The Mutual-Information Bound for Matrix Product Density Operators

2026-07-31

Verdict: false-source (open).

1 The assertion

Proposition 4.5 of Cirac–Pérez-García–Schuch–Verstraete states that, for any translationally invariant matrix product density operator of bond dimension D,

$$I_L \leq I_{L+1}, \qquad \lim_{L\to\infty} \lim_{N\to\infty} I_L = I_\infty < \infty.$$

The Appendix proves monotonicity from strong subadditivity and obtains boundedness from

$$I_L \leq 4 \log D,$$

which it attributes to Wolf–Verstraete–Hastings–Cirac. [1]

Here an MPDO means a matrix product operator whose periodic finite-chain operators are positive semidefinite. This does not include a local purification among its hypotheses.

2 The cited argument

The construction in [WVHC08] defines a mixed projected entangled-pair state by applying a completely positive map at each site to virtual entangled pairs. They purify every local map, obtaining a pure projected entangled-pair state with an additional local physical system. Monotonicity of mutual information under the local ancillary traces then gives

$$I(A:B)_\rho \leq I(A:B)_{|\Psi\rangle\langle\Psi|} = 2S(A)_{|\Psi\rangle\langle\Psi|} \leq 2|\partial A| \log D.$$

[1] Source locations: `Papers/1606.00608/MPDO-22-12-17-2.tex`:795–809 and 1316–1321.

For an interval in a periodic one-dimensional chain, $|\partial A| = 2$, and this gives $I(A : B) \leq 4 \log D$.

In that argument, D is the dimension of each virtual system on an edge of the locally completely positive construction. It remains the virtual bond dimension after the local maps are purified. It is not the operator bond dimension of an arbitrary MPO representation. If the density operator of a locally purified state is written as an MPO by pairing the ket and bra virtual indices, its natural operator bond has dimension at most D^2. Thus the two uses of the letter D refer to different representations and cannot be identified without an additional positive factorization.

Thus the cited proof applies to a locally purified tensor. It does not prove that an arbitrary family of positive periodic matrix product operators admits such a purification. The distinction is explicit in the source paper itself: immediately before Definition 4.3 it warns that an MPDO need not possess the displayed local purification, and after Theorem 4.4 it repeats that the purification definition does not apply to every MPDO. [2]

Positivity of all periodic finite-chain operators does not supply the missing factorization. For finite multipartite states, [DlCSPGC13] proves that the bond dimension D' of a matrix product purification cannot be bounded in terms of the MPO bond dimension D alone. More strongly, [DlCCC+16] constructs translationally invariant MPDOs that are positive for every chain length but admit no translationally invariant purification valid for every chain length; the construction already occurs for classical states. These results rule out extending the cited proof to every MPDO merely by choosing a bond-controlled local purification. They do not, by themselves, disprove the mutual-information estimate.

3 Why the marginal-rank argument does not repair the proof

Cutting a periodic MPO at the two boundary bonds gives an operator-Schmidt decomposition with at most D^2 summands. This does not imply that either reduced density matrix has rank at most D^2.

Indeed, let q be a full-rank one-site density matrix and let the bond dimension be one, with scalar tensor entries

$$M^{ij} = q_{ij}.$$

[2]Source locations: `Papers/1606.00608/MPDO-22-12-17-2.tex`:744 and 786.

The N-site operator is $q^{\otimes N}$, and the reduced state on L sites is $q^{\otimes L}$. Consequently

$$\operatorname{rank}(q^{\otimes L}) = \operatorname{rank}(q)^L > 1 = D^2$$

for $L > 0$. Its mutual information is nevertheless zero. The cancellation of the extensive marginal entropies in

$$I_L = S_L + S_{N-L} - S_N$$

is therefore essential. Bounding the two marginal entropies separately by $2 \log D$ is not a valid argument for mixed MPOs.

The operator-Schmidt decomposition across the cut therefore supplies no marginal-rank bound by itself. Nevertheless, a different argument using the supported-marginal quantum channel and a weighted Hilbert–Schmidt contraction does control the mutual information directly. This argument is given in the next section. It is not the proof cited by Cirac–Pérez-García–Schuch–Verstraete.

4 Mutual information and operator-Schmidt rank

Let ρ_{AB} be a finite-dimensional density operator and set

$$r = \operatorname{OSR}(\rho_{AB}).$$

Compress both tensor factors to the supports of ρ_A and ρ_B. Positivity implies that ρ_{AB} is supported on the tensor product of these two subspaces. The compression and its inverse isometric inclusion preserve both mutual information and operator-Schmidt rank. We may therefore write

$$\sigma = \rho_A > 0, \qquad \tau = \rho_B > 0.$$

Choose an eigenbasis of σ. The supported-marginal state–channel correspondence gives a trace-preserving completely positive map Φ such that

$$\rho_{AB} = (\mathrm{id} \otimes \Phi)(|\Omega_\sigma\rangle\langle\Omega_\sigma|), \qquad \tau = \Phi(\sigma), \qquad \dim_{\mathbb{C}} \operatorname{ran} \Phi = r,$$

where

$$|\Omega_\sigma\rangle = \sum_i \sqrt{p_i} |i, i\rangle \quad \text{for} \quad \sigma = \operatorname{diag}(p_i).$$

The last equality follows because the invertible scaling by $\sqrt{p_i p_j}$ does not change the span of the operator blocks.

