The Mutual-Information Bound for Matrix Product Density Operators

2026-07-31

Verdict: false-source (resolved).

1 The assertion

Proposition 4.5 of Cirac–Pérez-García–Schuch–Verstraete [CPGSV16] states that any translationally invariant matrix product density operator of bond dimension D satisfies

$$I_L \leq I_{L+1}, \tag{1}$$

$$\lim_{L\to\infty} \lim_{N\to\infty} I_L = I_\infty < \infty. \tag{2}$$

The Appendix derives (1) from strong subadditivity and obtains boundedness from

$$I_L \leq 4 \log D, \tag{3}$$

which it attributes to Wolf–Verstraete–Hastings–Cirac [WVHC08].[1]

Here an MPDO is a matrix product operator whose periodic finite-chain operators are positive semidefinite. This definition does not require a local purification.

2 The cited argument

Wolf–Verstraete–Hastings–Cirac construct a mixed projected entangled-pair state by applying a completely positive map at each site to virtual entangled

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

pairs [WVHC08]. 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

$$I(A:B)_\rho \le I(A:B)_{|\Psi\rangle\langle\Psi|}, \tag{4}$$

$$I(A:B)_{|\Psi\rangle\langle\Psi|} = 2S(A)_{|\Psi\rangle\langle\Psi|}, \tag{5}$$

$$2S(A)_{|\Psi\rangle\langle\Psi|} \le 2|\partial A| \log D. \tag{6}$$

For an interval in a periodic one-dimensional chain, $|\partial A| = 2$, so (4)–(6) imply

$$I(A:B)_\rho \le 4 \log D. \tag{7}$$

In this argument, D is the dimension of each virtual edge system in the locally completely positive construction. It remains the virtual bond dimension after purification of the local maps. 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 dimension is at most D^2. The two uses of D therefore refer to different representations and cannot be identified without an additional positive factorization.

The cited proof applies to a locally purified tensor. It does not show that an arbitrary family of positive periodic matrix product operators admits such a purification. Cirac–Pérez-García–Schuch–Verstraete explicitly warn immediately before Definition 4.3 that an MPDO need not possess the displayed local purification; after Theorem 4.4, they repeat that the purification definition does not apply to every MPDO [CPGSV16]. [2]

Positivity of all periodic finite-chain operators does not supply the missing factorization. De las Cuevas–Schuch–Pérez-García–Cirac prove that, for finite multipartite states, the bond dimension D' of a matrix product purification cannot be bounded in terms of the MPO bond dimension D alone [DlCSPGC13]. More strongly, de las Cuevas–Cubitt–Cirac–Wolf–Pérez-García construct translationally invariant MPDOs that are positive at every chain length but admit no translationally invariant purification valid at every chain length; their construction already applies to classical states [DlCCC+16]. These results prevent the cited proof from being extended to every MPDO by choosing a bond-controlled local purification. They do not, by themselves, disprove the mutual-information estimate.

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

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 decomposition does not imply that either reduced density matrix has rank at most D^2.

Let q be a full-rank one-site density matrix and consider the bond-dimension one tensor with scalar entries

$$M^{ij} = q_{ij}. \tag{8}$$

The N-site operator and its reduction to L sites are

$$\rho^{(N)} = q^{\otimes N}, \tag{9}$$

$$\rho_L = q^{\otimes L}. \tag{10}$$

Consequently,

$$\mathrm{rank}(q^{\otimes L}) = \mathrm{rank}(q)^L > 1 = D^2 \tag{11}$$

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

$$I_L = S_L + S_{N-L} - S_N \tag{12}$$

is therefore essential. Bounding the two marginal entropies separately by $2 \log D$ is invalid for mixed MPOs.

The operator-Schmidt decomposition alone supplies no marginal-rank bound. A different argument, based on the supported-marginal quantum channel and a weighted Hilbert–Schmidt contraction, controls the mutual information directly. The next section gives this argument; it is not the proof cited by Cirac–Pérez-García–Schuch–Verstraete [CPGSV16].

