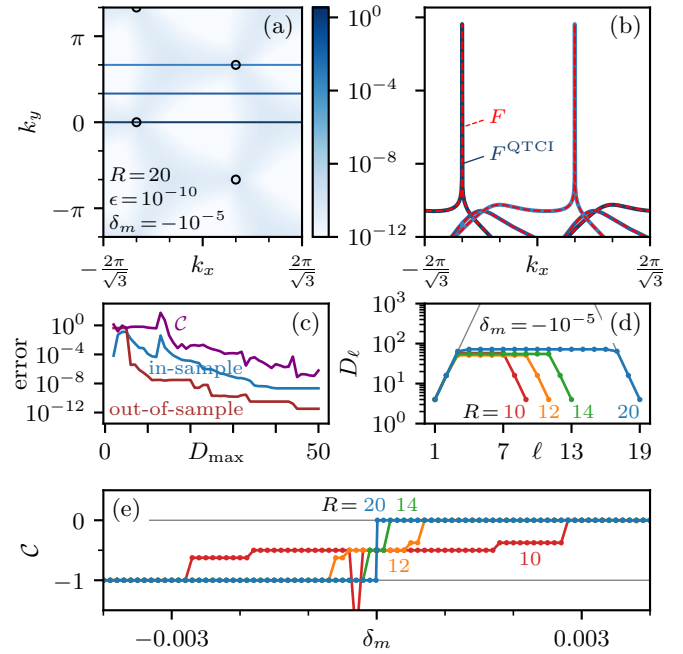


FIG. 3. Evaluation of the Chern number in the Haldane model using QTCI for the Berry flux $F(\mathbf{k})$, with error tolerance $\epsilon = 10^{-10}$ throughout. (a) QTCI of $F(\mathbf{k})$ on the integration domain (circles mark peak positions); (b) cuts through $F(\mathbf{k})$ along the colored lines shown in (a); and (c) errors for $F(\mathbf{k})$ and \mathcal{C} as functions of D_{max} , all computed for $\delta_m = -10^{-5}$, $R = 20$. (d) Bond dimension D_ℓ for QTTs of length $R = 10, 12, 14, 20$. (e) Chern number \mathcal{C} as a function of δ_m , for four choices of R .

good resolution of the sharp features at $\beta = 512$ requires $R > 20$ —well beyond the reach of SVD unfoldings.

The low computational cost of TCI allows us to investigate the β dependence of D_{max} , easily reaching $\beta = 2^{14} = 16384$. Figure 2(f) suggests $D_{\text{max}} \propto \beta^\alpha$ with $\alpha \approx 1/2$ for large β . Remarkably, this growth is slower than that, $D_{\text{max}} \propto \beta$, conjectured for a scheme based on SVD and patching [5]. A detailed analysis for general models and higher spatial dimensions is an interesting topic for future research.

Chern number.—Finally, we consider the Chern number, \mathcal{C} , for the Haldane model at $\mu = 0$ and $\beta = \infty$. To avoid cumbersome gauge-fixing procedures, we use the gauge-invariant method described in Ref. [34]. First, we discretize the Brillouin zone (BZ) into $2^R \times 2^R$ plaquettes. Then, the Chern number can be obtained from a sum over plaquettes, $\mathcal{C} \approx \frac{1}{2\pi i} \sum_{\mathbf{k} \in \text{BZ}} F(\mathbf{k})$, where $F(\mathbf{k}) \approx -i \arg(\langle \psi_{\mathbf{k}_1} | \psi_{\mathbf{k}_2} \rangle \langle \psi_{\mathbf{k}_2} | \psi_{\mathbf{k}_3} \rangle \langle \psi_{\mathbf{k}_3} | \psi_{\mathbf{k}_4} \rangle \langle \psi_{\mathbf{k}_4} | \psi_{\mathbf{k}_1} \rangle)$ is the Berry flux through the plaquette with corners $\mathbf{k}_1 \dots \mathbf{k}_4$, and $|\psi_{\mathbf{k}}\rangle$ are valence band wave functions.