For $\Gamma_\omega^z(X) = \omega^{z/2} X \omega^{z/2}$, define

$$L = \Gamma_\tau^{-1/2} \circ \Phi \circ \Gamma_\sigma^{1/2}.$$

Equation (18) of [Bei13, Theorem 6] implies, at exponent two, that L is a Hilbert–Schmidt contraction. Hence every singular value of L is at most one. Since the two modular conjugations are invertible,

$$\operatorname{rank} L = \operatorname{rank} \Phi = r.$$

The order-two sandwiched whitening of ρ_{AB} is the unnormalized Choi matrix of L:

$$W = (\sigma^T \otimes \tau)^{-1/4} \rho_{AB} (\sigma^T \otimes \tau)^{-1/4} = \sum_{i,j} E_{ij} \otimes L(E_{ij}).$$

The matrix units form a Hilbert–Schmidt orthonormal basis, so

$$\operatorname{Tr}(W^2) = \sum_{i,j} \|L(E_{ij})\|_2^2 = \sum_k s_k(L)^2 \le r.$$

A direct endpoint comparison gives

$$D(\rho\|\omega) \le \log Q_2(\rho, \omega)$$

for trace-one $\rho, \omega \ge 0$ with $\ker \omega \subseteq \ker \rho$. For a faithful reference, set $R = \sqrt{\rho}$, $T = \omega^{-1/2}$, $\Delta = R^T \otimes T$, and $v = \operatorname{vec}(R)$. Under column-stacking,

$$(B \otimes A) \operatorname{vec}(X) = \operatorname{vec}(AXB^T),$$

so the factor order in Δ gives $\Delta v = \operatorname{vec}(T\rho)$ and

$$m = \operatorname{Re}\langle v, \Delta v \rangle = \langle R, RTR \rangle_{\mathrm{HS}}.$$

The support projection of Δ fixes v. The support projectors in

$$\log(A \otimes B) = (\log A) \otimes P_{\operatorname{supp} B} + P_{\operatorname{supp} A} \otimes (\log B)$$

are essential when an eigenvalue vanishes: with the convention $\log 0 = 0$, the identity with ambient identities in their place is false. Together with $\log \sqrt{\rho} = \frac{1}{2} \log \rho$, $\log(\omega^{-1/2}) = -\frac{1}{2} \log \omega$, and transpose covariance, this gives

$$\operatorname{Re}\langle v, (\log \Delta) v \rangle = \frac{1}{2} D(\rho\|\omega).$$

Support-aware Jensen yields $\frac{1}{2}D(\rho\|\omega) \le \log m$, while Hilbert–Schmidt Cauchy–Schwarz gives

$$m^2 \le \|R\|_2^2\|RTR\|_2^2 = Q_2(\rho,\omega).$$

The auxiliary divergence and vector-state Jensen mechanism are adapted from [MLDS+13, arXiv:1306.3142v4, Definition 5 and Lemma 19, used in the proof of Theorem 7]. Only the direct order-two endpoint is used here, followed by the endpoint Cauchy–Schwarz estimate; the shortened proof is not Beigi's interpolation argument. Compression to the support of ω proves the singular-reference statement. Applying it to $\omega = \rho_A \otimes \rho_B$ therefore gives

$$I(A:B)_\rho = D(\rho_{AB}\|\rho_A \otimes \rho_B) \le \log \mathrm{Tr}(W^2) \le \log r.$$

Thus the coefficient-one estimate holds for every finite-dimensional bipartite density operator. It is sharp: the uniform perfectly correlated state on r classical symbols has operator-Schmidt rank r and mutual information $\log r$.

The supported-marginal state–channel correspondence and its exact range dimension are formalized.[3] The Choi-matrix estimate is also formalized.[4] The exponent-two weighted Hilbert–Schmidt contraction is formalized both for positive-definite output weights and for arbitrary positive-semidefinite output weights, with the negative quarter power taken on the output support.[5] The full-support eigenbasis calculation is formalized: the whitening is identified exactly with the rectangular Choi matrix, including the input transpose convention, the modular congruences preserve the range dimension, and the resulting operator is positive semidefinite with its trace square and squared Frobenius norm bounded by operator-Schmidt rank.[6] The single-reference support reduction is now formalized. Let

[3] Formal declarations: `Matrix.supportedMarginalMap_isKrausCPTP`, `Matrix.supportedMarginalReconstruction_eq`, and `Matrix.finrank_range_supportedMarginalMap` in `TNLean/Channel/SupportedMarginalChannel`.