4 Mutual information and operator-Schmidt rank

Let ρ_{AB} be a finite-dimensional density operator, and define its operator-Schmidt rank by

$$r = \mathrm{OSR}(\rho_{AB}). \tag{13}$$

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

$$\sigma = \rho_A > 0, \tag{14}$$

$$\tau = \rho_B > 0. \tag{15}$$

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

$$\rho_{AB} = (\mathrm{id} \otimes \Phi)(|\Omega_\sigma\rangle\langle\Omega_\sigma|), \tag{16}$$

$$\tau = \Phi(\sigma), \tag{17}$$

$$\dim_{\mathbb{C}} \operatorname{ran} \Phi = r, \tag{18}$$

where

$$|\Omega_\sigma\rangle = \sum_i \sqrt{p_i}|i,i\rangle, \tag{19}$$

$$\sigma = \operatorname{diag}(p_i). \tag{20}$$

The equality in (18) follows because the invertible scaling by $\sqrt{p_i p_j}$ preserves the span of the operator blocks.

For a positive-definite matrix ω, define

$$\Gamma_\omega^z(X) = \omega^{z/2} X \omega^{z/2}, \tag{21}$$

and set

$$L = \Gamma_\tau^{-1/2} \circ \Phi \circ \Gamma_\sigma^{1/2}. \tag{22}$$

Equation (18) in the proof of Theorem 6 of Ref. [Bei13] implies at exponent two that L is a Hilbert–Schmidt contraction. Every singular value of L is therefore at most one. Since the two modular congruences are invertible,

$$\operatorname{rank} L = \operatorname{rank} \Phi = r. \tag{23}$$

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

$$W = (\sigma^T \otimes \tau)^{-1/4} \rho_{AB} (\sigma^T \otimes \tau)^{-1/4}, \tag{24}$$

$$W = \sum_{i,j} E_{ij} \otimes L(E_{ij}). \tag{25}$$

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

$$\operatorname{tr}(W^2) = \sum_{i,j} \|L(E_{ij})\|_2^2, \tag{26}$$

$$\sum_{i,j} \|L(E_{ij})\|_2^2 = \sum_k s_k(L)^2, \tag{27}$$

$$\sum_k s_k(L)^2 \leq r. \tag{28}$$

For positive-semidefinite ρ and ω satisfying $\ker\omega \subseteq \ker\rho$, define the order-two sandwiched functional by

$$Q_2(\rho,\omega) = \operatorname{Re}\operatorname{tr}\left[\left(\omega^{-1/4}\rho\omega^{-1/4}\right)^2\right], \tag{29}$$

where the negative power vanishes on the kernel of ω. A direct endpoint comparison gives

$$D(\rho\|\omega) \leq \log Q_2(\rho,\omega) \tag{30}$$

for trace-one $\rho, \omega \geq 0$ satisfying the same support condition.

To prove the faithful-reference case, define

$$R = \sqrt{\rho}, \tag{31}$$

$$T = \omega^{-1/2}, \tag{32}$$

$$\Delta = R^T \otimes T, \tag{33}$$

$$v = \operatorname{vec}(R). \tag{34}$$

Under column stacking,

$$(B \otimes A)\operatorname{vec}(X) = \operatorname{vec}(AXB^T). \tag{35}$$

The factor order in Δ therefore gives

$$\Delta v = \operatorname{vec}(T\rho), \tag{36}$$

$$m = \operatorname{Re}\langle v, \Delta v\rangle, \tag{37}$$

$$m = \langle R, RTR\rangle_{\mathrm{HS}}. \tag{38}$$

The support projection of Δ fixes v. For positive-semidefinite operators A and B, the support-sensitive tensor logarithm is

$$\log(A \otimes B) = (\log A) \otimes P_{\operatorname{supp} B} + P_{\operatorname{supp} A} \otimes (\log B). \tag{39}$$