Close to the transition, for small $\delta_m = m - m_c$, the band gap is $2\delta_m$. This induces peaks of width $\sim \delta_m$ in the Berry flux $F(\mathbf{k})$, shown in Figs. 3(a, b) for $\delta_m = 10^{-5}$. There, we used a fused quantics representation with $R = 20$, ensuring a mesh spacing 2^{-R} well smaller than δ_m . Whereas a calculation of \mathcal{C} via direct summation or SVD unfolding would require $2^{2R} \approx 10^{12}$ function evaluations,

QTCI is much more efficient: for a relative tolerance of $\epsilon = 10^{-10}$, it needed only 4×10^5 samples (and 20 s runtime on a single core of an Apple M1 processor). It yielded a QTT with maximum bond dimension $D_{\max} = 50$, and a Chern number within 10^{-6} of the expected value $\mathcal{C} = -1$ (see Figs. 3(c,d)). When plotted as a function of δ_m , \mathcal{C} shows a sharp step from -1 to 0 at $\delta_m = 0$ if computed using $R = 20$ (Fig. 3(e)), beautifully demonstrating that the \mathbf{k} mesh is fine enough. For smaller R the mesh becomes too coarse, incorrectly yielding a plateau at $-1/2$ instead of a sharp step.

For benchmarking purposes, we deliberately chose a model that is analytically solvable. However, our prior knowledge of the peak positions of the Berry curvature was not made available to TCI. This demonstrates its reliability in finding sharp structures, provided enough quantic bits are provided to resolve them. Random sampling misses these sharp structures, which is why in Fig. 3(c) the out-of-sample error, obtained from 2000 random samples, lies well below the in-sample error.

Outlook.—We have shown that the combination of the quantics representation [1–11] with TCI [12–14, 17] is a powerful tool for uncovering low-rank structures in exponentially large, yet very common objects: functions of few variables resolved with high resolution. Numerical work with such objects always involves truncations—the radically new perspective opened up by QTCI is that they can be performed at *polynomial costs* by discarding weak entanglement between different scales. Once a low-rank QTT has been found, it may be further used within one of the many existing MPO/MPS algorithms [1, 5, 18–20].

We anticipate that the class of problems for which QTCI can be instrumental is actually very large, reaching well beyond the scope of physics. Intuitively speaking, the only requirement is that the functions should entail some degree of scale separation and not be too irregular (since random structures are not compressible). Thus, a large new research arena, potentially connecting numerous different branches of science, awaits exploration. Fruitful challenges: establish criteria for which types of multivariate functions admit low-rank QTT representations; develop improved algorithms for constructing low-rank approximations to tensors; and above all, explore the use of QTCI for any of the innumerable problems in science requiring high-resolution numerics. The initial diagnosis is easy: simply use SVDs or QTCI [27] to check whether the functions of interest are compressible or not.

We thank Takashi Koretsune, Ivan Oseledets and Björn Sbierski for inspiring discussions, and Jeongmin Shim for important help at the beginning of this work. We carried out part of the calculations using computer code based on ITensors.jl [35] written in Julia [36]. H. S. was supported by JSPS KAKENHI Grants No. 21H01041, and No. 21H01003, and JST PRESTO Grant No. JP-MJPR2012, Japan. X.W. acknowledges funding from the Plan France 2030 ANR-22-PETQ-0007 “EPIQ”; and






J. v. D. from the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy EXC-2111 (Project No. 390814868), and the Munich Quantum Valley, supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