[4] Formal declarations: `Matrix.rectangularChoi_frobenius_parseval` and `Matrix.rectangularChoi_frobenius_sq_le_finrank_range` in `TNLean/Algebra/RectangularChoi`.

[5] Formal declarations: `Matrix.weightedHilbertSchmidtMap_norm_le`, `Matrix.weightedHilbertSchmidtMap_isHilbertSchmidtContraction`, `Matrix.supportedWeightedHilbertSchmidtMap_norm_le`, and `Matrix.supportedWeightedHilbertSchmidtMap_isHilbertSchmidtContraction` in `TNLean/Channel/WeightedHilbertSchmidt`.

[6] Formal declarations: `Matrix.supportedMarginalWhitenedState_eq_rectangularChoi`, `Matrix.finrank_range_weightedHilbertSchmidtMap`, and `Matrix.supportedMarginalWhitenedState_trace_sq_re_le_operatorSchmidtRank` in `TNLean/Channel/WhitenedChoi`.

$\rho, \omega \geq 0$ satisfy $\ker \omega \subseteq \ker \rho$, and let $V : \mathbb{C}^k \to \mathbb{C}^D$ obey $V^\dagger V = I_k$ and $VV^\dagger = P_{\operatorname{supp} \omega}$. Then

$$D(\rho\|\omega) = D(V^\dagger \rho V \| V^\dagger \omega V), \tag{1}$$

$$Q_2(\rho, \omega) = Q_2(V^\dagger \rho V, V^\dagger \omega V). \tag{2}$$

Here the logarithm and negative quarter-power vanish on the zero eigenspace. The compressed reference is positive definite. If ρ and ω have trace one, isometric trace preservation gives the same normalization after compression. Consequently, the singular-reference inequality reduces to the faithful inequality

$$D(A\|B) \leq \log Q_2(A, B), \qquad A \geq 0, \quad B > 0, \quad \operatorname{tr} A = 1.$$

This faithful comparison is now proved; no normalization of B is required. The rectangular isometric-conjugation identities are formalized. [7] Faithful compression of a positive semidefinite matrix onto its support is formalized by `Matrix.PosSemidef.compression_on_support_posDef` in `TNLean/Analysis/SupportCompression`. The order-two functional is `TNLean.sandwichedRenyiTwoTrace` in `TNLean/Analysis/SandwichedRenyiTwo`. The reconstruction and entropy identities are formalized separately. [8]

A support-aware vector-state Jensen inequality for the logarithm is also formalized. If $A \geq 0$, $\langle v, v\rangle = 1$, and the support projection of A fixes v, then

$$\operatorname{Re}\langle v, (\log A)v\rangle \leq \log(\operatorname{Re}\langle v, Av\rangle)\,.$$

In an eigenbasis of A, the support condition makes every weight at a zero eigenvalue vanish. Scalar log concavity is therefore applied only to positive eigenvalues; neither concavity at zero nor positive definiteness of A is assumed. The proof uses shared spectral-weight formulas for quadratic forms

[7] Formal declarations: `Matrix.isometryConjNonUnitalStarAlgHom`, `Matrix.isometryConjNonUnitalStarAlgHom_apply`, `Matrix.cfc_conj_isometry_of_zero`, `Matrix.trace_mul_mul_conjTranspose_of_conjTranspose_mul_eq_one`, and `Matrix.conjTranspose_mul_mul_mul_conjTranspose_mul_of_isometry` in `TNLean/Analysis/IsometricCompression`.

[8] Formal declarations: `Matrix.PosSemidef.eq_expansion_compression_of_kernel_le_of_mul_conjTranspose_eq_supportProj`, `TNLean.quantumRelativeEntropy_support_compression`, `TNLean.sandwichedRenyiTwoTrace_support_compression`, and `TNLean.quantumRelativeEntropy_le_log_sandwichedRenyiTwoTrace_of_faithful` in `TNLean/Analysis/SupportCompressedEntropy`.

of a Hermitian matrix and its functional calculus. [9] This is an auxiliary analytic result, not a theorem of CPSV16. It does not alter the source-coverage count or complete Proposition 4.5.