The support projections in (39) are essential when an eigenvalue vanishes. With the convention $\log 0 = 0$, replacing them by ambient identity operators makes the formula false. Using

$$\log \sqrt{\rho} = \frac{1}{2} \log \rho, \tag{40}$$

$$\log(\omega^{-1/2}) = -\frac{1}{2} \log \omega, \tag{41}$$

together with transpose covariance gives

$$\mathrm{Re}\langle v, (\log \Delta) v \rangle = \frac{1}{2} D(\rho \| \omega). \tag{42}$$

Support-aware Jensen and Hilbert–Schmidt Cauchy–Schwarz yield

$$\frac{1}{2} D(\rho \| \omega) \le \log m, \tag{43}$$

$$m^2 \le \|R\|_2^2 \|RTR\|_2^2, \tag{44}$$

$$\|R\|_2^2 \|RTR\|_2^2 = Q_2(\rho, \omega). \tag{45}$$

Equations (43)–(45) prove (30) for faithful ω.

The auxiliary divergence and vector-state Jensen mechanism are adapted from Definition 5 and Lemma 19, as used in the proof of Theorem 7, of Ref. [MLDS+13]. Only the direct order-two endpoint and the endpoint Cauchy–Schwarz estimate are used here; this shortened proof is not Beigi's interpolation argument [Bei13]. Compression to the support of ω proves the singular-reference statement. Applying (30) with $\omega = \rho_A \otimes \rho_B$ gives

$$I(A : B)_\rho = D(\rho_{AB} \| \rho_A \otimes \rho_B), \tag{46}$$

$$D(\rho_{AB} \| \rho_A \otimes \rho_B) \le \log \mathrm{tr}(W^2), \tag{47}$$

$$\log \mathrm{tr}(W^2) \le \log r. \tag{48}$$

Thus every finite-dimensional bipartite density operator satisfies the sharp coefficient-one estimate

$$I(A : B)_\rho \le \log \mathrm{OSR}(\rho_{AB}). \tag{49}$$

The uniform perfectly correlated state on r classical symbols attains equality because its operator-Schmidt rank is r and its mutual information is $\log r$.

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 for both positive-definite and arbitrary positive-semidefinite output weights, with the negative quarter power restricted to the output support.[5] The full-support eigenbasis calculation is formalized: the whitening is identified 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 also formalized. Let $\rho, \omega \geq 0$ satisfy $\ker \omega \subseteq \ker \rho$, and let $V : \mathbb{C}^k \to \mathbb{C}^D$ satisfy

$$V^\dagger V = I_k, \tag{50}$$
$$VV^\dagger = P_{\operatorname{supp} \omega}. \tag{51}$$

Then

$$D(\rho \| \omega) = D(V^\dagger \rho V \| V^\dagger \omega V), \tag{52}$$
$$Q_2(\rho, \omega) = Q_2(V^\dagger \rho V, V^\dagger \omega V). \tag{53}$$

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 retains the same normalization after compression. The singular-reference inequality therefore reduces to

$$D(A \| B) \leq \log Q_2(A, B), \tag{54}$$
$$A \geq 0, \qquad B > 0, \qquad \operatorname{tr} A = 1. \tag{55}$$

[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`.

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

$$\mathrm{Re}\langle v, (\log A)v \rangle \leq \log(\mathrm{Re}\langle v, Av \rangle)\,. \tag{56}$$

In an eigenbasis of A, the support condition makes every weight at a zero eigenvalue vanish. Scalar logarithmic 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 of a Hermitian matrix and its functional calculus. [9] This auxiliary analytic result is not a theorem of Cirac–Pérez-García–Schuch–Verstraete [CPGSV16]. It does not alter the source-coverage count or complete Proposition 4.5.