-
- [1] J. J. García-Ripoll, Quantum-inspired algorithms for multivariate analysis: From interpolation to partial differential equations, *Quantum* **5**, 431 (2021).
 - [2] E. Ye and N. F. G. Loureiro, Quantum-inspired method for solving the Vlasov-Poisson equations, *Phys. Rev. E* **106**, 035208 (2022).
 - [3] N. Gourianov, M. Lubasch, S. Dolgov, Q. Y. van den Berg, H. Babae, P. Givi, M. Kiffner, and D. Jaksch, A quantum inspired approach to exploit turbulence structures, *Nat. Comp. Sci.* **2**, 30 (2022).
 - [4] N. Gourianov, *Exploiting the structure of turbulence with tensor networks*, Ph.D. thesis, University of Oxford (2022).
 - [5] H. Shinaoka, M. Wallerberger, Y. Murakami, K. Nogaki, R. Sakurai, P. Werner, and A. Kauch, Multiscale space-time ansatz for correlation functions of quantum systems based on quantics tensor trains, *Phys. Rev. X* **13**, 021015.
 - [6] S. Wiesner, Simulations of many-body quantum systems by a quantum computer, [arXiv:quant-ph/9603028](https://arxiv.org/abs/quant-ph/9603028) (1996).
 - [7] C. Zalka, Simulating quantum systems on a quantum computer, *Proc. R. Soc. A* **454**, 313 (1998).
 - [8] 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).
 - [9] J. I. Latorre, Image compression and entanglement, [arXiv:quant-ph/0510031](https://arxiv.org/abs/quant-ph/0510031) (2005).
 - [10] I. V. Oseledets, Approximation of matrices with logarithmic number of parameters, *Dokl. Math.* **80**, 653 (2009).
 - [11] B. N. Khoromskij, $\mathcal{O}(d \log n)$ -quantics approximation of $n - d$ tensors in high-dimensional numerical modeling, *Constr. Approximation* **34**, 257 (2011).
 - [12] Y. Núñez Fernández, M. Jeannin, P. T. Dumitrescu, T. Kloss, J. Kaye, O. Parcollet, and X. Waintal, Learning Feynman diagrams with tensor trains, *Phys. Rev. X* **12**, 041018 (2022).
 - [13] I. V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295 (2011).
 - [14] I. Oseledets and E. Tyrtyshnikov, TT-cross approximation for multidimensional arrays, *Linear Algebra Appl.* **432**, 70 (2010).
 - [15] D. Savostyanov and I. Oseledets, Fast adaptive interpolation of multi-dimensional arrays in tensor train format, in *The 2011 International Workshop on Multidimensional (nD) Systems* (IEEE, 2011) pp. 1–8.
 - [16] D. V. Savostyanov, Quasioptimality of maximum-volume cross interpolation of tensors, *Linear Algebra Appl.* **458**, 217 (2014).
 - [17] S. Dolgov and D. Savostyanov, Parallel cross interpolation for high-precision calculation of high-dimensional integrals, *Comput. Phys. Commun.* **246**, 106869 (2020).
 - [18] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (Amsterdam)* **326**, 96 (2011).
 - [19] M. Lubasch, P. Moinier, and D. Jaksch, Multigrid renormalization, *J. Comput. Phys.* **372**, 587 (2018).
 - [20] P. García-Molina, L. Tagliacozzo, and J. J. García-Ripoll,

- Global optimization of MPS in quantum-inspired numerical analysis, [arXiv:2303.09430](https://arxiv.org/abs/2303.09430) (2023).
- [21] See Supplemental Material, Sec. S-1 for a brief explanation of the QTT Fourier transform and Sec. S-2 for an overview of the sweeping algorithm. The supplemental material contains Refs. [5, 16, 22–27].
- [22] S. Dolgov, B. Khoromskij, and D. Savostyanov, Superfast Fourier transform using QTT approximation, *J. Fourier Anal. Appl.* **18**, 915 (2012).
- [23] M. Holzäpfel, T. Baumgratz, M. Cramer, and M. B. Plenio, Scalable reconstruction of unitary processes and Hamiltonians, *Phys. Rev. A* **91**, 042129 (2015).
- [24] J. Chen, E. Stoudenmire, and S. R. White, Quantum fourier transform has small entanglement, *PRX Quantum* **4**, 040318 (2023).
- [25] A. Cortinovis, D. Kressner, and S. Massei, On maximum volume submatrices and cross approximation for symmetric semidefinite and diagonally dominant matrices, *Linear Algebra Appl.* **593**, 251 (2020).
- [26] S. A. Goreinov and E. E. Tyrtyshnikov, Quasioptimality of skeleton approximation of a matrix in the Chebyshev norm, *Dokl. Math.* **83**, 374 (2011).
- [27] A ready-for-use QTCI toolbox will be published as open source library in the near future.
- [28] F. D. M. Haldane, Model for a quantum Hall effect without Landau levels: Condensed-matter realization of the “parity anomaly”, *Phys. Rev. Lett.* **61**, 2015 (1988).
- [29] For example, for $n = 2$, $R = 3$, the point $(u_1, u_2) = (\frac{5}{8}, \frac{4}{8})$ has the binary representation (101, 100). In the *interleaved* form, the bits are reordered such that $f(\frac{5}{8}, \frac{4}{8})$ is represented by $f_\sigma = f_{110010}$; in *fused* form, by $f_{\tilde{\sigma}} = f_{301}$.
- [30] I. V. Oseledets, Approximation of $2^d \times 2^d$ matrices using tensor decomposition, *SIAM J. on Matrix Anal. Appl.* **31**, 2130 (2010).
- [31] The in-sample error is $\max_{\sigma \in S} |f_\sigma^{\text{QTCI}} - f_\sigma|$, where S is the set of all σ evaluated during construction of the QTCI. For the out-of-sample error, S is instead a set of 2000 pseudorandom σ generated by Xoshiro256++ [32]. *Relative errors* are defined as $\max_{\sigma \in S} |f_\sigma^{\text{QTCI}} - f_\sigma| / |f_\sigma|$.
- [32] D. Blackman and S. Vigna, Scrambled linear pseudorandom number generators, *ACM Trans. on Math. Software* **47**, 36:1 (2021).
- [33] H. Watanabe, Y. Hatsugai, and H. Aoki, Manipulation of the Dirac cones and the anomaly in the graphene related quantum Hall effect, *J. Phys. Conf. Ser.* **334**, 012044 (2011).
- [34] T. Fukui, Y. Hatsugai, and H. Suzuki, Chern numbers in discretized Brillouin zone: Efficient method of computing (spin) Hall conductances, *J. Phys. Soc. Jpn.* **74**, 1674 (2005).
- [35] M. Fishman, S. White, and E. Stoudenmire, The ITensor software library for tensor network calculations, *SciPost Phys. Codebases* [10.21468/scipostphyscodeb.4](https://doi.org/10.21468/scipostphyscodeb.4) (2022).
- [36] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A fresh approach to numerical computing, *SIAM Rev.* **59**, 65 (2017).