The reusable functional-calculus identities and the entropy comparison are now formalized. Functional calculus commutes with transpose on Hermitian matrices; in particular, so do the logarithm and support projection, and every real power and the positive square root commute with transpose on positive-semidefinite matrices. These transpose identities include singular matrices and zero-dimensional matrix algebras. The tensor-product logarithm retains the two support projections, the logarithm of a positive-semidefinite square root is half the logarithm, and the logarithm of a real power of a positive-definite matrix is the exponent times the logarithm. The faithful theorem assumes only $\rho \geq 0$, $\omega > 0$, and $\operatorname{tr}\rho = 1$; the singular theorem assumes $\rho, \omega \geq 0$, both traces equal to one, and $\ker\omega \subseteq \ker\rho$. These functional-calculus identities are project-derived auxiliary results, not theorems of CPSV16; in particular, the tensor-product statement permits singular factors and zero-dimensional matrix algebras. [10] The one-sided output-support compression is also formalized: faithfulness of the input weight forces every channel output into the support of its image, the compressed channel is trace preserving, and the support-restricted negative quarter power agrees with the ordinary negative quarter power after compression. Simultaneous compression to both marginal supports is used in the arbitrary-support estimate below.

The unnormalized sandwiched Rényi trace term

$$\widetilde{Q}_\alpha(\rho\|\omega) = \operatorname{Re}\operatorname{Tr}\Big[\Big(\omega^{(1-\alpha)/(2\alpha)}\rho\,\omega^{(1-\alpha)/(2\alpha)}\Big)^{\alpha}\Big]$$

is now formalized. For trace-one ρ and $\alpha > 1$, this is the quantity inside the logarithm of the sandwiched Rényi divergence when $\rho, \omega \geq 0$ and

[9] Formal declarations: `Matrix.PosSemidef.re_dotProduct_cfc_log_mulVec_le_log` in `TNLean/Analysis/SupportLogJensen`, and `Matrix.IsHermitian.spectralWeight`, `Matrix.IsHermitian.spectralWeight_nonneg`, `Matrix.IsHermitian.sum_spectralWeight`, `Matrix.IsHermitian.re_dotProduct_cfc_mulVec_eq_sum`, and `Matrix.IsHermitian.re_dotProduct_mulVec_eq_sum` in `TNLean/Analysis/SpectralQuadraticForm`.

[10] Formal declarations: `Matrix.cfc_transpose`, `Matrix.IsHermitian.log_transpose`, `Matrix.PosSemidef.rpow_transpose`, `Matrix.PosSemidef.sqrt_transpose`, `Matrix.IsHermitian.supportProj_transpose`, `Matrix.log_kronecker_posSemidef`, `Matrix.PosSemidef.cfc_log_sqrt`, `Matrix.PosDef.cfc_log_rpow`, `TNLean.quantumRelativeEntropy_le_log_sandwichedRenyiTwoTrace_posDef`, and `TNLean.quantumRelativeEntropy_le_log_sandwichedRenyiTwoTrace` in `TNLean/Analysis/CfcConjugation`, `TNLean/Analysis/CfcKronecker`, `TNLean/Analysis/MatrixSqrt`, `TNLean/Analysis/SandwichedRenyiTwo`, and `TNLean/Analysis/SupportCompressedEntropy`.

$\ker \omega \subseteq \ker \rho$. This includes singular ω: the zero-on-kernel powers implement the generalized inverse on its support. In this regime, finiteness of the divergence requires the support inclusion. If the inclusion fails, the totalized continuous-functional-calculus value remains finite but is not the divergence trace term. The present application uses $1 < \alpha \leq 2$; the order-one case is the Umegaki endpoint. The formalized properties are nonnegativity for $\rho \geq 0$ and $\omega > 0$ and the exact identities at $\alpha = 1$ and $\alpha = 2$. [11] These are supporting analytic identities, not a theorem of Cirac–Pérez-García–Schuch–Verstraete. They prove neither continuity nor monotonicity in α, interpolation between orders one and two, differentiability or a limit at $\alpha = 1$, the logarithmic sandwiched Rényi divergence, nor a comparison with Umegaki relative entropy.

Simultaneous compression by the two marginal support isometries, exact reconstruction, preservation of ordinary operator-Schmidt rank, the exact formulas for the compressed marginals, and their positive definiteness are formalized. [12] The mathematical tasks tracked in #5209 and #5242 are complete: the faithful and singular relative-entropy comparisons are proved using the relative-modular half-moment, support-aware Jensen, and Hilbert–Schmidt Cauchy–Schwarz, without a general weighted interpolation theorem or an order-one limit.

The faithful marginal-whitening theorem now diagonalizes the first marginal, where the input transpose from the Choi convention agrees with the diagonal marginal, and proves

$$Q_2(\rho_{AB}, \rho_A \otimes \rho_B) \leq \mathrm{OSR}(\rho_{AB})$$

when both marginals are positive definite. The arbitrary-support theorem compresses simultaneously to the two marginal supports, applies the faithful theorem there, and returns using exact preservation of Q_2 and ordinary operator-Schmidt rank. It includes zero-dimensional support coordinates and requires no trace normalization. [13] Thus the mathematical work tracked

[11] Formal declarations: `TNLean.sandwichedRenyiTrace`, `TNLean.sandwichedRenyiTrace_nonneg`, `TNLean.sandwichedRenyiTrace_one`, and `TNLean.sandwichedRenyiTrace_two` in `TNLean/Analysis/SandwichedRenyi`.