The reusable functional-calculus identities and the entropy comparison are formalized. Functional calculus commutes with transpose on Hermitian matrices. In particular, the logarithm and support projection commute with transpose, and every real power and the positive square root commute with transpose on positive-semidefinite matrices. These identities include singular matrices and zero-dimensional matrix algebras. The

[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`.

[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`.

tensor-product logarithm retains both 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 Cirac–Pérez-García–Schuch–Verstraete [CPGSV16]; in particular, the tensor-product result 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. The arbitrary-support estimate below uses simultaneous compression to both marginal supports.

The unnormalized sandwiched Rényi trace term is

$$\widetilde{Q}_\alpha(\rho\|\omega) = \operatorname{Re}\operatorname{tr}\Big[\Big(\omega^{(1-\alpha)/(2\alpha)}\rho\,\omega^{(1-\alpha)/(2\alpha)}\Big)^{\alpha}\Big]. \tag{57}$$

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 $\ker\omega \subseteq \ker\rho$. This includes singular ω: powers that vanish on the kernel 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 supporting identities are not a theorem of Cirac–Pérez-García–Schuch–Verstraete

[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`.

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

[CPGSV16]. 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 through the relative-modular half-moment, support-aware Jensen inequality, and Hilbert–Schmidt Cauchy–Schwarz, without a general weighted interpolation theorem or an order-one limit.

The faithful marginal-whitening theorem diagonalizes the first marginal, so the input transpose from the Choi convention agrees with that marginal, and proves

$$Q_2(\rho_{AB}, \rho_A \otimes \rho_B) \leq \mathrm{OSR}(\rho_{AB}) \tag{58}$$

when both marginals are positive definite. The arbitrary-support theorem compresses simultaneously to the two marginal supports, applies (58), and returns through 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 by #5211 is complete.

The unrestricted coefficient-one inequality (49) is also formalized for every positive-semidefinite bipartite operator of trace one, using 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 using mono-

[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_operatorSchmidtRank_of_marginals_posDef` in `TNLean/Channel/FaithfulMarginalWhitenedChoi`, and `TNLean.sandwichedRenyiTwoTrace_product_marginals_le_operatorSchmidtRank` in `TNLean/Channel/MarginalSupportWhitenedChoi`.

tonicity of the logarithm.[14] Hence the mathematical work tracked by #5212 is complete.

The finite-chain MPO application 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 work does not formalize the false thermodynamic-limit clause of Proposition 4.5 of Ref. [CPGSV16] or alter its separate negative status.

5 Diagonal matrix product operators

No diagonal tensor can violate the proposed finite-chain estimate. Fix a chain of length $N = L + R$, and define its diagonal coefficient matrix by

$$P_{x,y} = \mathrm{tr}(M^{x_0x_0} \cdots M^{x_{L-1}x_{L-1}} M^{y_0y_0} \cdots M^{y_{R-1}y_{R-1}})\,, \tag{59}$$

where x and y label configurations on the two consecutive blocks. Set

$$A_x = M^{x_0x_0} \cdots M^{x_{L-1}x_{L-1}}, \tag{60}$$
$$B_y = M^{y_0y_0} \cdots M^{y_{R-1}y_{R-1}}. \tag{61}$$

Then

$$P_{x,y} = \mathrm{tr}(A_x B_y), \tag{62}$$
$$\mathrm{tr}(A_x B_y) = \sum_{a,b=0}^{D-1} (A_x)_{ab}(B_y)_{ba}. \tag{63}$$

The factorization in (63) passes through a complex vector space of dimension D^2, so

$$\mathrm{rank}_{\mathbb{C}}\, P \le D^2. \tag{64}$$

Suppose the finite-chain operator is positive semidefinite and diagonal, with positive trace

$$Z = \mathrm{tr}\,\rho^{(N)}(M) > 0. \tag{65}$$

[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`.