Supplemental Material:

Quantics Tensor Cross Interpolation for High-Resolution, Parsimonious Representations of Multivariate Functions in Physics and Beyond

Marc K. Ritter ¹ Yuriel Núñez Fernández ² Markus Wallerberger,³
Jan von Delft ¹ Hiroshi Shinaoka ⁴ and Xavier Waintal ²

¹*Arnold Sommerfeld Center for Theoretical Physics, Center for NanoScience, and Munich Center for Quantum Science and Technology, Ludwig-Maximilians-Universität München, 80333 Munich, Germany*

²*Université Grenoble Alpes, CEA, Grenoble INP, IRIG, Pheligs, F-38000 Grenoble, France*

³*Institute of Solid State Physics, TU Wien, 1040 Vienna, Austria*

⁴*Department of Physics, Saitama University, Saitama 338-8570, Japan*

This supplement offers more detail on the runtime complexity of calculating a Fourier transform using QTTs in Sec. S-1, and a brief review of the TCI algorithm in Sec. S-2.

S-1. FOURIER TRANSFORM OF A QTT

This section explains how to compute the Fourier transform of a function in QTT form to substantiate the claim that this can be done faster than fast Fourier transform (FFT). The algorithm is related to the so-called superfast Fourier transform [22] and quantum Fourier transform [23], and has been formulated in matrix product operator (MPO) form in Refs. [5, 24].

In quantics representation, the Fourier transform can be approximated as matrix product operator (MPO) with very small bond dimension, which can then be applied to a function in QTT form at a cost of $\mathcal{O}(D^3\mathcal{L})$ including recompression [5, 23, 24]. The cost of constructing the MPO is in most cases much smaller than the cost of applying it to the QTT. For QTT with moderate bond dimension D , the QTT Fourier transform can be orders of magnitude faster than FFT [22, 24]. For a generic non-compressible function, by contrast, the bond dimension will grow exponentially, which results in the same worst-case complexity as using FFT ($\mathcal{O}(N \log N)$ with $N = 2^\mathcal{L}$). A detailed analysis including numerical experiments can be found in Ref. [24].

S-2. TENSOR CROSS INTERPOLATION

This section gives an overview of the sweeping algorithm we used to construct a tensor cross interpolation (TCI) in practice. For details, see Ref. [12].