[12] Formal declarations: `Matrix.PosSemidef.eq_marginalSupport_expansion_compression`, `Matrix.PosSemidef.operatorSchmidtRank_marginalSupport_compression`, `Matrix.PosSemidef.partialTraceRight_marginalSupport_compression`, `Matrix.PosSemidef.partialTraceLeft_marginalSupport_compression`, and `Matrix.PosSemidef.marginalSupport_compression_marginals_posDef` in `TNLean/Channel/MarginalSupportAbsorption`.

[13] Formal declarations: `TNLean.sandwichedRenyiTwoTrace_product_marginals_le_op`

by #5211 is complete. The unrestricted coefficient-one inequality is also formalized:

$$I(A:B)_\rho \leq \log \mathrm{OSR}(\rho_{AB})$$

for every positive semidefinite bipartite operator of trace one, with ordinary operator-Schmidt rank and without positive-dimension or faithful-marginal assumptions. The proof identifies mutual information with relative entropy against the product of the marginals, applies the singular-reference relative-entropy comparison and the order-two whitening bound, and treats the zero value of the totalized order-two functional separately before applying monotonicity of the logarithm. [14] Hence the mathematical work tracked by #5212 is complete. Its application to the finite-chain MPO estimate is formalized by `MPOTensor.mutualInfoChain_le_two_log_bondDim`, and the source-facing coefficient-four consequence is formalized by `MPOTensor.IsMPDO.mutualInfoChain_le_four_log_bondDim`, both in `TNLean/MPS/MPDO/MutualInfoAreaLaw`. Thus the finite-chain work tracked by #5213 is complete. This does not formalize the false thermodynamic-limit clause of Proposition 4.5 or alter its separate negative status.

5 Diagonal matrix product operators

No diagonal tensor can violate the proposed estimate. Fix a chain of length $N = L + R$ and write its diagonal coefficient matrix as

$$P_{x,y} = \mathrm{tr}\big(M^{x_0x_0} \cdots M^{x_{L-1}x_{L-1}} M^{y_0y_0} \cdots M^{y_{R-1}y_{R-1}}\big),$$

where x and y label configurations on the two consecutive blocks. Set $A_x = M^{x_0x_0} \cdots M^{x_{L-1}x_{L-1}}$ and $B_y = M^{y_0y_0} \cdots M^{y_{R-1}y_{R-1}}$. Then

$$P_{x,y} = \mathrm{tr}(A_x B_y) = \sum_{a,b=0}^{D-1} (A_x)_{ab}(B_y)_{ba}.$$

`eratorSchmidtRank_of_marginals_posDef` in `TNLean/Channel/FaithfulMarginalWhitenedChoi`, and `TNLean.sandwichedRenyiTwoTrace_product_marginals_le_operatorSchmidtRank` in `TNLean/Channel/MarginalSupportWhitenedChoi`.

[14] Formal declarations: `Matrix.operatorSchmidtRank_pos_of_ne_zero`, `Matrix.PosSemidef.productMarginals_kernel_le`, `TNLean.sandwichedRenyiTwoTrace_submatrix_equiv`, and `Entropy.mutualInformation_le_log_operatorSchmidtRank` in `TNLean/Algebra/OperatorSchmidt`, `TNLean/Channel/MarginalSupportAbsorption`, `TNLean/Analysis/SandwichedRenyiTwo`, and `TNLean/Entropy/MutualInformationOperatorSchmidt`.

Consequently P factors through a complex vector space of dimension D^2, so $\operatorname{rank}_{\mathbb{C}} P \leq D^2$. Suppose that the finite-chain operator is positive semidefinite and diagonal, with positive trace $Z = \operatorname{tr} \rho^{(N)}(M)$. Then the joint probability matrix is

$$\widehat{P} = Z^{-1} P.$$

Multiplication by the nonzero real scalar Z^{-1} does not change either rank. Since $\widehat{P}$ is real, $\operatorname{rank}_{\mathbb{R}} \widehat{P} = \operatorname{rank}_{\mathbb{C}} \widehat{P} \leq D^2$.

Riazanov–Vyalyi prove for every finite joint probability matrix that [RV17, Theorem 4.1]

$$I(X:Y) \leq \log \operatorname{rank}_{\mathbb{R}} \widehat{P}.$$

Their rank is the ordinary real matrix rank, not the nonnegative rank. Their proof eliminates linearly dependent rows while preserving nonnegativity and the column marginals, choosing at each step the endpoint at which mutual information does not decrease. The final distribution has at most $\operatorname{rank}_{\mathbb{R}} P$ nonzero rows. It follows that every diagonal periodic MPO satisfies the stronger finite-chain estimate

$$I_L^{(N)} \leq 2 \log D.$$

This conclusion needs positivity and normalization only at the chain length under consideration; positivity at all other lengths is unnecessary. It gives an independent proof of the diagonal case. The coefficient-one theorem of the preceding section rules out both classical and genuinely quantum counterexamples to the unrestricted finite-chain estimate.