Its joint probability matrix is

$$\widehat{P} = Z^{-1} P. \tag{66}$$

Multiplication by the nonzero real scalar Z^{-1} preserves rank. Since $\widehat{P}$ is real,

$$\operatorname{rank}_{\mathbb{R}} \widehat{P} = \operatorname{rank}_{\mathbb{C}} \widehat{P}, \tag{67}$$

$$\operatorname{rank}_{\mathbb{C}} \widehat{P} \leq D^2. \tag{68}$$

Riazanov–Vyalyi prove for every finite joint probability matrix that

$$I(X:Y) \leq \log \operatorname{rank}_{\mathbb{R}} \widehat{P}. \tag{69}$$

This is Theorem 4.1 of Ref. [RV17]. 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. At each step it chooses an endpoint at which mutual information does not decrease. The final distribution has at most $\operatorname{rank}_{\mathbb{R}} P$ nonzero rows. Therefore every diagonal periodic MPO satisfies

$$I_L^{(N)} \leq 2 \log D. \tag{70}$$

This conclusion requires positivity and normalization only at the chain length under consideration; positivity at other lengths is unnecessary. It gives an independent proof of the diagonal case. The coefficient-one theorem (49) 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 Theorem 4.1, Section 4.2, of Ref. [RV17] cancels the factor $1 + \varepsilon\alpha_i$. In the source 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 in both $m_{ij}(1 + \varepsilon\alpha_i)$ and $p_i(1 + \varepsilon\alpha_i)$. The proof then chooses an endpoint of the admissible interval, 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 printed.

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

$$h(ab) = a\,h(b) + b\,h(a). \tag{71}$$

Equation (71) remains valid when either factor is zero. The additional terms cancel against the row-marginal entropy because each row marginal is the sum of its entries, while the row relation fixes the column marginal. Mutual information is therefore affine on the entire closed admissible interval, including both endpoints. This removes the endpoint singularity without strengthening the hypotheses or weakening the conclusion, so the discrepancy is a local correction.[15]

The rank factorization is formalized.[16] The information-versus-rank theorem of Riazanov–Vyalyi [RV17] 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`. It 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 does not by itself settle the unrestricted quantum question. The source-independent argument in the preceding section proves the unrestricted quantum estimate; 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 Result 1 and equations (10)–(14) of Ref. [DlCSPGC13]. It is also summarized in Section 4.3, equation (13), and Table 2 of Ref. [DlCDN19]. Their examples are diagonal states

$$\rho_t = \sum_{i,j} (M_t)_{ij} |i,j\rangle\langle i,j|, \tag{72}$$

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

This is a statement about the size of exact factorizations, not a lower bound on quantum mutual information. The same paper bounds the entanglement of purification above by the logarithm of the purification rank [DlCSPGC13]. A diverging purification rank makes this upper bound nonuniform but does not imply that the entanglement of purification diverges. No lower bound on entanglement of purification or quantum mutual information is proved for the rank-three family.

After normalizing M_t to a joint probability distribution, the quantum mutual information of the diagonal state equals the classical mutual information of that distribution. The ordinary-rank bound in Theorem 4.1 of Ref. [RV17] gives

$$I(X:Y) \le \log \operatorname{rank}(M_t), \tag{73}$$

$$\log \operatorname{rank}(M_t) = \log 3. \tag{74}$$

This theorem concerns ordinary matrix rank, not nonnegative rank. Its proof repeatedly eliminates a linearly dependent row while preserving nonnegativity and the column marginals, choosing an endpoint at which mutual information does not decrease. The final joint distribution has at most $\operatorname{rank}(M_t)$ nonzero rows, so its mutual information is at most the entropy of a random variable with that many values.