Given a tensor f_σ with indices $\sigma = (\sigma_1, \dots, \sigma_\mathcal{L})$ and $\sigma_\ell \in K_\ell := \{0, \dots, d_\ell - 1\}$, our goal is to construct a TCI approximation which fulfills a user-specified error tolerance ϵ with high efficiency. The TCI is assembled from three-leg tensors $[T_\ell]_{\alpha_\ell \beta_\ell}^{\sigma_\ell}$ and two-leg matrices $[P_\ell^{-1}]_{\beta_\ell \alpha_\ell}$

as

$$f_{\sigma_1 \dots \sigma_\mathcal{L}}^{\text{QTCI}} = [T_1]_{\alpha_0 \beta_1}^{\sigma_1} [P_1^{-1}]_{\beta_1 \alpha_1} [T_2]_{\alpha_1 \beta_2}^{\sigma_2} [P_2^{-1}]_{\beta_2 \alpha_2} \dots [T_\mathcal{L}]_{\alpha_{\mathcal{L}-1} \beta_\mathcal{L}}^{\sigma_\mathcal{L}}, \quad (\text{S1})$$

where $\alpha_0 = 1$ and $\beta_\mathcal{L} = 1$ are dummy indices, and repeated indices are summed over (see Fig. S-1). Using multi-index sets $\mathcal{I}_\ell \subseteq K_1 \otimes \dots \otimes K_\ell$ and $\mathcal{J}_\ell \subseteq K_\ell \otimes \dots \otimes K_\mathcal{L}$, all components of the TCI can be described as images of these multiindex sets:

$$T_\ell = f_{\mathcal{I}_{\ell-1} \otimes K_\ell \otimes \mathcal{J}_\ell}; \quad P_\ell^{-1} = (f_{\mathcal{I}_\ell \otimes \mathcal{J}_\ell})^{-1} \quad (\text{S2})$$

Then, the indices α_ℓ enumerate elements of \mathcal{I}_ℓ in Eq. (S1), and β_ℓ enumerate elements of \mathcal{J}_ℓ . The problem of finding a good TCI, i.e. good T_ℓ and P_ℓ , has now been reduced to optimizing \mathcal{I}_ℓ and \mathcal{J}_ℓ . The cost of the algorithm is determined by the *bond dimensions* $D_\ell = |\mathcal{I}_\ell| = |\mathcal{J}_\ell|$, which should optimally be kept as small as possible, while still large enough to approximate f_σ well.

We initialize the TCI starting from an arbitrary initial pivot point $(\bar{\sigma}_1, \dots, \bar{\sigma}_\mathcal{L})$ typically provided by the user. Now, set $\mathcal{I}_\ell = \{(\bar{\sigma}_1, \dots, \bar{\sigma}_\ell)\}$ and $\mathcal{J}_\ell = \{(\bar{\sigma}_\ell, \dots, \bar{\sigma}_\mathcal{L})\}$. Each T_ℓ now contains the full dependence of f_σ on a specific index σ_ℓ , but only takes a single value $\bar{\sigma}_{\ell'}$ of all other digits $\sigma_{\ell' \neq \ell}$ into account. This corresponds to exactly one *pivot* in MCI language. These T_ℓ and P_ℓ are now successively refined by adding pivots as follows.

The algorithm sweeps over the TCI repeatedly. During each sweep, two-site optimizations are applied to bonds between neighbouring tensors T_ℓ and $T_{\ell+1}$, where ℓ progresses from 1 to $\mathcal{L} - 1$ for forward sweeps, and vice versa for backward sweeps. For local optimization




FIG. S-1. Structure of the TCI during optimization in graphical notation. Squares denote 3-leg tensors T_ℓ , diamonds denote matrices P_ℓ^{-1} . The indices σ_ℓ are external, while the indices α_ℓ and β_ℓ are shared between neighbouring T_ℓ and P_ℓ .

of bond ℓ , the algorithm first evaluates a two-site tensor $[\Pi_\ell]_{\alpha\sigma_\ell;\sigma_{\ell+1}\beta} = f_{\alpha\sigma_\ell\sigma_{\ell+1}\beta}$, where $\alpha \in \mathcal{I}_{\ell-1}$, $\sigma_\ell \in K_\ell$, $\sigma_{\ell+1} \in K_{\ell+1}$ and $\beta \in J_{\ell+2}$. It describes the dependence of f_σ on the indices σ_ℓ and $\sigma_{\ell+1}$ exactly, but considers only one value of the other indices, namely the value specified by $\alpha \in \mathcal{I}_{\ell-1}$ and $\beta \in \mathcal{J}_{\ell+2}$. We fuse the index α with σ_ℓ and index $\sigma_{\ell+1}$ with β , such that Π_ℓ is reshaped into a matrix. Then, we perform an MCI