5.1 Local correction at the row-elimination endpoints

The displayed calculation in the proof of [RV17, Theorem 4.1, Section 4.2] cancels the factor $1 + \varepsilon\alpha_i$. In the source's notation, $m_{ij} = P_{x,y}$ is a joint-probability entry, $p_i = \sum_j m_{ij}$ is its row marginal, and $\alpha_i = \beta_x$ is the row-perturbation coefficient. The cancelled factor occurs between $m_{ij}(1 + \varepsilon\alpha_i)$ and $p_i(1+\varepsilon\alpha_i)$. The proof then chooses an endpoint of the admissible interval, precisely where this factor vanishes for at least one row. At such a row the displayed quotient has the indeterminate form $0/0$, so the cancellation does not justify the endpoint identity as written.

This is a local proof correction; the statement of the theorem is unchanged. Write mutual information as $H(X)+H(Y)-H(X,Y)$ and define $h(t) = -t \log t$ for $t > 0$, with $h(0) = 0$. Then

$$h(ab) = a\, h(b) + b\, h(a). \tag{3}$$

This identity is valid when either factor is zero. The extra terms cancel against the row-marginal entropy because each row marginal is the sum of its entries, while the column marginal is fixed by the row relation. Thus mutual information is affine on the whole closed admissible interval, including both endpoints. This removes the endpoint singularity without strengthening the hypotheses or weakening the conclusion, so the discrepancy is classified as a local correction. [15]

The rank factorization is formalized. [16] The information-versus-rank theorem of Riazanov–Vyalyi is formalized by `Entropy.classicalMutualInformation_le_log_rank` in `TNLean/Entropy/ClassicalMutualInformation`. Real scalar normalization preserves rank by `Matrix.rank_smul_of_mem_nonZeroDivisors`. Faithful scalar extension preserves matrix rank by `Matrix.rank_map_algebraMap` in `TNLean/Algebra/MatrixRankBaseChange`. The resulting finite-chain estimate `MPOTensor.diagonalCut_classicalMutualInformation_le_two_log` is in `TNLean/MPS/MPDO/DiagonalCutRank`. The estimate assumes a real coefficient matrix, its normalized joint distribution, and an identification of its complex scalar extension with the diagonal cut matrix. The fixed-length construction is formalized in `TNLean/MPS/MPDO/DiagonalFiniteChain`. [17] For a positive diagonal operator at the chosen length, the real diagonal coefficients are nonnegative, their scalar extension is the cut matrix, and division by their positive total mass gives a joint distribution. Diagonality identifies these normalized coefficients with the full classical finite-chain distribution. This formal conclusion is classical and therefore does not by itself settle the unrestricted quantum question. The source-independent argument in the preceding section proves the unrestricted quantum estimate mathematically; its Lean formalization remains separate from this diagonal special case.

[15] The corrected calculation is used in the proof of `Entropy.classicalMutualInformation_le_log_rank` in `TNLean/Entropy/ClassicalMutualInformation`. Its zero-valid product identity is `Real.negMulLog_mul`.

[16] Formal declarations: `MPOTensor.diagonalCutMatrix` and `MPOTensor.diagonalCutMatrix_rank_le` in `TNLean/MPS/MPDO/DiagonalCutRank`.

[17] Formal declarations: `MPOTensor.IsDiagonalAt`, `MPOTensor.IsDiagonal`, `MPOTensor.diagonalCutRealMatrix`, `MPOTensor.diagonalCutMass`, `MPOTensor.normalizedDiagonalCutDistribution`, `MPOTensor.diagonalCutRealMatrix_map_ofReal_of_posSemidef`, `MPOTensor.diagonalCutRealMatrix_nonneg_of_posSemidef`, `MPOTensor.diagonalCutMass_eq_trace_re`, `MPOTensor.normalizedDiagonalCutDistribution_isJointDistribution`, and `MPOTensor.diagonalFiniteChain_classicalMutualInformation_le_two_log`. The global diagonal-MPDO corollary is `MPOTensor.diagonalMPDO_classicalMutualInformation_le_two_log`.

6 What the rank-three separation establishes

De las Cuevas–Schuch–Pérez-García–Cirac give a rank-three separation in [DlCSPGC13, Result 1 and equations (10)–(14)]; it is also summarized in [DlCDN19, Section 4.3, equation (13), and Table 2]. The examples are diagonal states

$$\rho_t = \sum_{i,j} (M_t)_{ij} |i,j\rangle\langle i,j|,$$

where the nonnegative matrix M_t has ordinary rank three while its positive semidefinite rank diverges. Under their correspondence, this says that ρ_t has operator-Schmidt rank three while its purification rank, and hence also its separable rank, diverges.