The rank-three examples therefore have uniformly bounded mutual information. They show that a small MPO bond dimension 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 a tensor with physical dimension two and bond dimension three whose only nonzero physical slices are

$$M^{00} = \mathrm{diag}(1, 0, 0), \tag{75}$$

$$M^{11} = \mathrm{diag}(0, 1, -1). \tag{76}$$

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

$$\rho^{(N)}(M) = |0^N\rangle\langle 0^N| + (1 + (-1)^N)|1^N\rangle\langle 1^N|. \tag{77}$$

The operator in (77) is positive semidefinite for every $N \geq 1$, including odd lengths, and

$$\mathrm{tr}\, \rho^{(N)}(M) = 2 + (-1)^N \neq 0. \tag{78}$$

This tensor is therefore an MPDO under exactly the hypotheses of Proposition 4.5 of Ref. [CPGSV16].

At odd lengths, the normalized state is

$$\widehat{\rho}^{(N)}(M) = |0^N\rangle\langle 0^N|. \tag{79}$$

At positive even lengths, it is

$$\widehat{\rho}^{(N)}(M) = \frac{1}{3}|0^N\rangle\langle 0^N| + \frac{2}{3}|1^N\rangle\langle 1^N|. \tag{80}$$

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

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

$$I_L^{(N)} = 0 \qquad \text{for odd } N. \tag{81}$$

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

$$h_2(1/3) = -\frac{1}{3}\log\frac{1}{3} - \frac{2}{3}\log\frac{2}{3} > 0. \tag{82}$$

The even-length mutual information is

$$I_L^{(N)} = h_2(1/3) + h_2(1/3) - h_2(1/3), \tag{83}$$

$$I_L^{(N)} = h_2(1/3) \qquad \text{for even } N. \tag{84}$$

Equations (81) and (84) show that $\lim_{N\to\infty} I_L^{(N)}$ does not exist. The eigenvalue -1 of the physical-trace transfer operator is the peripheral sector responsible for the parity oscillation. Primitivity or another aperiodicity hypothesis would exclude this example, but Proposition 4.5 of Ref. [CPGSV16] contains no such hypothesis.

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

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

$$I_L^{(N)} \le I_{L+1}^{(N)} \qquad (2L+1 \le N). \tag{85}$$

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

$$I_L \le 4 \log D'. \tag{86}$$

[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`, `MPOTensor.ThermodynamicLimitCounterexample.not_exists_tendsto_parityMutualInfoAfterCut`, and the source-facing conjunction `MPOTensor.ThermodynamicLimitCounterexample.proposition45_limit_counterexample` in `TNLean/MPS/MPDO/ThermodynamicLimitCounterexample`.

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

The local-purification step is also complete.[20] For explicit matrices $A^{(i,k)}$ witnessing the locally purified form, the two blockwise ancillary traces of the 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 (86), 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 terms of D', not the MPO bond dimension D.

This result does not produce a local purification for an arbitrary positive MPO. Cirac–Pérez-García–Schuch–Verstraete warn that such a representation need not exist [CPGSV16], and the no-go results of Refs. [DlCSPGC13, DlCCC+16] show that MPO positivity and bond dimension do not recover it uniformly.

The finite-chain estimate is nevertheless 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

$$\mathrm{OSR}(\widehat{\rho}^{(N)}_{L:N-L}) \leq D^2. \tag{87}$$

Combining (87) with (49) gives the stronger conclusion

$$I_L^{(N)} \leq 2 \log D. \tag{88}$$

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 of Ref. [CPGSV16] is formally refuted at its stated hypothesis set by `MPOTensor.ThermodynamicLimitCounterexample.proposition45_limit_counterexample`.

[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`.

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

The finite-chain monotonicity and bond-dimension bounds remain valid as separate results. Primitivity, aperiodicity, or another convergence condition would exclude the parity obstruction, but such a condition is a genuine boundary of a corrected theorem and is absent from the printed proposition. This resolves the false-source limit clause tracked by #7070.

References

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

[CPGSV16] J. Ignacio Cirac, David Pérez-García, Norbert Schuch, and Frank Verstraete. Matrix product density operators: Renormalization fixed points and boundary theories. Preprint; published version: Annals of Physics 378 (2017), 100–149, 2016.

[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.