$$[\Pi_\ell]_{\alpha_{\ell-1}\sigma_\ell;\sigma_{\ell+1}\beta_{\ell+2}} \approx C_{\alpha_{\ell-1}\sigma_\ell;j} [P^{-1}]_{ji} R_{i;\sigma_{\ell+1}\beta_{\ell+1}}, \quad (\text{S3})$$

where C contains D_ℓ columns of Π_ℓ , R contains D_ℓ rows, and P their intersection points:

$$\Pi_\ell = \left(\begin{array}{cccccccc} \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \end{array} \right) \approx \left(\begin{array}{ccc} \circ & \bullet & \circ \\ \bullet & \circ & \bullet \\ \circ & \bullet & \circ \end{array} \right) \left(\begin{array}{ccc} \circ & \bullet & \circ \\ \bullet & \circ & \bullet \\ \circ & \bullet & \circ \end{array} \right)^{-1} \left(\begin{array}{cccccccc} \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \\ \bullet & \circ & \bullet & \circ & \bullet & \circ & \bullet & \circ \end{array} \right)$$

For a given bond dimension D_ℓ , the error is quasi-optimal if P is chosen such that $|\det P|$ is maximal (maximum volume principle) [25, 26]. We construct C, P and R iteratively, where in each iteration we add the row and column that maximize $|\det P|$. The iteration continues until the difference between the left- and right-hand side of Eq. (S3), $\varepsilon_\ell = \|\Pi_\ell - CP^{-1}R\|_\infty$, is smaller than a user-specified tolerance ϵ_{MCI} . This is equivalent to a fully pivoted rank-revealing LU decomposition, which is used to construct the MCI in practice [25].

Since Π_ℓ is a subset of the f_σ tensor, the choice of rows and columns in the MCI corresponds to choosing a subset of important multi-indices at bond ℓ . Therefore, we update \mathcal{I}_ℓ by including the set of rows that form R , and \mathcal{J}_ℓ by including the set of columns that form C , possibly increasing the number of elements in those sets. Concretely, the tensors $T_\ell, T_{\ell+1}$ are updated as

$$\begin{aligned} [T_\ell]_{\alpha_{\ell-1}\beta_\ell}^{\sigma_\ell} &\leftarrow C_{\alpha_{\ell-1}\sigma_\ell;\beta_\ell}, & [P_\ell^{-1}]_{\beta_\ell\alpha_\ell} &\leftarrow [P^{-1}]_{\beta_\ell\alpha_\ell}, \\ [T_{\ell+1}]_{\alpha_\ell\beta_{\ell+1}}^{\sigma_{\ell+1}} &\leftarrow R_{\alpha_\ell;\sigma_{\ell+1}\beta_{\ell+1}}. \end{aligned}$$

This preserves Eq. (S2), now with updated \mathcal{I}_ℓ and \mathcal{J}_ℓ .

Local optimizations are performed successively for each bond during the sweep, such that one full sweep corresponds to optimization of all \mathcal{I}_ℓ and \mathcal{J}_ℓ . The algorithm continues sweeping back and forth until the bond dimension $D_{\text{max}} = \max_\ell D_\ell = \max_\ell |\mathcal{I}_\ell|$ is converged, and a global error estimate $\varepsilon \approx \max_\sigma |f_\sigma^{\text{QTCI}} - f_\sigma|$ is smaller than a user-specified tolerance ϵ . The global error estimate is given by $\varepsilon = \max_\ell \varepsilon_\ell$, where ε_ℓ is the maximum absolute difference between the left- and right-hand-side of Eq. (S3). These local errors ε_ℓ are a good estimate of the global error because the multi-indices in \mathcal{I}_ℓ and \mathcal{J}_ℓ are *nested*, i.e. $\mathcal{I}_\ell \subseteq \mathcal{I}_{\ell-1} \otimes K_\ell$ and $\mathcal{J}_\ell \subseteq K_\ell \otimes \mathcal{J}_{\ell+1}$ [12, 16].

In practice, it is advantageous to choose the MCI tolerance ϵ_{MCI} slightly smaller than the TCI tolerance ϵ , and the initial pivot $\bar{\sigma}_\ell$ close to the largest structures in f_σ . An optimized library that was used for the examples in the main text will be released in the near future [27].