This is a statement about the size of exact factorizations, not a lower bound on quantum mutual information. In the same paper, the entanglement of purification is bounded above by the logarithm of the purification rank. A diverging purification rank makes this upper bound nonuniform, but does not imply that the entanglement of purification itself diverges. No lower bound on the entanglement of purification or on quantum mutual information is proved for the rank-three family.

Indeed, after normalizing M_t to a joint probability distribution, the quantum mutual information of the diagonal state is exactly the classical mutual information of that distribution. The ordinary-rank bound [RV17, Theorem 4.1] gives

$$I(X : Y) \leq \log \operatorname{rank}(M_t) = \log 3.$$

This theorem concerns ordinary matrix rank, rather than nonnegative rank. Its proof repeatedly eliminates a linearly dependent row while preserving nonnegativity and the column marginals, choosing the endpoint for which mutual information does not decrease. The final joint distribution has at most $\operatorname{rank}(M_t)$ nonzero rows, and hence its mutual information is at most the entropy of a random variable with that many values. Thus the rank-three examples have uniformly bounded mutual information. They show that a small MPO bond does not control the bond dimension of an exact local purification; they neither prove nor disprove the unrestricted $4 \log D$ mutual-information estimate.

7 The iterated limit is false without an aperiodicity hypothesis

The normalized periodic states need not converge on local algebras. Consider the physical dimension two, bond dimension three tensor whose only nonzero physical slices are

$$M^{00} = \mathrm{diag}(1,0,0), \qquad M^{11} = \mathrm{diag}(0,1,-1).$$

For every positive chain length, contraction of the virtual trace gives

$$\rho^{(N)}(M) = |0^N\rangle\langle 0^N| + \big(1 + (-1)^N\big)|1^N\rangle\langle 1^N|.$$

Thus the operator is positive semidefinite for every $N \geq 1$, including odd lengths, and

$$\operatorname{tr}\rho^{(N)}(M) = 2 + (-1)^N \neq 0.$$

This is therefore an MPDO at exactly the hypothesis set of Proposition 4.5.

At odd lengths the normalized state is

$$\widehat{\rho}^{(N)}(M) = |0^N\rangle\langle 0^N|,$$

whereas at positive even lengths it is

$$\widehat{\rho}^{(N)}(M) = \frac{1}{3}|0^N\rangle\langle 0^N| + \frac{2}{3}|1^N\rangle\langle 1^N|.$$

In particular, the all-zero entry of the one-site reduced state alternates between 1 and 1/3, so the normalized local states have no thermodynamic limit.

The mutual information itself also oscillates. For any fixed $L \geq 1$ and all $N > L$, the odd-length state is a product state and hence

$$I_L^{(N)} = 0.$$

For even N, both nonempty marginals and the full state have the two nonzero eigenvalues 1/3 and 2/3. Writing

$$h_2(1/3) = -\frac{1}{3}\log\frac{1}{3} - \frac{2}{3}\log\frac{2}{3} > 0,$$

one obtains

$$I_L^{(N)} = h_2(1/3) + h_2(1/3) - h_2(1/3) = h_2(1/3).$$

Hence $\lim_{N\to\infty} I_L^{(N)}$ does not exist. The eigenvalue -1 of the physical-trace transfer is the peripheral sector responsible for the parity oscillation. Primitivity or another aperiodicity assumption would exclude this example, but no such assumption occurs in Proposition 4.5.

The finite-chain form, positivity, trace formula, alternating normalized states and marginals, their exact entropies and mutual informations, and the nonconvergence of both the one-site entry and mutual information are formalized. [18]

Accordingly, the precise unconditional result that remains from the cited argument is the finite-chain monotonicity

$$I_L^{(N)} \le I_{L+1}^{(N)} \qquad (2L+1 \le N).$$

8 Resolution

The finite-chain analytic implication used by the cited proof is now formalized.[19] If the normalized state across a cut is the image of a normalized pure matrix product state of bond dimension D' under independent local trace-preserving completely positive maps, then data processing and the pure-state boundary estimate give

$$I_L \le 4 \log D'.$$

The local-purification step is now complete.[20] For explicit matrices $A^{(i,k)}$ witnessing the locally purified form, the two blockwise ancillary traces of the

[18] Formal declarations: `MPOTensor.ThermodynamicLimitCounterexample.mpo_parityTensor_eq_diagonal`, `MPOTensor.ThermodynamicLimitCounterexample.parityTensor_isMPDO`, `MPOTensor.ThermodynamicLimitCounterexample.trace_mpo_parityTensor`, `MPOTensor.ThermodynamicLimitCounterexample.reducedBlockState_one_zero_zero_of_odd`, `MPOTensor.ThermodynamicLimitCounterexample.reducedBlockState_one_zero_zero_of_even`, `MPOTensor.ThermodynamicLimitCounterexample.reducedBlockState_parityTensor_of_even`, `MPOTensor.ThermodynamicLimitCounterexample.blockEntropy_parityTensor_of_even`, `MPOTensor.ThermodynamicLimitCounterexample.mutualInfoChain_parityTensor_of_odd`, `MPOTensor.ThermodynamicLimitCounterexample.mutualInfoChain_parityTensor_of_even`, `MPOTensor.ThermodynamicLimitCounterexample.not_exists_tendsto_reducedBlockState_one_zero_zero`, and `MPOTensor.ThermodynamicLimitCounterexample.not_exists_tendsto_parityMutualInfoAfterCut` in `TNLean/MPS/MPDO/ThermodynamicLimitCounterexample`.

[19] Formal declaration: `MPOTensor.mutualInfoChain_le_of_bipartitioned_channel_image` in `TNLean/MPS/MPDO/LocalPurificationAreaLaw`.

[20] Formal declarations: `MPOTensor.bipartitionedNormalizedMPO_eq_tensorMapBoth_blockAncillaryTraceMap`, `MPOTensor.mutualInfoChain_le_of_lpdoWitness`, and `MPOTensor.IsLPDO.exists_mutualInfoChain_le_purifyingBond` in `TNLean/MPS/MPDO/LocalPurificationAreaLaw`.

normalized purifying pure state equal the normalized operator generated by M. Hence, on every nonempty chain, a locally purified tensor with nonzero finite-chain trace satisfies

$$I_L \leq 4 \log D',$$

where D' is the bond dimension of the purifying matrix product state. The bond-space identification in the locally purified form identifies D with $(D')^2$ at the level of cardinalities, but the proved estimate is stated in D', not in the MPO bond dimension D.

This result does not supply a local purification for an arbitrary positive MPO. The source itself warns at the purification step that such a representation need not exist, and the no-go results above show that this missing representation cannot be recovered uniformly from MPO positivity and bond dimension.

The finite-chain estimate is mathematically valid for arbitrary positive periodic MPO families of bond dimension D at every chain length for which the periodic operator has positive trace. After normalization, opening the two virtual boundary bonds gives

$$\mathrm{OSR}(\widehat{\rho}^{(N)}_{L:N-L}) \leq D^2,$$

and the coefficient-one theorem above gives the stronger conclusion

$$I_L^{(N)} \leq 2 \log D.$$

The operator-Schmidt rank estimate is formalized.[21] The coefficient-one entropy implication is formalized by `Entropy.mutualInformation_le_log_operatorSchmidtRank` in `TNLean/Entropy/MutualInformationOperatorSchmidt`. The resulting finite-chain estimate and its coefficient-four source-facing consequence are formalized by `MPOTensor.mutualInfoChain_le_two_log_bondDim` and `MPOTensor.IsMPDO.mutualInfoChain_le_four_log_bondDim` in `TNLean/MPS/MPDO/MutualInfoAreaLaw`.

The thermodynamic-limit clause of Proposition 4.5 is false at its stated hypothesis set, while the finite-chain bound for every positive periodic operator of nonzero trace is formalized. The remaining negative conclusion is tracked by issue #4169 under the MPDO coverage audit #2380.

[21] Formal declaration: `MPOTensor.operatorSchmidtRank_bipartitionedNormalizedMPO_le` in `TNLean/MPS/MPDO/MutualInfoBridge`.

References

[Bei13] Salman Beigi. Sandwiched rényi divergence satisfies data processing inequality. *Journal of Mathematical Physics*, 54(12):122202, 2013.

[DlCCC+16] Gemma De las Cuevas, Toby S. Cubitt, J. Ignacio Cirac, Michael M. Wolf, and David Pérez-García. Fundamental limitations in the purifications of tensor networks. *Journal of Mathematical Physics*, 57(7):071902, 2016.

[DlCDN19] Gemma De las Cuevas, Tom Drescher, and Tim Netzer. Separability for mixed states with operator schmidt rank two. *Quantum*, 3:203, 2019.

[DlCSPGC13] Gemma De las Cuevas, Norbert Schuch, David Pérez-García, and J. Ignacio Cirac. Purifications of multipartite states: Limitations and constructive methods. *New Journal of Physics*, 15:123021, 2013.

[MLDS+13] Martin Müller-Lennert, Frédéric Dupuis, Oleg Szehr, Serge Fehr, and Marco Tomamichel. On quantum rényi entropies: A new generalization and some properties. *Journal of Mathematical Physics*, 54(12):122203, 2013.

[RV17] Andrii Riazanov and Mikhail Vyalyi. Exploring the bounds on the positive semidefinite rank. 2017.

[WVHC08] Michael M. Wolf, Frank Verstraete, Matthew B. Hastings, and J. Ignacio Cirac. Area laws in quantum systems: Mutual information and correlations. *Physical Review Letters*, 100(7):070502, 2008.